

Abstract of “Bounds and Applications of Concentration of Measure in Fair Machine Learning and Data Science” by Cyrus Cousins, Ph.D., Brown University, May 2021.

I introduce novel concentration-of-measure bounds for the *supremum deviation*, several *variance concepts*, and a family of *game-theoretic welfare functions*. I then apply these bounds to analyze machinelearning and data science problems, with deep theoretical consequences and impactful practical ramifications. Parts I and II assume an *independently and identically distributed* (i.i.d.) setting, where we observe many statistically independent occurrences, and part III extends some of my methods and themes to study non-i.i.d. mean-estimation problems. Naturally, some conclusions are weaker in these relaxed settings, but I find that many of the same data-dependent and variance-sensitivity themes apply, and give efficient algorithms for real-world problems where the i.i.d. assumption is prohibitive.

This thesis is split into three parts, each representing a significant advancement in its respective field; the first in *statistical learning theory*, with an emphasis on *algorithms* and *generalization guarantees* making efficient use of training data; the second in the *axiomatic underpinnings, practice, and statistical elements* of *fair machine learning*; and the third in showing that themes and bounds from *sampling problems* in standard i.i.d. settings translate well into non-i.i.d. settings. Special attention is paid to keep each part independently readable, however the reader will better appreciate thematic connections, applications, and technical synergy when they are considered as a whole. Concretely, I note that the methods of part I and part III are not mutually incompatible, opening the door to strong *uniform convergence bounds* in *non-i.i.d.* settings, and furthermore, the *fair learning setting* of part II benefits from the statistical bounds of part I, and similar analysis is possible in non-i.i.d. settings. Indeed, in a connected and interdependent world, it may be the case that practical fair systems need to consider the intricacies of non-i.i.d. learning. Ultimately, the importance of philosophically grounded and statistically rigorous fair-learning systems, operating on real-world data with all the messy dependence structures that may entail, seems to be not only of immense academic interest, but also a necessary step in joining the theory and practice of machine learning and its societal impact.

Bounds and Applications of Concentration of Measure in Fair Machine Learning and Data Science

by
Cyrus Cousins
Sc.M., Brown University, 2017
Sc.B., Tufts University, 2015

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Computer Science at Brown University

Providence, Rhode Island
May 2021

© Copyright MMXXI by Cyrus Cousins

This dissertation by Cyrus Cousins is accepted in its present form by
the Department of Computer Science as satisfying the dissertation requirement
for the degree of Doctor of Philosophy.

Date _____

Eli Upfal, Director

Recommended to the Graduate Council

Date _____

Michael Littman, Reader

Date _____

Matteo Riondato, Reader

Approved by the Graduate Council

Date _____

Andrew G. Campbell
Dean of the Graduate School

Curriculum Vitae

A native of Providence, RI, Cyrus showed an interest in computers and mathematics from an early age. He pursued his baccalaureate studies at Tufts University, and whilst there, he served as a teaching assistant in various courses, ranging from computational geometry to bioinformatics. He also took his first steps as a researcher, working with professors Donna Slonim, Norman Ramsey, and Mitchell Wand, on a variety of competition biology and inference problems. In 2015, Cyrus graduated, earning his Bachelor’s degree in computer science, mathematics, and biology, and also earning highest honors on his senior thesis, on the subject of information-theoretic anomaly detection in biological systems. With an infectious thirst for knowledge, he immediately matriculated to Brown University to pursue his doctoral studies the next autumn. At Brown University, Cyrus served as a graduate teaching assistant for CS1810: *Computational Molecular Biology*, DATA2040: *Deep Learning and Special Topics in Data Science*, and CS1450: *Probability for Computing and Data Analysis*. He earned his master’s degree in computer science from Brown University in 2017 en route to his Ph.D. In the summer of 2018, he took a position as research intern at Two Sigma labs, working with Matteo Riondato on statistical machine learning methods, and returned the following summer to work with Larry Rudolph on problems in multi-agent reinforcement learning. Throughout his time at Brown University, he published with myriad professors on a gamut of problems, ranging from empirical game theory to statistical learning theory, though a dedication to rigorous mathematical analysis of practically-relevant sampling problems echoes throughout his body of work. In 2021, Cyrus was awarded the Dean’s Faculty Fellowship at Brown University.

Conference and Journal Publications

1. Cyrus Cousins and Matteo Riondato. Sharp uniform convergence bounds through empirical centralization. In *Advances in Neural Information Processing Systems*, 2020
2. Leonardo Pellegrina, Cyrus Cousins, Fabio Vandin, and Matteo Riondato. MCRapper: Monte-Carlo Rademacher averages for POSET families and approximate pattern mining. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2165–2174, 2020
3. Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. Improved algorithms for learning equilibria in simulation-based games. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 79–87, 2020
4. Enrique Areyan Viqueira, Cyrus Cousins, Yasser Mohammad, and Amy Greenwald. Empirical mechanism design: Designing mechanisms from data. In *Uncertainty in Artificial Intelligence*, pages 1094–1104. PMLR, 2020
5. Cyrus Cousins and Matteo Riondato. CaDET: interpretable parametric conditional density estimation with decision trees and forests. *Machine Learning*, 108(8-9):1613–1634, 2019
6. Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. Learning simulation-based games from data. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019
7. Cyrus Cousins and Eli Upfal. The k -nearest representatives classifier: A distance-based classifier with strong generalization bounds. In *2017 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10. IEEE, 2017

Workshop Papers, Notable Preprints, and Extended Abstracts

1. Cyrus Cousins. An axiomatic theory of provably-fair welfare-centric machine learning. *arXiv preprint arXiv:2104.14504*, 2021
2. Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. Learning competitive equilibria in noisy combinatorial markets. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021
3. Cyrus Cousins, Shahrzad Haddadan, and Eli Upfal. Making mean-estimation more efficient using an MCMC trace variance approach: DynaMITE. *arXiv preprint arXiv:2011.11129*, 2020
4. Carsten Binnig, Benedetto Buratti, Yeounoh Chung, Cyrus Cousins, Tim Kraska, Zeyuan Shang, Eli Upfal, Robert Zeleznik, and Emanuel Zraggen. Towards interactive curation & automatic tuning of ML pipelines. In *Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*, pages 1–4, 2018
5. Clayton Sanford, Cyrus Cousins, and Eli Upfal. Uniform convergence bounds for codec selection. *arXiv preprint arXiv:1812.07568*, 2018
6. Cyrus Cousins, Chirstopher M Pietras, and Donna K Slonim. Scalable FRaC variants: Anomaly detection for precision medicine. In *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 253–262. IEEE, 2017

Preface and Acknowledgements

I dedicate this thesis to my parents, my coauthors, my colleagues, and my mentors, without whom I would not have been able to complete this work. In particular, Shahrzad Haddadan, Amy Greenwald, Enrique Areyan Viqueira, Alessio Mazzetto, Clayton Sanford, Eli Upfal, Larry Rudolph, and Matteo Riondato have all worked tirelessly on many submissions, which have made their way into this thesis in one form or another. I would also like to thank Amir Dib, Kavosh Asadi, and Cristina Menghini for their insightful discussion and collaborative effort in exploring many topics, though our work has not made it into the thesis. Finally I wish to thank Eli Upfal, Shahrzad Haddadan, and Amir Dib a second time, for helping to connect my work with receptive audiences, for tolerating my flaws, and for always pushing me towards betterment.

Contents

The Prelude	1
I Concentration of Measure and Uniform Convergence in Machine Learning	5
1 Empirically Centralized Rademacher Averages	8
1.1 Introduction	8
1.2 Empirical centralization	9
1.3 Uniform convergence bounds	10
1.4 Experimental evaluation	15
1.5 Conclusions	16
2 Adversarial Learning with Weak Supervision	18
2.1 Introduction	18
2.2 Related Work	20
2.3 Preliminaries	20
2.4 Adversial Labeling Method	22
2.4.1 Risk Guarantees and Statistical Learning	24
2.4.2 The Centerpoint Strategy: Bounds without Adversarial Learning	25
2.4.3 Applications	26
2.4.4 Constraining the feasible set	27
2.5 Experiments	28
2.5.1 Setup	29
2.5.2 Baselines and algorithms	30
2.5.3 Results	30
2.6 Conclusion	31
II Fairness with Population Means: Malfare, Welfare, and Fair PAC Learning	32
3 Quantifying Population Sentiment with Welfare and Malfare	35
3.1 Introduction	35
3.2 Related Work	36
3.3 Quantifying Population-Level Sentiment	37
3.3.1 Axioms of Cardinal Welfare and Malfare	38
3.3.2 The Power Mean	39
3.3.3 Properties of Welfare and Malfare Functions	40
3.3.4 A Comparison with the Additively Separable Form	41
3.3.5 Relating Power Means and Inequality Indices	41
3.4 Statistical Estimation of Malfare Values	43
4 <i>Ewo</i> <i>k</i> s: an Algorithm for Fair Codec Selection	45
4.1 Introduction	45
4.2 The <i>Ewo</i> <i>k</i> s Algorithm	46
4.2.1 Welfare-Optimal k codec Selection	47
4.2.2 Empirical Welfare Optimization	48

4.2.3	Generalization Analysis	49
4.2.4	Linear Loss Families	50
4.3	Experimental Evaluation of EwoKs	51
4.3.1	Data-Dependence and Pareto Optimality	52
4.3.2	Fairness and Welfare-Optimality	54
4.3.3	Uniform Convergence Bounds	55
4.4	Discussion	56
5	Fair Probably Approximately Correct Learning	57
5.1	Statistical and Computational Efficiency Learning Guarantees	58
5.2	Characterizing Fair Statistical Learnability with FPAC-Learners	61
5.3	Characterizing Computational Learnability with Efficient FPAC Learners	64
5.3.1	Efficient FPAC Learning with Convex Optimization	65
5.3.2	Uniform Convergence and Efficient Covering	66
5.4	Conclusion	68
5.4.1	Contrasting Malfare and Welfare	68
5.4.2	FPAC Learning: Contributions and Open Questions	68
III Sample-Efficient Mean-Estimation with Dependent Sequential Data		70
6	Mean Estimation in Block Databases	74
6.1	Preliminaries	74
6.1.1	Variance and Concentration Inequalities	75
6.1.2	Prior Work	76
6.2	Algorithms	77
6.2.1	A Simple Block Sampling Algorithm	77
6.2.2	A Refined Block Sampling Algorithm	79
6.2.3	Setting Parameters	81
6.3	Experiments	83
6.3.1	Frequency Estimation Tasks	83
6.3.2	Sigmoid Functions	84
6.3.3	Block Size and Sample Complexity	85
6.4	Conclusion	86
7	Mean Estimation and the Markov-Chain Monte-Carlo Method	87
7.1	Introduction and Related Work	87
7.2	Preliminaries	89
7.2.1	The Inter-Trace Variance	89
7.2.2	Classic MCMC-Mean Estimators	91
7.3	DYNAMITE	92
Conclusion and Discussion		97
Bibliography		99
Appendices and Supplementary Material		105
A	Supplementary Material for Chapter 1	107
A.1	Proofs	107
A.2	Details on the Experimental Evaluation	117
A.2.1	Data Generation	119
A.2.2	Supplementary Plots	119

B	Supplementary Material for Chapter 2	121
B.1	Deferred proofs	121
B.2	Additional Experimental Details	124
B.2.1	Weak Supervision Sources	124
B.2.2	Algorithm Hyperparameters	124
B.3	Additional Figures	125
B.3.1	Animals with Attributes	125
B.3.2	DomainNet	125
B.4	Experiments on Synthetic Data	126
C	Supplementary Material for Chapter 3	129
C.1	Welfare and Malfare	129
D	Supplementary Material for Chapter 4	132
D.1	Pseudocode and EWOks Algorithm Details	132
D.2	Supplementary Experiments and Methods	133
D.2.1	Experimental Setup	133
D.2.2	Video Experiments	134
D.2.3	Supplementary Audio Experiments	135
D.3	Proof Compendium	137
D.3.1	Proof of Theorem 4.2.2	138
D.3.2	Proof of Theorem 4.2.3	142
D.3.3	Justifications of Equations	144
E	Supplementary Material for Chapter 5	146
E.1	Efficient FPAC-Learning	146
F	Supplementary Material for Chapter 7	153
F.1	A Compendium of Missing Proofs	153
F.1.1	Correctness and Efficiency of DYNAMITE	154

The Prelude

In this thesis, I introduce novel concentration-of-measure bounds for the *supremum deviation*, several *variance concepts*, and a family of *game-theoretic welfare functions*. It is divided into three parts, the first focusing on statistical methods and machine learning, followed by fair machine learning, and concluding with non-i.i.d. statistical estimation tasks. In particular, I introduce *empirically centralized Rademacher averages* to probabilistically bound the deviation between the *empirical* and *true* means over a *family of functions*, with applications to *multiple comparisons problems* in statistical and scientific settings (smaller p -values and tighter confidence intervals), and to various supervised and unsupervised machine learning settings (reduced sample complexity and sharper generalization bounds). I then show applications of these bounds to various machine-learning, fair machine-learning, and data-science settings, with deep theoretical implications and impactful practical consequences. Parts I and II assume an *independently and identically distributed* (i.i.d.) setting, where we observe many statistically independent occurrences before drawing conclusions, but in closing, part III extends some of my methods and themes to study non-i.i.d. mean-estimation problems. Naturally, some conclusions are weaker in these relaxed settings, but I find that many of the same data-dependent and variance-sensitivity themes apply, and give practical algorithms for realistic problems where the i.i.d. assumption is prohibitive.

Special attention is made throughout this thesis to connect complicated bounds and ideas to simple and intuitive concepts (like the central limit theorem or linear regression), while explaining and rigorously justifying all finite-sample probabilistic guarantees, as well as assumptions and proof techniques. Consequently, at the heart of this thesis is the marriage of simple ideas about intuitive random processes to sophisticated techniques and bounds from the theory of concentration of measure and probabilistic methods, applied in complicated settings to a variety of learning and statistical estimation problems. Topics are chosen for both their theoretical relevance and how well they illustrate and connect such ideas, but also for their practical relevance and applicability to high-impact real-world problems. Note also that, except when sufficiently succinct and pedagogically pertinent, all proofs are relegated to the appendices.

Part I: Concentration of Measure and Uniform Convergence in Machine Learning Part I deals primarily with statistical estimation guarantees in standard machine-learning settings. Chapter 1 introduces the *empirically centralized Rademacher average*, which yields probabilistic data-dependent bounds on the *supremum deviation* (SD) of empirical means of functions in a family \mathcal{F} from their expectations (i.e., $\sup_{f \in \mathcal{F}} |\hat{\mathbb{E}}[f] - \mathbb{E}[f]|$), with optimal dependence on the *supremum variance* and the function ranges. Such bounds are impactful in machine learning, as they bound the *generalization gap* between training and test error, and thus control overfitting and quantify the bias-complexity tradeoff (i.e., the tension at the heart of machine learning between *better fit* and *increased bias* as model complexity grows). More generally, such uniform convergence bounds have applications to *multiple comparisons problems* in statistical and scientific settings (smaller p -values and tighter confidence intervals), and to various supervised and unsupervised machine learning settings (reduced sample complexity and sharper generalization bounds). Empirical centralization yields *data-dependent* supremum deviation bounds that improve the dependence of *non-centralized* (standard)

Rademacher averages on *raw variances* to *centralized variances*, thus matching (asymptotically) known lower-bounds for mean estimation. To compute the bounds in practice, I develop novel tightly-concentrated Monte-Carlo estimators for the empirical Rademacher average of the empirically-centralized family, and show novel concentration results for the empirical wimpy variance. My experimental evaluation shows that these bounds greatly outperform their non-centralized counterparts, and are extremely practical, even at small sample sizes.

A major issue, ubiquitous in supervised learning settings, is the cost of obtaining labeled training data. The methods of chapter 1 may be viewed as a mechanism to get more out of *limited labeled data*, i.e., by showing that far less of it is required than previously thought, but chapter 2 takes a complimentary and orthogonal approach. In chapter 2, I adopt a *transductive learning* setting, wherein we have a vast set of points, a tiny fraction of which are labeled, and wish to learn to predict the remaining labels. I then describe a transductive learning algorithm that uses these data, alongside additional knowledge in the form of *weak labelers*, which are arbitrarily inaccurate predictors for either the target task or some related task. In particular, I show that accurate learning is possible with a small labeled training set when (a subset of) the family of weak labelers is somewhat well-aligned with the target task. The strategy here is to estimate and (probabilistically) bound appropriately-chosen statistics of the weak labelers using the supervised data, and then consider a *feasible set of possible labelings* for the *unlabeled data* that respect these statistics. When sufficient labeled data are available to well-estimate the chosen family of statistics, and the feasible label space is of low diameter, learning the *minimax-optimal* classifier over this feasible set then yields strong generalization guarantees. In particular, if the weak-labeler statistics give sufficient information about the target task, and we have sufficient labeled data to accurately constrain their statistics, then all *feasible labelings* are reasonably accurate, and if sufficiently many *unlabeled data* are available, then we can learn a complicated model; these conditions neatly factor the labeled and unlabeled sample complexities, and describe conditions under which weak-labelers are sufficient to supplement a small labeled dataset.

Part II: Fairness with Population Means: Malfare, Welfare, and Fair PAC Learning In part II, I begin with an axiomatic justification to the *power mean* family of welfare functions, which summarize societal wellbeing, while making tradeoffs between the needs of the overall population and of marginalized groups (cf., *utilitarianism* versus *egalitarianism*). From the same axiomatization, starting with a measure of *discontentment* (loss) rather than *contentment* (utility), I derive the parallel concept of *malfare*. For linear welfare functions, malfare acts as the *negative welfare* of the *negative utility*, however *nonlinear* welfare functions (e.g., the egalitarian welfare, or the Nash social welfare) may be *undefined* for negative utilities, whereas nonlinear malfare functions may be applied directly to loss or risk values. This is crucial to fairness in ML, where we generally seek to minimize loss (rather than maximize utility), and we require nonlinear power-mean malfare functions to specify fairness tradeoffs. These arguments are strongly grounded in the economic theory of *cardinal welfare*, but from them I show statistical estimation and learning guarantees more characteristic of the computer science literature. In particular, malfare is a natural target in machine learning problems where we *minimize* (negatively connoted) loss, rather than *maximize* (positively connoted) utility. As an application, I cast a *streaming-media codec-selection* problem as a *fairness-sensitive* learning problem, wherein we seek to *efficiently select a small set* of media-encoders that can mutually satisfy a userbase with *diverse preferences* (e.g., quality versus bandwidth consumption). Optimizing welfare objectives maximize the impact of each selected codec, whereas without considering fairness and welfare objectives, it's easy to optimize only for a target demographic, or under invalid assumptions on users, thus potentially discriminating against some groups of users. This is an important *accessibility issue*, as these types of considerations ensure that audiovisual streaming and telecommunications services effectively serve populations that are often sidelined by digital services, including those with limited internet access, as well as those with various audiovisual perception conditions. I explore various *welfare* and *Pareto* optimality concepts, and how the bias-complexity tradeoff manifests in multivariate settings and with fairness issues.

From a more theoretical angle, I also show *statistical estimation guarantees* for welfare and malfare, and from the *social planning problem*, develop a theory of fair machine learning, based on the *probably approximately correct* (PAC) learning framework, termed *fair-PAC* (FPAC) learning. An FPAC-learner is an algorithm that learns an ε - δ malfare-optimal model with bounded sample complexity, for *any data distribution*, and for *any* (axiomatically justified) malfare concept. We show broad conditions under which, with appropriate modifications, many standard PAC-learners may be converted to FPAC learners. This places FPAC learning on firm theoretical ground, as it yields statistical, and in some cases computational, efficiency guarantees for many well-studied machine-learning models. FPAC-learning is also practically relevant, as it democratizes fair machine learning, by providing concrete training algorithms and rigorous generalization guarantees for these models.

Part III: Sample-Efficient Mean-Estimation with Dependent Sequential Data Finally, part III extends the methods and themes of the previous parts, where I assume an *independently and identically distributed* (i.i.d.) sample, into more general *weakly dependent* settings. As in the *data-dependent* guarantees with *empirically centralized Rademacher averages*, the goal here is to get the *strongest guarantees* under the *weakest assumptions*: in particular, this means my algorithms must be sensitive to structure found *in the data*. While notions of *approximate independence* can be difficult to rigorously bound from data, I find that appropriate *variance concepts* are easy to bound, and are sufficient to obtain asymptotically near-optimal sample-complexity guarantees for *mean estimation* in two dependent settings. In particular, I first examine *block databases*, where the key assumption is that *contiguous blocks of records* can be accessed nearly as efficiently as *individual records*, and second, I examine *Markov chains*, where each step of the chain is a *memoryless* random variable (i.e., conditionally independent from its history given the previous step), and it is as easy to collect a *trace of dependent* (often correlated) samples as to collect a pair of near-independent samples.

In both cases, I give *data-dependent* algorithms for ε - δ mean-estimation that avoid *worst-case* sample-complexity behavior. In particular, by considering not just independent or near-independent samples (as guaranteed via structural assumptions), but instead all of the dependent samples (available at no extra cost), I find that often the *variance* of the mean estimate decreases; particularly so when the data are less dependent than indicated by *a priori* structural assumptions. This decrease in variance implies more rapid convergence of empirical mean estimates by various central limit theorems, though this work undertakes the significantly more challenging endeavor of showing commensurate improvements to *finite-sample guarantees* on convergence rates. Furthermore, as I do not assume *a priori* knowledge of the appropriate variance concepts, sufficient sample sizes are not known *a priori*, and thus I employ a *progressive sampling strategy* to avoid drawing too many samples. The details of variance-estimation and necessary multiple-comparisons corrections entailed are subtle, but intuitively, this acts as a “guess and check” method, wherein we optimistically select an initial small sample size, which I use to estimate variances and means, and increase it until a sufficiently large sample has been drawn to provide the desired guarantee. Surprisingly, despite being variance-oblivious, these strategies are asymptotically optimal, up to logarithmic factors, in both the block-database and Markov-chain settings. In both cases, *a priori* guarantees on the amount of dependence do appear in our bounds, but only transiently, and as the additive error is taken to 0, terms involving only variance, which is estimated entirely from the data, come to dominate. This calls into question the importance of difficult-to-bound quantities measuring the degree of dependence.

This Work, as a Whole Read separately, each part represents a significant advancement in its respective field; the first in *statistical learning theory*, with an emphasis on *algorithms* and *generalization guarantees* requiring less *labeled data* than previous methods; the second in the *axiomatic philosophy, practice*, and *statistical learning guarantees* of *fair machine learning*; and the third in showing that themes and bound forms from *sampling problems* in standard i.i.d. settings translate well into non i.i.d. settings (with some

additional work). Special attention is paid to keep each part readable outside of the greater context of this work, however the reader will better appreciate thematic connections, applications, and technical synergy when they are considered as a whole.

To better appreciate the connections and common themes running through this thesis, the preface of each part contains a preview of the concentration inequalities derived within. In particular, I develop a leitmotif of analogues of the Hoeffding and Bennett inequalities for bounded random variables, in the i.i.d. mean-estimation, malfare estimation, and non-i.i.d. mean-estimation settings, laying bare the common threads between each part. In their original forms, both results characterize the rate (with exact constants) that i.i.d. sums of *bounded random variables* converges to its mean. The Bennett inequality yields sub-Poisson (Poisson-like) or sub-gamma tails, with asymptotic behavior dominated by the *variance* of the random variables, and the Hoeffding inequality yields similar sub-Gaussian tails, in terms of the *worst-case* variance bound $r^2/4$, where r is the range of the random variables in question. Where appropriate, I also discuss more sophisticated uniform-convergence bounds, and how they manifest in each setting. Such bounds are important in many settings, as they allow us to study the simultaneous convergence of not just one or a handful of functions to their means, but rather large and even infinite sets of functions, without relying on union bounds (a.k.a. the *Bonferroni correction*). In the first part, these bounds are straightforward, as we consider risk (mean loss) in machine-learning settings with i.i.d. training data. The second part generalizes these results to consider *malfare*, which is a *nonlinear function* of risk over multiple groups, but we still obtain similar range, variance, and Rademacher average dependent tail bounds in this setting. Finally, in part three, we relinquish the i.i.d. assumption, which complicates matters greatly, but ultimately we attain similar bounds, now with additional terms capturing the *degree of (approximate) independence*.

To further solidify the connections between the three main parts, I note that the methods of part I and part III are not mutually incompatible, opening the door to strong *uniform convergence bounds* in *non-i.i.d.* settings, and furthermore, the *fair learning setting* of part II obviously benefits from the statistical bounds of part I, but similar analysis is certainly possible in non-i.i.d. settings. Indeed, in a connected and interdependent world, it may be the case that practical fair systems need to consider the intricacies of non-i.i.d. learning, and ultimately the importance of philosophically grounded and statistically rigorous fair-learning systems, operating on real-world data with all the messy dependence structures that may entail, just may be exactly what is needed, not only to bring new deep and interesting problems to the computer science community, but also to solve problems of algorithmic and natural injustice and unfairness in the world at large.

Part I

Concentration of Measure and Uniform Convergence in Machine Learning

The highly theoretical work of this part is key to the success in the data science and fair machine learning domains. Furthermore, this work in *concentration of measure* and *uniform convergence* is key to understanding and assessing the quality of our solutions in the pattern mining, data science, and fairness settings, as the basic statistical estimation themes echo throughout in various forms. Throughout the work we maintain a running metaphor of *sub-Gaussian* and *sub-gamma* bounds, and describe how, in most cases we can't expect to beat the sub-Gaussian bound, which are *asymptotically equivalent* to sub-gamma bounds.

Tail Bounds and Concentration of Measure For some function f , we are interested in the convergence of the *empirical mean* $\hat{\mathbb{E}}[f]$ on m samples to its expectation $\mathbb{E}[f]$. The Central Limit Theorem (CLT) tells us that $\hat{\mathbb{E}}[f]$ is asymptotically Gaussian, thus by the Chernoff bound for the Gaussian distribution, we have

$$\lim_{m \rightarrow \infty} \mathbb{P} \left(\left| \mathbb{E}[f] - \hat{\mathbb{E}}[f] \right| \geq \sqrt{\frac{2 \mathbb{V}[f] \ln \frac{2}{\delta}}{m}} \right) \leq \delta .$$

[Devroye et al., 2016] show matching finite-sample lower-bounds for *any mean estimator*, thus the best bounds we can hope to achieve should be on par with the above. This class of bound is generally called *sub-Gaussian*, with $\mathbb{V}[f]$ replaced by various *variance proxies*. If we restrict our attention to *bounded* f , i.e., $f \in \mathcal{X} \rightarrow [a, b]$, where $r \doteq b - a$, Hoeffding's inequality [Hoeffding, 1963] states

$$\mathbb{P} \left(\left| \mathbb{E}[f] - \hat{\mathbb{E}}[f] \right| \geq \underbrace{\sqrt{\frac{r^2 \ln \frac{2}{\delta}}{2m}}}_{\text{VARIANCE TERM}} \right) \leq \delta . \quad (1)$$

Here the *sub-Gaussian variance proxy* is $\frac{r^2}{4}$, which by Popoviciu's inequality, is the *worst-case* variance of any range r random variable. Bennett's inequality [Bennett, 1962] tells us that

$$\mathbb{P} \left(\left| \mathbb{E}[f] - \hat{\mathbb{E}}[f] \right| \geq \underbrace{\frac{r \ln \frac{2}{\delta}}{3m}}_{\text{SCALE TERM}} + \underbrace{\sqrt{\frac{2 \mathbb{V}[f] \ln \frac{2}{\delta}}{m}}}_{\text{VARIANCE TERM}} \right) \leq \delta , \quad (2)$$

which yields the desired *variance dependence*, at the cost of a fast-decaying *scale term*. As such, the bound is *sub-gamma* [see Boucheron et al., 2013, chapter 2], rather than sub-Gaussian. We may thus view the *scale term* as a *finite-sample correction* to a sub-Gaussian bound.

Our uniform convergence bounds then apply simultaneously to *large sets of functions*, which translate to various guarantees on generalization error, estimation quality, and fairness. The natural comparison is to Hoeffding's inequality with a union bound over \mathcal{F} , which tells us that

$$\mathbb{P} \left(\sup_{f \in \mathcal{F}} \left| \mathbb{E}[f] - \hat{\mathbb{E}}[f] \right| \geq \sqrt{\frac{r \ln \frac{2|\mathcal{F}|}{\delta}}{2m}} \right) \leq \delta .$$

Despite its weaknesses, this result powers many sample complexity guarantees in computer science. The principal aim of chapter 1 is to show that in many applications, we can improve over Hoeffding's inequality, and may nearly match variance-sensitive lower-bounds or Bennett's inequality.

Semisupervised Learning with Weak Labelers In (semi)supervised learning settings, the need for sharp tail bounds and generalization inequalities is often motivated by a dearth of *labeled training data*. Empirically centralized Rademacher averages are a mechanism to get more out of *limited labeled data*, but in chapter 2, I investigate a *weakly semisupervised setting*, where a small number of training labels are augmented

with the output of *weak labelers*, which are assumed to be loosely correlated with the target task. I then describe a transductive learning algorithm for this setting, with strong data-dependent learning guarantees when the family of weak labelers is somewhat well-aligned with the target task.

In particular, the strategy is to estimate and bound appropriately-chosen statistics of the weak labelers using *labeled data* and applying *concentration inequalities*, e.g., Hoeffding, Bennett, or (empirically centralized) Rademacher bounds. The algorithm then considers a *feasible set* \mathbb{Y} of *possible labelings* for the *unlabeled data* that respect these statistics. When sufficient labeled data are available to sharply-bound the chosen family of statistics, and the feasible label space is of *low diameter*, learning the *minimax-optimal* classifier over this feasible set then yields strong generalization guarantees. In particular, if the weak-labeler statistics give sufficient information about the target task, and we have sufficient labeled data to constrain their statistics, then all *feasible labelings* are reasonably accurate. Furthermore, if sufficiently *unlabeled data* are available, then we can learn a *complicated model* (i.e., a model of high Rademacher complexity); these conditions neatly factor the labeled and unlabeled sample complexities, and describe conditions under which weak-labelers are sufficient to supplement a small labeled training set.

In this setting, generalization error bounds may be composed into three major terms, with some variations in their exact form, depending on technical details. The first is the *minimax empirical risk* term, representing both the fundamental difficulty of the (possibly unrealizable) learning task with the given hypothesis class, which may be improved by learning a more complicated model. Second, we have *uncertainty over the feasible set* \mathbb{Y} , which may be quantified, e.g., as the *diameter* of \mathbb{Y} , and may be improved by constraining additional statistics, adding weak labelers, or adding more labeled or unlabeled points. Then, thirdly we have an *overfitting term*, which describes overfitting of a model to the minimax guarantee, which may be controlled *solely by adding more unlabeled training data*. To make this concrete: the first two terms are *independent* of the complexity of the model we are learning, so if we wish to, e.g., fit a sophisticated deep neural network to a small amount of *labeled training data*, we may do so, offsetting the model complexity only by increasing the size of the *unlabeled training set*. Surprising as this may be, the reader should consider that this term addresses only *overfitting*, and good performance also requires that the *diameter* of the feasible set remain small, which requires subtle conditions on the weak-labeler statistics: essentially it must be possible to reconstruct the true labels of the target task with reasonable accuracy from them. Thus the method does not perform the impossible alchemy of synthesizing an accurate model with generalization guarantees from thin air, but rather *provides a mechanism* by which to *utilize and quantify* the information provided by weak-labelers.

Chapter 1

Empirically Centralized Rademacher Averages

I introduce the use of *empirical centralization* to derive novel practical, probabilistic, sample-dependent bounds to the Supremum Deviation (SD) of empirical means of functions in a family from their expectations. My bounds have optimal dependence on the maximum (i.e., wimpy) variance and the function ranges, and the same dependence on the number of samples as existing SD bounds. To compute the bounds in practice, I develop novel tightly-concentrated Monte-Carlo estimators of the empirical Rademacher average of the empirically-centralized family, and we show novel concentration results for the empirical wimpy variance. Experimental evaluation shows that our bounds greatly outperform non-centralized bounds and are extremely practical even at small sample sizes. This chapter is adapted from Cousins and Riondato [2020].

1.1 Introduction

The *supremum deviation* of the empirical means of functions in a family $\mathcal{F} \subseteq \mathcal{X} \rightarrow [a, b] \subset \mathbb{R}$ from their expectations is a key object in the study of empirical processes [Pollard, 1984]. Formally, let \mathcal{D} be a distribution on the domain \mathcal{X} and $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ be a collection of m independent samples from \mathcal{D} . The *Supremum Deviation (SD)* of \mathcal{F} on \mathbf{x} is the quantity

$$\text{SD}(\mathcal{F}, \mathbf{x}) \doteq \sup_{f \in \mathcal{F}} \left| \hat{\mathbb{E}}_{\mathbf{x}}[f] - \mathbb{E}_{\mathcal{D}}[f] \right|, \text{ where } \hat{\mathbb{E}}_{\mathbf{x}}[f] \doteq \frac{1}{m} \sum_{i=1}^m f(\mathbf{x}_i) .$$

The sample-dependent *Empirical Rademacher Average (ERA)* $\hat{\mathfrak{R}}_m(\mathcal{F}, \mathbf{x})$ of \mathcal{F} on \mathbf{x} and its expectation, the *Rademacher Average (RA)* $\mathfrak{R}_m(\mathcal{F}, \mathcal{D})$ of \mathcal{F} [Koltchinskii, 2001, Bartlett and Mendelson, 2002], imply *bidirectional bounds* on the SD (see (4)). Let $\boldsymbol{\sigma}$ be a collection of m independent Rademacher variables (i.e., uniform on $\{-1, 1\}$). These two quantities are defined as

$$\hat{\mathfrak{R}}_m(\mathcal{F}, \mathbf{x}) \doteq \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i f(\mathbf{x}_i) \right| \right], \text{ and } \mathfrak{R}_m(\mathcal{F}, \mathcal{D}) \doteq \mathbb{E}_{\mathbf{x}} [\hat{\mathfrak{R}}_m(\mathcal{F}, \mathbf{x})] . \quad (3)$$

The RA controls the finite-sample *expected* SD as [Van der Vaart and Wellner, 1996]

$$\frac{1}{2} \mathfrak{R}_m(\mathcal{F}, \mathcal{D}) - \frac{1}{\sqrt{m}} \sup_{f \in \mathcal{F}} \|f\|_{\infty} \leq \mathbb{E}_{\mathbf{x}} [\text{SD}(\mathcal{F}, \mathbf{x})] \leq 2 \mathfrak{R}_m(\mathcal{F}, \mathcal{D}) . \quad (4)$$

Probabilistic deviation bounds can be obtained by studying the convergence properties of the SD, and sample-dependent versions use the ERA and its deviation from the RA (see also theorem 1.3.2). The dependence on the *maximum* $q \doteq \sup_{f \in \mathcal{F}} \|f\|_\infty$ of \mathcal{F} makes the lower bound unsatisfactory, as this quantity can be very large. This downside is particularly evident at relatively small sample sizes, which are actually the most interesting in practice. As uniform convergence bounds are now used not “just” for the theoretical analysis of the performance of learning, but also to develop randomized approximation algorithms for many tasks [Riondato and Upfal, 2015, 2018, Pellegrina et al., 2019, Areyan Viqueira et al., 2019, Pellegrina et al., 2020], we believe it is extremely important to derive *practical* bounds to the SD that are optimized not just in terms of the number of samples, but also of other important parameters, such as the maximum and the *wimpy variance* (see (10)). In this work, we use various forms of *centralization* to develop such practical bounds. Define the *distributional centralization* $C_{\mathcal{D}}(\mathcal{F})$ w.r.t. \mathcal{D} as the family

$$C_{\mathcal{D}}(\mathcal{F}) \doteq \{x \mapsto f(x) - \mathbb{E}[f], f \in \mathcal{F}\} . \quad (5)$$

$C_{\mathcal{D}}(\mathcal{F})$ contains one function g for each $f \in \mathcal{F}$, such that g is f shifted by its expectation w.r.t. \mathcal{D} , thus, $\mathbb{E}_{\mathcal{D}}[g] = 0$ for each $g \in C_{\mathcal{D}}(\mathcal{F})$. The Rademacher Average $\mathfrak{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D})$ of $C_{\mathcal{D}}(\mathcal{F})$ *sharply* controls the finite-sample expected SD as [Boucheron et al., 2013, Lemma 11.4]

$$\frac{1}{2} \mathfrak{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \mathbb{E}_{\mathbf{x}}[\text{SD}(\mathcal{F}, \mathbf{x})] \leq 2 \mathfrak{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) . \quad (6)$$

Comparing (4) and (6), it is evident that the RA could be an *arbitrary large* multiplicative factor away from the expected SD, especially at small-sample regimes or when the maximum q of \mathcal{F} is large. The RA of the distributional centralization instead is *always* at most a multiplicative factor *two* away in both directions. Distributional centralization is therefore already known to be beneficial in the *expected case*, but can this gain be generalized to the probabilistic case, possibly using only sample-dependent quantities?

Contributions. In this work we introduce the use of *empirical centralization* to derive *practical, probabilistic* bounds to the SD. Our bounds exhibit a better or no worse dependence on important parameters such as the wimpy variance, the range (see (8)), and the sample size m (see theorem 1.3.3). We also show that the dependence on the wimpy variance that we obtain is optimal (lemma 1.3.4 and corollary 1.3.5). We introduce a novel empirical counterpart to the RA of the distributional centralization which uses *empirical centralization* to bound the SD. We analyze the bias of this quantity (lemma 1.2.1) and derive its concentration properties (theorem 1.2.2) using tail bounds for *self-bounding functions* [Boucheron et al., 2000, 2009]. In order to obtain fully-sample-dependent bounds, we introduce a Monte-Carlo estimator with sharp deviation bounds (theorem 1.3.7), and we also develop novel tight bounds for the empirical wimpy variance (theorem 1.3.1), which we believe to be of independent interest (e.g., in *matrix concentration inequalities* for estimating *eigenvalues of covariance matrices*). The results of our experimental evaluation show the advantages of centralization: the computed bounds to the SD are much smaller than those computed without centralization, even at small sample sizes. All proofs are relegated to appendix A.1.

1.2 Empirical centralization

We define the *empirical centralization* $\hat{C}_{\mathbf{x}}(\mathcal{F})$ of \mathcal{F} w.r.t. the sample $\mathbf{x} \in \mathcal{X}^m$ as

$$\hat{C}_{\mathbf{x}}(\mathcal{F}) \doteq \{x \mapsto f(x) - \hat{\mathbb{E}}_{\mathbf{x}}[f], f \in \mathcal{F}\} .$$

This quantity is an empirical counterpart to the distributional centralization $C_{\mathcal{D}}(\mathcal{F})$ of \mathcal{F} (see (5)). The key quantity that we use to derive the sample-dependent probabilistic bounds to the SD (section 1.3) is the ERA of the empirical centralization of \mathcal{F} , i.e., the quantity

$$\hat{\mathfrak{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) .$$

This quantity is completely dependent on the realized \mathbf{x} , even more, in some sense, than a “standard” ERA (see (3)), because the considered family $\hat{C}_{\mathbf{x}}(\mathcal{F})$ is also a function of \mathbf{x} , i.e., it is *sample-dependent*. We now derive its important properties: bias and concentration.

Bias The expectation w.r.t. \mathbf{x} of $\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$ is *not* the RA of the distributional centralization of \mathcal{F} (i.e., $\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D})$), but we now show that the bias decreases rapidly in m , i.e., $\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \in \Theta(\mathbb{E}_{\mathbf{x}}[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})])$. For ease of notation, let

$$b(m) \doteq \mathbb{E}_{\boldsymbol{\sigma}} \left[\left| \frac{1}{m} \sum_{i=1}^m \sigma_i \right| \right] \left(\text{which is } \Theta \left(\frac{1}{\sqrt{m}} \right) \right). \quad (7)$$

Lemma 1.2.1. Suppose $m \geq 4$. Then

$$\frac{\mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right]}{1 + 2b(m)} \leq \mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \frac{\mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right]}{1 - 2b(m)}.$$

Concentration We now show that $\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$ is tightly concentrated around its expectation because it is a *self-bounding function* [Boucheron et al., 2000, 2009] (see also definition A.1.1 in the supplementary material). We call the *widest range of \mathcal{F}* the quantity

$$r \doteq \sup_{f \in \mathcal{F}} \left(\max_{x \in \mathcal{X}} f(x) - \min_{y \in \mathcal{X}} f(y) \right) \quad (\leq b - a). \quad (8)$$

It is possible that $r \ll b - a$, for example, when \mathcal{F} contains a function f and a function $g = f + c$ for some $c \in \mathbb{R}$. The widest range of the empirical and distributional centralizations of \mathcal{F} is the same as the widest range of \mathcal{F} .

Theorem 1.2.2. Suppose $m \geq 1$, and let $\chi \doteq 1 + 2b(m)$. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the choice of \mathbf{x} , it holds that

$$\mathbb{E}_{\mathbf{x}}[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})] \leq \hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) + \frac{2r\chi \ln \frac{1}{\delta}}{3m} + \sqrt{\frac{\left(r\chi \ln \frac{1}{\delta} \right)^2}{(\sqrt{3}m)} + \frac{2r\chi(\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) + rb(m)) \ln \frac{1}{\delta}}{m}}. \quad (9)$$

The ERA of \mathcal{F} is a self-bounding function [Boucheron et al., 2003, Sect. 5.1], but proving this fact for the ERA of the empirical centralization $\hat{C}_{\mathbf{x}}(\mathcal{F})$ of \mathcal{F} is more challenging (see proof in the supplementary material), because the empirical centralization $\hat{C}_{\mathbf{x}}(\mathcal{F})$ itself depends on the sample \mathbf{x} . This result, together with lemma 1.2.1, enables us to use the ERA of the empirical centralization, and Monte-Carlo estimations of it, to derive practical sharp upper-bounds to the SD.

1.3 Uniform convergence bounds

We now introduce novel bounds to the SD using the ERA of the empirical centralization. Before doing so, we must introduce an important technical concept.

Wimpy variance The *raw* (i.e., non-centralized) *wimpy variance* $W^r(\mathcal{F})$ of \mathcal{F} and the (centralized) *wimpy variance* $W(\mathcal{F})$ of \mathcal{F} are key quantities in the study of probabilistic tail bounds to the SD [Boucheron et al., 2013, Ch. 11]. They are defined as

$$W^r(\mathcal{F}) \doteq \sup_{f \in \mathcal{F}} \mathbb{E}_{x \sim \mathcal{D}} \left[(f(x))^2 \right], \text{ and } W(\mathcal{F}) \doteq \sup_{f \in \mathcal{F}} \mathbb{E}_{x \sim \mathcal{D}} \left[\left(f(x) - \mathbb{E}_{\mathcal{D}}[f] \right)^2 \right]. \quad (10)$$

Naturally, the raw wimpy variance is always greater or equal to its centralized counterpart, and potentially much larger. A key identity that we use throughout this work is

$$W(\mathcal{F}) = W^r(C_{\mathcal{D}}(\mathcal{F})) = W(C_{\mathcal{D}}(\mathcal{F})) .$$

Empirical estimators on \mathbf{x} for the raw wimpy variance and for the wimpy variance are

$$\widehat{W}_{\mathbf{x}}^r(\mathcal{F}) \doteq \sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m (f(x_i))^2, \text{ and } \widehat{W}_{\mathbf{x}}(\mathcal{F}) \doteq \sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \left(f(x_i) - \widehat{\mathbb{E}}_{\mathbf{x}}[f] \right)^2 .$$

To compute the *sample-dependent* bounds to the SD that we introduce later in this section, we develop novel tail bounds to these estimators, which we believe to be of independent interest. Most prior work assumed known *a priori* bounds to the wimpy variances, but we show that they can be replaced by *empirical* bounds. Maurer and Pontil [2009] prove that the sample variance (i.e., when \mathcal{F} is a singleton) is a *weakly self-bounding function* [McDiarmid and Reed, 2006]. Our result holds for general \mathcal{F} , and is stronger, as we show that the wimpy variance is a (*strongly*) *self-bounding function* [Boucheron et al., 2000, 2009] (see also definition A.1.1 in the supplementary material).

Theorem 1.3.1. Suppose $m \geq 2$. Let $\delta \in (0, 1)$. With probability $\geq 1 - \delta$ over the choice of \mathbf{x} ,

$$W(\mathcal{F}) \leq \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F}) + \frac{r^2 \ln \frac{1}{\delta}}{m-1} + \sqrt{\left(\frac{r^2 \ln \frac{1}{\delta}}{m-1} \right)^2 + \frac{2r^2 \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F}) \ln \frac{1}{\delta}}{m-1}} . \quad (11)$$

Bounds to the SD Bousquet [2002, Thm. 2.3 (presented here for clarity in a slightly weaker form)] uses the wimpy variance to derive concentration bounds for the SD.

Theorem 1.3.2 (Bousquet, 2002, Thm. 2.3). Let $\delta \in (0, 1)$. With probability $\geq 1 - \delta$ over the choice of \mathbf{x} ,

$$\text{SD}(\mathcal{F}, \mathbf{x}) \leq \mathbb{E}_{\mathbf{x}} [\text{SD}(\mathcal{F}, \mathbf{x})] + \frac{2r \ln \frac{1}{\delta}}{3m} + \sqrt{\frac{2 \left(W(\mathcal{F}) + 4r \mathbb{E}_{\mathbf{x}} [\text{SD}(\mathcal{F}, \mathbf{x})] \right) \ln \frac{1}{\delta}}{m}} . \quad (12)$$

By plugging the r.h.s. of the symmetrization inequalities (4) and (6) in the r.h.s. of (12), one can obtain bounds that depend on the RA of \mathcal{F} or on the RA of the *distributional* centralization $C_{\mathcal{D}}(\mathcal{F})$. Neither of these bounds are sample-dependent. Such a bound can be obtained, for example, by using the ERA of \mathcal{F} and a tail bound (e.g., McDiarmid [1989]’s inequality or a tail bound for self-bounding functions [Boucheron et al., 2009]) on the deviation of the ERA from the RA. The following result states our sample-dependent bound to the SD using the *empirical centralization* $\widehat{C}_{\mathbf{x}}(\mathcal{F})$ and tail bounds to the wimpy variance, obtained by combining lemma 1.2.1 and theorems 1.3.1 to 1.3.2.

Theorem 1.3.3. Assume $m \geq 4$, and let $\eta \in (0, 1)$. Take ν to be the r.h.s. of (11) computed with $\delta = \eta/3$, so $\mathbb{P}(W(\mathcal{F}) > \nu) \leq \eta/3$, and take λ to be the r.h.s. of (9) computed with $\delta = \eta/3$, so $\mathbb{P}(\mathbb{E}_{\mathbf{x}}[\widehat{\mathfrak{R}}_m(\widehat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})] > \lambda) \leq \eta/3$. With probability $\geq 1 - \eta$ over the choice of \mathbf{x} , it holds that

$$\text{SD}(\mathcal{F}, \mathbf{x}) \leq \frac{2\lambda}{1-2b(m)} + \frac{2r \ln \frac{3}{\eta}}{3m} + \sqrt{\frac{2(\nu + 8r\lambda/(1-2b(m))) \ln \frac{3}{\eta}}{m}} .$$

The r.h.s. is

$$\frac{2}{1-2b(m)} \widehat{\mathfrak{R}}_m \left(\widehat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x} \right) + \mathbf{O} \left(\frac{r \ln \frac{1}{\eta}}{m} + \sqrt{\frac{(W(\mathcal{F}) + r \widehat{\mathfrak{R}}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) + r^2/\sqrt{m}) \ln \frac{1}{\eta}}{m}} \right) .$$

Is there any advantage in using this bound, i.e., in using empirical centralization, rather than using a bound involving the ERA of \mathcal{F} ? I.e., how does it compare to the standard bound

$$\text{SD}(\mathcal{F}, \mathbf{x}) \leq 2\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) + \mathbf{O}\left(\frac{q \ln \frac{1}{\eta}}{m} + \sqrt{\frac{(\mathbf{W}(\mathcal{F}) + q\hat{\mathbf{R}}_m(\mathcal{F}, \mathcal{D}) + q\sqrt{\mathbf{W}^r(\mathcal{F})/\sqrt{m}}) \ln \frac{1}{\eta}}{m}}\right) ?$$

We shall see that the r^2/\sqrt{m} and $q\sqrt{\mathbf{W}^r(\mathcal{F})/\sqrt{m}}$ terms are incomparable, though both appear only in transient $\mathbf{O}(m^{-3/4})$ terms, and the remaining differences all favor centralization. Most previous studies focused on the behavior of SD bounds as functions of the sample size m , but we believe that efficient SD bounds for practical applications (e.g., [Riondato and Upfal, 2018, 2015, Pellegrina et al., 2019, Areyan Viqueira et al., 2019, Pellegrina et al., 2020]), must improve the dependence also on the other parameters, the *wimpy variance* being the most important. Indeed, developing such bounds is the goal of this work.

First of all, we remark that the dependence on the wimpy variance shown in (12) cannot be improved: any bound to the SD of \mathcal{F} must be $\Omega\sqrt{\mathbf{W}(\mathcal{F}) \ln \frac{1}{\delta}/m}$, as can be shown using minimax lower bounds and median-of-means bounds [Devroye et al., 2016, Lugosi and Mendelson, 2019]. The question is thus whether the *complexity terms*, i.e., $\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x})$ and $\hat{\mathbf{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$, can match this lower bound. Lemma 1.3.4 answers this question in the *negative* for $\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x})$, and in the *positive* for $\hat{\mathbf{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$: the ERA of \mathcal{F} is controlled (in part) by the empirical *raw* wimpy variance, whereas the ERA of $\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F})$ has corresponding dependence on the empirical (centralized) wimpy variance. As with ordinary function variances, the raw wimpy variance can be unboundedly larger than the (centralized) wimpy variance, e.g., in the *constant function family* $\mathcal{F} \doteq \{x \mapsto c\}$.

Lemma 1.3.4. For any $\mathbf{x} \in \mathcal{X}^m$, it holds

$$\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) \geq \sqrt{\frac{\widehat{\mathbf{W}}_{\mathbf{x}}^r(\mathcal{F})}{2m}} \text{ and } \hat{\mathbf{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \geq \sqrt{\frac{\widehat{\mathbf{W}}_{\mathbf{x}}(\mathcal{F})}{2m}} .$$

Furthermore, it holds

$$\lim_{m \rightarrow \infty} \sqrt{m} \hat{\mathbf{R}}_m(\mathcal{F}, \mathcal{D}) \geq \sqrt{\frac{2}{\pi} \mathbf{W}^r(\mathcal{F})} \text{ and } \lim_{m \rightarrow \infty} \sqrt{m} \hat{\mathbf{R}}_m(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \geq \sqrt{\frac{2}{\pi} \mathbf{W}(\mathcal{F})} .$$

To make the result concrete, consider that as soon as \mathcal{F} contains a function f and a “ c -shifted” version of it $f + c$, for some $c \in \mathbb{R}^+$, then $\sup_{g \in \mathcal{F}} |\hat{\mathbb{E}}_{\mathbf{x}}[g]| \geq c/2$, thus $\widehat{\mathbf{W}}_{\mathbf{x}}^r(\mathcal{F}) \geq c^2/4$, and from the above lemma, $\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) \geq c/\sqrt{8m}$, but $\hat{\mathbf{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$ does not suffer from this issue.

The significance of lemma 1.3.4 is that a dependence on the (*centralized*) wimpy variance *cannot* be obtained *without* empirical centralization. One must settle for dependence on the *raw* wimpy variance, which can be unboundedly larger than its centralized counterpart. The result also tells us that a dependence on the (centralized) wimpy variance *may* be attained with empirical centralization. We show next that such is indeed the case.

Optimal dependence on wimpy variance The quantity $\hat{\mathbf{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$ is an ERA, thus it can be upper-bounded using Massart’s finite-class lemma [Massart, 2000, lemma 5.2]. We now apply this celebrated result to bound the ERA under *empirical centralization* while including the *absolute value* (absent from some presentations) inside the supremum of the ERA.

Corollary 1.3.5. Assume that \mathcal{F} is finite. Let $\mathcal{F}_{\pm} \doteq \mathcal{F} \cup \{-f, f \in \mathcal{F}\}$. It holds

$$\hat{\mathbf{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \leq \sqrt{\frac{2\widehat{\mathbf{W}}_{\mathbf{x}}(\mathcal{F}) \ln |\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}_{\pm})|}{m}} . \tag{13}$$

The use of \mathcal{F}_\pm is needed¹ to handle the absolute value in our definition of the ERA (see (3)). Without empirical centralization, the dependence would be on the raw wimpy variance, which equals the squared ℓ_2 norm in “classic” presentations of Massart’s lemma. Corollary 1.3.5 shows that empirical centralization enables *optimal* dependence on the *centralized* wimpy variance, which *cannot be obtained without empirical centralization*, as shown in lemma 1.3.4.

Monte-Carlo estimation The quantity $\hat{\mathbf{R}}_m(\hat{C}_x(\mathcal{F}), \mathbf{x})$ is an ERA, so it “suffers” from the usual issue of how to actually compute or bound it in order to bound the SD via theorem 1.3.3. While analytical methods (e.g., Massart’s lemma) yield (generally loose) bounds, Monte-Carlo estimation with proper tail bounds gives better results in practice, and it was proposed almost concurrently with the introduction of the ERA [Bartlett and Mendelson, 2002].

Definition 1.3.6. Let $\boldsymbol{\sigma} \in (\pm 1)^{n \times m}$ be a matrix of i.i.d. Rademacher r.v.’s. The *Monte-Carlo ERA* $\hat{\mathbf{R}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma})$ of \mathcal{F} on \mathbf{x} w.r.t. $\boldsymbol{\sigma}$ is the quantity

$$\hat{\mathbf{R}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) \doteq \frac{1}{n} \sum_{i=1}^n \sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{j=1}^m \sigma_{i,j} f(\mathbf{x}_i) \right| .$$

It clearly holds $\mathbb{E}_\sigma[\hat{\mathbf{R}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma})] = \hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x})$. Bartlett and Mendelson [2002, Thm. 11] show that the MC-ERA with $n = 1$ is concentrated about the ERA as

$$\mathbb{P}_\sigma \left(\left| \hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) - \hat{\mathbf{R}}_m^1(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) \right| \geq \varepsilon \right) \leq 2 \exp \left(\frac{-2m\varepsilon^2}{q^2} \right) .$$

The r.h.s. can be used in theorem 1.3.3 inside the definition of λ (with the needed adjustment of the confidence parameter δ using a union bound), thus obtaining an upper bound to the SD using the MC-ERA. The leitmotif of this work is to obtain strong, practical, sample-dependent bounds to the SD, so we derive a novel tail bound to the MC-ERA (theorem 1.3.7) *for general n* , where the strong dependence on q^2 of the above bound is replaced by a much weaker dependence, primarily on $W(\mathcal{F})$. This change is similar to how theorem 1.3.2 improves over textbook bounds to the SD that use McDiarmid’s bounded difference inequality. Our improved variance-sensitive bound uses a transportation-method inequality due to Samson [2007] to *upper bound the expectation of suprema of empirical processes*. This result is, to our knowledge, novel, and is worst-case asymptotically equivalent to the McDiarmid bounds, and improves over it when the wimpy variance is small. The bound uses the *empirical maximum* $\hat{q}_\mathcal{F}(\mathbf{x})$ of \mathcal{F} on \mathbf{x} , defined as

$$\hat{q}_\mathcal{F}(\mathbf{x}) \doteq \sup_{f \in \mathcal{F}, \mathbf{x} \in \mathbf{x}} |f(\mathbf{x})| \quad (\leq q) .$$

Theorem 1.3.7. Let $\boldsymbol{\sigma} \in (\pm 1)^{n \times m}$ be a matrix of i.i.d. Rademacher r.v.’s. Let $\delta \in (0, 1)$. With probability at least $1 - \delta$ over the choice of $\boldsymbol{\sigma}$, it holds

$$\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) \leq \hat{\mathbf{R}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \frac{2\hat{q}_\mathcal{F}(\mathbf{x}) \ln \frac{1}{\delta}}{3nm} + \sqrt{\frac{4\widehat{W}_\mathbf{x}^r(\mathcal{F}) \ln \frac{1}{\delta}}{nm}} . \quad (14)$$

Empirical centralization obtains a dependence on the empirical wimpy variance of \mathcal{F} , rather than on the raw (i.e., non-centralized) wimpy variance. This advantage propagates when using the MC-ERA of the empirical centralization to bound the SD of \mathcal{F} . The dependence on the empirical maximum changes from $\hat{q}_\mathcal{F}(\mathbf{x})$ to $\hat{q}_{\hat{C}_x(\mathcal{F})}(\mathbf{x})$, which can be a large improvement (and $\hat{q}_{\hat{C}_x(\mathcal{F})}(\mathbf{x}) < 2\hat{q}_\mathcal{F}(\mathbf{x})$ at most).

¹Since $|\hat{C}_x(\mathcal{F}_\pm)| \leq 2|\hat{C}_x(\mathcal{F})|$, the bound can be reformulated as function of $|\hat{C}_x(\mathcal{F})|$ only.

Corollary 1.3.8. Let $\sigma \in (\pm 1)^{n \times m}$ be a matrix of i.i.d. Rademacher r.v.'s. Let $\delta \in (0, 1)$. With probability at least $1 - \delta$ over the choice of σ , it holds

$$\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}), \mathbf{x}) \leq \hat{\mathfrak{R}}_m^n(\hat{C}_x(\mathcal{F}), \mathbf{x}, \sigma) + \frac{2\hat{q}_{\hat{C}_x(\mathcal{F})}(\mathbf{x}) \ln \frac{1}{\delta}}{3nm} + \sqrt{\frac{4\widehat{W}_x(\mathcal{F}) \ln \frac{1}{\delta}}{nm}}.$$

Although $n = 1$ Monte-Carlo trials are sufficient to match the convergence rate of theorem 1.3.2, the Monte-Carlo estimation error term can still be a significant portion of the total SD bound. For practical usage, particularly with small sample sizes, or when extremely tight bounds are needed, more Monte-Carlo trials (i.e., larger n) rapidly reduce the Monte-Carlo estimation error, and this error is soon dominated by the tail bound terms of theorem 1.3.2.

Example: Batch panel of experts Consider now the *batch panel of experts* problem, where \mathcal{F} is a *finite family of experts*, and the task is to select the (approximately) *most accurate* among them, given a sample of *labeled instances*. With the Monte-Carlo method, we may sharply bound the SD whenever evaluating the requisite suprema is computationally feasible, e.g., via *enumeration* of \mathcal{F} . Furthermore, we automatically benefit from *data-dependent* and *distribution-dependent* structure, e.g., highly correlated or anticorrelated experts, and *low wimpy variance* over uniformly accurate \mathcal{F} . This example immediately extends to *model selection via structural risk minimization* if, e.g., the experts are organized into *concentric groups* by some *a priori confidence* or *quality* estimate.

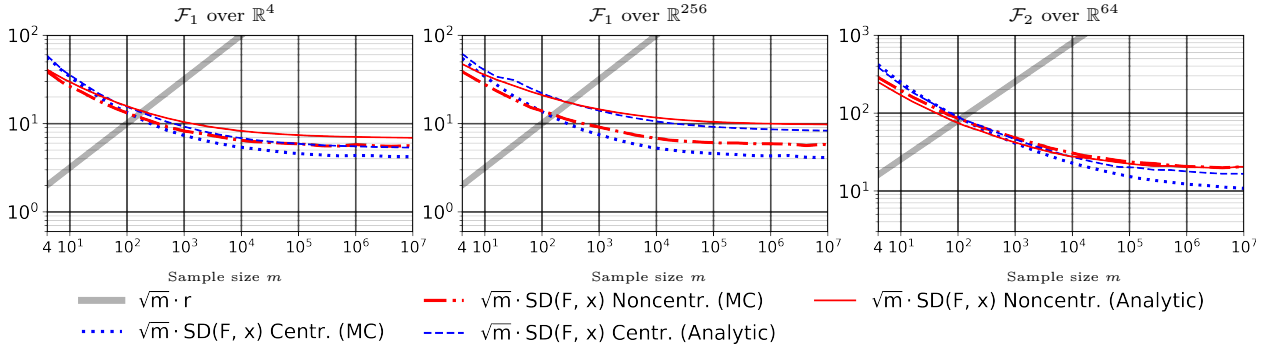


Figure 1.1: Comparison of SD bounds as functions of the sample size m . See the main text for an explanation of the results.

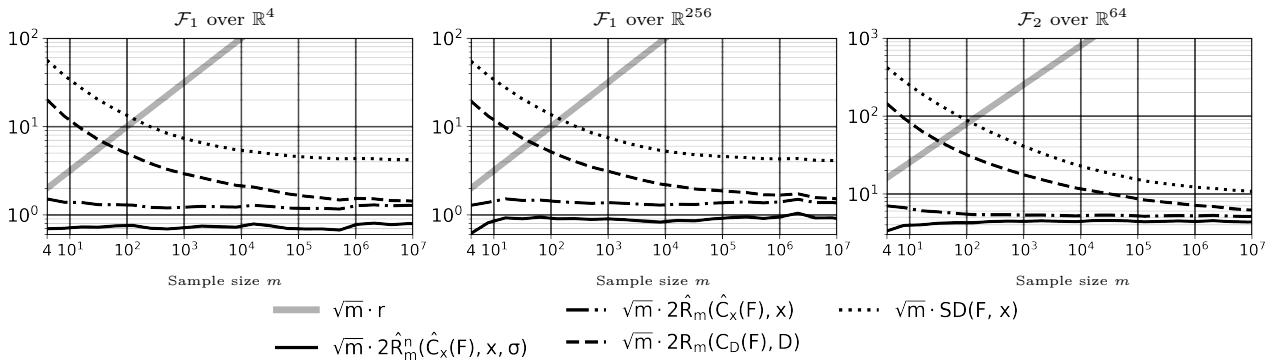


Figure 1.2: Upper bounds to complexity measures and SD as functions of the sample size m . See the main text for details.

1.4 Experimental evaluation

We performed experiments to evaluate the various bounds presented in the previous sections and compare the bounds to the SD using empirical centralization to those without centralization. The code is included in the supplementary material.

Function families We consider the function families \mathcal{F}_p , for any $p \geq 1$, containing all unit ℓ_p -norm-constrained linear functions in \mathbb{R}^d , i.e.,

$$\mathcal{F}_p \doteq \{\mathbf{x} \mapsto w \cdot \mathbf{x}, w \in \mathbb{R}^d \text{ s.t. } \|w\|_p \leq 1\} .$$

These families are of immediate interest in many machine learning settings, such as the analysis of support vector machines and neural networks, as both consist of Lipschitz loss and/or activation functions applied to one or more linear functions (see, e.g., [Bartlett and Mendelson, 2002] for analysis). Additionally, a bound on the SD of \mathcal{F}_p over distribution \mathcal{D} over \mathbb{R}^d corresponds to the radius of the $\ell_{p/p-1}$ (Hölder dual norm) ball about $\hat{\mathbb{E}}[\mathbf{x}]$ in which $\mathbb{E}_{\mathcal{D}}[\mathbf{x}] \in \mathbb{R}^d$ falls. Such balls can be used to estimate *covariance matrices*, high-dimensional *sufficient statistics* in graphical models [Bradley and Guestrin, 2012], and to learn *equilibria* in simulation-based games [Areyan Viqueira et al., 2019, 2020].

Analytical bounds to the ERA of \mathcal{F}_p on \mathbf{x} and Monte-Carlo estimates of it (see definition 1.3.6) are relatively straightforward [Shalev-Shwartz and Ben-David, 2014, Lemmas 26.10, 26.11] (see lemma A.2.1 in the supplementary material). The following lemma extends these results to the empirical centralization.

Lemma 1.4.1. Let $\bar{\mathbf{x}} \doteq \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i \in \mathbb{R}^d$. For the ℓ_1 norm, it holds

$$\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}_1), \mathbf{x}) = \mathbb{E}_{\sigma} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_{\infty} \right] \leq \max_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|_{\infty} \sqrt{\frac{2 \ln(2d)}{m}},$$

while for the ℓ_2 norm, it holds

$$\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}_2), \mathbf{x}) = \mathbb{E}_{\sigma} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_2 \right] \leq \max_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|_2 \frac{1}{\sqrt{m}} .$$

Similar bounds are possible for other values of p ; e.g., by linearity, the case of $p = \infty$ is trivial. Note that in addition to computing MC-ERAs from $\ell_{p/p-1}$ dual norms, we may also compute (raw) empirical wimpy variances from *operator norms* of (raw) *covariance matrices* of \mathbf{x} . In particular, for \mathcal{F}_1 , it is easy to show that the wimpy variance is simply the *largest variance* along any *standard basis vector*. Similarly, for \mathcal{F}_2 , the wimpy variance is simply the *maximum variance* along any *unit vector*, i.e., the *spectral norm* of the covariance matrix.

Data generation and parameter values We generated the samples \mathbf{x} for our experiments from random distributions over \mathbb{R}^d . The ERA of the family \mathcal{F}_1 is susceptible to the value of d (see lemma 1.4.1 and lemma A.2.1 in the supplementary material), so we use $d = 4$ and $d = 256$, while in the case of \mathcal{F}_2 the ERA is independent of d , so we use $d = 64$. Details of the distributions are in the supplementary material. Range-like quantities (i.e., q , \hat{q} , or r) can be computed from the data and/or known *a priori* bounds: $r = 1$ for our \mathcal{F}_1 experiments and $r = 8$ for the \mathcal{F}_2 case. (Raw) wimpy variances correspond to norms of the (raw) covariance matrices used for data generation (see the supplementary material for details). In all experiments, we used $\delta = 0.01$ and $n = 32$ (we comment on this choice below). The sample size m varied from 4 (the minimum possible, due to lemma 1.2.1) to 10^7 .

A note on results visualization We present all of our results in plots with *log-log axes*, so that convergence rates are clearly visible as slopes, and constant factors as vertical offsets. The x -axis is the sample size m . Since we expect asymptotic convergence rates $\propto C/\sqrt{m}$, where C depends on the (possibly raw) wimpy variance of \mathcal{F} , r , and δ , we plot all quantities *multiplied by \sqrt{m}* . This transformation allows to clearly visualize $\Theta(C/\sqrt{m})$ behaviors as *straight horizontal lines*, and $\mathfrak{o}(C/\sqrt{m})$ behaviors as (transient) downward slopes. For completeness, we show plots without the scaling by \sqrt{m} in the supplementary material.

Results Figure 1.1 compares four bounds to the SD: using the Monte-Carlo estimate $\hat{\mathfrak{R}}_m^n(\hat{C}_x(\mathcal{F}_p), \mathbf{x}, \sigma)$ for the ERA $\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}_p), \mathbf{x})$ of the empirical centralization of \mathcal{F}_p on \mathbf{x} , using the Monte-Carlo estimate $\hat{\mathfrak{R}}_m^n(\mathcal{F}_p, \mathbf{x}, \sigma)$ for the ERA $\hat{\mathfrak{R}}_m(\mathcal{F}_p, \mathbf{x})$ of the non-centralized \mathcal{F}_p , using analytical bounds to $\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}_p), \mathbf{x})$ from lemma 1.4.1, and using analytical bounds to $\hat{\mathfrak{R}}_m(\mathcal{F}_p, \mathbf{x})$ from lemma A.2.1 (in the supplementary material). The thicker grey line is the quantity \sqrt{mr} ; bounds above this line are *vacuous*.

At very small sample sizes (when all bounds are *vacuous*), the bounds obtained without centralization are sharper than the bounds with empirical centralization, due to the bias-correction of lemma 1.2.1 (see ξ in theorem 1.3.3) and the (fast-decaying) $\Theta(r/m^{3/4})$ term of theorem 1.2.2. Before $m \approx 200$, when bounds become non-vacuous, the advantages of empirical centralization become clear, and increase with the sample size. Recall that each bound is scaled by \sqrt{m} , thus all are *asymptotically horizontal*, as $\Theta(C/\sqrt{m})$ terms eventually dominate the bound to the SD, where C varies greatly between bounds and methods. Thus without empirical centralization, obtaining the same bound to the SD would require a larger sample size m than with empirical centralization (this effect can be better observed in the non- \sqrt{m} -scaled plots in figure A.1 in the supplementary material.) The Monte-Carlo estimate, despite using only $n = 32$ Monte-Carlo trials, gives better bounds to the SD than an analytical approach.

In figure 1.2, we drill down on the SD bounds using the Monte-Carlo estimate $\hat{\mathfrak{R}}_m^n(\hat{C}_x(\mathcal{F}_p), \mathbf{x}, \sigma)$ for the ERA $\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}_p), \mathbf{x})$ of the empirical centralization of \mathcal{F}_p on \mathbf{x} , showing this quantity, together with the *upper bounds* to other intermediate quantities, that eventually lead to the SD bound: the ERA $\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}_p), \mathbf{x})$ (obtained by applying theorem 1.3.7 to the MC-ERA), the RA $\mathfrak{R}_m(\mathcal{F}_p, \mathcal{D})$ (obtained by applying theorem 1.2.2 and lemma 1.2.1 to the bound on the ERA), and SD (obtained by applying the r.h.s. of (6) and theorem 1.3.2 to the bound on the RA).

At small sample sizes, the fast-decaying terms dominate the bounds to the RA and SD, but, true to their nature, quickly become negligible: all bounds are asymptotically $\Theta(C/\sqrt{m})$, where C , which in the plots in figure 1.2 appear as the vertical offset of each curve at high sample sizes, depends mostly on the wimpy variance of \mathcal{F} and the range r . The bounds that decay as $\Theta(\sqrt{W(\mathcal{F})/m})$ (i.e., the MC-ERA \rightarrow ERA and RA \rightarrow SD bounds) introduce constant factor terms, manifest as asymptotic vertical gaps, whereas the remaining bounds entirely vanish asymptotically. The gap from the MC-ERA to the ERA would disappear as the number n of Monte-Carlo trials (which we fixed at $n = 32$) increases.

The range and wimpy variances are approximately the same in both \mathcal{F}_1 experiments but the MC-ERA are much larger when $d = 256$ because here the RA is essentially the expected largest distance traveled over d random walks, which increases with d (see also lemma 1.4.1).

In conclusion, the results confirm the advantages of empirical centralization to obtain tighter bounds to the SD with optimal dependence on the wimpy variance, while still maintaining the same behavior in terms of the number of samples as bounds not using centralization.

1.5 Conclusions

We develop practical, sharp, sample-dependent probabilistic bounds to the SD through *empirical centralization*, together with novel results on the concentration of the wimpy variance and of Monte-Carlo estimates of the ERA. Our bounds exhibit optimal dependence on the wimpy variance and the same dependence on the

number of samples as bounds not using centralization. The results of our experimental evaluation show that the advantage is significant even at small sample sizes, and remains so as the sample size grows. In future work, we will explore the important relationship between centralization and localization [Koltchinskii, 2006, Giné and Koltchinskii, 2006].

Chapter 2

Adversarial Learning with Weak Supervision

In this work, I develop a rigorous approach for using a set of possibly correlated weak supervision sources in order to solve a multi class classification task, when only a very small set of labeled data is available. The algorithm optimally learns a model that has minimum expected risk among all the feasible solutions for a set of unlabeled data, where the feasibility of a labeling is computed through constraints defined by estimated statistics of the weak supervision sources. I provide theoretical guarantees for this approach that depend on the information provided by the weak supervision sources. Notably, this method does not require the weak supervision sources to have the same labeling space as the multi-class task. Finally, I demonstrate the effectiveness of the approach with experiments on various image classification tasks. This chapter is based on joint work with Alessio Mazzeo, Dylan Sam, Stephen Bach, and Eli Upfal.

2.1 Introduction

In the last decade, deep neural networks have been applied to accurately solve a wide range of classification tasks in different domains, but the supervised learning of these models requires a considerable amount of labeled data. An alternative strategy is to learn from *weak supervision*, i.e., sources of labels that are noisy or heuristic. Examples include hand-written rules [Ratner et al., 2017, Wu et al., 2018, Safranchik et al., 2020] and classifiers trained for related tasks [Varma et al., 2017, Bach et al., 2019, Chen et al., 2019]. Even if these sources of information are noisy, results show that they can lead to high-quality models, particularly when the outputs from many sources are combined.

A key technical challenge in such work is how to combine multiple sources of weak supervision, since they might conflict with each other. We assume access to little to no ground truth data. Instead, much prior work on aggregating noisy labels [Dawid and Skene, 1979, Zhang et al., 2016, Gao and Zhou, 2013, Karger et al., 2014, Ghosh et al., 2011, Dalvi et al., 2013, Ratner et al., 2016, 2019] assumes that the sources make independent errors, which is a very strong assumption. Some recent works [Bach et al., 2017, Varma et al., 2019] attempt to learn more sophisticated distributions but still rely on parametric assumptions that make conditional independence assumptions. Independence assumptions in models of weak supervision sources are hard to verify and limiting in practice. Many useful weak supervision sources, particularly ones learned from related datasets, can be arbitrarily correlated, as there are systematic differences between the target classification task and the mildly related tasks used to learn them. For example, if all the labelers are fine-tuned from the same pretrained model, they are likely to inherit some of the same biases.

Recently, two pieces of work have addressed the problem of combining weak labelers without distribution assumptions by taking an adversarial approach. Adversarial label learning (ALL) [Arachie and Huang, 2019] first formulated the problem as a minimax optimization, learning a model that minimizes risk using the worst-case assignment to the unknown ground truth labels. ALL is restricted to binary classification, it does not solve the minimax optimization problem optimally, and provides no theoretical guarantees for the models it learns. A later work, performance guaranteed majority vote (PGMV) [Mazzetto et al., 2021], takes an alternative approach for the binary classification setting. It uses a small amount of labeled data and a large amount of unlabeled data to empirically estimate properties of the labelers and define a constrained optimization over the assignments to unknown labels with a worst-case error bound. However, this approach is also limited to binary classification because it exploits the fact that when two labelers disagree, one must be correct. Further, it requires performing the worst case optimization over each subset of labelers considered. The resulting combinatorial blowup limits the number of labelers that can be considered simultaneously.

In this work, we address the limitations of ALL and PGMV by providing a computationally efficient framework for multi class weak supervision with performance guarantees. Similar to ALL, we formulate the search for ground truth as a search over a set of ground truth labels that satisfy statistical constraints on the weak supervision sources. Unlike ALL, we optimally solve the optimization problem, and we show a novel analysis that factors in the information provided by the weak supervision sources with respect to the target classification, geometrically represented as the ℓ_1 diameter of the region of feasible labelings, to evaluate the quality of the learnt model.

Contributions. We introduce a novel method to use the information provided by a set of arbitrarily correlated weak supervision sources in order to learn a classifier for a target classification task of interest. Inspired by previous work, we use a little amount of labeled data to compute statistics on the weak supervision sources, and we formulate an optimization problem to find the prediction model that achieves the lowest risk among all the labelings of an unlabeled dataset that agree with those statistics.

Our main contributions are the following:

1. We develop the first rigorous method for learning multi class classification from weak supervision sources (Section 4).
2. We provide theoretical analysis of our method, proving approximation guarantees on the quality of our solution, and complexity bounds for the algorithm’s run-time (Section 4).
3. We evaluate the quality of the solution provided by our method using a geometrical quantity that represents the aggregate information provided by the weak supervision sources with respect to the target classification task (Section 4.2).
4. While the presentation of our method is general, we demonstrate the applicability of our approach through two practical instances of prediction model and loss function: convex combination of the weak supervision sources and multinomial logistic regression (Section 4.1)
5. We show how to extend our method for the case when the weak supervision sources do not have the same labeling space as the multi-classification task. This is useful, for example, when learning with attributes. In many weak supervision tasks we have information that partially constrains the correct label, such as knowing that if a labeler detects stripes on an animal, it narrows down the possible animal types (Section 4.3).
6. We conduct extensive experiments demonstrating the effectiveness of our novel approach for multi class classification tasks. We also show that the performance of our approach compares favorably with the best published adversarial weak learning algorithms for binary classification, ALL [Arachie and Huang, 2019] and PGMV [Mazzetto et al., 2021].

2.2 Related Work

The problem of learning from multiple, possibly conflicting weak labelers with little to no ground truth data has received considerable attention recently [Ratner et al., 2016, Bach et al., 2017, Ratner et al., 2017, Varma et al., 2019, Arachie and Huang, 2019, Mazzetto et al., 2021]. This setting is distinct from much work on ensemble learning [Zhang and Ma, 2012], such as boosting [Schapire, 1990, Freund, 1995], where abundant labeled examples are used to learn to combine ensemble members. Other ensemble methods, such as bagging [Breiman, 1996] take an unweighted vote of ensemble members but rely on the assumption that each member is trained on labeled data sampled from the target distribution. Unlike these methods, in weak supervision, the goal is to use other statistical properties of the labelers, such as their agreements and disagreements, to learn to combine them. In this way, the labelers can potentially be improved without increasing the need for labeled training data.

This work has its roots in crowdsourcing, where the “labelers” are people with different, unknown levels of reliability. Dawid and Skene’s seminal work [Dawid and Skene, 1979] showed how the accuracy of each labeler could be estimated with expectation maximization by assuming a naive Bayes distribution over the labelers’ votes and the latent ground truth. Since then, much work has provided theoretically guaranteed algorithms for learning under these assumptions [Zhang et al., 2016, Gao and Zhou, 2013, Karger et al., 2014, Ghosh et al., 2011, Dalvi et al., 2013]. When the labelers are humans working without coordination, the independence assumption is a reasonable one.

Recently, frameworks for weakly supervised machine learning like Snorkel [Ratner et al., 2016, Bach et al., 2017, Ratner et al., 2017] have used and extended these learning techniques to the setting in which the labelers are programmed rules, weak classifiers, or other heuristics. As described in the introduction, learned and programmed labelers can have heavily correlated errors because of common elements in the heuristics they use. This potential problem has motivated attempts to relax the independence assumption. One line of work [Bach et al., 2017, Varma et al., 2019] has tried to learn more sophisticated parametric models of the labelers, but they are still limited by how correct their assumptions are, which are hard to verify in practice. In this work, we therefore focus on methods for learning from weak supervision without assumptions on the distribution of labeler outputs and ground truth.

2.3 Preliminaries

We denote scalar and generic items as lowercase symbols, vectors as lowercase bold symbols, and matrices are bold uppercase symbols. The i -th column of a matrix \mathbf{A} as denoted by the corresponding lowercase symbol \mathbf{a}_i , i.e., $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n]$. Due to space constraints, the proofs are deferred to the Appendix.

In multi class learning, we have a domain \mathcal{X} and a classifier function h that maps each $x \in \mathcal{X}$ to one of k possible labels (classes). Since we will work later with distributions over the k classes, it is convenient to represent label i , $1 \leq i \leq k$ as a k dimension vector \mathbf{e}_i with all the components set to 0 except for the i -th component that is set to 1. Thus, $h : \mathcal{X} \rightarrow \mathcal{Y} = \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$. A classifier (e.g., softmax layer of a neural network) may output a probability distribution vector $\mathbf{y} \in \mathbb{R}_{\geq 0}^k$ over the k classes, where y_c is the probability that the item belongs to class c , and $\sum_c y_c = 1$. We call $\mathcal{Y}^\circ \supset \mathcal{Y}$ the set of all possible probability vectors. A *loss function* $\ell : \mathcal{Y}^\circ \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ quantifies the error of the classifier’s output $h(x)$ with respect to the true label \mathbf{y} . Let $p_{\mathcal{X}\mathcal{Y}}$ be the probability distribution over $\mathcal{X} \times \mathcal{Y}$. Given a classifier h , its *risk* is defined as

$$R(h) \doteq \mathbb{E}_{x \sim p_{\mathcal{X}\mathcal{Y}}} \ell(h(x), \mathbf{y})$$

In standard supervised learning, we are given *labeled samples* from $p_{\mathcal{X}\mathcal{Y}}$, and we find a classifier with low risk among a set of classifiers \mathcal{H} , which is also called a *hypothesis class*. The amount of labeled data required to guarantee that we can find (or *train*) such a classifier is referred to *sample complexity*, and it is proportional

to a measure of richness of \mathcal{H} . For many classification tasks of interest, there could be low availability of labeled data, and this is a critical problem for a wide range of domains, where the most successful hypothesis classes have high dimension (e.g., convolutional neural networks for images).

In this work we assume access to m_L i.i.d. labeled samples $(\tilde{x}_1, \tilde{\mathbf{y}}_1), \dots, (\tilde{x}_{m_L}, \tilde{\mathbf{y}}_{m_L})$ from $p_{\mathcal{X}\mathcal{Y}}$, where the amount m_L is insufficient for the direct supervised learning of \mathcal{H} . To circumvent the lack of sufficient training data, we assume access to a set of weak labelers (classifiers) $\varphi_1, \dots, \varphi_n$, also called weak supervision sources. These labelers are *weak* in that they can be inaccurate with respect to the target classification task. For example, the weak labelers could be trained for classification tasks that are only mildly related to the target classification task. For example, we could train a labeler to detect stripes on zebras and horses, and then attempt to use it to label data as either tigers or lions. Moreover, we add no further assumptions on the properties of those classifiers, and their output could be arbitrarily correlated. We also assume access to m unlabeled data points $X = \{x_1, \dots, x_m\}$ sampled independently from the marginal distribution $p_{\mathcal{X}}$. Our method uses the weak supervision sources $\varphi_1, \dots, \varphi_n$ to constraint the space of the possible labels that can be given to those unlabeled data points. We use the limited labeled data to compute statistics about the weak classifier, and then consider possible labeling of the unlabeled data X that satisfy those statistics.

As an example, suppose that we use the m_L labeled data points to compute as a statistic the *empirical risk* of each weak supervision source, i.e., we take

$$\hat{\mu}_i \doteq \frac{1}{m} \sum_{j=1}^m \ell(\varphi_i(\tilde{x}_j), \tilde{\mathbf{y}}_j) ,$$

for $i \in 1, \dots, n$. In section 2.4 we will use related statistics in order to prove strong theoretical guarantees. If we were to assign a labeling to the unlabeled data points X , a reasonable approach would be to find a labeling such that the empirical risk of the weak supervision source i computed with respect of those labels is equal to $\hat{\mu}_i$. However, this is a computationally hard problem, as we have to assign a discrete label (from \mathcal{Y}) to each item, and each label affects the empirical risk of all the weak supervision sources. Moreover, there is no guarantee that we can find such a labeling for the unlabeled data, and it is unclear which labeling to choose in case there are multiple solutions.

To address the computational issues with a discrete selection of the labels, we assign a probability vector from \mathcal{Y}° to each unlabeled data point. In other words, for each unlabeled item x_j , we assign a probability vector \mathbf{y} , where the c -th component of that vector represents the probability that item x_j belongs to class c . Given a classifier h , we compute the loss of the classifier on item x with respect to the probability vector $\mathbf{y} \in \mathcal{Y}^\circ$ as *expected loss*. Abusing notation, let $\mathbf{e} \sim \mathbf{y}$ denote that $\mathbf{e} = \mathbf{e}_c \in \mathcal{Y}$ with probability y_c , and define

$$\ell^\circ(h(x), \mathbf{y}) \doteq \mathbb{E}_{\mathbf{e} \sim \mathbf{y}} \ell(h(x), \mathbf{e}) = \sum_{c=1}^k y_c \cdot \ell(h(x), \mathbf{e}_c) . \quad (15)$$

We observe that this definition of loss generalizes the one computed with respect to a discrete labeling, in fact for each $\mathbf{e} \in \mathcal{Y}$, we have that $\ell(h(x), \mathbf{e}) = \ell^\circ(h(x), \mathbf{e})$. The loss (15) is computationally easier to work with, as it is *linear* with respect to the labeling \mathbf{y} . Let $\mathbf{Y} \in \mathbb{R}^{k \times m}$ be a matrix that describes a possible labeling of the unlabeled data points; in particular the j th column of the matrix \mathbf{Y} is $\mathbf{y}_j \in \mathcal{Y}^\circ$, and it denotes the probability vector of the labeling of the item x_j . The empirical risk of a classifier h on the unlabeled data X with labeling \mathbf{Y} is then defined as

$$\hat{\mathbf{R}}(h; X, \mathbf{Y}) \doteq \frac{1}{m} \sum_{j=1}^m \ell^\circ(h(x_j), \mathbf{y}_j) .$$

Finding a labeling \mathbf{Y} for which $\hat{\mathbf{R}}(h; X, \mathbf{Y}) = \hat{\mu}_i$ for $i \in 1, \dots, n$ is equivalent to the computationally easy task of solving a linear system with $\mathbf{O}(n + m)$ constraints (the n constraints on the empirical risk equality and m constraints on probability vectors summing to 1) and $\mathbf{O}(mk)$ variables. However, there still could be

multiple solutions to such a linear system. The core idea of the method presented in section 2.4 is to find a model that has the lowest risk across all those possible solutions.

2.4 Adversial Labeling Method

Let $\mathcal{H} = \{\mathbf{h}_\theta : \theta \in \Theta \subseteq \mathbb{R}^d\}$ be the hypothesis class that we will use to find the classifier for the classification task of interest, where each classifier $h \in \mathcal{H}$ is parametrized by a vector of weights θ .

Let \mathbf{Y}^* be the (unknown) true labelings of the unlabeled data X . For each weak supervision source i , we use the labeled data to compute a interval Δ_i such that with high probability we have that $\hat{R}(\varphi_i(x); X, \mathbf{Y}^*) \in \Delta_i$ for $i \in 1, \dots, n$. This is a crucial property that we will need to show our theoretical bound (theorem 2.4.8), and one may refer to lemma 2.4.1 to see how to build such intervals.

Let \mathbb{Y}^\diamond be the set of all possible labeling matrices \mathbf{Y} such that the empirical risk of φ_i computed with respect to the labeling \mathbf{Y} of the unlabeled data X belongs to the corresponding interval Δ_i for each weak supervision source. Formally, the set \mathbb{Y}^\diamond is defined as

$$\mathbb{Y}^\diamond \doteq \left\{ \mathbf{Y} \in \mathbb{R}^{k \times m} \left| \begin{array}{ll} \mathbf{y}_j & \in \mathcal{Y}^\diamond \quad \forall j \in 1, \dots, m \\ \hat{R}(\varphi_i; X, \mathbf{Y}) & \in \Delta_i \quad \forall i \in 1, \dots, n \end{array} \right. \right\} .$$

We will refer to \mathbb{Y}^\diamond as the set of feasible labelings. The next lemma shows how to build the intervals Δ_i to guarantee that with high probability the true labeling \mathbf{Y}^* is feasible.

Lemma 2.4.1 (Two-Sample Bounded Difference Interval Statistics). Suppose that the codomain of the loss function ℓ is contained in the interval $[0, B]$. Let $\hat{\mu}_1, \dots, \hat{\mu}_n$ be the empirical risks of $\varphi_1, \dots, \varphi_n$ computed with respect to the m_L labeled samples. Fix a value $\delta \in (0, 1)$ and let γ be defined as

$$\gamma \doteq B \sqrt{\frac{(m_L + m) \ln \frac{2n}{\delta}}{2m_L m}} .$$

If we set $\Delta_i = [\mu_i - \gamma, \mu_i + \gamma]$, then with probability at least $1 - \delta$ it holds that $\mathbf{Y}^* \in \mathbb{Y}^\diamond$.

We want to find the classifier that achieves the lowest empirical risk among the feasible labelings of the unlabeled data points. That is, we choose the classifier $\mathbf{h}_{\hat{\theta}} \in \mathcal{H}$, where $\hat{\theta}$ is the solution of the following minimax problem

$$\hat{\theta} \doteq \operatorname{argmin}_{\theta \in \Theta} \max_{\mathbf{Y} \in \mathbb{Y}^\diamond} \hat{R}(\mathbf{h}_\theta; X, \mathbf{Y}) . \quad (16)$$

The optimization problem above has some nice properties. The set \mathbb{Y}^\diamond is specified by linear constraints in \mathbf{Y} . Moreover, the objective of the minimax (16) problem is also linear in \mathbf{Y} . Hence, it is easy to see that for a given $\theta \in \Theta$, it is possible to solve the maximization problem

$$f(\theta) \doteq \max_{\mathbf{Y} \in \mathbb{Y}^\diamond} \hat{R}(\mathbf{h}_\theta; X, \mathbf{Y}) \quad (17)$$

through a *linear program* with $\mathbf{O}(mk)$ variables and $\mathbf{O}(m + n)$ linear constraints.

In order to solve the minimax problem (16), we will introduce a few assumptions on the loss function and the model choice \mathcal{H} , which are satisfied by many classic machine learning settings. In particular, we would like the function $f(\theta)$ to be convex, so that we can solve the minimization problem $\min_{\theta \in \Theta} f(\theta)$. Even if $f(\theta)$ is convex, we may not be able to apply a gradient-based optimization method as $f(\theta)$ involves a maximization, hence it is not differentiable everywhere. To solve this issue, we will use the concept of subgradient, which is a generalization of the concept of gradient. This will require the loss function to be Lipschitz. A function $g : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2}$ is said to be L -Lipschitz if for any $\mathbf{a}, \mathbf{a}' \in \mathbb{R}^{d_1}$, it holds that $\|g(\mathbf{a}) - g(\mathbf{a}')\|_2 \leq L \|\mathbf{a} - \mathbf{a}'\|_2$.

Definition 2.4.2 (The Subgradient). Let $\mathcal{A} \subseteq \mathbb{R}^b$ be the domain of a function g . A vector $\mathbf{v} \in \mathbb{R}^b$ is a subgradient for a function g at $\mathbf{x} \in \mathcal{A}$ if for any $\mathbf{y} \in \mathcal{A}$ we have that

$$g(\mathbf{a}') - g(\mathbf{a}) \geq \mathbf{v}^T \cdot (\mathbf{a}' - \mathbf{a}) .$$

For each $\mathbf{a} \in \mathcal{A}$, we define

$$\partial g(\mathbf{a}) \doteq \{ \mathbf{v} : \mathbf{v} \text{ is a subgradient of } g \text{ at } \mathbf{a} \} .$$

If a function is differentiable in a point, then its subgradient with respect to that point is unique, and it is equal to the gradient.

Fact 2.4.3. If a function g is convex, then there exists at least one subgradient for each point of its domain.

The following intermediate result that immediately follows from the definition of ℓ° will prove useful throughout this discussion.

Lemma 2.4.4 (Lipschitz Constants under Convex Combination). Let $\ell(\mathbf{h}_\theta(x), \mathbf{e})$ be convex and L -Lipschitz with respect to θ for any $(x, \mathbf{e}) \in \mathcal{X} \times \mathcal{Y}$. Then, for any probability vector $\mathbf{y} \in \mathcal{Y}^\circ$, the function $\ell^\circ(\mathbf{h}_\theta(x), \mathbf{y})$ is also convex and L -Lipschitz with respect to θ .

The next Lemma shows that under some conditions, it is possible to compute the subgradient of the function f .

Lemma 2.4.5 (The Subgradient at Nondifferentiable Points). Fix a value $\theta \in \Theta$, let $\mathbf{Y}' \doteq \operatorname{argmax}_{\mathbf{Y} \in \mathbb{Y}^\circ} \hat{\mathbf{R}}(\mathbf{h}_\theta; X, \mathbf{Y})$, and assume that $\ell(\mathbf{h}_\theta(x), \mathbf{e})$ is convex with respect to θ for any $x \in \mathcal{X}$ and $\mathbf{e} \in \mathcal{Y}$. Then

$$\emptyset \neq \partial \mathbf{R}(\mathbf{h}_\theta; X, \mathbf{Y}') \subseteq \partial f(\theta) .$$

A subgradient-based optimization approach [Shor, 2012] is similar to gradient descent, however at each iteration we use a subgradient instead of the gradient, and we memorize the best solution found among all the iterations.

The subgradient-based optimization algorithm used to solve the optimization problem (16) is presented in algorithm 1.

As observed before, \mathbf{Y}' as defined in the algorithm can be computed by solving a linear program. The projection step depends on the set of parameters Θ . While this is not a requirement for our approach, if the loss function $\ell(\mathbf{h}_\theta(x), \mathbf{y})$, is differentiable with respect to θ , then we can compute the gradient of the empirical risk instead of a subgradient.

Theorem 2.4.6 (Optimality Guarantees of the Subgradient Method). Suppose that for any $(x, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$, $\ell(\mathbf{h}_\theta(x), \mathbf{y})$ is L -Lipschitz continuous and convex with respect to θ . Let *step size* $h > 0$, and *iteration count* $T \in \mathbb{N}$, and $\tilde{\theta}$ as returned by algorithm 1. Then, we have that

$$f(\tilde{\theta}) - f(\hat{\theta}) \leq \frac{\varphi^2(\Theta) + L^2 h^2 T}{2hT} ,$$

where $\varphi(\cdot)$ is computed with respect to the ℓ_2 -norm, i.e., $\varphi^2(\Theta) \doteq \max_{\theta_1, \theta_2 \in \Theta} \|\theta_1 - \theta_2\|_2^2$, and $\hat{\theta}$ is defined as in (16).

Corollary 2.4.7 (Convergence Rates of the Adversarial Subgradient Algorithm). Suppose that the same assumptions of theorem 2.4.6 hold. For any $\varepsilon > 0$, then if $h = \varepsilon/L^2$ and $T \geq \frac{L^2 \varphi^2(\Theta)}{\varepsilon^2}$, we have that

$$f(\tilde{\theta}) - f(\hat{\theta}) \leq \varepsilon .$$

Therefore, we can compute a solution within ε -additive error of (16) in $\mathbf{O}(\frac{L^2 \varphi^2(\Theta)}{\varepsilon^2})$ steps of the subgradient algorithm.

Algorithm 1 Adversarial Subgradient Algorithm

Input: Number of iterations T , step size h , \mathcal{H} , X , $\varphi_1, \dots, \varphi_n$

Output: Approximate solution $\tilde{\theta}$ of (16) (See theorem 2.4.6)

$\tilde{\theta}^{(0)} = \theta^{(0)} \leftarrow$ arbitrary point $\theta \in \Theta$

for $t \in 1, \dots, T$ **do**

$Y' \leftarrow \operatorname{argmax}_{Y \in \mathbb{Y}} \hat{R}(h_{\theta^{(t-1)}}, Y)$

$v \leftarrow$ arbitrary vector from $\partial \hat{R}(h_{\theta}, Y')$

$\theta^{(t)} \leftarrow \operatorname{Proj}_{\Theta}(\theta^{(t-1)} - hv)$ ($\operatorname{Proj}_{\Theta}(\cdot)$ denotes projection onto Θ)

$\tilde{\theta}^{(t)} \leftarrow \operatorname{argmin}\{f(\tilde{\theta}^{(t-1)}), f(\theta^{(t)})\}$

end for

Return $\tilde{\theta}^{(T)}$

2.4.1 Risk Guarantees and Statistical Learning

In this section, we develop a bound on the true risk of the classifier $h_{\hat{\theta}}$ that is a solution of the optimization problem (16). The bounds are expressed in function of the Rademacher complexity of the function family $\mathcal{L} \doteq \{\ell^\circ \circ h : h \in \mathcal{H}\}$ that describes the loss of each function $h \in \mathcal{H}$, the risk minimizer $\theta^* = \operatorname{argmin}_{\theta \in \Theta} R(h_\theta)$, and the normalized ℓ_1 diameter $D_{\mathbb{Y}^\circ}$ of the feasible set of solutions \mathbb{Y}° .

$$D_{\mathbb{Y}^\circ} \doteq \sup_{Y', Y'' \in \mathbb{Y}^\circ} \frac{1}{m} \sum_{j=1}^m \|\mathbf{y}'_j - \mathbf{y}''_j\|_1. \quad (18)$$

The quantity $D_{\mathbb{Y}^\circ}$ characterizes the information given by the classifiers $\varphi_1, \dots, \varphi_n$ on the classification task. In particular, a weak supervision source provides useful information on the classification task of interest only if it reduces the size of the feasible set, and it provably improves the performance of our algorithm if it decreases the average diameter $D_{\mathbb{Y}^\circ}$.

Given a function family \mathcal{L} , we define the empirical Rademacher average [see Mitzenmacher and Upfal, 2017] of the unlabeled items X and a possible labeling Y of those items as

$$\hat{\mathfrak{R}}(\mathcal{L}; X, Y) \doteq \mathbb{E}_{\sigma} \left[\sup_{\ell^\circ \circ h \in \mathcal{L}} \frac{1}{m} \sum_{i=1}^m \sigma_i \ell^\circ(h(x_i), \mathbf{y}_i) \right],$$

where $\sigma_1, \dots, \sigma_m$ are independent random variables from the Rademacher distribution, i.e., $\mathbb{P}(\sigma_i = 1) = \mathbb{P}(\sigma_i = -1) = \frac{1}{2}$. Intuitively, this quantity measures the *capacity* of \mathcal{H} to *overfit*, and under mild conditions, approaches 0 as sample size m tends to infinity, in which case overfitting becomes impossible.

Theorem 2.4.8 (Optimality of Adversarial Solutions). Let $h_{\hat{\theta}}$ be the solution of (16). Let $\theta^* \in \operatorname{argmin}_{\theta \in \Theta} R(h_\theta)$. Suppose that the codomain of the loss function ℓ is contained in the interval $[0, B]$. Let Y^* be the true (unknown) labeling of the unlabeled data X , and assume that $Y^* \in \mathbb{Y}^\circ$. Then, with probability $1 - \delta$ it holds that

$$R(h_{\hat{\theta}}) \leq R(h_{\theta^*}) + BD_{\mathbb{Y}^\circ} + \sup_{Y \in \mathbb{Y}^\circ} 4\hat{\mathfrak{R}}(\mathcal{L}; X, Y) + \mathbf{O} \left(B \sqrt{\frac{\ln \frac{1}{\delta}}{m}} \right). \quad (19)$$

The bound of theorem 2.4.8 requires the computation of the diameter of the set \mathbb{Y}^\diamond with respect to the ℓ_1 -norm. This set is a polytope $\mathbb{Y}^\diamond \subseteq \mathbb{R}^{km}$, though due to the *equality constraints* for each unit k -simplex on each of the m unlabeled points, it may be embedded and expressed as a standard H -polytope in $\mathbb{R}^{(k-1)m}$, as it is specified by *linear inequality constraints* on the possible labelings. While this bound provides insight on how the information provided by the weak supervision sources affect the quality of our solution, computing the diameter is hard. Brieden et al. [2001] show that computing the ℓ_1 diameter is NP-hard, as is computing even a multiplicative $\sqrt{(k-1)m}$ approximation. Furthermore, particularly in the case of $k=2$, it seems straightforward to construct problem instances with a great degree of control over the constraints, so there is no reason to think our problem is in general easier than the general case.

2.4.2 The Centerpoint Strategy: Bounds without Adversarial Learning

Although we have shown a polynomial-time algorithm for (convex) minimax optimization, it is worth noting that solving a linear program at each subgradient step adds significant overhead to training time, and this strategy also precludes using our method with out-of-the-box optimizers and machine learning algorithms. We also note that the minimax gap is necessarily small when the feasible set \mathbb{Y}^\diamond is small, as captured by the *diameter term* of theorem 2.4.8. Taking this intuition into account, we know that if the feasible set is small, then there is little difference between a *minimax-optimal* model, and a model trained to optimize *any point* in the *feasible set*.

In this subsection, we adopt a methodology based on this idea, where a point in the feasible set, termed a *centerpoint*, is selected, a model is optimised for that (fixed) labelling, using any black-box optimizer, and generalization bounds are obtained for said model. We derive these bounds as corollaries of a minimax bound, so they are not as statistically efficient, however they may be computationally more efficient, and are clearly more modular, and thus this method is easier to incorporate into existing systems. Finally, while any point in the feasible set may be used as a centerpoint, we know that particular choices may yield desirable properties and intuitive interpretations, so we discuss such choices. The next theorem quantifies the performance of the centerpoint method.

Theorem 2.4.9 (Generalization Bounds for Feasibility-Constrained Learning). Fix $\theta \in \Theta$, and let $\mathbf{Y} \in \mathbb{Y}^\diamond$ be a feasible labeling. Suppose that the codomain of the loss function ℓ is contained in the interval $[0, B]$. Let \mathbf{Y}^* be the true (unknown) labeling of the unlabeled data X , and assume that $\mathbf{Y}^* \in \mathbb{Y}^\diamond$. Then, with probability $1 - \delta$, it holds that

$$R(\mathbf{h}_\theta) \leq \sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \hat{R}(\mathbf{h}_\theta; (X, \mathbf{Y}')) + 2 \sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \hat{\mathbf{r}}_m(\mathcal{L}; (X, \mathbf{Y}')) + 3B \sqrt{\frac{\ln \frac{1}{\delta}}{2m}} \quad (20)$$

$$\leq \hat{R}(\mathbf{h}_\theta; (X, \mathbf{Y})) + \frac{B}{2} \sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \|\mathbf{Y} - \mathbf{Y}'\|_{1,1} + 2 \sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \hat{\mathbf{r}}_m(\mathcal{L}; (X, \mathbf{Y}')) + 3B \sqrt{\frac{\ln \frac{1}{\delta}}{2m}} \quad (21)$$

$$\leq \hat{R}(\mathbf{h}_\theta; (X, \mathbf{Y})) + BD_{\mathbb{Y}^\diamond} + 2 \sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \hat{\mathbf{r}}_m(\mathcal{L}; (X, \mathbf{Y}')) + 3B \sqrt{\frac{\ln \frac{1}{\delta}}{2m}} . \quad (22)$$

Theorem 2.4.9 provides generalization bounds with respect to a fixed model $\theta \in \Theta$. The inequality (20) is a standard minimax bound, and it can be observed that the right-hand side is minimized by choosing θ equal to the solution of (16). Naturally, the question arises, how should θ be selected to optimize the other forms? In particular, we seek reasonable choices of \mathbf{Y} such that θ trained to optimize \mathbf{Y} perform well. For this purpose, we will ignore the $\hat{R}(\mathbf{h}_\theta; (X, \mathbf{Y}))$ term, as this can't be computed without knowing \mathbf{Y} , and also the $2 \sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \hat{\mathbf{r}}_m(\mathcal{L}; (X, \mathbf{Y}'))$ term, as this term can generally be bounded independently of \mathbf{Y} .

The next bound (21) shows that the generalization bound is minimized by picking the point \mathbf{Y} that minimizes the *radius term* $\sup_{\mathbf{Y}' \in \mathbb{Y}^\diamond} \|\mathbf{Y} - \mathbf{Y}'\|_{1,1}$. Such a minimizer is known as an ℓ_1 *outer Chebyshev center*, and for

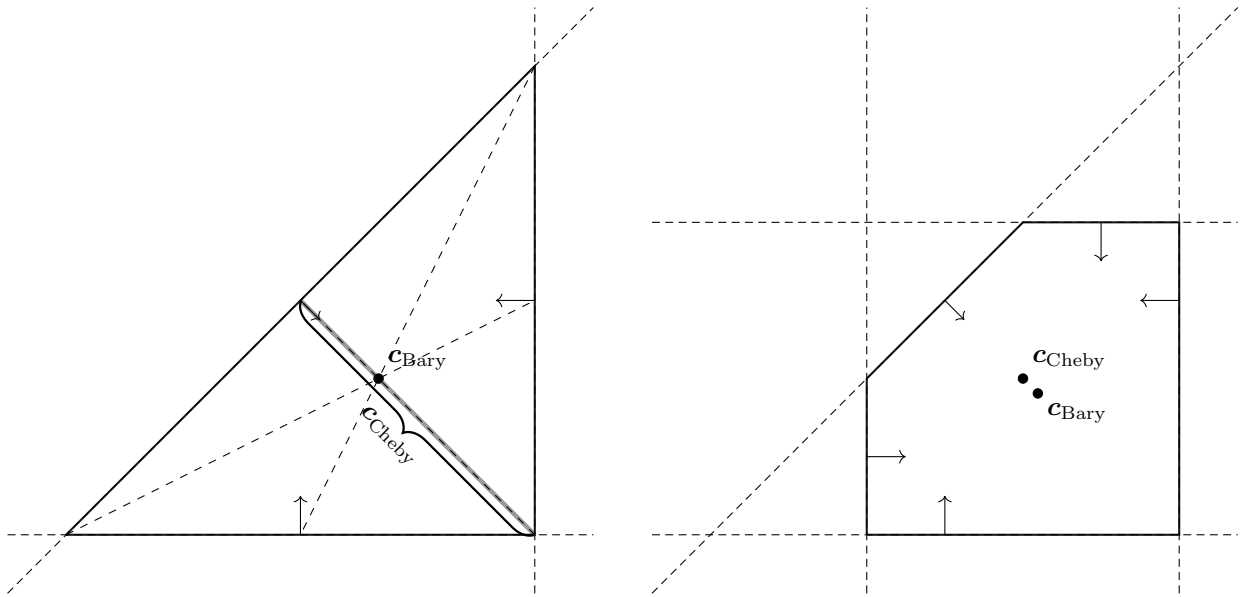


Figure 2.1: Two polytopes in \mathbb{R}^2 , with their ℓ_1 outer Chebyshev centers and barycenters. The barycenter is always uniquely defined, but in the right isosceles triangle, the Chebyshev center is not unique, but rather consists of the entire perpendicular bisector of the hypotenuse.

this reason is a natural choice of centerpoint. Unfortunately, Brieden et al. [2001] show that computing the ℓ_1 outer Chebyshev center of a polytope is as hard as computing the diameter ℓ_1 , making this centerpoint choice difficult to work with. In a related problem, Eldar et al. [2008] obtain promising results with a relaxation of the ℓ_2 outer Chebyshev center, and we are hopeful that similar techniques may lead to an effective and easily computed centerpoint choice in this setting.

Another intuitive choice is the *arithmetic mean* (a.k.a. the *centroid* or *barycenter*) of the feasible region \mathbb{Y}^\diamond . This is rather elegant, as the hypothesis h_θ that minimizes the empirical risk *on average* over \mathbb{Y}^\diamond also minimizes empirical risk of the *average labeling* in \mathbb{Y}^\diamond , i.e., we have

$$\mathbf{Y} \doteq \min_{\theta \in \Theta} \frac{1}{Z} \int_{\mathbf{y} \in \mathbb{Y}^\diamond} \hat{\mathbf{R}}(h_\theta; X, Y) dY = \min_{\theta \in \Theta} \hat{\mathbf{R}} \left(h_\theta; X, \frac{1}{Z} \int_{\mathbb{Y}^\diamond} Y dY \right),$$

for normalizing constant Z . This gives us a nice interpretation as an *average-case optimal solution* over the uncertainty over feasible labelings. Similarly, we can also consider only the *integer solutions* (i.e., integrate and normalize w.r.t. \mathbb{Y} instead of \mathbb{Y}^\diamond). This is an instance of the *vertex centroid problem*, which Elbassioni and Tiwary [2008] show to be $\#P$ hard, however they also show it can be approximated efficiently via sampling methods. The ℓ_1 outer Chebyshev center and barycenter are contrasted graphically on example polytopes in figure 2.1.

2.4.3 Applications

In order to feature the generality of our framework, we show two examples of different instantiations of the optimization problem (16) for different choices of loss function and prediction models for which we can apply corollary 2.4.7.

Convex combination of the weak supervision sources. Let $\Theta = \{\theta = (\theta_1, \dots, \theta_n) \in \mathbb{R}_+^n : \sum_{i=1}^n \theta_i = 1\}$. Our prediction model is a convex combination of the output of the weak classifiers $\varphi, \dots, \varphi_n$. In particular,

given $\theta \in \Theta$, the classifier \mathbf{h}_θ is defined as $\mathbf{h}_\theta(x) = \sum_{i=1}^n \theta_i \varphi_i(x)$ for any $x \in \mathcal{X}$. It is easy to see that $\varnothing(\Theta) \leq \sqrt{2}$. Given an arbitrary vector $\mathbf{v} \in \mathbb{R}^n$, the projection step to Θ can be done efficiently by using for example the algorithm in Wang and Carreira-Perpinán [2013].

Let ℓ be the Brier loss, defined as for any $(x, \mathbf{e}) \in \mathcal{X} \times \mathcal{Y}$ is defined as:

$$\begin{aligned} \ell(\mathbf{h}_\theta(x), \mathbf{e}) &\doteq \sum_{c=1}^k (\mathbf{h}_\theta(x)_c - e_c)^2 \\ &= \|\mathbf{h}_\theta(x)\|_2^2 - 2\mathbf{h}_\theta(x)^T \cdot \mathbf{e} + 1 . \end{aligned}$$

It is easy to see that the function $\ell(\mathbf{h}_\theta(x), \mathbf{e})$ is convex and differentiable with respect to θ . Moreover, the loss function has codomain $[0, 2]$.

Lemma 2.4.10. The loss $\ell(\mathbf{h}_\theta(x), \mathbf{e})$ of a prediction model \mathbf{h}_θ defined as in this subsection is $2\sqrt{n}$ -Lipschitz continuous with respect to θ .

softMax (multinomial logistic regression). Suppose that each item is a vector in \mathbb{R}^b , i.e., $\mathcal{X} \subseteq \mathbb{R}^b$, and assume that $\|\mathbf{x}\|_2 \leq B_x$ for any $x \in \mathcal{X}$. Let $\Theta \doteq \{\theta = (\mathbf{w}_1^T \dots \mathbf{w}_k^T) \in \mathbb{R}^{k \times b} : \mathbf{w}_c \in \mathbb{R}^b \wedge \|\mathbf{w}_c\|_2 \leq B_w \text{ for } c \in 1, \dots, k\}$. That is, θ is the concatenation of k vectors with bounded norm. If we are given θ Observe that with this definition of Θ , we have that $\varnothing(\Theta) \leq \sqrt{2k}B_w$. Given a vector $\theta = (\mathbf{w}_1^T \dots \mathbf{w}_k^T)$, the projection step to Θ is simply $\tilde{\theta} = (\tilde{\mathbf{w}}_1^T \dots \tilde{\mathbf{w}}_k^T)$, where $\tilde{\mathbf{w}}_c = \mathbf{w}_c / \min(B_w / \|\mathbf{w}_c\|_2, 1)$ for $c \in 1, \dots, k$.

Given $\theta = (\mathbf{w}_1^T \dots \mathbf{w}_k^T) \in \Theta$, and $\mathbf{x} \in \mathcal{X}$, we define

$$\mathbf{h}_\theta(\mathbf{x}) \doteq \left(\frac{\exp(\mathbf{w}_1^T \cdot \mathbf{x})}{\sum_{c=1}^k \exp(\mathbf{w}_c^T \cdot \mathbf{x})}, \dots, \frac{\exp(\mathbf{w}_k^T \cdot \mathbf{x})}{\sum_{c=1}^k \exp(\mathbf{w}_c^T \cdot \mathbf{x})} \right)^T .$$

This classifier is a particular instantiation of softMax combined with a linear model. For a vector $\mathbf{v} = (v_1, \dots, v_d)^T$, define $\ln \mathbf{v} \doteq (\ln v_1, \dots, \ln v_d)^T$. Given $(x, \mathbf{e}) \in \mathcal{X} \times \mathcal{Y}$, we define the cross-entropy loss ℓ of the prediction model \mathbf{h}_θ as

$$\ell(\mathbf{h}_\theta(\mathbf{x}), \mathbf{e}) \doteq -\mathbf{e}^T \cdot \ln(\mathbf{h}_\theta(\mathbf{x})) .$$

This combination of prediction model and loss function is also known as multinomial logistic regression. It is easy to see that the loss function is differentiable with respect to θ , and it is a known result that $\ell(\mathbf{h}_\theta(\mathbf{x}), \mathbf{e})$ is convex with respect to θ for any $(\mathbf{x}, \mathbf{e}) \in \mathcal{X} \times \mathcal{Y}$ [Böhning, 1992]. Moreover, the loss function is bounded, and Lipschitz continuous, as shown by the next two lemmas.

Lemma 2.4.11. For any $(x, \mathbf{e}) \in \mathcal{X} \times \mathcal{Y}$, and $\theta \in \Theta$, we have that $\ell(\mathbf{h}_\theta(\mathbf{x}), \mathbf{e}) \in [0, B_w B_x + \ln k]$.

Lemma 2.4.12. For any (x, \mathbf{e}) , the loss $\ell(\mathbf{h}_\theta(x), \mathbf{e})$ of a prediction model \mathbf{h}_θ defined as in this subsection is (kB_x) -Lipschitz continuous with respect to θ .

2.4.4 Constraining the feasible set

The presentation of section 2.4 implicitly assumes an alignment between the output classes of the weak supervision sources $\varphi_1, \dots, \varphi_n$ and the target classification task. In fact, as seen in lemma 2.4.1, we compute the intervals Δ_i based on the empirical risk of the weak supervision sources using labeled data of the target classification task. However, for many application of interests, the weak supervision sources could output to a different codomain, potentially with a unequal number of classes. As an example, suppose that we would like to distinguish between images of {cat, dog, rabbit, bear}. A binary classifier that tells us if the animal represented in an image has a tail or not still provides a useful clue with respect to the target classification task, and we would like to use that information.

In this section, we will show how to constrain the feasible set of labelings \mathbb{Y}^\diamond in a more general setting, where the weak supervision source φ_i is a classifier that maps elements from the domain \mathcal{X} to soft labels over k_i classes, i.e., $\varphi_i : \mathcal{X} \rightarrow \mathcal{Y}_{k_i}^\diamond$, where $\mathcal{Y}_{k_i}^\diamond = \{\mathbf{v} \in \mathbb{R}_{\geq 0}^{k_i} : \sum_c v_c = 1\}$.

Consider the weak supervision source φ_i . For each $c \in 1, \dots, k$ and $\tilde{c} \in 1, \dots, k_i$, we use the m_L labeled data $(\tilde{x}_1, \tilde{\mathbf{y}}_1), \dots, (\tilde{x}_{m_L}, \tilde{\mathbf{y}}_{m_L})$ to compute the statistic

$$\hat{\mu}_{i,c,\tilde{c}}(\tilde{X}, \tilde{Y}) \doteq \frac{1}{|\tilde{X}|} \sum_{j=1}^{|\tilde{X}|} y_{j,c} [\varphi_i(x_j)]_{\tilde{c}} .$$

It is clear that the function $\hat{\mu}_{i,c,\tilde{c}}(\tilde{X}, \tilde{Y})$ is linear in \tilde{Y} . For each weak supervision source φ_i , true class $c \in 1, \dots, k$, and weak supervision source's output class $\tilde{c} \in 1, \dots, k_i$, based on the value $\hat{\mu}_{i,c,\tilde{c}}(\tilde{X}, \tilde{Y})$, we compute an interval $\Delta_{i,c,\tilde{c}}$, defined as

$$\Delta_{i,c,\tilde{c}} \doteq [\hat{\mu}_{i,c,\tilde{c}}(\tilde{X}, \tilde{Y}) - \gamma, \hat{\mu}_{i,c,\tilde{c}}(\tilde{X}, \tilde{Y}) + \gamma] .$$

where the value γ is specified in lemma 2.4.13.

Given a labeling \mathbf{Y} of the unlabeled dataset X , we say that \mathbf{Y} is a feasible solution if for each i, c and \tilde{c} , it holds that

$$\hat{\mu}_{i,c,\tilde{c}}(X, \mathbf{Y}) \in \Delta_{i,c,\tilde{c}} . \quad (23)$$

That is, the set of all the feasible solutions \mathbb{Y}^\diamond is defined as

$$\mathbb{Y}^\diamond \doteq \left\{ \mathbf{Y} \in \mathbb{R}^{k \times m} \mid \begin{array}{ll} \mathbf{y}_j & \in \mathcal{Y}^\diamond \quad \forall j \in 1, \dots, m \\ \hat{\mu}_{i,c,\tilde{c}} & \in \Delta_{i,c,\tilde{c}} \quad \forall i \in 1, \dots, n, \forall c, \tilde{c} \end{array} \right\} .$$

Notice that the constraints specified in \mathbb{Y}^\diamond are still linear in \mathbf{Y} , therefore we can still compute the value $f(\boldsymbol{\theta})$ (as in (17)) by solving a linear program, and all discussion of empirical-risk based constraints still applies.

In order to be able to give the theoretical bound of theorem 2.4.8, we need to guarantee that the true labeling \mathbf{Y}^* of the unlabeled data X is feasible. This is possible by choosing a suitable value γ when defining the intervals $\Delta_{i,c,\tilde{c}}$, in the following result, which closely resembles lemma 2.4.1.

Lemma 2.4.13 (Confusion Matrix Interval Risk Bounds). Suppose failure probability $\delta \in (0, 1)$. For every $i \in 1, \dots, n$, $c \in 1, \dots, k$, and $\tilde{c} \in 1, \dots, k_i$ let $\Delta_{i,c,\tilde{c}}$ be computed as in (23), and let $K \doteq k \sum_{i=1}^n k_i$. Now, if we use the value

$$\gamma \doteq \sqrt{\frac{(m_L + m) \ln \frac{2nK}{\delta}}{2m_L m}}$$

as the radius of these intervals, then with probability at least $1 - \delta$ it holds that $\mathbf{Y}^* \in \mathbb{Y}^\diamond$.

2.5 Experiments

We demonstrate the applicability and performance of our method on image multi class classification tasks derived from the DomainNet [Peng et al., 2019] dataset. We also provide experiments on image binary classification tasks derived from the Animals with Attributes 2 [Xian et al., 2018] dataset in order to compare our methods with additional baselines.

DomainNet contains images from 6 different domains $P = \{\text{clipart, infograph, painting, quickdraw, real, sketch}\}$ and features 345 different classes. Animals with Attributes 2 contains natural images of 50 types of animals. Associated with the dataset is a list of 85 attributes for each animal class, which we use to create weak supervision sources. Animals with Attributes 2 is divided into 40 ‘‘seen’’ classes and 10 ‘‘unseen’’ classes,

where the seen classes can be used to train attribute classifiers without leaking information about the unseen classes.

We refer to our algorithms by using the acronyms **AMCL-CC** and **AMCL-LR**, where AMCL stands for **A**dversial **M**ulti **C**lass **L**earning. AMCL-CC is an implementation of our method that uses a convex combination of the weak supervision sources as the prediction model, whereas AMCL-LR uses multinomial logistic regression (see Section 4.1). For every image, we compute the output of a pretrained ResNet-18 and use it as inputs for AMCL-LR.

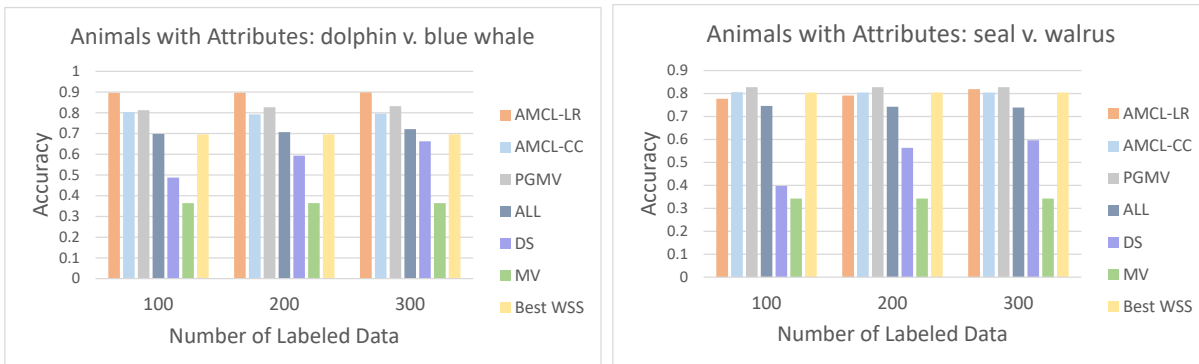


Figure 2.2: Experimental results on Animals with Attributes for the binary classification tasks of dolphin v. blue whale (left) and seal v. walrus (right) as we vary the amount of labeled data. Each method uses 560 unlabeled data for dolphin v. blue whale and 602 unlabeled data for seal v. walrus.

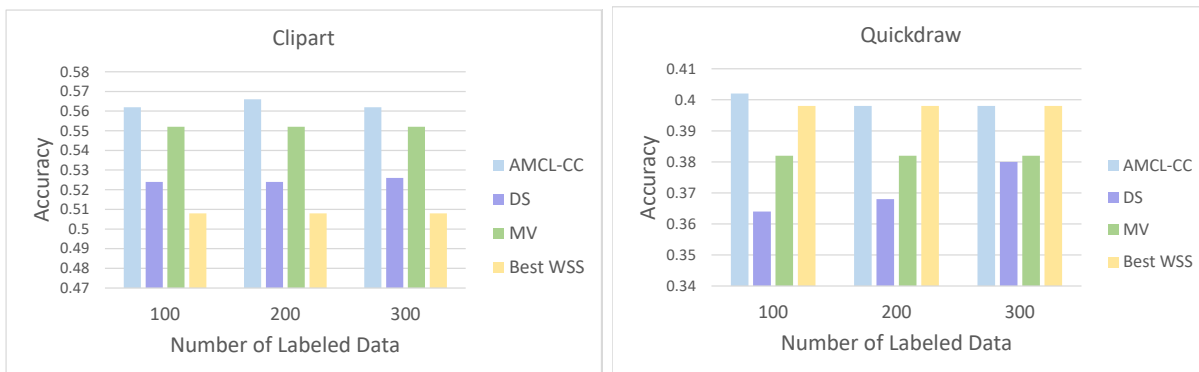


Figure 2.3: Experimental results on Domain Net for the clipart and quickdraw domains as we vary the amount of labeled data. Each method uses 500 unlabeled data. Results are listed for the 5 classes of {sea turtle, vase, whale, bird, violin }.

2.5.1 Setup

From DomainNet, we select $k = 5$ random classes from the 25 classes with the largest number of data. Then, for each domain $p \in P$, we learn a multi class classifier φ_p for those k classes in domain p . The classifier φ_p is trained by fine-tuning a pretrained ResNet-18 network [He et al., 2016] using 60% of the labeled data for that domain. For each domain p , we consider the classifiers trained in domains other than p as weak supervision sources, i.e., the classifiers $\{\varphi_q\}$ for $q \in P \setminus \{p\}$. We remark that these weak supervision sources never have access to samples from domain p .

From Animals With Attributes, we create binary classification tasks by selecting pairs of unseen classes.

Following Mazzetto et al. [2021], we create weak supervision sources by using the seen classes to train classifiers for the attributes that distinguish them. Similarly to Domain Net, these classifiers are learned by fine-tuning a pretrained ResNet-18 network using labeled data from the seen classes. In order to focus on the most challenging tasks, where the weak supervision sources are not highly accurate, we select the 4 class pairs among the unseen classes with the lowest majority vote accuracy.

We remark that all algorithms that require unlabeled data are evaluated in a transductive setting: the unlabeled data used by the algorithms is also used to evaluate the final learnt prediction models.

2.5.2 Baselines and algorithms

Following the example of [Mazzetto et al., 2021], we compare our method with the following baselines and algorithms:

Best Weak Supervision Source (Best WSS): We report the accuracy of the best weak supervision source.

Majority Vote (MV): We consider a simple approach to combining the weak supervision sources: we average their output and select the most voted class. This approach requires no learning, but is suboptimal when weak supervision sources’ errors are not independent and/or have unequal levels of accuracy.

Semi-Supervised Dawid-Skene Estimator (DS): We also consider a semi-supervised extension to the standard crowdsourcing algorithm [Dawid and Skene, 1979] that finds the optimal aggregation of the outputs of independent weak supervision sources. The Dawid-Skene estimator is also the default aggregation method for the Snorkel system [Ratner et al., 2017]. Here, we use a semi-supervised version of this algorithm, for a fair comparison with our work. We simply optimize the marginal likelihood of the weak supervision sources’ outputs for unlabeled data and the joint likelihood with the label when it is observed.

Adversarial Label Learning (ALL): This algorithm [Arachie and Huang, 2019] learns a prediction model that has the highest expected accuracy with respect to an adversarial labeling of an unlabeled dataset, where this labeling must satisfy error constraints on the weak supervision sources. This approach shares similarities with our method; however, it fails to provide theoretical guarantees on the learning of the prediction model. For a fair comparison to our method, we use logistic regression as the prediction model, and use the same features given to AMCL-LR.

Performance-guaranteed Majority Vote (PGMV): This method finds a subset of weak supervision sources whose majority vote achieves high accuracy with respect to the worst-case distribution of the output the weak supervision sources. Again, this worst-case distribution is constrained by using statistics computed on the weak supervision sources (individual error rates and pairwise differences).

Due to the limitations of PGMV and ALL, we can run those algorithms only for binary classification tasks.

2.5.3 Results

Animals With Attributes (binary classification): In figure 2.2, we report the results on the Animals With Attributes dataset for two binary classification tasks.

In the binary setting, our methods match or outperform the state-of-the-art methods PGMV and ALL over all amounts of labeled data. We note that even though AMCL-LR and ALL use the same inputs and train the same prediction model, our method achieves overall higher accuracies in addition to providing theoretical guarantees on the learning of the prediction model.

Domain Net (multi-class classification): In figure 2.3, we report the accuracies of the different algorithms on the Domain Net dataset for the clipart and quickdraw domains. We remark that ALL and PGMV cannot be used in this setting, as they are restricted to binary classification.

In the multi class setting, our methods again match or outperform the baselines over all amounts of labeled data. We note that in the quickdraw domain, the weak supervision sources are overall very inaccurate, and it is difficult to recover useful information from them. However, differently from the baselines DS and MV, AMCL-CC can still recover and improve upon the best weak supervision source.

Again, as noted by the Best WSS column, the weak supervision sources are quite inaccurate in this dataset. Therefore, we do not report the results for the AMCL-LR algorithm as the weak supervision sources do not constrain the feasible set of solutions enough for our method to learn a more complex model such as multinomial logistic regression.

Due to space constraints, additional plots and experimental details for both datasets are reported in the Appendix.

2.6 Conclusion

We develop the first rigorous method that can use information provided by arbitrarily correlated weak supervision sources in order to learn a prediction model for a multi class classification task. In many practical settings, our method can optimally learn the model that achieves the smallest risk with respect to an adversarial feasible labeling of a unlabeled dataset, and we provide theoretical guarantees on the quality of the learnt model based on a measure of the information provided by the weak supervision sources. Finally, we provide experiments that illustrate the applicability of our approach and its advantages over existing methods.

Part II

Fairness with Population Means: Malfare, Welfare, and Fair PAC Learning

The rise of machine learning in industry and government has brought about a commensurate rise in algorithmic bias and fairness issues, with increasing recognition by academics that such issues must be explicitly considered in practical fairness-sensitive machine learning settings. It is now well-known that historic data reflect historic biases and injustices, and that, due to availability bias, machine learning models are frequently trained on data from prominent, privileged, or majority groups while minorities remain understudied and underdocumented, and these issues have well-studied effects on the efficacy and (un)fairness of many contemporary machine learning systems. While part I addresses this issue only implicitly, by tailoring machine learning methods and generalization guarantees to the *small sample* setting (and thus democratizing such methods to operate not just on well-studied large groups, but also on smaller populations), here I take a more direct approach to confronting algorithmic bias and fairness issues in machine learning.

For context, traditional constraint-based notions of algorithmic fairness have risen to prominence in machine learning, with the potential to correct for some forms of data and/or algorithmic bias by ensuring *demographic-parity* concepts, e.g., *equality of opportunity*, *equality of outcome*, and other such desiderata. Fairness by demographic-parity constraint has several prominent flaws: most notably, many parity constraints are *mutually unsatisfiable*, and by nature, there is inherent tension between *minimizing loss* and *constraining for fairness*, where additional *tolerance parameters* (which quantify the degree to which fairness constraints impact a solution) must be specified to strike a balance between *accuracy* and *fairness*. Perhaps in response to these issues, some recent work has trended toward welfare-based fairness-concepts, wherein both *accuracy* and *fairness* are encoded in a *welfare function* defined on a *group of subpopulations*. Welfare metrics quantify *overall wellbeing* across such populations, and welfare-based objectives and constraints incentivize *fair machine learning methods* to produce satisfactory solutions that consider the diverse needs of multiple groups. Unfortunately, welfare based approaches require a notion of (positive) per-group utility, and I argue that this is not natural to many machine learning tasks, which instead seek to *minimize* some *negatively connoted* loss value.

Introducing Malfare In this part, I derive an equivalently-axiomatically justified alternative to welfare, termed malfare, and propose its use as a cardinal objective in *fair machine learning tasks*. In particular, malfare measures overall *societal harm* (rather than wellbeing), and *malfare minimization* naturally generalizes *risk minimization* to (nonlinearly) consider model performance across multiple protected groups, and is thus ideal for tasks like ensuring *facial recognition* or *speech recognition* systems are *accurate* and *accessible* across a diverse population. I then cast fair machine learning as a direct *malfare minimization* problem, where a group’s dissatisfaction is simply their risk (expected loss). Surprisingly, except in two trivial cases, the axioms of cardinal welfare (malfare) dictate that this is not equivalent to simply defining utility as negative loss. The framework is established in chapter 3, and it is instantiated in a practical learning setting in chapter 4.

Building upon these concepts, in chapter 5 I define *fair-PAC learning*, where a fair PAC-learner is an algorithm that learns an ϵ - δ malfare-optimal model with bounded sample complexity, for *any data distribution*, and for *any* (axiomatically justified) malfare concept. We show broad conditions under which, with appropriate modifications, many standard PAC-learners may be converted to fair-PAC learners. This places fair-PAC learning on firm theoretical ground, as it yields *statistical efficiency* guarantees for many well-studied machine-learning models, and is also practically relevant, as it democratizes fair machine learning by providing concrete training algorithms and rigorous generalization guarantees for these models. In particular, I show via a constructive polynomial reduction that *realizable fair PAC-learning* reduces to *realizable PAC-learning*, and furthermore, I show, non-constructively, that for learning problems where PAC-learnability implies uniform convergence, it is equivalent to fair-PAC-learnability. In addition to *statistical guarantees*, I also show *computational guarantees*: in particular, I show that when training is possible via *convex optimization* or efficient-enumeration of an approximate cover of the hypothesis space, then ϵ - δ training malfare-optimal models, as with risk-optimal models, requires polynomial time.

A Preview of Technical Results As for statistical guarantees in malfare estimation, corollary 3.4.2 shows that for any fair \mathbf{w} -weighted malfare function $\Lambda(\cdot; \mathbf{w})$ (precisely defined in chapter 3), empirical malfare is tightly concentrated about true malfare. In particular, assume g groups, any *loss function* $\ell : \mathcal{X} \rightarrow [0, r]$, risk values $\mathcal{S} \in [0, r]^g$ s.t. $\mathcal{S}_i = \mathbb{E}_{\mathcal{D}_i}[\ell]$, samples $\mathbf{x}_i \sim \mathcal{D}_i^m$ and empirical risk values $\hat{\mathcal{S}}_i = \frac{1}{m} \sum_{j=1}^m \ell(\mathbf{x}_{i,j})$. Then with probability at least $1 - \delta$ over choice of \mathbf{x} , we have

$$\left| M_p(\mathcal{S}; \mathbf{w}) - M_p(\hat{\mathcal{S}}; \mathbf{w}) \right| \leq r \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}} . \quad (24)$$

Malfare estimation is more challenging than risk estimation, as the empirical malfare is in general a *biased estimator* of true malfare. Still, this result is clearly analogous to Hoeffding's inequality for the risk, suffering only a $\sqrt{\log g}$ dependence on the number of groups g . Alternatively, again with probability at least $1 - \delta$ over choice of \mathbf{x} , we have a (variance-dependent) Bennett's inequality analogue of

$$\left| M_p(\mathcal{S}; \mathbf{w}) - M_p(\hat{\mathcal{S}}; \mathbf{w}) \right| \leq \frac{r \ln \frac{2g}{\delta}}{3m} + \max_{i \in \{1, \dots, g\}} \sqrt{\frac{2 \mathbb{V}_{\mathcal{D}_i}[\ell] \ln \frac{2g}{\delta}}{m}} . \quad (25)$$

These results should be contrasted with inequalities (1) and (2) for mean-estimation, as they generalize said results to malfare estimation. It is a straightforward matter to extend them to depend on *empirical variances*, and I further generalize to give *uniform convergence guarantees over function families* in chapter 4.

Chapter 3

Quantifying Population Sentiment with Welfare and Malfare

In this chapter, I address an inherent difficulty in welfare-theoretic fair machine learning by proposing an equivalently-axiomatically justified alternative and studying the resulting computational and statistical learning questions. Welfare metrics quantify *overall wellbeing* across a population consisting of one or more groups, and welfare-based objectives and constraints have recently been proposed to incentivize *fair machine learning methods* to produce satisfactory solutions that consider the diverse (often mutually incompatible) needs of multiple groups. Unfortunately, many machine-learning problems are more naturally cast as *loss minimization*, rather than *utility maximization* tasks, which complicates direct application of welfare-centric methods to fair machine learning tasks. I thus define a complementary measure to *welfare*, termed *malfare*, measuring overall *societal harm* (rather than wellbeing), with axiomatic justification via the standard axioms of cardinal welfare.

With this definition in hand, I then cast fair machine learning as a direct *malfare minimization* problem, where a group’s malfare is their risk (expected loss). Surprisingly, except in two trivial cases, the axioms of cardinal welfare (malfare) dictate that this is not equivalent to simply defining utility as negative loss. Previous welfare based methods have generally defined welfare in an *ad-hoc* manner, and I show that, in particular, seemingly intuitive definitions of welfare in terms of loss lead to absurdities or non-trivial and uninterpretable hyperparameters. Furthermore, I show that malfare exhibits rapidly-converging finite-sample concentration-of-measure guarantees, and thus may be estimated with great accuracy from a small sample. This chapter is adapted from the first half of Cousins [2021].

3.1 Introduction

It is now well-understood that contemporary ML systems for tasks like facial recognition [Buolamwini and Geburu, 2018, Cook et al., 2019, Cavazos et al., 2020], medical settings [Mac Namee et al., 2002, Ashraf et al., 2018], and many others exhibit differential accuracy across gender, race, and other protected-group membership. This immediately yields *accessibility issues* to users of such systems, and can lead to *direct discrimination*, e.g., facial recognition in policing yielding disproportionate false-arrest rates, and ML in medical technology yielding disproportionate health outcomes, exacerbating existing structural and societal inequalities impacting many minority groups. In welfare-centric ML methods, both *accuracy* and *fairness* are encoded in a single *welfare function* defined on a *group of subpopulations*. Welfare is then directly optimized [Rolf et al., 2020] or constrained [Speicher et al., 2018, Heidari et al., 2018] to promote *fair learning* across *all groups*. This addresses differential performance and bias issues across groups by ensuring that (1),

each group is *seen* and *considered* during training, and (2), an outcome is incentivized that is desirable overall, ideally according to some mutually-agreed-upon welfare function. Unfortunately, welfare based metrics require a notion of (positive) utility, and we argue that this is not natural to many machine learning tasks, where we instead *minimize* some *negatively connoted* loss value (e.g., in decision-theory or with proper scoring rules). We thus define a complementary measure to *welfare*, termed *malfare*, measuring societal harm (rather than wellbeing). In particular, malfare arises naturally when one applies the standard *axioms of cardinal welfare* (with appropriate modifications) to *loss* rather than *utility*. With this framework, we then cast fair machine learning as a direct *malfare minimization* problem, where a group’s malfare is their risk (expected loss).

Perhaps surprisingly, defining and minimizing a *malfare function* is *not equivalent* to defining and maximizing some *welfare function*, while taking utility to be negative loss (except in the trivial cases of egalitarian and utilitarian malfare). This is essentially because nearly every function satisfying the standard axioms of cardinal welfare requires *nonnegative* inputs, and it is not in general possible to contort a loss function into a utility function while satisfying this requirement. For example, while minimizing the 0-1 loss, which simply counts the number of mistakes a classifier makes, is isomorphic to maximizing the 1-0 gain, which counts number of correct classifications, minimizing some *malfare function* defined on 0-1 loss over groups *is not* in general equivalent to maximizing any *welfare function* defined on 1-0 gain. More strikingly, for problems like minimizing square error (i.e., in regression), it is in general not even possible to define a complementary nonnegative gain function without changing the optimal solution, *even for a single group*.

We briefly summarize our contributions as follows.

1. We derive in section 3.3 the *malfare concept*, extending welfare to measure *negatively-connoted* attributes, and show that *malfare-minimization* naturally generalizes *risk-minimization* to produce *fairness-sensitive* machine-learning objectives that consider multiple protected groups.
2. We show in section 3.4 that in many cases, while empirical estimates of welfare and malfare are *statistically biased*, they may be *sharply estimated* using standard finite-sample concentration-of-measure bounds.

3.2 Related Work

Constraint-based notions of algorithmic fairness have risen to prominence in machine learning, with the potential to ensure demographic-parity (e.g., equality of opportunity, equality of outcome, or equalized odds) thus correcting for some forms of data or algorithmic bias. While noble in intent and intuitive by design, fairness by demographic-parity constraints has several prominent flaws: most notably, several popular parity constraints are mutually unsatisfiable [Kleinberg et al., 2017], and their constraint-based formulation inherently puts *accuracy* and *fairness* at odds, where additional *tolerance parameters* are required to strike a balance between the two. Furthermore, recent works [Hu and Chen, 2020, Kasy and Abebe, 2021] have shown that *welfare* and even *disadvantaged group utility* can decrease even as fairness constraints are tightened, calling into question whether demographic parity constraints are even beneficial to those they purport to aid. Without deeper axiomatic underpinnings, it is unclear why one should choose one notion of fairness over another, and it is unsatisfying that we do not have a way of comparing and selecting between them without appealing to informal arguments.

Perhaps in response to these issues, some recent work has trended toward welfare-based fairness-concepts [Hu and Chen, 2020, Rolf et al., 2020], wherein both *accuracy* and *fairness* are encoded in a *welfare function* defined on a *group of subpopulations*. Welfare is then directly optimized [Rolf et al., 2020] or constrained [Speicher et al., 2018, Heidari et al., 2018] to promote *fair learning* across *all groups*. Perhaps the most similar to our work is a method of [Hu and Chen, 2020], wherein they *directly maximize* empirical welfare over linear (halfspace) classifiers; however as with other previous works, an appropriate utility function must be selected, which we avoid by instead using malfare. We argue that *empirical welfare maximization* is an effective strategy when an appropriate and natural measure of *utility* is available, but in machine learning contexts

like this, there is no “correct” or clearly neutral way to convert loss to utility. Our strategy avoids this issue by working directly in terms of malfare and loss.

The most poignant contrast to existing work we can make is to the *Seldonian learner* [Thomas et al., 2019] framework, which can be thought of as extending PAC-learning to learning problems with both *constraints* and *arbitrary nonlinear objectives*. We argue that this generality is harmful to the utility of the concept as a mathematical or practical object, as nearly any ML problem can be posed as a constrained nonlinear optimization task. The utility in fair PAC learning is that it is sophisticated enough to handle fairness issues, with a particular axiomatically justified objective, but remains simple enough to study as a mathematical object; in particular reductions between various PAC and FPAC learnable classes are of great value in understanding FPAC learning, which would not be possible in a more general framework.

3.3 Quantifying Population-Level Sentiment

A generic *population mean* function $M(\mathcal{S}; \mathbf{w})$ quantifies some *sentiment value* \mathcal{S} , across a population Ω weighted by \mathbf{w} . In particular, $\mathcal{S} : \Omega \rightarrow \mathbb{R}_{0+}$ describes the *values* over which we take the mean, and \mathbf{w} , a probability measure over Ω , describes their *weights*. When \mathcal{S} measures a *desirable quantity*, generally termed *utility*, the population mean is a measure of *cardinal welfare* [Moulin, 2004], and thus quantifies overall *well-being*. We also consider the inverse-notion of *ill-being*, termed *malfare*, in terms of an *undesirable* \mathcal{S} , generally *loss* or *risk*, which naturally extends the concept. We show an equivalent *axiomatic justification* for malfare, and argue that its use is more natural in many situations, particularly when considering or optimizing *loss functions*.

Definition 3.3.1 (Population Means: Welfare and Malfare). A *population mean* function $M(\mathcal{S}; \mathbf{w})$ measures the *overall sentiment* of population Ω , measured by *sentiment function* $\mathcal{S} : \Omega \rightarrow \mathbb{R}_{0+}$, weighted by *probability measure* \mathbf{w} over Ω . If \mathcal{S} denotes a desirable quantity (e.g., utility), we call $M(\mathcal{S}; \mathbf{w})$ a *welfare function*, written $W(\mathcal{S}; \mathbf{w})$, and inversely, if it is undesirable (e.g., disutility, loss, or risk), we call $M(\mathcal{S}; \mathbf{w})$ a *malfare function*, written $\Lambda(\mathcal{S}; \mathbf{w})$.

For now, think of the term *population mean* as signifying that an entire population, with diverse and subjective desiderata, is considered and summarized, rather than a single objective viewpoint. As we introduce axioms and show consequent properties, the appropriateness of the term shall become more apparent. Note that we use the term *sentiment* to refer to \mathcal{S} with neutral connotation, but when discussing welfare or malfare, we often refer to \mathcal{S} as *utility* or *risk*, respectively, as in these cases, \mathcal{S} describes a well-understood pre-existing concept. Coarsely speaking, the three notions are identical, all being functions of the form $(\Omega \rightarrow \mathbb{R}_{0+}) \times \text{MEASURE}(\Omega, 1) \rightarrow \mathbb{R}_{0+}$, however we shall see that in order to promote fairness, the desirable axioms of malfare and welfare functions differ slightly. The notation reflects this; $M(\mathcal{S}; \mathbf{w})$ is an M for *mean*, whereas $W(\mathcal{S}; \mathbf{w})$ is a W for *welfare*, and $\Lambda(\mathcal{S}; \mathbf{w})$ is an Λ (*inverted W*), to emphasize its inverted nature.

Often we are interested in *unweighted* population means, given sentiments as a *vector* $\mathcal{S} \in \mathbb{R}_{0+}^g$. Unweighted means may be defined in terms of weighted means, as

$$M(\mathcal{S}) \doteq M\left(i \mapsto \mathcal{S}_i; \left\{i \mapsto \frac{1}{g}\right\}\right) \quad ,$$

abusing notation to concisely express the uniform measure. Indeed, it may seem antithetical to fairness to allow for weights in malfare and welfare definitions; consider however that weights can represent *differential population sizes*, and thus ensure that the welfare or malfare of *weight-preserving decompositions* of groups into subgroups with equal risk or utility remains constant.

Example 3.3.2 (Utilitarian Welfare). Suppose individuals reside in some space \mathcal{X} , where *distributions* $\mathcal{D}_{1:g}$ over domain \mathcal{X} describe the distribution over individuals in *each group*. Suppose also *utility function*

$U(x) : \mathcal{X} \rightarrow \mathbb{R}_{0+}$, describing the *level of satisfaction* of an individual, w.r.t., e.g., some *situation, allocation, or classifier*. We now take the *sentiment function* to be the *mean utility* (per-group), i.e.,

$$\mathcal{S}(\omega_i) \doteq \mathbb{E}_{x \sim \mathcal{D}_i} [U(x)] = \mathbb{E}_{\mathcal{D}_i} [U] .$$

Now, given a *weights vector* \mathbf{w} , describing the *relative frequencies* of membership in each of the g groups, we define the *utilitarian welfare* as

$$W_1(\mathcal{S}; \mathbf{w}) \doteq \sum_{i=1}^g \mathbf{w}(\omega_i) \mathcal{S}(\omega_i) = \mathbb{E}_{\omega \sim \mathbf{w}} [\mathcal{S}(\omega)] = \mathbb{E}_{\mathbf{w}} [\mathcal{S}] .$$

Of course, in statistical, sampling, and machine learning contexts, $\mathcal{D}_{1:g}$ and \mathbf{w} may be unknown, so we now discuss an *empirical analogue*. Section 3.4 is then devoted to showing how and when empirical population means well-approximate their true counterparts.

Example 3.3.3 (Empirical Utilitarian Welfare). Now suppose $\mathcal{D}_{1:g}$ are unknown, but instead, we are given a *sample* $\mathbf{x}_{1:g,1:m} \in \mathcal{X}^{g \times m}$, where $\mathbf{x}_{i,1:m} \sim \mathcal{D}_i$. We define an *empirical analogue* of the utilitarian welfare as in example 3.3.2, instead taking

$$\hat{\mathcal{S}}(\omega_i) \doteq \hat{\mathbb{E}}_{x \in \mathbf{x}_i} [U(x)] \quad \& \quad \hat{W}_1(\hat{\mathcal{S}}, \mathbf{w}) \doteq \mathbb{E}_{\mathbf{w}} [\hat{\mathcal{S}}] .$$

Similarly, if \mathbf{w} is unknown, but we may sample from some \mathcal{D} over $\Omega \times \mathcal{X}$, we can use *empirical frequencies* $\hat{\mathbf{w}}$ in place of *true frequencies* \mathbf{w} , and define $\hat{\mathcal{S}}(\omega_i)$ as *conditional averages* over the subsample associated with group i .

3.3.1 Axioms of Cardinal Welfare and Malfare

Definition 3.3.4 (Axioms of Cardinal Welfare and Malfare). We define the *population-mean axioms* for population-mean function $M(\mathcal{S}; \mathbf{w})$ below. For each item, assume (if necessary) that the axiom applies $\forall \mathcal{S}, \mathcal{S}' \in \Omega \rightarrow \mathbb{R}_{0+}$, scalars $\alpha, \beta \in \mathbb{R}_{0+}$, and probability measures \mathbf{w} over Ω .

1. (Strict) Monotonicity: $\forall \varepsilon : \Omega \rightarrow \mathbb{R}_{0+}$ s.t. $\int_{\mathbf{w}} \varepsilon(\omega) d(\omega) > 0$: $M(\mathcal{S}; \mathbf{w}) < M(\mathcal{S} + \varepsilon; \mathbf{w})$.

2. Symmetry: \forall permutations π over Ω : $M(\mathcal{S}; \mathbf{w}) = M(\pi(\mathcal{S}); \pi(\mathbf{w}))$.

3. Continuity: $\{\mathcal{S}' \mid M(\mathcal{S}'; \mathbf{w}) \leq M(\mathcal{S}; \mathbf{w})\}$ and $\{\mathcal{S}' \mid M(\mathcal{S}'; \mathbf{w}) \geq M(\mathcal{S}; \mathbf{w})\}$ are closed sets.

4. Independence of unconcerned agents: Suppose subpopulation $\Omega' \subseteq \Omega$. Then

$$M \left(\begin{array}{l} p \in \Omega' : \alpha \\ p \notin \Omega' : \mathcal{S}(p) \end{array} ; \mathbf{w} \right) \leq M \left(\begin{array}{l} p \in \Omega' : \alpha \\ p \notin \Omega' : \mathcal{S}'(p) \end{array} ; \mathbf{w} \right) \Rightarrow M \left(\begin{array}{l} p \in \Omega' : \beta \\ p \notin \Omega' : \mathcal{S}(p) \end{array} ; \mathbf{w} \right) \leq M \left(\begin{array}{l} p \in \Omega' : \beta \\ p \notin \Omega' : \mathcal{S}'(p) \end{array} ; \mathbf{w} \right) .$$

5. Independence of common scale: $M(\mathcal{S}; \mathbf{w}) \leq M(\mathcal{S}'; \mathbf{w}) \Rightarrow M(\alpha \mathcal{S}; \mathbf{w}) \leq M(\alpha \mathcal{S}'; \mathbf{w})$.

6. Multiplicative linearity: $M(\alpha \mathcal{S}; \mathbf{w}) = \alpha M(\mathcal{S}; \mathbf{w})$.

7. Unit scale: $M(\mathbf{1}; \mathbf{w}) = M(1, \dots, 1; \mathbf{w}) = 1$.

8. Pigou-Dalton transfer principle: Suppose $\mu = \mathbb{E}_{\mathbf{w}}[\mathcal{S}] = \mathbb{E}_{\mathbf{w}}[\mathcal{S}']$, and $\forall p \in \Omega : |\mu - \mathcal{S}'(p)| \leq |\mu - \mathcal{S}(p)|$. Then $W(\mathcal{S}'; \mathbf{w}) \geq W(\mathcal{S}; \mathbf{w})$.

9. Anti-Pigou-Dalton transfer principle: Suppose as in 8, and conclude $M(\mathcal{S}'; \mathbf{w}) \leq M(\mathcal{S}; \mathbf{w})$.

We take a moment to comment on each of these axioms, to preview their purpose and assure the reader of their necessity. Axioms 1-5 are the standard *axioms of cardinal welfarism* (1-4 are discussed by Sen [1977], Roberts [1980], and 5 by Debreu [1959], Gorman [1968]). Together, they imply (via the Debreu-Gorman theorem) that any population-mean can be decomposed as a *monotonic function* of a *sum* (over groups) of *log*

or *power* functions. Axiom 6 is a natural and useful property, and ensures that *dimensional analysis* on mean functions is possible: in particular, the *units* of mean functions match those of sentiment. Note that axiom 6 implies axiom 4, and it is thus a simple strengthening of a traditional cardinal welfare axiom. This axiom also ensures that *units* of population means preserve the *units* of \mathcal{S} , making dimensional analysis greatly more convenient; we will also see that it is essential to show convenient *statistical* and *learnability* properties. Axiom 7 furthers this theme, as it ensures that not only do *units* of means match those of \mathcal{S} , but scale does as well (making comparisons like \mathcal{S}_i is *above* the population-welfare meaningful), and also enabling comparison *across populations* (in the sense that comparing *averages* is more meaningful than *sums*). Finally, axiom 8 (the *Pigou-Dalton transfer principle* [see Pigou, 1912, Dalton, 1920]) is also standard in cardinal welfare theory as it ensures fairness, in the sense that welfare is higher when utility values are more uniform, i.e., incentivizing *equitable redistribution of “wealth”* in welfare. Its antithesis, axiom 9, encourages the opposite; in the context of welfare, this perversely incentivizes an expansion of inequality, but for malfare, which we generally wish to *minimize*, the opposite occurs, thus this axiom characterizes *fairness* in the context of *malfare*.

Axioms 6-4 are novel to this work, and are key in strengthening the Debreu-Gorman theorem to ensure that all welfare and malfare functions are *power means* in the sequel. Axiom 9 is also novel, as it is necessary to flip the inequality of axiom 8 when the sense of the population mean is inverted from welfare to malfare; in particular, the semantic meaning shifts from requiring that “redistribution of utility is desirable” to “redistribution of disutility is not undesirable.”

For context, we present an additional axiom; that of *additive separability*. For simplicity, we present it only in the unweighted discrete case, as there is some subtlety to an equivalent measure-theoretic formulation, and we derive no benefit from assuming this axiom, as it is largely incompatible with our assumptions, and presented only for comparison purposes.

Definition 3.3.5 (Additive Separability). Population-mean $M(\mathcal{S}_{1:g})$ is *additively separable* if there exist functions $f_{1:g}$ s.t. each $f_i \in \mathbb{R}_{0+} \rightarrow \mathbb{R}_{0+}$, and $M(\mathcal{S}_{1:g})$ may be decomposed as

$$M(\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_g) = \sum_{i=1}^g f_i(\mathcal{S}_i) .$$

While seemingly quite important to early welfare theorists (e.g., the Debreu-Gorman theorem is generally presented in additively-separable form), and it gives rise to some convenient interpretability and computational properties, we argue that these are far-outstripped by those stemming from the *multiplicative linearity* and *unit scale* axioms, with no real difference in generality.¹ Furthermore, unlike these and other standard cardinal welfare axioms, additive separability seems a bit-heavy handed, assuming something very specific that is supposedly convenient for the economist, with little justification as to why and how it serves as a fundamental property of *cardinal welfare* itself. These properties are discussed in section 3.3.3, and directly compared with the additively-separable form in section 3.3.4.

3.3.2 The Power Mean

We now define the *p-power mean* $M_p(\cdot)$, for any $p \in \mathbb{R} \cup \pm\infty$, which we shall use to quantify both malfare and welfare. Power means arise often when analyzing population means obeying the various axioms of definition 3.3.4, and as we shall see in theorem 3.3.7, are a particularly important class of population means.

Definition 3.3.6 (Power-Mean Welfare and Malfare). Suppose $p \in \mathbb{R} \cup \pm\infty$. We first define the *unweighted*

¹In this sense, the *additive welfare functionals* are somewhat like the *natural sufficient statistics* of the exponential family, in that, while additivity is sometimes convenient, it comes only with the sacrifice of other desiderata.

power mean of sentiment vector $\mathcal{S} \in \mathbb{R}_{0+}^g$ as

$$M_p(\mathcal{S}) \doteq \begin{cases} p \in \mathbb{R} \setminus \{0\} & \sqrt[p]{\frac{1}{g} \sum_{i=1}^g \mathcal{S}_i^p} \\ p = -\infty & \min_{i \in \{1, \dots, g\}} \mathcal{S}_i \\ p = 0 & \sqrt[g]{\prod_{i=1}^g \mathcal{S}_i} = \exp\left(\frac{1}{g} \sum_{i=1}^g \ln(\mathcal{S}_i)\right) \\ p = \infty & \max_{i \in \{1, \dots, g\}} \mathcal{S}_i \end{cases}.$$

We now define the *weighted power mean*, given *sentiment value function* $\mathcal{S} : \Omega \rightarrow \mathbb{R}_{0+}$ and probability measure \mathbf{w} over Ω , as

$$M_p(\mathcal{S}; \mathbf{w}) \doteq \begin{cases} p \in \mathbb{R} \setminus \{0\} & \sqrt[p]{\int_{\mathbf{w}} \mathcal{S}^p(\omega) d(\omega)} = \sqrt[p]{\mathbb{E}_{\omega \sim \mathbf{w}}[\mathcal{S}^p(\omega)]} \\ p = -\infty & \min_{\omega \in \text{Support}(\mathbf{w})} \mathcal{S}(\omega) \\ p = 0 & \exp\left(\int_{\mathbf{w}} \ln \mathcal{S}(\omega) d(\omega)\right) = \exp\left(\mathbb{E}_{\omega \sim \mathbf{w}}[\ln \mathcal{S}(\omega)]\right) \\ p = \infty & \max_{\omega \in \text{Support}(\mathbf{w})} \mathcal{S}(\omega) \end{cases}.$$

In both the weighted and unweighted cases, $p \in \{-\infty, 0, \infty\}$ resolve as their (unique) limits, and for all $p \in \mathbb{R}$, note that *power means* are special cases of the (weighted) *generalized mean*, defined for strictly monotonic f as $M_f(\mathcal{S}; \mathbf{w}) \doteq f^{-1}(\mathbb{E}_{\omega \sim \mathbf{w}}[\ln \mathcal{S}(\omega)])$.

Theorem 3.3.7 (Properties of the Power-Mean). Suppose \mathcal{S}, ε are sentiment values $\Omega \rightarrow \mathbb{R}_{0+}$, and \mathbf{w} is a probability measure over Ω . Then

1. Monotonicity: $M_p(\mathcal{S}; \mathbf{w})$ is weakly-monotonically-increasing in p , and strictly if \mathcal{S} attains distinct $a, b \in \mathbb{R}$ with nonnegligible probability.
2. Subadditivity: $\forall p \geq 1 : M_p(\mathcal{S} + \varepsilon; \mathbf{w}) \leq M_p(\mathcal{S}; \mathbf{w}) + M_p(\varepsilon; \mathbf{w})$.
3. Contraction: $\forall p \geq 1 : |M_p(\mathcal{S}; \mathbf{w}) - M_p(\mathcal{S}'; \mathbf{w})| \leq M_p(|\mathcal{S} - \mathcal{S}'|; \mathbf{w}) \leq \|\mathcal{S} - \mathcal{S}'\|_{\infty}$.
4. Curvature: $M_p(\mathcal{S}; \mathbf{w})$ is concave in \mathcal{S} for $p \in [-\infty, 1]$ and convex for $p \in [1, \infty]$.

3.3.3 Properties of Welfare and Malfare Functions

We now show that the axioms of definition 3.3.4 are sufficient to characterize many properties of welfare and malfare.

Theorem 3.3.8 (Population Mean Properties). Suppose population-mean function $M(\mathcal{S}; \mathbf{w})$. If $M(\cdot; \cdot)$ satisfies (subsets of) the population-mean axioms (see definition 3.3.4), we have that $M(\cdot; \cdot)$ exhibits the following properties. For each, assume arbitrary sentiment-value function $\mathcal{S} : \Omega \rightarrow \mathbb{R}_{0+}$ and weights measure \mathbf{w} over Ω . The following then hold.

1. *Identity*: Axioms 6 & 7 imply $M(\omega \mapsto \alpha; \mathbf{w}) = \alpha$.
2. Axioms 1-5 imply $\exists p \in \mathbb{R}$, strictly-monotonically-increasing continuous $F : \mathbb{R} \rightarrow \mathbb{R}_{0+}$ s.t.

$$M(\mathcal{S}; \mathbf{w}) = F\left(\int_{\mathbf{w}} f_p(\mathcal{S}(\omega)) d(\omega)\right) = F\left(\mathbb{E}_{\omega \sim \mathbf{w}}[f_p(\mathcal{S}(\omega))]\right), \quad \text{with } \begin{cases} p = 0 & f_0(x) \doteq \ln(x) \\ p \neq 0 & f_p(x) \doteq \text{sgn}(p)x^p \end{cases}.$$

3. Axioms 1-7 imply $F(x) = f_p^{-1}(x)$, thus $M(\mathcal{S}; \mathbf{w}) = M_p(\mathcal{S}; \mathbf{w})$.
4. Axioms 1-5 and 8 imply $p \in (-\infty, 1]$.
5. Axioms 1-5 and 9 imply $p \in [1, \infty)$.

Taken together, the items of theorem 3.3.8 tell us that the mild conditions of axioms 1-4 (generally assumed for welfare), along with *multiplicative linearity*, imply that welfare and utility, or malfare and loss, are measured in the *same units* (e.g., *nats* or *bits* for *cross-entropy loss*, square- \mathcal{V} -units for *square error*, or *dollars* for *income utility*), and power-mean malfare is effectively the only reasonable choice of welfare or malfare function. Even without multiplicative linearity, axioms 1-5 imply population mean functions are still *monotonic transformations* of power-mean. Furthermore, the entirely milquetoast *unit scale* axiom implies that sentiment values and population means have the same *scale*, making comparisons like “the risk of group i is above (or below) the population malfare” meaningful. Finally, we also have that $p \in [-\infty, 1)$ incentivizes redistribution of utility from better-off groups to worse-off groups, and similarly $p \in [1, \infty)$ incentivizes redistribution of harm² from worse-off groups to better-off groups.

3.3.4 A Comparison with the Additively Separable Form

In particular, assuming axioms 1-5, all population means are *monotonic transformations* of the power mean, thus whether we assume *additive separability*, and get the additively-separable form

$$M(\mathcal{S}) = c \sum_{i=1}^g f_p(\mathcal{S}_i) = cgM_p^c(\mathcal{S}) \text{ ,}$$

for $p \in \mathbb{R} \setminus \{0\}$ (where usually $c = 1$ is taken as the canonical form), or we assume axioms 6 & 7, and get the power-mean, there is no real loss of descriptiveness, as all such population-means remain *isomorphic* under the binary comparison operator (\leq). With additive separability, it is straightforward to compute the welfare of a population from the welfares of *subpopulations*, but it is still computationally trivial to do this with power means. Furthermore, the limiting cases of $p \in \pm\infty$ become undefined in the additively separable form, and we lose monotonicity in p (see theorem 3.3.7), both of which are remedied with power means.

With additive separability, welfare summarizes population sentiment by intuitively generalizing the idea of *summation*. In contrast, in our setting, we prefer to think of welfare as a generalized *average*. This yields desirable statistical estimation and learnability properties (shown in the sequel), but is also useful in and of itself, as it allows us to, for instance compare individual sentiment values to welfare, as both the *units* and *scale* match.

Another reason to prefer the power-mean over the additively separable form is the potential for direct comparisons between *group sentiments*, population means of *subgroups*, and *overall population means*. In particular, the *dimensional analysis* properties of the *power mean* are convenient, as these comparisons agree in *units* (due to axiom 6) and *scale* (due to axiom 7). No such dimensional analysis is possible with the additively separable form, as, e.g., if \mathcal{S} is measured in dollars, then $W_2^2(\mathcal{S}; \mathbf{w})$ is measured in *square dollars* (a rather unintuitive unit).

3.3.5 Relating Power Means and Inequality Indices

We now discuss and define *relative inequality indices* $I(\mathcal{S}; \mathbf{w})$, which have been employed in the literature [Sen et al., 1997] to construct *welfare functions* of the form

$$W(\mathcal{S}; \mathbf{w}) = W_1(\mathcal{S}; \mathbf{w})(1 - I(\mathcal{S}; \mathbf{w})) \text{ .}$$

²Note that, mathematically speaking, it is entirely *valid* to quantify welfare with $p > 1$ or malfare with $p < 1$, and indeed such characterizations may arise in the analysis of unfair systems; however we generally advocate against intentionally creating such unfair systems.

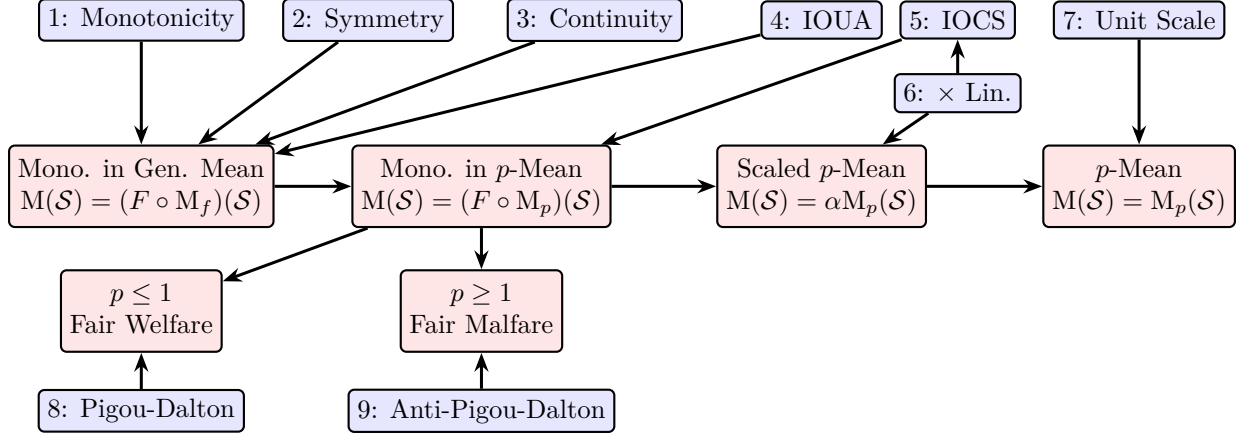


Figure 3.1: Relationships between population-mean axioms and properties. Assumptions and axioms shown in pastel blue, and properties shown in pastel red. Equivalent properties hold for weighted population mean functions.

This characterization intuitively starts with the *utilitarian welfare*, which measures *overall satisfaction* and then *downweights* based on how unfairly distributed utility is amongst the population. The “relative” in relative inequality indices connotes the fact that they are restricted to domain $[0, 1]$, thus the welfare metric matches the utilitarian under no inequality, and is 0 under maximal inequality.

We show that a large class of such functions are actually power means, which both gives them axiomatic justification, and shows prior support in the literature for the power mean. In particular, we first consider the *Atkinson index* relative inequality measure family [Atkinson et al., 1970].

Definition 3.3.9 (Atkinson Index). For all $\varepsilon \in \mathbb{R}$, we define the Atkinson index as

$$\text{Atk}_\varepsilon(\mathcal{S}; \mathbf{w}) \doteq 1 - \frac{M_{1-\varepsilon}(\mathcal{S}; \mathbf{w})}{M_1(\mathcal{S}; \mathbf{w})} .$$

Note that often the Atkinson index is restricted to $\varepsilon \in [0, 1]$; outside this range, it may exceed 1. Furthermore, the Atkinson index is generally stated without weights, and in a mathematically equivalent form, in which the resemblance to the power mean is less obvious, but for our purposes the above form is clearer. From it, we immediately have the following lemma.

Lemma 3.3.10 (Relating Atkinson’s Indices and Power Means). Suppose some $\varepsilon \in \mathbb{R}$, and take $p = 1 - \varepsilon$. It then holds that

$$M_p(\mathcal{S}; \mathbf{w}) = M_1(\mathcal{S}; \mathbf{w})(1 - \text{Atk}_\varepsilon(\mathcal{S}; \mathbf{w})) .$$

Proof. This is a direct consequence of definition 3.3.9, noting $p = 1 - \varepsilon \Leftrightarrow \varepsilon = 1 - p$. □

This is not particularly surprising in light of the welfare-centric derivation of [Atkinson et al., 1970], but nonetheless it yields a valuable alternative way to think about power means and inequality-weighted welfare functions. In particular, it gives a direct *axiomatic justification* of the welfare function $W(\mathcal{S}; \mathbf{w}) = W_1(\mathcal{S}; \mathbf{w})(1 - \text{Atk}_\varepsilon(\mathcal{S}; \mathbf{w}))$ (see theorem 3.3.8), and also gives an alternative intuitive interpretation of power-mean welfare (as inequality-weighted utilitarian welfare).

Note that similar properties may be shown for the isomorphic inequality measures of *generalized entropy indices* and *Theil indices* [Theil, 1967], though in this context their forms are generally less pleasing.

3.4 Statistical Estimation of Malfare Values

We first illustrate the ease with which p -power means can be estimated, in contrast to the standard additive welfare formulations. Perhaps surprisingly, we find that empirical welfare and malfare are *biased estimators*, yet they admit much sharper finite-sample tail bounds than the additive welfare formulations, which are unbiased.

Lemma 3.4.1 (Statistical Estimation). Suppose probability distribution \mathcal{D} , population-mean M obeying *monotonicity*, sentiment value function \mathcal{S} such that, given functions f_ω , we have $\mathcal{S}(\omega) = \mathbb{E}_{x \sim \mathcal{D}}[f_\omega(x)]$, sample $\mathbf{x} \sim \mathcal{D}^m$, and *empirical sentiment value estimate* $\hat{\mathcal{S}} \doteq \mathbb{E}_{\mathbf{x} \in \mathcal{X}}[f_\omega(x)]$. If it holds with probability at least $1 - \delta$ that $\forall \omega : \mathcal{S}'(\omega) - \varepsilon(\omega) \leq \mathcal{S}(\omega) \leq \mathcal{S}'(\omega) + \varepsilon(\omega)$, then with said probability, we have

$$M_p(\mathbf{0} \vee (\hat{\mathcal{S}} - \varepsilon); \mathbf{w}) \leq M_p(\hat{\mathcal{S}}; \mathbf{w}) \leq M_p(\hat{\mathcal{S}} + \varepsilon; \mathbf{w}) ,$$

where $\mathbf{a} \vee \mathbf{b}$ denotes the (elementwise) minimum.

Proof. This result follows from the assumption, and the *monotonicity* axiom (i.e., adding or subtracting ε can not decrease or increase the power mean, respectively). The minimum with 0 on the LHS is required simply because by definition, sentiment values are nonnegative, and M_p is in general undefined with negative inputs. \square

We now reify this result, applying the well-known Hoeffding [1963] and Bennett [1962] bounds to show concentration and derive an explicit form for ε .

Corollary 3.4.2 (Statistical Estimation with Hoeffding and Bennett Bounds). Suppose fair power-mean malfare $\Lambda(\cdot; \cdot)$ (i.e., $p \geq 1$), *loss function* $\ell : \mathcal{X} \rightarrow [0, r]$, $\mathcal{S} \in [0, r]^g$ s.t. $\mathcal{S}_i = \mathbb{E}_{\mathcal{D}_i}[\ell]$, samples $\mathbf{x}_i \sim \mathcal{D}_i^m$, and take $\hat{\mathcal{S}}_i \doteq \frac{1}{m} \sum_{j=1}^m \ell(\mathbf{x}_{i,j})$. Then with probability at least $1 - \delta$ over choice of \mathbf{x} ,

$$\left| M_p(\mathcal{S}; \mathbf{w}) - M_p(\hat{\mathcal{S}}; \mathbf{w}) \right| \leq r \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}} .$$

Alternatively, again with probability at least $1 - \delta$ over choice of \mathbf{x} , we have

$$\left| M_p(\mathcal{S}; \mathbf{w}) - M_p(\hat{\mathcal{S}}; \mathbf{w}) \right| \leq \frac{r \ln \frac{2g}{\delta}}{3m} + \max_{i \in \{1, \dots, g\}} \sqrt{\frac{2 \mathbb{V}_{\mathcal{D}_i}[\ell] \ln \frac{2g}{\delta}}{m}} .$$

As corollary 3.4.2 follows directly from lemma 3.4.1, with Hoeffding and Bennett inequalities applied to derive ε bounds, similar results are immediately possible with arbitrary concentration inequalities. In particular, similar *data-dependent* bounds may be shown, e.g., with *empirical Bennett bounds*, removing dependence on *a priori known variance*. In particular, we may apply theorem 1.3.3 to bound ε , either elementwise (over groups) with a *union bound*, or jointly, if the loss functions for each group may be combined into a single function family.

Furthermore, note that while these bounds may be used for *evaluating* the welfare or malfare of a *particular* classifier or mechanism (through \mathcal{S} and $\hat{\mathcal{S}}$), they immediately extend to *learning* over a *finite family* via the union bound. Much like in standard statistical learning theory, the exponential tail bounds of corollary 3.4.2 allow the family to grow *exponentially*, at *linear* cost to sample complexity. Also as in standard uniform convergence analysis (generally discussed in the context of *empirical risk minimization*), we can easily handle infinite hypothesis classes and obtain much sharper bounds by considering *uniform convergence bounds* over the family, e.g., with Rademacher averages and appropriate concentration-of-measure bounds.

In learning contexts, minimizing the *malfare* among all groups immediately generalizes minimizing *risk* of a single group. These statistical estimation bounds immediately imply that the *empirical malfare-optimal* solution is a reasonable proxy for the true malfare-optimal solution, as we now formalize.

Definition 3.4.3 (The Empirical Malfare Minimization (EMM) Principle). Suppose hypothesis class $\mathcal{H} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$, training samples $\mathbf{z}_{1:g}$ drawn from distributions $\mathcal{D}_{1:g}$ over $\mathcal{X} \times \mathcal{Y}$, loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{0+}$, malfare function \mathbb{M} , group weights \mathbf{w} . The *empirical malfare minimizer* is then defined as

$$\hat{h} \doteq \operatorname{argmin}_{h \in \mathcal{H}} \mathbb{M} \left(i \mapsto \hat{R}(h; \ell, \mathbf{z}_i); \mathbf{w} \right) ,$$

and the EMM principle states that \hat{h} is a reasonable proxy for the true malfare minimizer

$$h^* \doteq \operatorname{argmin}_{h^* \in \mathcal{H}} \mathbb{M} \left(i \mapsto R(h^*; \ell, \mathcal{D}_i); \mathbf{w} \right) .$$

In all cases, these bounds provide further justification for using the *power-mean* over the *additively separable* form, as it is substantially harder to achieve comparable bounds on $M_p^p(\mathcal{S}; \mathbf{w})$, due to the increased difficulty of controlling the *range* and *Lipschitz constant* of these quantities. Of course, as power-mean bounds imply bounds on the additively-separable form (and vice-versa), we recommend working with power means, and then converting back to the additively separable form (if so desired).

Chapter 4

Ewoks: an Algorithm for Fair Codec Selection

In this chapter, I frame the problem of selecting *sets of streaming media codecs* as a *supervised learning task*, which I address with the *Empirical Welfare-Optimal k codec Selection* (EWOKS) algorithm. I consider multivariate optimality concepts (over, e.g., *audio quality* versus *bandwidth consumption*), such as *Pareto optimality*, *minimax-optimality*, and various *welfare* concepts. Welfare solutions perform well *in aggregate* over *sets* or *distributions* of users with *varied preferences*, quantifying and providing fine-grained control over *algorithmic fairness* and *bias* issues. I prove that EWOKS-solutions are approximately welfare-optimal over the underlying media distribution via rigorous statistical learning guarantees, and show empirically that EWOKS solutions are fairer (as measured by welfare) than single-objective solutions. This chapter is based on joint work with Clayton Sanford and Eli Upfal.

4.1 Introduction

Streaming video now accounts for an estimated 58% of all internet traffic [Sandvine, 2018], thus there is strong user demand and economic incentive to optimize stream quality and resource utilization. Most streamed audio and video data is *compressed* with a *codec* (coder-decoder), with complicated, subjective, multivariate (often conflicting) objectives; for example *resource utilization metrics*, (e.g., *compression ratio*, *encode/decode CPU time*), *fidelity / distortion (quality) metrics*, and *fault tolerance*, are all important, to varying degrees for various users, applications, and media types. Furthermore, while codecs are generally developed and tested on particular use cases (e.g., lossless audio, low-bandwidth speech data, various qualities of music data), and relative performance on each attribute varies depending on the media being encoded.

While many codecs purport efficacy for various use cases, little effort has been put into rigorous multivariate comparative statistical analysis. In practice, sets of streaming media codecs are generally selected *ad-hoc* by domain experts to target particular bitrates or quality thresholds, which may be suboptimal for users with different preferences, or over different media distributions. For example, low-quality mp3 encodings may be used for a low-bitrate music stream, and lossless FLAC encodings for a high-quality stream, but for a speech stream, there exist specialized codecs that will achieve higher quality than a general-purpose codec at a given compression ratio. Furthermore, relative codec performance across media is inconsistent, and usually selecting the best encoding among a set of codecs outperforms any individual codec *on average*.

The multivariate aspect of the codec-selection problem induces both deep theoretical quandries and impactful real-world fairness and algorithmic bias issues. New technologies often fail to adequately serve global users,

resulting in products with varying degrees of success across markets, harming both potential users and developers. We model the preferences of a user with a *loss function* ℓ mapping *codec attributes* to *user dissatisfaction*, and model *user diversity* by considering a *set* or *distribution* \mathcal{L} over loss functions.

We address all of these issues with a novel *fair media streaming strategy*, employing a *codec set* \mathbf{c} , s.t. when a user with loss function ℓ requests media x , we return the encoding $c(x)$ for $c \in \mathbf{c}$ that minimizes $(\ell \circ c)(x)$, thus optimizing for *data-dependent user preference*. We then propose the *Empirical Welfare-Optimal k codec Selection* (EWoks) algorithm, which *learns* a data-dependent set of k codecs \mathbf{c} that is *welfare-optimal* over \mathcal{L} . Data-dependence is key, as EWoks considers each codec in an unbiased way, allowing empirical performance to dictate whether it is selected, which automates the process of selecting domain-appropriate codecs, eliminating bias due to distribution-shift.

While various *welfare concepts* characterize different *notions of fairness*, EWoks is not tied to any particular welfare function. Rather, we show how to optimize and bound generalization error across a *broad class of welfare concepts*, where both *utilitarian* and *egalitarian* welfare arise as special cases.

It may be computationally intractable to optimize welfare over large datasets, and one may wish to approximate with a *sample* [see e.g. Riondato and Upfal, 2015, 2016, 2018], and in dynamic settings, new media are continuously uploaded, and thus are unavailable during training. We show that EWoks solutions are not only *empirically optimal*, but also *approximately optimal* over the *underlying media distribution* of the training data, thus mitigating overfitting concerns, by posing codec selection as a *statistical learning problem* and showing data-dependent uniform-convergence bounds.

We show novel finite-sample data-dependent welfare-generalization bounds with Monte-Carlo Rademacher averages [Koltchinskii, 2001, Bartlett and Mendelson, 2002] and modern entropy-method concentration inequalities [Boucheron et al., 2003, 2013], which outperform standard Martingale-based Rademacher bounds [Mitzenmacher and Upfal, 2017], particularly in the small-sample setting. Our bounds are sensitive to the *empirical variances* of codec performances, and can substantially improve upon the generalization guarantees of variance-insensitive methods. Such welfare-generalization guarantees are particularly important, as unlike with risk and loss functions, *empirical welfare* is *not* an unbiased estimator of *welfare*.

We now summarize our contributions:

1. We introduce *k codec selection*, reframing traditionally manual codec selection as a *supervised learning problem*, and propose the EWoks algorithm it.
2. We control for *algorithmic bias* and *fairness* issues by learning *distribution-dependent* models and robustly optimizing over a broad class of *welfare concepts*.
3. We show welfare generalization bounds through *uniform convergence* theory and novel sharp *variance-sensitive* tail bounds.
4. We validate experimentally that EWoks robustly controls for overfitting and unfairness due to optimizing for specific user models or data distributions.

4.2 The EWoks Algorithm

In the codec selection framework, we assume media exist in some space \mathcal{X} , and a family \mathcal{C} of *codecs*, which act as functions from \mathcal{X} onto a compressed \mathcal{X}' and back, which we abstract to functions from \mathcal{X} onto the space of *attribute values* \mathbb{R}_{0+}^a . Attributes represent various *objectively measurable* properties such as *quality* and *resource utilization* metrics. For example, mp3, with fixed encoder hyperparameters, is a codec, and the quality / distortion metrics and compression ratio are *attributes* of an encoding. We generally assume monotonically minimizing each objective is desirable (i.e., decreasing some attribute while others remain constant is never bad), though often this may be relaxed.

We also assume a family $\mathcal{L} \subseteq \mathbb{R}_{0+}^a \rightarrow \mathbb{R}_{0+}$ of *loss functions*, such that each loss $\ell \in \mathcal{L}$, codec $c \in \mathcal{C}$, media sample $x \in \mathcal{X}$, $(\ell \circ c)(x)$ quantifies the *subjective dissatisfaction* of ℓ on encoding $c(x)$. Therefore \mathcal{L} represents the *space of user preferences* for various options and tradeoffs, for instance for video *framerate*, *quality*, *resolution*, and *compression ratio*. In other words, the *family \mathcal{L} objectively captures the subjective desiderata* of users, each represented by some $\ell \in \mathcal{L}$. While each $\ell \in \mathcal{L}$ is an objective metric, by considering all of \mathcal{L} simultaneously, we remain sensitive to the *subjective nature* of the problem.

Given a set \mathbf{c} of k codecs, when a user with some loss function ℓ requests media $x \in \mathcal{X}$, we serve the request by selecting the $c \in \mathbf{c}$ that *minimizes* $(\ell \circ c)(x)$, and returning $(c, c(x))$. Consequently, it's generally sufficient to select a codec set such that *on average* (over the media distribution), *at least one* codec is satisfactory for *each user*. We quantify the aversion of a user with loss function ℓ to codec set \mathbf{c} over media distribution \mathcal{D} with the *risk*

$$R(\mathbf{c}; \ell, \mathcal{D}) \doteq \mathbb{E}_{x \sim \mathcal{D}} \left[\min_{c \in \mathbf{c}} (\ell \circ c)(x) \right]. \quad (26)$$

The brief explanation above, and the optimality concepts discussed in the sequel, are sufficient to understand and analyze EWOKS.

4.2.1 Welfare-Optimal k codec Selection

Given a large codec family \mathcal{C} , the computation and storage costs of serving each query with the optimal $c \in \mathcal{C}$ become prohibitive. Depending on the data distribution \mathcal{D} and user needs (represented by \mathcal{L}), many codecs may contribute little marginal benefit, hence a subset of \mathcal{C} is often sufficient. For a single loss function ℓ , ℓ -optimality encourages selecting a *diverse codec set \mathbf{c}* , as there is little marginal benefit to adding a high-performing codec c to \mathbf{c} if its performance is highly correlated with some $c' \in \mathbf{c}$ (as may occur with *similar encoding algorithms* or *parameters*). Furthermore, in general, EWOKS considers a *loss family \mathcal{L} in aggregate*, and welfare-objectives incentivize selecting a *diverse codec set* that performs well for all $\ell \in \mathcal{L}$. In short, our goal is to select some $\mathbf{c} \subseteq \mathcal{C}$ from among the set of size- k subsets of \mathcal{C} , henceforth written $\binom{\mathcal{C}}{k}$, that performs well in aggregate over \mathcal{L} w.r.t. \mathcal{D} .

With a specific user or objective use-case, represented by loss function ℓ , *risk* is the natural way to quantify performance. We then define the *risk minimizer \mathbf{c}_ℓ^** as

$$\mathbf{c}_\ell^* \doteq \operatorname{argmin}_{\mathbf{c} \in \binom{\mathcal{C}}{k}} R(\mathbf{c}; \ell, \mathcal{D}). \quad (27)$$

With a *family* or *distribution \mathcal{L}* over loss functions, we face tradeoffs as to the degree to which each $\ell \in \mathcal{L}$ may be satisfied. One approach to algorithmic fairness and robustness against arbitrary use-cases is to select a set of codecs \mathbf{c} that optimizes worst-case performance (minimax-optimality) over \mathcal{L} . This corresponds to the *egalitarian-welfare optimal* solution concept¹ over \mathcal{L} . Welfare concepts capture various aspects of multivariate optimality in multi-agent systems [Gibbons, 1992], and egalitarian welfare ensures fairness, because it always incentivizes selecting \mathbf{c} to benefit the *least-satisfied* users. Alternatively, *utilitarian welfare* considers the *mean loss* across a *distribution* over users, which is fair in a different sense. Utilitarian welfare will not ignore the desires of large groups in favor of improving small worst-case groups, though in ML systems it may create positive feedback loops, magnifying preexisting inequity. Rather than argue for a particular welfare concept, a deep philosophical choice which depends on many extrinsic factors, we show that our statistical methods are sufficient to control for generalization error *within a broad class* of welfare concepts.

In general, take *welfare concept $M(\mathbf{c}; \mathcal{L}, \mathcal{D})$* to be a function that takes a *codec set \mathbf{c}* , a *distribution over loss functions \mathcal{L}* , and a *media distribution \mathcal{D}* over \mathcal{X} , and yields a real-valued measure of overall well-being. In

¹Although welfare is usually defined in terms of (desirable) *utility $u(\cdot)$* , we instead consider (undesirable) *loss $\ell(\cdot)$* , and thus in some sense seek to *minimize* a notion of *antiwelfare* over loss. Our methods can be reposed to work in either direction, e.g., by defining utility $u(\mathbf{a}) = 1 - \ell(\mathbf{a})$.

this section, we show all results in full generality, but for intuition, recall the *power mean* definition 3.3.6. We also refer to both *welfare* and *malware* as *welfare*, as whether we wish to minimize or maximize depends on whether codec attributes are positive or negative (e.g., *positive quality* or *negative distortion*). Recall that $M_p(\dots)$ is *monotonic* in p , 1 & ∞ correspond to *utilitarian* and *egalitarian* welfare, respectively, and intermediate p are of interest as *compromises thereof* (e.g., the $p = 2$ *quadratic welfare* lies between *egalitarian* and *utilitarian* welfare, and weighs high-risk users more heavily than low-risk groups). We require only that $M(\cdot)$ is λ -*Lipschitz continuous* in the ℓ_∞ -norm of $\ell \mapsto \mathbb{E}_{x \sim \mathcal{D}}[\min_{c \in \mathcal{C}}(\ell \circ c)(x)]$ (i.e., an ε -perturbation to *all* risk values changes $M(\cdot)$ by $\leq \lambda\varepsilon$). We also assume that $M(\cdot)$ is *monotonic* in the risk of each user, making *Pareto-optimality* a meaningful optimality concept.

Given *loss function family* or *distribution* \mathcal{L} , media distribution \mathcal{D} over \mathcal{X} , and welfare concept $M(\cdot)$, we now quantify the performance of some k codec selection $\mathbf{c} \in \binom{\mathcal{C}}{k}$ as its *welfare* w.r.t. $M(\cdot)$. The M -optimal codec set \mathbf{c}_M^* is then defined as

$$\mathbf{c}_M^* \doteq \underset{\mathbf{c} \in \binom{\mathcal{C}}{k}}{\operatorname{argmin}} M(\mathbf{c}; \mathcal{L}, \mathcal{D}) . \quad (28)$$

4.2.2 Empirical Welfare Optimization

We now consider the task of *approximating* the risk-optimal \mathbf{c}_ℓ^* or welfare-optimal \mathbf{c}_M^* , given only a *sample* $\mathbf{x} \sim \mathcal{D}^m$. We define the *empirical risk* $\hat{R}(\mathbf{c}; \ell, \mathbf{x})$ as the *sample average loss*

$$\hat{R}(\mathbf{c}; \ell, \mathbf{x}) \doteq \hat{\mathbb{E}}_{x \in \mathbf{x}} \left[\min_{c \in \mathcal{C}} (\ell \circ c)(x) \right] , \quad (29)$$

and similarly, define the *empirical risk minimizer* $\hat{\mathbf{c}}_\ell$ as

$$\hat{\mathbf{c}}_\ell \doteq \underset{\mathbf{c} \in \binom{\mathcal{C}}{k}}{\operatorname{argmin}} \hat{R}(\mathbf{c}; \ell, \mathbf{x}) , \quad (30)$$

corresponding to eqs. 26 and 27, respectively.

Similarly, for welfare $M(\cdot)$, we take the *empirical welfare* $\hat{M}(\mathbf{c}; \mathcal{L}, \mathbf{x})$ to be welfare computed with *empirical risk* $\hat{R}(\mathbf{c}; \ell, \mathbf{x})$ instead of *risk* $R(\mathbf{c}; \ell, \mathcal{D})$, and define the *empirical welfare minimizer*, by analogy with equation 28, as

$$\hat{\mathbf{c}}_M \doteq \underset{\mathbf{c} \in \binom{\mathcal{C}}{k}}{\operatorname{argmin}} \hat{M}(\mathbf{c}; \mathcal{L}, \mathbf{x}) . \quad (31)$$

The *Empirical Welfare-Optimal k codec Selection* (EWOKS) algorithm computes $\hat{\mathbf{c}}_M$ as a proxy for \mathbf{c}_M^* . We assume the combinatorial optimization aspect of this task is tractable via enumeration, or by exploiting its inherent *submodularity* (diminishing returns in adding codecs), as usually the computational bottleneck is encoding each \mathbf{x}_i with each $c \in \mathcal{C}$. Although EWOKS is straightforward, the difficulty arises in determining the *minimum sample size* that guarantees that the welfare (over \mathcal{D}) of $\hat{\mathbf{c}}_M$ approaches that of \mathbf{c}_M^* .

In the statistical parlance, the *empirical risk* is an *unbiased estimator* of *risk*, i.e., $\mathbb{E}[\hat{R}(\mathbf{c}; \ell, \mathbf{x})] = R(\mathbf{c}; \ell, \mathcal{D})$, and $\hat{\mathbf{c}}_\ell$ is an *M -estimator* [Huber, 1964]. The *Empirical Risk Minimization* (ERM) paradigm [Vapnik, 1992] states that $\hat{\mathbf{c}}_\ell$ is a reasonable proxy for \mathbf{c}_ℓ^* , and much of *statistical learning theory* quantifies the finite-sample *selection bias* (overfitting) of $\hat{\mathbf{c}}_\ell$. Similarly, with insufficient data, $\hat{\mathbf{c}}_M$ can overfit, in which case the *estimation error* $M(\hat{\mathbf{c}}_M; \mathcal{L}, \mathcal{D}) - M(\mathbf{c}_M^*; \mathcal{L}, \mathcal{D}) \gg 0$, and $\hat{\mathbf{c}}_M$ is a poor proxy for \mathbf{c}_M^* . In general, the *empirical welfare* is a *biased estimator* of *welfare* (i.e., $\mathbb{E}[\hat{M}(\mathbf{c}; \mathcal{L}, \mathbf{x})] \neq M(\mathbf{c}, \mathcal{L}, \mathcal{D})$), however in section 4.2.3 we show tail bounds on the welfare-estimation error, thus controlling for overfitting and providing methods to guarantee that EWOKS solutions are statistically significant.

Note that k is a fixed hyperparameter of the EWOKS algorithm, and should be selected to balance the tradeoff as k increases between the *cost of managing more codecs* and *increased overfitting* against empirical-welfare improvement. As model complexity increases in k , optimizing k to balance the bias-variance tradeoff is statistically-efficient with *structural risk minimization* [Koltchinskii, 2001].

4.2.3 Generalization Analysis

We now show exponential tail bounds on the *welfare optimality-gap* $M(\hat{\mathbf{c}}_M; \mathcal{L}, \mathcal{D}) - M(\mathbf{c}_M^*; \mathcal{L}, \mathcal{D})$. We bound the *supremum deviation* over all *risk values* with novel *variance-sensitive tail bounds* which we then use to bound the welfare optimality-gap. Here it is convenient to discuss a generic function family, so given codec family $\mathbf{c} \subseteq \mathcal{X} \rightarrow \mathbb{R}_{0+}^a$, loss family $\mathcal{L} \subseteq \mathbb{R}_{0+}^a \rightarrow \mathbb{R}_{0+}$, and $k \in \mathbb{N}$, we define

$$\mathcal{F} \doteq \min \circ \mathcal{L} \circ \binom{\mathcal{C}}{k} = \left\{ x \mapsto \min_{\mathbf{c} \in \binom{\mathcal{C}}{k}} (\ell \circ \mathbf{c})(x) \mid \mathbf{c} \in \binom{\mathcal{C}}{k}, \ell \in \mathcal{L} \right\} .$$

\mathcal{F} captures the k codec selection-specific details of the problem, thus we seek to show *uniform convergence* over \mathcal{F} , which we will use to bound the gap between *empirical* and *expected* welfare of EWoks solutions.

Definition 4.2.1 (Rademacher Averages). Suppose *distribution* \mathcal{D} over \mathcal{X} , *sample* $\mathbf{x} \in \mathcal{X}^m$, n *Monte-Carlo trials*, and Rademacher matrix $\boldsymbol{\sigma} \in (\pm 1)^{n \times m}$. We define

1. for fixed $\boldsymbol{\sigma}$ and \mathbf{x} , the *n-Monte-Carlo Empirical Rademacher Average* (*n-MCERA*):

$$\hat{\mathfrak{R}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) \doteq \frac{1}{n} \sum_{j=1}^n \sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_{j,i} f(\mathbf{x}_i) \right| ;$$

2. averaging over $\boldsymbol{\sigma} \sim$ i.i.d. Rademacher (uniform on ± 1), the *Empirical Rademacher Average* (ERA):

$$\hat{\mathfrak{R}}_m(\mathcal{F}, \mathbf{x}) \doteq \mathbb{E}_{\boldsymbol{\sigma}} [\hat{\mathfrak{R}}_m^1(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma})] ; \&$$

3. averaging over $\mathbf{x} \sim \mathcal{D}^m$, the *Rademacher Avg.* (RA):

$$\mathfrak{R}_m(\mathcal{F}, \mathcal{D}) \doteq \mathbb{E}_{\mathbf{x}} [\hat{\mathfrak{R}}_m(\mathcal{F}, \mathbf{x})] .$$

The *n-MCERA* is a *Monte-Carlo* estimate (w.r.t. $\boldsymbol{\sigma}$) of the ERA, and the ERA is a *sample estimate* (w.r.t. \mathbf{x}) of the RA. Recall that through symmetrization, we have

$$\mathbb{E}_{\mathbf{x}} \left[\underbrace{\sup_{f \in \mathcal{F}} \left| \mathbb{E}_{\mathcal{D}}[f] - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right|}_{\text{Supremum Deviation}} \right] \leq 2\mathfrak{R}_m(\mathcal{F}, \mathcal{D}) , \quad (32)$$

where the *supremum deviation* (SD) *uniformly bounds* the gap between *empirical* and *expected* values of each $f \in \mathcal{F}$ *simultaneously*.

While RAs control the *expected* SD, and boundedness is sufficient to obtain sharp tail bounds on the convergence of *n-MCERA* and ERA to the RA, sharp bounds on the SD require also that the *variance* of each $f \in \mathcal{F}$ is controlled. We define the *wimpy variance* (see Boucheron et al. [2013, ch. 11]), and our novel estimate, the *empirical wimpy variance* as

$$v \doteq \sup_{f \in \mathcal{F}} \mathbb{V}_{\mathcal{D}}[f] \quad \& \quad \hat{v} \doteq \sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathcal{D}}[f])^2 ,$$

respectively. We now show bounds that quantify the relationships between the SD, *n-MCERA*, v , and \hat{v} in the following novel theorem, which enjoys the asymptotic improvements of *variance-sensitive bounds* [Bousquet, 2002] without assuming *a priori* variance knowledge. In particular, our bound excels in the small-sample setting for bounded functions of uniformly low variance, as occurs in the codec selection problem.

Theorem 4.2.2 (Variance-Sensitive Bounds). Suppose *distribution* \mathcal{D} over \mathcal{X} , *training sample* $\mathbf{x} \sim \mathcal{D}^m$, *Monte-Carlo trial count* n , i.i.d. Rademacher matrix $\boldsymbol{\sigma} \in (\pm 1)^{n \times m}$, codec family \mathcal{C} , loss family $\mathcal{L} \subseteq \mathbb{R}_{0+}^a \rightarrow [-r, r]$, λ -Lipschitz *welfare* $M(\cdot)$, and codec selection size k . $\forall \delta \in (0, 1)$, let $\mathcal{F} \doteq \min \circ \mathcal{L} \circ \binom{\mathcal{C}}{k}$, take \hat{v} to be the *empirical wimpy variance* of \mathcal{F} over \mathbf{x} and \hat{v}^{raw} to be the *empirical raw wimpy variance*, and take

$$\begin{aligned}
1. \quad \varepsilon_{\text{ERA}} &\doteq \frac{4r^{\frac{4}{\delta}}}{3nm} + \sqrt{\frac{4\hat{v}^{\text{raw}} \ln \frac{4}{\delta}}{nm}}; & 3. \quad \tilde{\mathfrak{K}} &\doteq \hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}} + \varepsilon_{\text{RA}}; \\
2. \quad \varepsilon_{\text{RA}} &\doteq \frac{2r \ln \frac{4}{\delta}}{m} + \sqrt{\frac{2r(\hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}) \ln \frac{4}{\delta}}{m}}; & 4. \quad \tilde{v} &\doteq \frac{m}{m-1} \hat{v} + \frac{4r^2 \ln \frac{4}{\delta}}{m} + \sqrt{\frac{2r^2 \hat{v} \ln \frac{4}{\delta}}{m-1}}; \& \\
& & 5. \quad \varepsilon_{\text{SD}} &\doteq \frac{r \ln \frac{4}{\delta}}{3m} + \sqrt{\frac{2(\tilde{v} + 4r\tilde{\mathfrak{K}}) \ln \frac{4}{\delta}}{m}}.
\end{aligned}$$

It then holds with with pr. $\geq 1 - \delta$ over $\mathbf{x}, \boldsymbol{\sigma}$ that

$$\begin{aligned}
1. \quad & \left| \text{M}(\hat{\mathbf{c}}_{\text{M}}, \mathcal{L}, \mathcal{D}) - \hat{\text{M}}(\hat{\mathbf{c}}_{\text{M}}, \mathcal{L}, \mathbf{x}) \right| && \leq 2\lambda\tilde{\mathfrak{K}} + \lambda\varepsilon_{\text{SD}} \quad \& \\
2. \quad & \text{M}(\hat{\mathbf{c}}_{\text{M}}, \mathcal{L}, \mathcal{D}) - \text{M}(\mathbf{c}_{\text{M}}^*, \mathcal{L}, \mathcal{D}) && \leq 4\lambda\tilde{\mathfrak{K}} + 2\lambda\varepsilon_{\text{SD}} \quad .
\end{aligned}$$

The details of the proof are rather cumbersome, but the key idea is that we take a union bound over 4 probabilistic upper-tail-bounds; in particular, w.h.p., $v \leq \tilde{v}$ (1 tail bound), and $\mathfrak{K}_m(\mathcal{F}, \mathcal{D}) \leq \tilde{\mathfrak{K}}$ (2 tail bounds for n -MCERA \rightarrow ERA \rightarrow RA). The fourth tail bound is on the SD of \mathcal{F} , which by Lipschitz properties yields bounds on welfare-gaps. Note that if we assume $\lambda = 1$, $4\tilde{\mathfrak{K}} + 2\varepsilon_{\text{SD}}$ is

$$4\hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \Theta\left(\frac{r \ln \frac{1}{\delta}}{m} + \sqrt{\frac{(v + r\mathfrak{K}_m(\mathcal{F}, \mathcal{D})) \ln \frac{1}{\delta}}{m}}\right). \quad (33)$$

With worst-case variance $v = r^2/4$, we recover the slow-decaying $\Theta\sqrt{r^2 \ln \frac{1}{\delta} / m}$ McDiarmid terms of standard methods, however generally we get asymptotic *mixed-rate convergence*, where the $\Theta(r \ln \frac{1}{\delta} / m)$ term depends on the scale r , but *decays quickly* in m , and the $\Theta\sqrt{(v + r\mathfrak{K}_m(\mathcal{F}, \mathcal{D})) \ln \frac{1}{\delta} / m}$ term depends on \sqrt{v} instead of r , but *decays slowly* in m . Note that the latter term simplifies to $\Theta\sqrt{v \ln \frac{1}{\delta} / m}$ when we consider the limiting behavior as $m \rightarrow \infty$ and assume $\mathfrak{K}_m(\mathcal{F}, \mathcal{D})$ tends to 0, thus the essence of the bound is *fast-decay* as r/m plus *slow decay* as $\sqrt{v/m}$.

This dichotomy is key to understanding the bound's strong performance in the small sample setting: we see initial *rapid decay* while the r term is dominant, followed by *slow decay* as the \sqrt{v} term comes to dominate. Where the transition occurs is primarily dictated by the relative values of r and v , but for sufficiently small v , we see fast-decaying tail bounds and excellent performance in the small-sample regime (see section 4.3.3).

4.2.4 Linear Loss Families

For large \mathcal{C} and k , it can become computationally intractable to compute the wimpy variance and n -MCERA, and in general they are rather opaque, due to the supremum over codec combinations and the minimum over selected codecs. We now show simple *compositional bounds* for *linear loss families* mapping *attribute vectors* $\mathbf{a} \in \mathbb{R}_{0+}^a$ onto \mathbb{R}_{0+} , taking

$$\mathcal{L}_p \doteq \{ \ell(\mathbf{a}) \doteq \mathbf{w} \cdot \mathbf{a} \mid \mathbf{w} \in \mathbb{R}_{0+}^a \text{ s.t. } \|\mathbf{w}\|_p \leq 1 \} \quad (34)$$

for $p \geq 1$. To control the *range* of \mathcal{L}_p , we assume also a *bounded domain* \mathcal{X} , in particular we assume bounded q -norm for $q \doteq \frac{p-1}{p}$, or $q = 1$ for $p = \infty$, thus $\|\cdot\|_p$ and $\|\cdot\|_q$ are *dual norms* (i.e., by Hölder's inequality, $|\mathbf{w} \cdot c(x)| \leq \|\mathbf{w}\|_p \|c(x)\|_q$). We now show that both the variance and ERA of the (nonlinear) k codec selection problem can be bounded with their (linear) counterparts in 1 codec selection.

Theorem 4.2.3 (Bounds for Linear Loss Families). Suppose p, q , s.t. $|p^{-1}| + |q^{-1}| = 1$, *codec family* \mathcal{C} , \mathcal{L}_p as in (34), and assume $\sup_{c \in \mathcal{C}, x \in \mathcal{X}} \|c(x)\|_q \leq s$. For all $k \in 1, \dots, |\mathcal{C}|$, take $\mathcal{F}_k \doteq \min \circ \mathcal{L}_p \circ \binom{\mathcal{C}}{k}$, and \hat{v}_k to be the *empirical wimpy variance* of \mathcal{F}_k on *sample* \mathbf{x} . Taking $\hat{\mathbb{C}}_{\mathbf{x}}[c] \in \mathbb{R}^{a \times a}$ to denote the *empirical covariance matrix* of the attributes of codec c over \mathbf{x} , and $\|\cdot\|_{p \rightarrow q}$ to be the ℓ_p - ℓ_q *operator norm*, it holds that

$$\hat{v}_1 = \sup_{c \in \mathcal{C}} \sup_{\ell \in \mathcal{L}} \hat{\mathbb{V}}[\ell \circ c] \leq \sup_{c \in \mathcal{C}} \|\hat{\mathbb{C}}_{\mathbf{x}}[c]\|_{q \rightarrow p}, \quad \&$$

$$\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) \leq \frac{1}{n} \sum_{j=1}^n \sup_{c \in \mathcal{C}} \left\| \frac{1}{m} \sum_{i=1}^m \boldsymbol{\sigma}_{j,i} c(\mathbf{x}_i) \right\|_q .$$

Furthermore, for any $a, b \in \mathbb{N}$ s.t. $k = a + b$, we have

$$\hat{v}_k \leq \hat{v}_a + \hat{v}_b \leq k \hat{v}_1 , \ \&$$

$$\hat{\mathbf{R}}_m(\mathcal{F}_k, \mathbf{x}) \leq \hat{\mathbf{R}}_m(\mathcal{F}_a, \mathbf{x}) + \hat{\mathbf{R}}_m(\mathcal{F}_b, \mathbf{x}) \leq k \hat{\mathbf{R}}_m(\mathcal{F}_1, \mathbf{x}) .$$

Here we see that linearity of \mathcal{L}_p is convenient, though it *does not* in general convert the EWoks problem to a linear problem, due to the *minimum* over selected codecs. However, \mathcal{F}_1 is completely linear, thus for $k = 1$ it is easy to compute Rademacher averages and variances. With sufficient computational resources, $\hat{\mathbf{R}}_m(\mathcal{F}_k, \mathbf{x})$ and \hat{v}_k can be computed as precisely as required, and theorem 4.2.3 allows us to trade *statistical efficiency* for *computational efficiency* by truncating the computation at some selection size $j < k$, and then loosely bounding $\hat{\mathbf{R}}_m(\mathcal{F}_k, \mathbf{x})$ and \hat{v}_k in terms of $\hat{\mathbf{R}}_m(\mathcal{F}_j, \mathbf{x})$ and \hat{v}_j , with particularly efficient computation for $j = 1$. We now illustrate the use of theorem 4.2.3 with a simple corollary for the ℓ_2 linear loss family.

Corollary 4.2.4 (Bounds for Linear Loss Families). Suppose as in theorem 4.2.3, and also $p = q = 2$, and all attribute values are contained by the unit sphere. Then

$$\hat{v}_k \leq k \sup_{c \in \mathcal{C}} \left\| \hat{\mathbb{C}}[c] \right\|_{\mathbf{x}} \Big|_{2 \rightarrow 2} ,$$

where $\|\cdot\|_{2 \rightarrow 2}$ is the *spectral norm* (largest eigenvalue). Additionally, with pr. $\geq 1 - \delta$ over $\mathbf{x}, \boldsymbol{\sigma}$, it holds that

$$\begin{aligned} \hat{\mathbf{R}}_m(\mathcal{F}_k, \mathbf{x}) &\leq k \left(\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \sqrt{\frac{\ln \frac{1}{\delta}}{2nm}} \right) \\ &\leq k \left(\frac{1}{n} \sum_{i=1}^n \sup_{c \in \mathcal{C}} \left\| \frac{1}{m} \sum_{j=1}^m \boldsymbol{\sigma}_{i,j} c(\mathbf{x}_j) \right\|_2 + \sqrt{\frac{\ln \frac{1}{\delta}}{2nm}} \right) . \end{aligned}$$

Consequently, theorem 4.2.2 holds with

$$\begin{aligned} \varepsilon_{\text{RA}} &\leq \frac{\ln \frac{4}{\delta}}{3m} + \sqrt{\frac{2k(\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}) \ln \frac{4}{\delta}}{m}} ; \\ \hat{\mathbf{R}} &\leq k(\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}) + \varepsilon_{\text{RA}} ; \ \& \\ \tilde{v} &\leq \frac{m}{m-1} k \hat{v}_1 + \frac{2 \ln \frac{4}{\delta}}{m} + \sqrt{\frac{2k \hat{v}_1 \ln \frac{4}{\delta}}{m-1}} . \end{aligned}$$

Here we see that, for linear loss family \mathcal{L}_2 , we can bound the wimpy variance by computing the *maximum variance* along any *unit vector* (much like in PCA or SVD). Similarly, bounding the ERA is straightforward as well, as each trial of the n -MCERA essentially corresponds to measuring the maximum (over codecs) *distance traveled* (ℓ_2 norm) in *attribute space* over a *random walk* (directions dictated by each $\boldsymbol{\sigma}_{i,j}$). Note that together, corollary D.3.6 & theorem 4.2.2 imply that for 1-Lipschitz welfare $M(\cdot)$, $r \in \Theta(1)$, with pr. $\geq 1 - \delta$, we have $M(\hat{c}_M, \mathcal{L}, \mathcal{D}) - M(c_M^*, \mathcal{L}, \mathcal{D}) \leq$

$$4k \hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \Theta \left(\frac{\ln \frac{1}{\delta}}{m} + \sqrt{\frac{k(v_1 + \hat{\mathbf{R}}_m(\mathcal{F}_1, \mathcal{D})) \ln \frac{1}{\delta}}{m}} \right) . \quad (35)$$

4.3 Experimental Evaluation of Ewoks

We study various notions of multivariate optimality on four datasets, spanning speech and several music genres. The *Debussy* dataset consists of the complete orchestral works of Claude Debussy, conducted by Yan

Type	Codec #	Parameter	Distortion
VBR	c ₀ -c ₉	0, 1, ..., 9	Low-Med
ABR	c ₁₀ -c ₁₃	320, 256, 128, 64	Low-Med
VBR	c ₁₄ -c ₁₆	9.9, 9.99, 9.999	Med-High
ABR	c ₁₇ -c ₁₉	32, 16, 8	Med-High
WAV	c ₂₀	—	None

Table 4.1: We employ *Variable Bit Rate* (VBR) mp3, parameterized by *quality*, *Average Bit Rate* (ABR) mp3, parameterized by *bitrate*, and uncompressed wav codecs.

Pascal Tortelier; the *Explosions* dataset is the 2000–2013 discography of progressive rock band *Explosions in the Sky*; the *Zeppelin* dataset is the complete discography of hard rock band *Led Zeppelin*; and the *LibriVox* dataset contains public-domain LibriVox audiobooks, all cut into short nonoverlapping audio samples. In some experiments, we pool all music datasets or all four datasets to evaluate EWOKS on *heterogeneous data*. Table 4.1 describes the codecs c₀, c₁, ..., c₂₀ used in all experiments.

We quantify distortion in the audio domain with the *Perceptual Evaluation of Audio Quality* (PEAQ) [Thiede et al., 2000] metric, which is an *objective* measure of how well one audio sample approximates another, based on psychoacoustic principles, that aims to measure the degree of human-perceived difference. In all experiments, we consider two attributes, the first being PEAQ distortion, normalized to [0, 1], and the second being the *compression ratio*, also in [0, 1].

4.3.1 Data-Dependence and Pareto Optimality

A codec set $\mathbf{c} \in \binom{\mathcal{C}}{k}$ is said to be *Pareto-optimal* among $\binom{\mathcal{C}}{k}$ if there exists a *linear loss function* (nonnegative linear combination of attributes; see equation 34) $\ell \in \mathcal{L}_2$ for which the risk (equation 26) of \mathbf{c} is minimal among $\binom{\mathcal{C}}{k}$, i.e.,

$$\text{Pareto}(\binom{\mathcal{C}}{k}, \mathcal{D}) \doteq \bigcup_{\ell \in \mathcal{L}_2} \underset{\mathbf{c} \in \binom{\mathcal{C}}{k}}{\text{argmin}} \mathbb{R}(\mathbf{c}; \ell, \mathcal{D}) . \quad (36)$$

For $k = 1$, the *Pareto-optimal frontier* of \mathcal{C} in \mathbb{R}_{0+}^d is then the *lower convex-hull*² LCH(\cdot) of the set of *mean attribute values* of all Pareto-optimal codecs, i.e.,

$$\text{PFrnt}(\mathcal{C}, \mathcal{D}) \doteq \text{LCH} \left\{ \mathbb{E}_{x \sim \mathcal{D}} [c(x)] \mid c \in \mathcal{C} \right\} . \quad (37)$$

This notion is easily visualized (see figure 4.1), and it concisely represents the set of *mean attribute values* (and thus risk values) of $\mathbf{c}_\ell^* \in \text{CH}(\mathcal{C})$ for any $\ell \in \mathcal{L}_2$.

To generalize the *Pareto-optimal frontier* to $k > 1$, first observe that the *mean attribute values* of a k codec set depend on $\ell \in \mathcal{L}_2$, as EWOKS takes a *minimum* over k codecs. Consequently, each $\mathbf{c} \in \binom{\mathcal{C}}{k}$ is associated with a *mean-attribute set*, and we take the Pareto-optimal frontier of $\binom{\mathcal{C}}{k}$ to be the *lower convex hull* of the *union* of the mean-attribute sets of each $\mathbf{c} \in \binom{\mathcal{C}}{k}$. Algebraically, $\text{PFrnt}(\binom{\mathcal{C}}{k}, \mathcal{D})$ is the LCH of

$$\bigcup_{\mathbf{c} \in \binom{\mathcal{C}}{k}} \left\{ \mathbb{E}_{x \sim \mathcal{D}} \left[\underset{c(x) \in \mathbf{c}(x)}{\text{argmin}} \ell(c(x)) \right] \mid \ell \in \mathcal{L}_2 \right\} . \quad (38)$$

Observe that equation 38 generalizes equation 37, as for $k = 1$ and $\mathcal{C} = \{c\}$, we have $\underset{c(x) \in \mathbf{c}(x)}{\text{argmin}} \ell(c(x)) = c(x)$. Note that the *subadditive nonlinearity* of the argmin (i.e., the improvement from selecting the *best* among k codecs) improves $\text{PFrnt}(\binom{\mathcal{C}}{k}, \mathcal{D})$ *monotonically* in k . Again the Pareto-optimal frontier concisely visualizes the possible *mean attribute values* (and thus *risks*) of the optimal $\mathbf{c} \in \binom{\mathcal{C}}{k}$ for any $\ell \in \mathcal{L}_2$.

Note that, like the welfare concepts of section 4.2.1, Pareto-optimality is a notion of multivariate optimality. As with welfare, we may define the *empirical Pareto-optimal set* and the *empirical Pareto-optimal front*,

²NB the *convex hull* $\text{CH}(\mathbf{a})$ of pointset $\mathbf{a} \subseteq \mathbb{R}^a$ is the *convex closure* of \mathbf{a} , and the *lower convex hull* is $\text{LCH}(\mathbf{x}) \doteq \{\mathbf{a}' \in \text{CH}(\mathbf{a}) \mid \forall \varepsilon > \mathbf{0}: \mathbf{a}' - \varepsilon \notin \text{CH}(\mathbf{a})\}$.

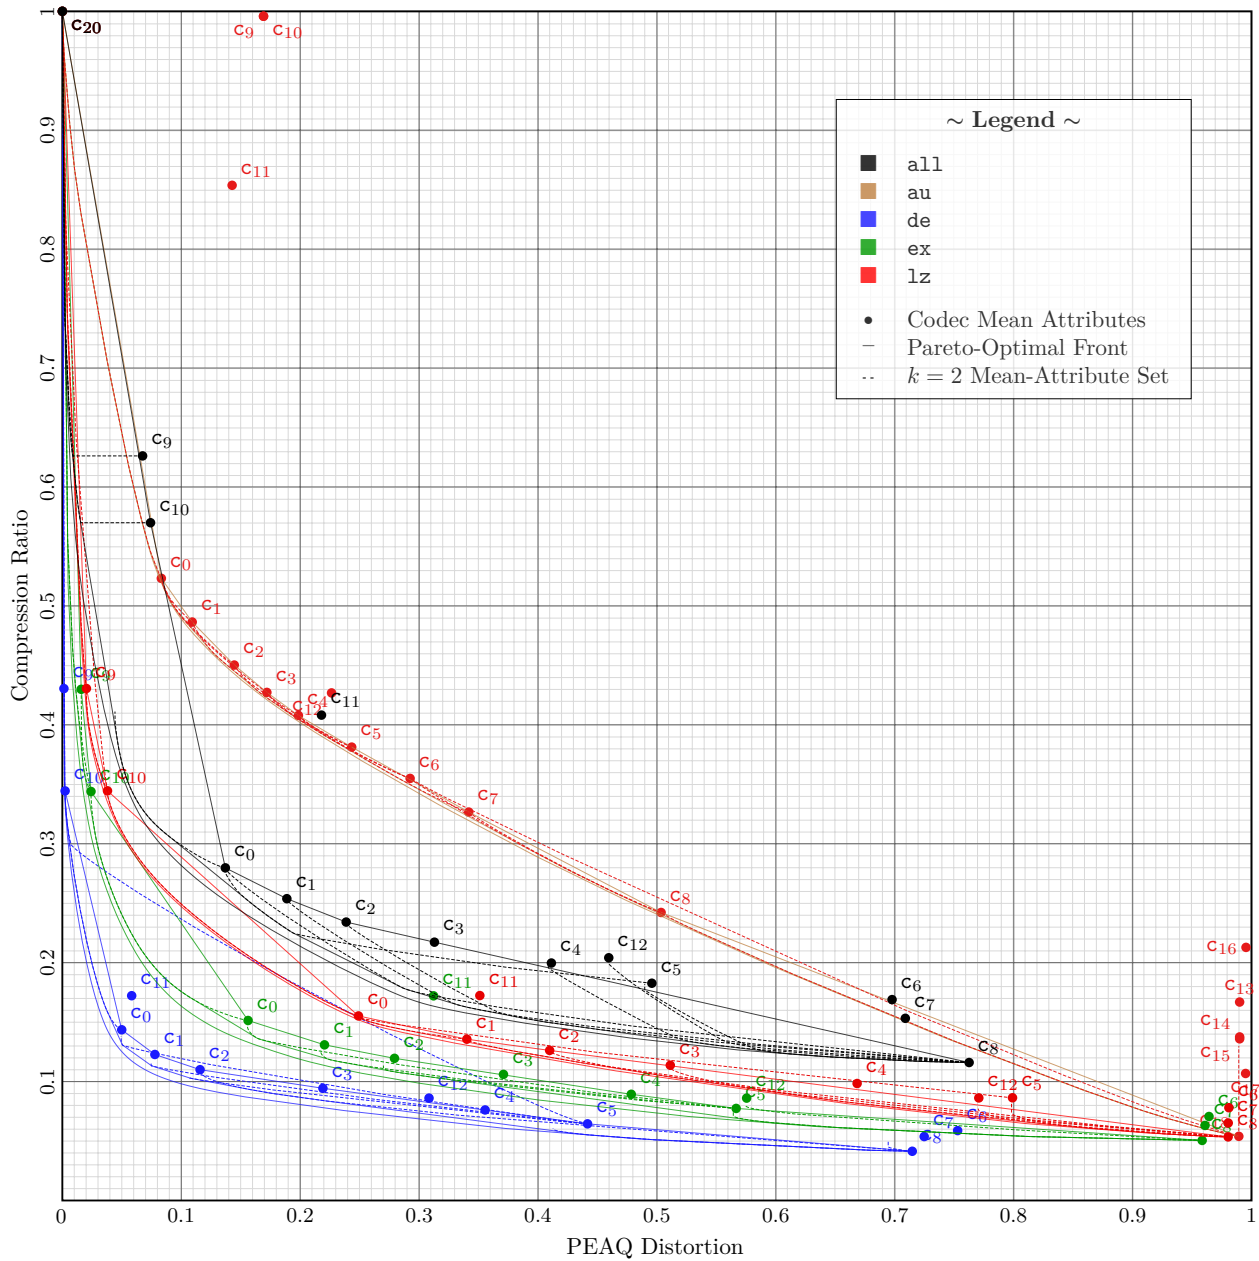


Figure 4.1: Pareto-optimal frontiers for all datasets. Mean distortion (x axis) and compression ratio (y axis) shown for each codec \bullet , with Pareto-optimal frontiers for codec families of sizes $k \in \{1, 2, |\mathcal{C}|\}$, color-coded by dataset. Curves for each k are unambiguous, since k increases from top-right to bottom-left. Mean-attribute sets for Pareto-optimal $k = 2$ pairs are shown in the inset as dotted curves.

Table 4.2: Experiments optimizing $\ell_{1/2}$ loss, and $M_1(\cdot)$, $M_\infty(\cdot)$, and $M_2(\cdot)$ welfares on au. **a)** *Objective values*, and **b)** *regret* for each row-objective, when optimizing \hat{c}_C for each column-objective o_C .

a) Objective Values:					b) $k = 3$ Regret Matrix:				
k	$\ell_{1/2}$	M_1	M_∞	M_2		$\ell_{1/2}$	M_1	M_∞	M_2
1	.297	.297	.341	.309	$\ell_{1/2}$	0	.0098	.0097	.0015
2	.294	.217	.328	.234	M_1	.1343	0	.1221	.0189
3	.293	.208	.327	.229	M_∞	.0811	.0374	0	.0103
$ \mathcal{C} $.293	.202	.326	.223	M_2	.0607	.0035	.0454	0

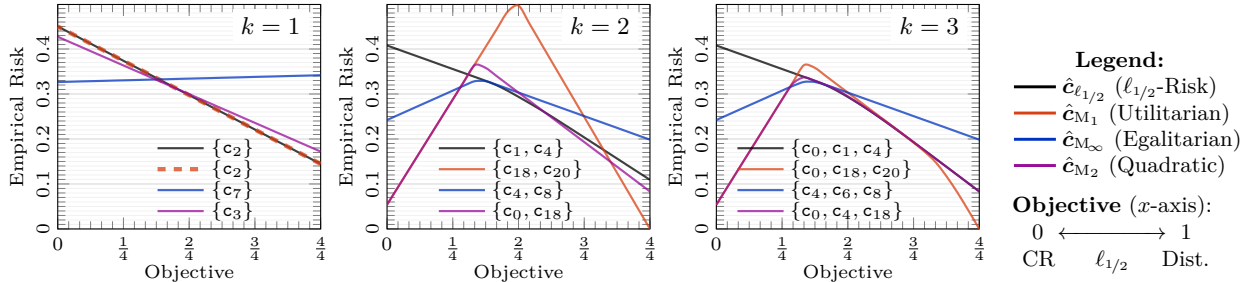


Figure 4.2: Empirical risk (y -axis) plotted against linear objective family \mathcal{L}_1 (x -axis), with *distortion weight* $\mathbf{w}_1 = x$ and *compression ratio weight* $\mathbf{w}_2 = 1 - x$, on au, for $k \in \{1, 2, 3\}$. Empirically-optimal codec sets $\hat{c}_o \in \binom{\mathcal{C}}{k}$ for each objective o are also given.

w.r.t. sample \mathbf{x} , by replacing $\mathbb{E}_{\mathcal{D}}[\cdot]$ with $\hat{\mathbb{E}}_{\mathbf{x}}[\cdot]$. Furthermore, by its inherently linear nature, it is trivial to adapt the uniform convergence guarantees of theorem 4.2.3 to bound the gap between empirical and true Pareto-optimality concepts.

In figure 4.1, we plot the empirical mean PEAQ distortion and compression ratios of codecs, and empirical Pareto-optimal frontiers of codec sets of sizes $k \in \{1, 2, |\mathcal{C}|\}$ on the music, speech, and mixed datasets. Immediately we see that the Pareto-optimal frontiers differ between datasets, illustrating the importance of data-dependent codec selection. We see that on the speech and music datasets, the $k = 1$ (top-right most) Pareto-optimal frontier usually performs similarly to the $k = |\mathcal{C}|$, which indicates little variability in codec performance over these distributions, thus a static codec for each objective performs reasonably well. However, for the *mixed dataset*, we see a massive improvement in moving from $k = 1$ to $k = 2$ codecs, and modest diminishing returns for $k > 2$. This makes sense, as for fixed objectives, different codecs are often optimal for the LibriVox and music datasets, thus combining them is optimal in the mixed dataset.

4.3.2 Fairness and Welfare-Optimality

The above Pareto-optimality experiments indicate that *for any fixed objective*, there exists a *small codec set* that is near-optimal (i.e., $\text{PFrnt}(\binom{\mathcal{C}}{2}, \mathcal{D}) \approx \text{PFrnt}(\binom{\mathcal{C}}{|\mathcal{C}|}, \mathcal{D})$) However, we see in the inset that *individual pairs* of codecs (dotted lines) only achieve Pareto-optimality for *small ranges* of \mathcal{L} , thus $k > 2$ codecs may be required to well-optimize many loss functions *simultaneously*; consequently, we expect various *welfare concepts* to continue improving beyond $k = 2$. In table 4.2 & figure 4.2, we study the effect of increasing k on the *welfare* of various EWOKS-optimal solution concepts, which is sensitive to performance *across all of* \mathcal{L} .

In this experiment, we consider the 2-dimensional ℓ_1 -linear loss family $\mathcal{L}_1 = \{\mathbf{a} \mapsto \mathbf{w} \cdot \mathbf{a} : \mathbf{w}_1 \in [0, 1], \mathbf{w}_2 = 1 - \mathbf{w}_1\}$ (see eq 34), where \mathbf{w}_1 and \mathbf{w}_2 represent the penalty-weight a user places on *distortion* and *compression ratio*, respectively. We optimize the neutral single linear objective (see eq 34) $\ell_{1/2}(\mathbf{a}) \doteq (1/2, 1/2) \cdot \mathbf{a}$, *utilitarian welfare* $M_1(\cdot)$ on the loss distribution over \mathcal{L}_1 with density $\propto |\mathbf{w}_1 - 1/2|^2$ (i.e., the user distribution is biased towards the extremes), as well as *egalitarian welfare* $M_\infty(\cdot)$ over \mathcal{L}_1 and *quadratic welfare* $M_2(\cdot)$ with uniform density over \mathcal{L}_1 .

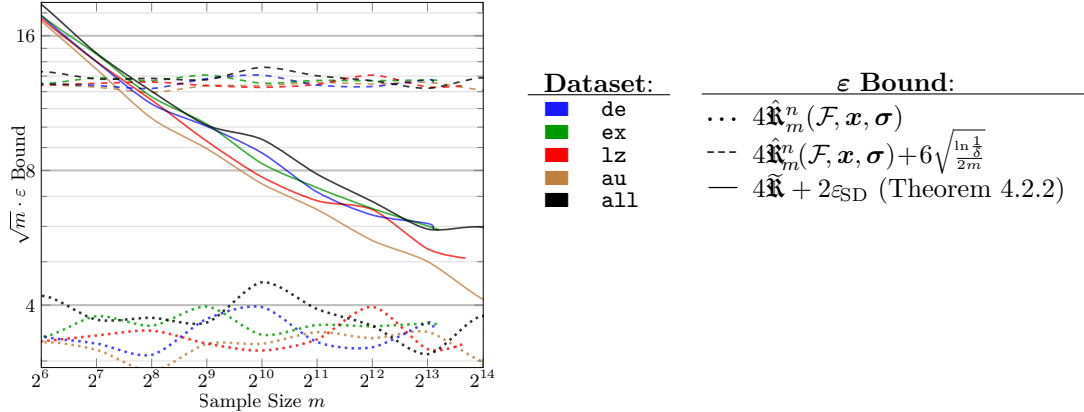


Figure 4.3: Log-log plots of EWOks \sqrt{m} -scaled $M(\cdot)$ -optimality gap bound $\sqrt{m}\varepsilon$ (x -axis), versus sample size m (y -axis), in expectation with $4\hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma})$, and w.h.p. with McDiarmid and theorem 4.2.2.

Table 4.2a shows the raw empirical objective values of each optimality concept, for $k \in \{1, 2, 3, |\mathcal{C}|\}$, and table 4.2b shows the *regret* of (row) objective o_R on $\hat{\mathbf{c}}_C$ optimized for each (column) objective o_C , which is the amount by which objective o_R would *improve* if $\hat{\mathbf{c}}_C$ were optimized for o_R instead of o_C .

To better understand the tradeoffs made by EWOks under various objectives, for each objective o , figure 4.2 plots the *empirical risk* (see equation 29) of each o -optimal k codec selection $\hat{\mathbf{c}}_o$ w.r.t. each $\ell \in \mathcal{L}_1$ as a function of *distortion weight* \mathbf{w}_1 . In particular, each plot corresponds to a *selection size* k , and each line represents the *empirical risk* w.r.t. the linear loss function with $\mathbf{w}_1 = x$, i.e., from left to right we interpolate from *pure distortion* to *pure compression ratio* losses.

From the *objective values* matrix in table 4.2, we see that the single-objective $\ell_{1/2}$ experiences rapidly diminishing returns as k increases, improving by only 0.004 from $k = 1$ to $k = |\mathcal{C}|$, whereas the welfare objectives gradually improve with increasing k . This is as expected, because a sufficiently diverse set of codecs must be selected so as to perform reasonably well *across all* $\ell \in \mathcal{L}$, rather than *for some fixed* ℓ .

This experiment also clearly illustrates the *fairness impact* of objective choice, as we see that *improving utilitarian welfare* $M_1(\cdot)$ can *decrease egalitarian welfare* $M_\infty(\cdot)$, as occurs for $k = 2$, when $\hat{\mathbf{c}}_{M_1} = \{\mathbf{c}_{18}, \mathbf{c}_{20}\}$, which are the second-lowest bitrate and uncompressed codecs, respectively, resulting in high regret (0.1343) for $M_\infty(\cdot)$ (see table 4.2a). This is unsurprising, as it is easy to produce low distortion *or* low compression ratio encodings; the difficulty lies in optimizing both *simultaneously*. The regret matrix (table 4.2b) also shows fairness tradeoffs; here we see that the quadratic-welfare-optimal $\hat{\mathbf{c}}_{M_2}$ is much better on-average than $\hat{\mathbf{c}}_{M_\infty}$ (utilitarian regret 0.0189 versus 0.1121), despite worst-case performance only 0.01 worse than $\hat{\mathbf{c}}_{M_\infty}$. This confirms that while optimizing the *worst-case* $M_\infty(\cdot)$ is at-odds with optimizing the *average-case* $M_1(\cdot)$, the quadratic $M_2(\cdot)$ is a reasonable compromise.

4.3.3 Uniform Convergence Bounds

We now show that our uniform convergence bounds yield valuable conclusions as to the approximate optimality of EWOks solutions, and outperform McDiarmid bounds on real data. In figure 4.3, we plot EWOks optimality-gap bounds against sample size m , assuming $M(\cdot)$ is 1-Lipschitz, and using selection-size $k = 3$, loss family \mathcal{L}_1 (see (34)), $n = 100$ Monte-Carlo trials, and tail bounds with failure probability $\delta = 0.01$. Here all bounds are scaled by \sqrt{m} to visualize the *deviation* of rates from $\Theta\sqrt{1/m}$, which appears flat under \sqrt{m} -scaling.

We immediately see that our variance-sensitive bounds (theorem 4.2.2) are significantly sharper than the McDiarmid bounds, yielding valuable conclusions after only $m \approx 1000$ samples. The *mixed convergence rates* (see theorem 4.2.2) of these bounds are visible as initial *fast-decay* (*negative slopes*), straightening out as

$m \rightarrow \infty$ and the asymptotic $\Theta \sqrt{v/m}$ slow rate is reached. In contrast, the McDiarmid bounds, and the n -MCERAs themselves, all exhibit slow $\Theta \sqrt{1/m}$ rates (even slopes). In particular, the poor performance of the McDiarmid bounds occurs because they are always dominated by their concentration term $6\sqrt{\ln \frac{1}{\delta}}/2m$, visible as a large parallel gap between them and the n -MCERAs.

We see also from their relative order that **au** is the easiest to learn, and **all** the most difficult. This is unsurprising; in music and mixed data, we expect more *variability* (thus *variance*), and also more benefit (overfitting) with $k > 1$ (thus higher \mathfrak{K}).

4.4 Discussion

We propose a novel media streaming strategy, where a small set of codecs \mathbf{c} is selected such that the set includes a “good” solution for each $\ell \in \mathcal{L}$ over \mathcal{D} . We develop the *Empirical Welfare-Optimal k codec Selection* (EwoKs) algorithm, which *learns* a data-dependent set of k codecs \mathbf{c} that is *welfare-optimal* over \mathcal{L} . The *selection size k* balances between the cost of *compressing / storing* multiple encodings of media, and the improvement to user experience and fairness.

Our experiments show the importance of considering both the objective *and* the data-distribution. We see that while k codec selection yields modest improvement for *single-objectives* (section 4.3.1), particularly for heterogenous (mixed speech / music) data, it greatly improves *welfare-objectives* over *loss families* (section 4.3.2). Our experiments confirm that bias-issues exist in codec selection, and exhibit several surprising phenomena (e.g., optimizing a single neutral loss function can yield bias issues, learning a more sophisticated [higher k] model for one welfare-concept can harm other welfare concepts), further motivating rigorous analysis and provable methods for controlling algorithmic bias.

We rigorously analyze the generalization error of EwoKs (section 4.2.3), applying *Rademacher averages* with a novel variance-sensitive bound for welfare-generalization. We handle arbitrary Lipschitz welfare functions, whereas previous applications gave weaker bounds limited to simple cases, such as *utilitarian welfare* in *auctions* [Hoy et al., 2017]. Since general welfare-concepts depend nonlinearly on many loss values *simultaneously*, uniform convergence is an ideal tool to analyze them. Our experiments (section 4.3.3) also show our bounds significant improvement traditional methods in the small-sample domain.

Our data-driven codec-selection approach follows trends in machine learning and databases. We replace *fixed codecs* with *learned codecs*, in much the same way that autoML systems replace fixed models with learned pipelines [Hutter et al., 2011], and database systems replace static indices with learned indices [Kraska et al., 2018]. As this trend continues, and more critical services integrate learned components, the importance of fairness continues to grow; we are confident that our welfare analysis methods for provable fairness can be applied beyond the domain of codec selection.

Chapter 5

Fair Probably Approximately Correct Learning

Building upon the concept of malfare of chapter 3, I define *fair-PAC learning*, where a fair PAC-learner is an algorithm that learns an ε - δ malfare-optimal model with bounded sample complexity, for *any data distribution*, and for *any* (axiomatically justified) malfare concept, i.e., power-mean malfare with $p \in [1, \infty)$. This definition extends Valiant’s [Valiant, 1984] classic PAC-learning formalization of machine learning, and I show that, with appropriate modifications, many (standard) PAC-learners may be converted to fair-PAC learners, and argue that fair-PAC-learners are intuitive and easy to use, as one must only select an axiomatically justified malfare concept (encoding their desired fairness concept), hypothesis class, and confidence guarantees, and then receives a provably ε - δ optimal model. I show broad conditions under which, with appropriate modifications, many standard PAC-learners may be converted to fair-PAC learners. I also argue that fair-PAC-learners are intuitive and easy to use, as one must only select a malfare concept (encoding their desired fairness concept), hypothesis class, and confidence guarantees, and then one receives a provably (probabilistically) near-optimal model.

This is in contrast to similar frameworks, e.g., the *Seldonian learner* of [Thomas et al., 2019], where (1), the arbitrarily statistically infeasible task of constrained optimization must be addressed, and (2), learning guarantees are in terms of less-interpretable *constrained objectives* (e.g., specified objectives, parity constraints, and tolerance values), rather than singular cardinal *malfare objectives*, which directly encode an interpretable single-parameter fairness concept that ranges between *average case* ($p = 1$) and *worst case* ($p = \infty$) performance across groups. This places fair-PAC learning on firm theoretical ground, as it yields statistical, and in some cases computational, efficiency guarantees for many well-studied machine-learning models, and is also practically relevant, as it democratizes fair machine learning by providing concrete training algorithms and rigorous generalization guarantees for these models.

In this framework, both PAC and FPAC learning are parameterized by a *learning task* (model space and loss function), and I explore the rich hierarchy of learnability under variations of these concepts. In particular, I show that for many loss functions, PAC and FPAC learning are *statistically equivalent* (i.e., PAC-learnability implies FPAC-learnability) in section 5.2. Furthermore, in section 5.3 I show that standard *convexity conditions* sufficient for PAC-learnability are also sufficient for FPAC-learnability, as are certain *coverability conditions* that ensure the hypothesis space is sufficiently small so as to limit overfitting, and may be efficiently approximately searched. I explore the basic relationships between various learnability classes, but many open questions remain, and I hope that future work will further characterize these practically interesting and theoretically deep problems. This chapter is adapted from the latter half of Cousins [2021].

5.1 Statistical and Computational Efficiency Learning Guarantees

In this section, we define a formal notion of fair-learnability, termed *fair PAC-learning*, where a loss function and hypothesis class are fair PAC-learnable essentially if any distribution can be approximately learned from a polynomially-sized sample (w.h.p.). We then construct various fair-PAC learners, and relate the concept to standard PAC learning [Valiant, 1984], with the understanding that this allows the vast breadth of research of PAC-learning algorithms, and quite saliently, necessary and sufficient conditions, to be applied to fair-PAC learning. In particular, we show a hierarchy of fair-learnability via generic statistical and computational learning theoretic bounds and reductions.

Hypothesis Classes and Sequences We now define *hypothesis class sequences*, required to discuss nontrivial *computational complexity* in learning. This definition is adapted from [Shalev-Shwartz and Ben-David, 2014, Def. 8.1], which treats only *binary classification*.

Definition 5.1.1 (Hypothesis Class Sequence). A *hypothesis class sequence* $\mathcal{H} = (\mathcal{H}_1, \mathcal{H}_2, \dots)$ is a *nested sequence of hypothesis classes* mapping $\mathcal{X} \rightarrow \mathcal{Y}$. In other words, $\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \dots$.

Usually, \mathcal{H}_i is easily derived from \mathcal{H}_j . For instance, *linear classifiers* naturally form a sequence of families using their *dimension*:

$$\mathcal{H}_d \doteq \left\{ \mathbf{x} \mapsto \text{sgn}(\mathbf{x} \cdot (\mathbf{w} \circ \mathbf{0})) \mid \mathbf{w} \in \mathbb{R}^d \right\} .$$

Here each \mathcal{H}_i is defined over domain \mathbb{R}^∞ , but it is more natural to discuss them as objects over \mathbb{R}^d . Similarly, unit-scale *univariate polynomial regression* naturally decomposes as

$$\mathcal{H}_d \doteq \left\{ x \mapsto (x, x^2, \dots, x^d) \cdot \mathbf{w} \mid \mathbf{w} \in [-1, 1]^d \right\} .$$

The *hypothesis class sequence* concept is necessary, as it allows us to distinguish statistically easy problems, like learning hyperplanes in finite-dimensional space, from the statistically challenging problem of learning hyperplanes in \mathbb{R}^∞ . It is also important to studying computational-hardness of statistically-easy learning problems, as this essentially boils down to selecting a hypothesis class sequence such that learning *time complexity* grows *superpolynomially* (in d).

For context, we first present a generalized notion of PAC-learnability, which we then generalize to fair-PAC-learnability. Standard presentations consider only classification under 0-1 loss, but following the generalized learning setting of Vapnik [2013], some authors consider generalized notions for other learning problems (see, e.g., Shalev-Shwartz and Ben-David [2014, Definition 3.4]).

Definition 5.1.2 (PAC-Learnability). Suppose *nested hypothesis class sequence* $\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \dots$, all over $\mathcal{X} \rightarrow \mathcal{Y}$, and *loss function* $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{0+}$. We say \mathcal{H} is *PAC-learnable* w.r.t. ℓ if \exists a (randomized) algorithm \mathcal{A} , such that \forall :

1. sequence index d ;
2. instance distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$;
3. additive approximation error $\varepsilon > 0$; and
4. failure probability $\delta \in (0, 1)$;

it holds that \mathcal{A} can identify a hypothesis $\hat{h} \in \mathcal{H}$, i.e., $\hat{h} \leftarrow \mathcal{A}(\mathcal{D}, \varepsilon, \delta, d)$, such that

1. there exists some *sample complexity* function $m(\varepsilon, \delta, d) : (\mathbb{R}_+ \times (0, 1) \times \mathbb{N}) \rightarrow \mathbb{N}$ s.t. $\mathcal{A}(\mathcal{D}, \varepsilon, \delta, d)$ consumes no more than $m(\varepsilon, \delta, d)$ samples from \mathcal{D} (i.e., has finite sample complexity); and

2. with probability at least $1 - \delta$ (over randomness of \mathcal{A}), \hat{h} obeys

$$R(\hat{h}; \ell, \mathcal{D}) \leq \inf_{h^* \in \mathcal{H}} R(h^*; \ell, \mathcal{D}) + \varepsilon .$$

The class of such learning problems is denoted PAC, thus we write $(\mathcal{H}, \ell) \in \text{PAC}$ to denote PAC-learnability. Furthermore, if for all d , the space of \mathcal{D} is restricted such that

$$\exists h \in \mathcal{H}_d \text{ s.t. } R(h; \ell, \mathcal{D}) = 0 ,$$

then (\mathcal{H}, ℓ) is *realizable-PAC-learnable*, written $(\mathcal{H}, \ell) \in \text{PAC}^{\text{Rlz}}$.

Observation 5.1.3 (On Realizable Learning). Our definition of realizability appears to differ from the standard form, in which \mathcal{D} is a distribution over only \mathcal{X} , and y is simply computed as $h^*(x)$, for some $h^* \in \mathcal{H}$. We instead constrain \mathcal{D} such that there exists a 0-risk $h^* \in \mathcal{H}$, which is equivalent for any loss function ℓ such that $\ell(y, \hat{y}) = 0 \Leftrightarrow y = \hat{y}$, e.g., 0-1 classification loss, or absolute or square regression loss. With our definition, it is much clearer that realizable learning is a special case of agnostic learning, and furthermore, we handle a much broader class of problems, for which there may be some amount of *noise*, or wherein a ground truth may not even exist.

For instance, in a *recommender system*, y may represent the *set of items* that x will like, and h may predict a singleton set, and thus we take $\ell(y, \{\hat{y}\}) = \mathbb{1}_y(\hat{y})$. There is no ground-truth here, but rather we seek a *compatible solution* that recommends appropriate items to everyone. Similarly, in high-dimensional classification, we often treat classifier output as a *ranked list* of predictions, and the *top- k loss* is taken to be $\ell(y, \hat{y}) = \mathbb{1}_{\hat{y}_{1:k}}(y)$. Again, no ground truth exists, but 0-risk learning is still possible if there is not “too much” ambiguity (e.g., foxes and dogs can be confused, as long as they are separated from horses and zebras).

We now generalize this concept to *fair-PAC-learnability*. In particular, we replace the *univariate risk-minimization* task with a *multivariate malfare-minimization* task. Following the theory of section 3.3.3, we do not commit to any particular objective, but instead require that a fair-PAC-learner is able to maximize *any* fair malfare function satisfying the standard axioms. Furthermore, here problem instances grow not just in problem complexity d , but also in the number of groups g .

Definition 5.1.4 (Fair PAC (FPAC)-Learnability). Suppose *nested hypothesis class sequence* $\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \dots \subseteq \mathcal{X} \rightarrow \mathcal{Y}$, and *loss function* $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{0+}$. We say \mathcal{H} is *fair PAC-learnable* w.r.t. *group count* g and *loss function* ℓ if \exists a (randomized) algorithm \mathcal{A} , such that \forall :

1. sequence index d ;
2. group count g ;
3. instance distribution $\mathcal{D}_{1:g}$ over $(\mathcal{X} \times \mathcal{Y})^g$;
4. group weights measure \mathbf{w} over $\{1, \dots, g\}$;
5. malfare concept $\mathbb{M}(\cdot; \cdot)$ satisfying axioms 1-7 and 9;
6. additive approximation error $\varepsilon > 0$; and
7. failure probability $\delta \in (0, 1)$;

it holds that \mathcal{A} can identify a hypothesis $\hat{h} \in \mathcal{H}$, i.e., $\hat{h} \leftarrow \mathcal{A}(\mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}, \varepsilon, \delta, d)$, such that

1. there exists some *sample complexity* function $m(\varepsilon, \delta, d, g) : (\mathbb{R}_+ \times (0, 1) \times \mathbb{N} \times \mathbb{N}) \rightarrow \mathbb{N}$ s.t. $\mathcal{A}(\mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}, \varepsilon, \delta, d)$ consumes no more than $m(\varepsilon, \delta, d, g)$ samples (finite sample complexity); and
2. with probability at least $1 - \delta$ (over randomness of \mathcal{A}), \hat{h} obeys

$$\mathbb{M}(i \mapsto R(\hat{h}; \ell, \mathcal{D}_i); \mathbf{w}) \leq \operatorname{argmin}_{h^* \in \mathcal{H}} \mathbb{M}(i \mapsto R(h^*; \ell, \mathcal{D}_i); \mathbf{w}) + \varepsilon .$$

The class of such fair-learning problems is denoted FPAC, thus we write $(\mathcal{H}, \ell) \in \text{FPAC}$ to denote fair-PAC-learnability.

Finally, if for all d , the space of \mathcal{D} is restricted such that

$$\exists h \in \mathcal{H}_d \text{ s.t. } \max_{i \in \{1, \dots, g\}} R(h; \ell, \mathcal{D}_i) = 0 \text{ ,}$$

then (\mathcal{H}, ℓ) is *realizable-fair-PAC-learnable*, written $(\mathcal{H}, \ell) \in \text{FPAC}^{\text{Rlz}}$.

We now observe that a few special cases are familiar learning problems, though we argue that all cases are of interest, and simply represent different ideals of fairness which may be situationally appropriate.

Observation 5.1.5 (Malfare Functions and Special Cases). By assumption, $\mathbb{M}(\cdot; \cdot)$ must be $\mathbb{M}_p(\cdot; \cdot)$ for $p \in [1, \infty)$. Taking $g = 1$ implies $\mathbf{w} = (1)$, and $\mathbb{M}_p(\mathcal{S}; \mathbf{w}) = \mathcal{S}_1$, thus reducing the problem to standard PAC-learning. Similarly, taking $p = 1$ converts the problem to *weighted loss minimization* (weights determined by \mathbf{w}), and $p = \infty$ yields a *minimax optimization problem*, where the max is over groups, commonly encountered in adversarial and robust learning settings.

An aside on the flexibility of fair-PAC-learnability Note that the generalized definition of (fair) PAC-learnability is sufficiently broad so as to include many *supervised*, *semi-supervised*, and *unsupervised* learning problems. While this is not immediately apparent, consider that, for instance, *k-means clustering* can be expressed as a *learning problem* where the task is to identify a set of k vectors in \mathbb{R}^d where the hypothesis class maps a vector \mathbf{x} on to the nearest cluster center, and the loss function is the square distance. This is a surprisingly natural fairness issue when cast as a *resource allocation problem*; if, for instance each cluster center represents a cellphone tower, then we seek to place towers to serve *all groups*, and to avoid serving one or more groups particularly well at the expense of others.

On Computational Efficiency Some authors consider not just the *statistical* but also the *computational* performance of learning, generally requiring that \mathcal{A} have *polynomial* time complexity (and thus implicitly sample complexity). In other words, they require that $\mathcal{A}(\mathcal{D}, \varepsilon, \delta, d)$ terminates in $m(\varepsilon, \delta, d) \in \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d)$ steps. The main focus of this manuscript is sample complexity, but for completeness, we note that a similar concept of *polynomial-time fair-PAC learnability* is equally interesting, where here we assume $\mathcal{A}(\mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}, \varepsilon, \delta, d)$ may be computed by a Turing machine (with access to *sampling* and *entropy* oracles) in $m(\varepsilon, \delta, d, g) \in \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$ steps, and thus implicitly the same bound on sample complexity. We denote these concepts $\text{PAC}_{\text{Poly}}^{\text{Agn}}$, $\text{PAC}_{\text{Poly}}^{\text{Rlz}}$, $\text{FPAC}_{\text{Poly}}^{\text{Agn}}$, and $\text{FPAC}_{\text{Poly}}^{\text{Rlz}}$.

Some trivial reductions We first observe (immediately from definitions 5.1.2 and 5.1.4) that PAC-learning is a special case of fair-PAC-learning. In particular, taking $g = 1$ implies $\mathbb{M}_p(\mathcal{S}) = \mathbb{M}_1(\mathcal{S}) = \mathcal{S}_1$, thus *malfare-minimization* coincides with *risk minimization*. The more interesting question, which we seek to answer in the remainder of this document, is *when* and *whether* the *converse* holds. Furthermore, when possible, we would like to show practical, sample-and-compute-efficient *constructive reductions*.

Realizability We first show that in the *realizable case*, *PAC-learnability* implies *fair PAC-learnability*. In particular, we employ a simple and practical *constructive* polynomial-time reduction. More efficient reductions are possible for particular values of p, g . Our reduction simply takes a sufficiently number of samples from the *uniform mixture distribution* over all g groups, and PAC-learns on this distribution. As the reduction is constructive, (and efficiency-preserving) this gives us generic algorithms for (efficient) FPAC-learning in terms of algorithms for PAC-learning.

Theorem 5.1.6 (Realizable Reductions). Suppose loss function ℓ and hypothesis class \mathcal{H} . Then

1. $(\mathcal{H}, \ell) \in \text{PAC}^{\text{Rlz}} \Rightarrow (\mathcal{H}, \ell) \in \text{FPAC}^{\text{Rlz}}$; and
2. $(\mathcal{H}, \ell) \in \text{PAC}_{\text{Poly}}^{\text{Rlz}} \Rightarrow (\mathcal{H}, \ell) \in \text{FPAC}_{\text{Poly}}^{\text{Rlz}}$.

We construct a(n) (efficient) FPAC-learner for (\mathcal{H}, ℓ) by noting that there exists some \mathcal{A}' with sample-complexity $m_{\mathcal{A}'}(\varepsilon, \delta, d)$ and time complexity $t_{\mathcal{A}'}(\varepsilon, \delta, d)$ to PAC-learn (\mathcal{H}, ℓ) , and taking $\mathcal{A}(\mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}, \varepsilon, \delta, d) \doteq \mathcal{A}'(\text{mix}(\mathcal{D}_{1:g}), \frac{\varepsilon}{g}, \delta, d)$. Then \mathcal{A} FPAC-learns (\mathcal{H}, ℓ) , with sample-complexity $m_{\mathcal{A}}(\varepsilon, \delta, d, g) = m_{\mathcal{A}'}(\frac{\varepsilon}{g}, \delta, d)$, and time-complexity $t_{\mathcal{A}}(\varepsilon, \delta, d, g) = t_{\mathcal{A}'}(\frac{\varepsilon}{g}, \delta, d)$.

Proof. We first show the sample complexity result. Suppose $\hat{h} \leftarrow \mathcal{A}'(p, \mathbf{w}, \mathcal{D}_{1:g}, \varepsilon, \delta, d)$. Then, with probability at least $1 - \delta$ (by the guarantee of \mathcal{A}), we have

$$\begin{aligned} \mathbb{M}_p(i \mapsto \text{R}(h; \ell, \mathcal{D}_i)), \mathbf{w} &\leq \mathbb{M}_\infty(i \mapsto \text{R}(h; \ell, \mathcal{D}_i), i \mapsto \frac{1}{g}) \\ &\leq g \mathbb{M}_1(i \mapsto \text{R}(h; \ell, \mathcal{D}_i), i \mapsto \frac{1}{g}) \\ &= g \text{R}(h; \ell, \text{mix}(\mathcal{D}_{1:g})) \leq g \frac{\varepsilon}{g} = \varepsilon . \end{aligned}$$

We thus may conclude that (\mathcal{H}, ℓ) is efficiently *realizable-PAC-learnable* by \mathcal{A}' , with *sample complexity* $m_{\mathcal{A}}(\varepsilon, \delta, d, g) = m_{\mathcal{A}'}(\frac{\varepsilon}{g}, \delta, d)$. Similarly, if \mathcal{A} has polynomial runtime, then so too does \mathcal{A}' , thus we may also conclude efficiency. \square

While mathematically correct, if somewhat trivial, unfortunately, this argument does not extend to the agnostic case. This is essentially because in general it is not possible to simultaneously satisfy all groups, and even with reweighting, the malfare of ERM solutions are highly unstable in the weights. To see this, suppose $\mathcal{Y} \doteq \{a, b\}$, group A always wants a , and group B always wants b , with symmetric preferences, and we wish to optimize egalitarian malfare. For any reweighting, the utility-optimal solution is always to produce all a or all b , except with equal weights, when all solutions are equally good. Thus for no reweighting do all reweighted-risk solutions even approximate the egalitarian-optimal solution (which is evenly split between a and b). We thus conclude that simple constructive reductions using PAC-learners as subroutines is likely to solve the general problem.

In addition to the argument being inextensible to the agnostic case, we note that, philosophically speaking, realizable FPAC learning is rather uninteresting, essentially because in a world where all parties may be satisfied completely, the obvious solution is to do so (and this solution is in fact an equilibrium). Thus unfairness and bias issues logically only arise in a world of *conflict* (e.g., in zero-sum settings, or under limited resources constraints, which foster *competition* between groups). Thus while the realizable case is straightforward and solvable, we argue that in practice, the agnostic learning setting is far more relevant.

5.2 Characterizing Fair Statistical Learnability with FPAC-Learners

We first consider only questions of *statistical learning*. In other words, we ignore computation for now, and show only that *there exist* fair-PAC-learning algorithms of unbounded runtime. In particular, we show a generalization of the *fundamental theorem of statistical learning* to fair learning problems. The aforementioned result relates *uniform convergence* and *PAC-learnability*, and is generally stated for binary classification only. We define a natural generalization of uniform convergence to arbitrary learning problems within our framework, and then show conditions under which a generalized fundamental theorem of (fair) statistical learning holds. In particular, we show that, neglecting computational concerns, PAC-learnability and fair-PAC learnability are equivalent for learning problems with a particular *no-free-lunch* guarantee.

We first define a *generalized notion of uniform convergence* as defined by [Shalev-Shwartz and Ben-David, 2014]. In particular, our definition applies to *any bounded loss function*,¹ thus greatly generalizes the standard notion for binary classification.

Definition 5.2.1 (Uniform Convergence). Suppose $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, r] \subseteq \mathbb{R}$ and hypothesis class $\mathcal{H} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$. We say $(\mathcal{H}, \ell) \in \text{UC}$ if

$$\lim_{m \rightarrow \infty} \sup_{\mathcal{D}} \sup_{\text{over } \mathcal{X} \times \mathcal{Y}} \mathbb{E}_{z \sim \mathcal{D}^m} \left[\sup_{h \in \mathcal{H}} \left| \hat{R}(h; \ell, z) - R(h; \ell, \mathcal{D}) \right| \right] = 0 .$$

Observation 5.2.2. We stress that this definition is both uniform over ℓ composed with the *hypothesis class* \mathcal{H} and uniform over *all possible distributions* \mathcal{D} . The classical definition of *uniform convergence in probability* applies to a singular \mathcal{D} , however it is standard in PAC-learning and VC theory to assume uniformity over \mathcal{D} , so we adopt this convention.

Standard uniform convergence definitions consider only the convergence of *empirical frequencies* of events to their *true frequencies*, whereas we generalize to consider uniform convergence of the *empirical means* of functions to their *expected values*.

In discussing uniform convergence, it is often necessary to consider not the loss function or hypothesis class *in isolation*, but their *composition*, defined as

$$\forall h \in \mathcal{H} : (\ell \circ h)(x, y) \doteq \ell(y, h(x)) \quad \& \quad \ell \circ \mathcal{H} \doteq \{\ell \circ h \mid h \in \mathcal{H}\} .$$

It is also helpful to consider the *sample complexity* of ε - δ uniform-convergence, where we take

$$m_{\text{UC}}(\ell \circ \mathcal{H}, \varepsilon, \delta) \doteq \operatorname{argmin} \left\{ m \mid \sup_{\mathcal{D}} \sup_{\text{over } \mathcal{X} \times \mathcal{Y}} \mathbb{P} \left(\sup_{h \in \mathcal{H}} \left| \mathbb{E}_{\mathcal{D}}[\ell \circ h] - \hat{\mathbb{E}}_{z \sim \mathcal{D}^m}[\ell \circ h] \right| > \varepsilon \right) \leq \delta \right\} ,$$

i.e., the minimum sufficient sample size to ensure ε - δ uniform-convergence over the *loss family* $\ell \circ \mathcal{H}$.

It is in general true that *uniform convergence* implies *PAC-learnability*; this is well-known for binary classification, but we show the generalized result for completeness. The converse is true for some learning problems, but not for others, which we shall use as a powerful tool to characterize when PAC-learnability implies fair-PAC-learnability. We first state the fundamental theorem of statistical learning (for classification) [Shalev-Shwartz and Ben-David, 2014, theorems 6.2, 29.3].

Theorem 5.2.3 (Fundamental Theorem of Statistical Learning (Classification)). Suppose ℓ is the 0-1 loss for k -class classification, where $k < \infty$. Then the following are equivalent.

1. $\forall d \in \mathbb{N}$: \mathcal{H}_d has finite Natarajan-dimension (= VC dimension for $k = 2$ classes).
2. $\forall d \in \mathbb{N}$: (ℓ, \mathcal{H}_d) has the uniform convergence property.
3. Any ERM rule is a successful agnostic-PAC learner for \mathcal{H} .
4. \mathcal{H} is agnostic-PAC learnable.
5. Any ERM rule is a successful realizable-PAC learner for \mathcal{H} .
6. \mathcal{H} is realizable-PAC learnable.

It is somewhat subtle to generalize this result to arbitrary learning problems. In particular, there are PAC-learnable problems for which uniform convergence *does not hold*. However, Alon et al. [1997] show similar results for various regression problems, with the (scale-sensitive) γ -*fat-shattering dimension* playing the role of the Vapnik-Chervonenkis or Natarajan dimension in classification.

¹Boundedness should not be strictly necessary for learnability even uniform convergence, but vastly simplifies all aspects of the analysis. In many cases, it can be relaxed to moment-conditions, such as *sub-Gaussian* or *sub-exponential* assumptions.

Theorem 5.2.4 (Fundamental Theorem of Fair Statistical Learning). Suppose ℓ such that $\forall \mathcal{H} : (\mathcal{H}, \ell) \in \text{PAC}^{\text{Rlz}} \Rightarrow (\mathcal{H}, \ell) \in \text{UC}$. Then, for any hypothesis class sequence \mathcal{H} , the following are equivalent:

1. $\forall d \in \mathbb{N}$: (ℓ, \mathcal{H}_d) has the (generalized) uniform convergence property.
2. Any EMM rule is a successful agnostic-fair-PAC learner for (ℓ, \mathcal{H}) .
3. (ℓ, \mathcal{H}) is agnostic-fair-PAC learnable.
4. Any EMM rule is a successful realizable-fair-PAC learner for (ℓ, \mathcal{H}) .
5. (ℓ, \mathcal{H}) is realizable-fair-PAC learnable.

Proof. First note that $1 \Rightarrow 2$ is a rather straightforward consequence of the definition of uniform convergence and the contraction property of fair malfare functions (theorem 3.3.7 item 3). In particular, take $m \doteq m_{\text{UC}}(\ell \circ \mathcal{H}_d, \frac{\varepsilon}{2}, \frac{\delta}{g})$. By union bound, this implies that with probability at least $1 - \delta$, taking samples $\mathbf{z}_{1:g, 1:m} \sim \mathcal{D}_1^m \times \cdots \times \mathcal{D}_g^m$, we have

$$\forall i \in \{1, \dots, g\} : \sup_{h \in \mathcal{H}_d} \left| \mathbb{R}(h; \ell, \mathcal{D}_i) - \hat{\mathbb{R}}(h; \ell, \mathbf{z}_i) \right| \leq \frac{\varepsilon}{2} .$$

Consequently, as $\mathbb{M}(\cdot; \mathbf{w})$ is $1 - \|\cdot\|_{\infty}$ in risk (see lemma 3.4.1), it holds with probability at least $1 - \delta$ that

$$\forall h \in \mathcal{H}_d : \left| \mathbb{M}(i \mapsto \hat{\mathbb{R}}(h; \ell, \mathbf{z}_i); \mathbf{w}) - \mathbb{M}(i \mapsto \mathbb{R}(h; \ell, \mathcal{D}_i); \mathbf{w}) \right| \leq \frac{\varepsilon}{2} .$$

Now, for EMM-optimal \hat{h} , and malfare-optimal h^* , we apply this result twice to get

$$\begin{aligned} \mathbb{M}(i \mapsto \mathbb{R}(\hat{h}; \ell, \mathcal{D}_i); \mathbf{w}) &\leq \mathbb{M}(i \mapsto \hat{\mathbb{R}}(\hat{h}; \ell, \mathbf{z}_i); \mathbf{w}) + \frac{\varepsilon}{2} \\ &\leq \mathbb{M}(i \mapsto \hat{\mathbb{R}}(h^*; \ell, \mathbf{z}_i); \mathbf{w}) + \frac{\varepsilon}{2} \\ &\leq \mathbb{M}(i \mapsto \mathbb{R}(h^*; \ell, \mathcal{D}_i); \mathbf{w}) + \varepsilon . \end{aligned}$$

Therefore, under uniform convergence, the EMM algorithm agnostic fair-PAC learns (\mathcal{H}, ℓ) with finite sample complexity $m_{\mathcal{A}}(\varepsilon, \delta, d, g) = g \cdot m_{\text{UC}}(\ell \circ \mathcal{H}_d, \frac{\varepsilon}{2}, \frac{\delta}{g})$, completing $1 \Rightarrow 2$.

Now, observe that $2 \Rightarrow 3$ and $4 \Rightarrow 5$ are almost tautological: the existence of (agnostic / realizable) fair-PAC learning algorithms imply (agnostic / realizable) fair-PAC learnability.

Now, $2 \Rightarrow 4$ and $3 \Rightarrow 5$ hold, as realizable learning is a special case of agnostic learning.

As 1 implies 2-4, which in turn each imply 5, it remains only to show that $5 \Rightarrow 1$, i.e., if \mathcal{H} is realizable fair PAC learnable, then \mathcal{H} has the uniform convergence property. In general, the question is rather subtle, but here the assumption “suppose ℓ such that $(\mathcal{H}, \ell) \in \text{PAC}^{\text{Rlz}} \Rightarrow (\mathcal{H}, \ell) \in \text{UC}$ ” does most of the work. In particular, as PAC-learning is a special case of FPAC-learning, we have

$$(\mathcal{H}, \ell) \in \text{FPAC}^{\text{Agn}} \Rightarrow (\mathcal{H}, \ell) \in \text{PAC}^{\text{Agn}} ,$$

then applying the assumption yields $(\mathcal{H}, \ell) \in \text{UC}$. □

The reductions and equivalences that compose this result are graphically depicted in figure 5.1.

Observation 5.2.5 (The Gap Between Uniform Convergence and (fair) PAC-Learnability). Note that the assumption “suppose ℓ such that $(\mathcal{H}, \ell) \in \text{PAC}^{\text{Rlz}} \Rightarrow (\mathcal{H}, \ell) \in \text{UC}$ ” does not in general hold. In many cases of interest, it is known to hold, e.g., finite-class classification under 0-1 loss, and bounded regression under square and absolute loss.

In general, verifying this condition is a rather subtle task, and until more general techniques are developed, must be repeated for each learning problem (loss function). Regardless, this lies beyond the scope of this manuscript, as we have reduced the question of *fair*-PAC learnability to one of (generic) PAC-learnability.

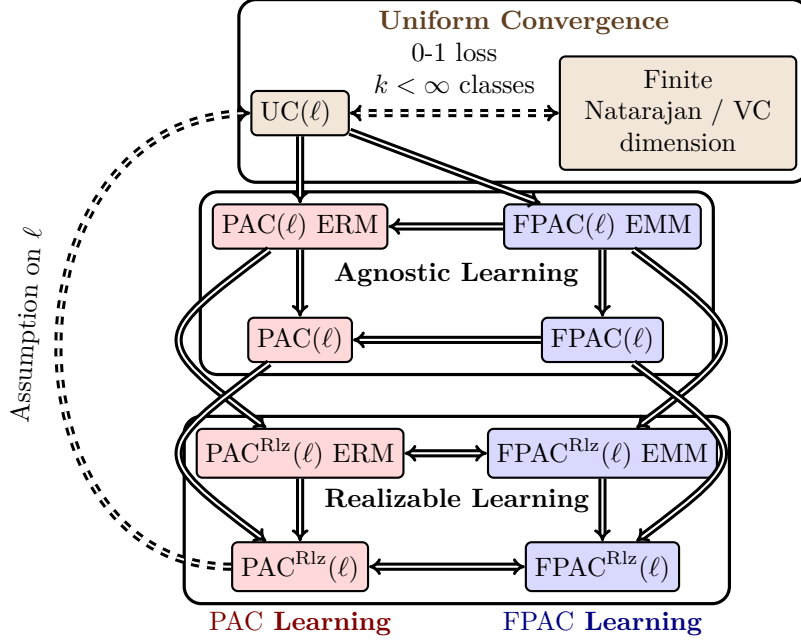


Figure 5.1: Implications between membership in PAC and fair-PAC classes. In particular, for arbitrary fixed ℓ , implication denotes *implication of membership* of some \mathcal{H} (i.e., containment); see theorems 5.2.4 and 5.1.6. Dashed implication arrows hold conditionally on ℓ . Note that when the assumption on ℓ (see theorem 5.2.4) holds, the hierarchy collapses, and in general, under realizability, some classes are known to coincide.

5.3 Characterizing Computational Learnability with Efficient FPAC Learners

In this section, we consider the more granular question of whether FPAC learning is computationally harder than PAC learning. In other words, where previously we showed conditions under which $\text{PAC} = \text{FPAC}$, here we focus on the subset of models with polynomial time training efficiency guarantees, i.e., $\text{PAC}_{\text{Poly}} = \text{FPAC}_{\text{Poly}}$. Theorem 5.1.6 has already characterized the computational complexity of *realizable* FPAC-learning, so we now focus on the agnostic case. In the *agnostic case*, we show neither a generic reduction or non-constructive proof that $\text{PAC}_{\text{Poly}} = \text{FPAC}_{\text{Poly}}$, nor do we show a counterexample; rather we leave this question for future work. We do, however, show that under conditions commonly leveraged as sufficient for efficient PAC-learning, so too is efficient FPAC-learning possible. In particular, section 5.3.1 provides an efficient constructive reduction (i.e., an algorithm) for efficient FPAC-learning under standard *convex optimization* settings, and section 5.3.2 shows the same when a small cover of \mathcal{H} exists and may be efficiently enumerated. The results of this section are summarized graphically in figure 5.2.

In both the convex optimization and efficient enumeration settings, the proofs take the same general form: we show that ε -approximate EMM on m total samples is computationally efficient (in $\text{Poly}(m, \varepsilon, d)$ time), and then argue that so long as *sample complexity* of FPAC-learning via EMM is polynomial, then we may construct an FPAC learner using ε -approximate EMM with polynomial *time complexity*. In particular, the proofs simply account for *optimization* and *sampling* error, and in both cases construct efficient FPAC-learners.

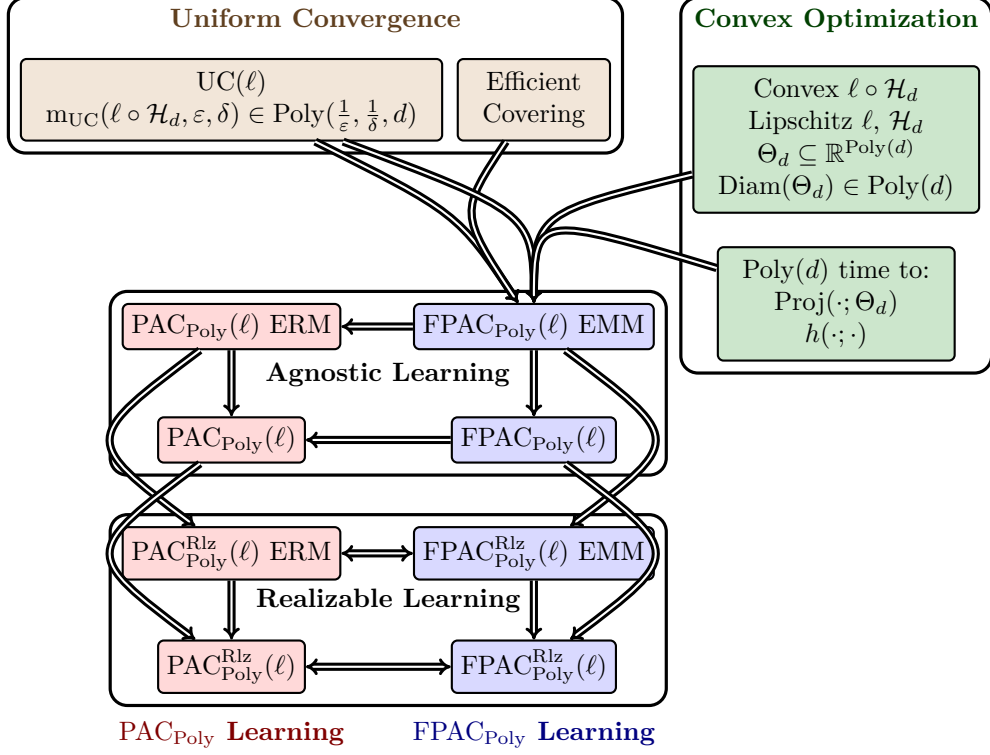


Figure 5.2: Implications between membership in various poly-time PAC and fair-PAC classes. In particular, for arbitrary but fixed ℓ , implication denotes *implication of membership* of some \mathcal{H} (i.e., containment). See theorems 5.3.1 and 5.3.2.

5.3.1 Efficient FPAC Learning with Convex Optimization

Here we show concretely and constructively the existence of fair-PAC-learners under standard convex optimization assumptions via the *subgradient method* [Shor, 2012], with constants fully derived. Sharper analyses are of course possible, and potential improvements are discussed subsequently, but our result is immediately practical and can be applied verbatim to problems like *generalized linear models* [Nelder and Wedderburn, 1972] and many *kernel methods* with little analytical effort.

Theorem 5.3.1 (Efficient FPAC Learning via Convex Optimization). Suppose each hypothesis space $\mathcal{H}_d \in \mathcal{H}$ is indexed by $\Theta_d \subseteq \mathbb{R}^{\text{Poly}(d)}$, i.e., $\mathcal{H}_d = \{h(\cdot; \theta) \mid \theta \in \Theta_d\}$, s.t. (Euclidean) $\text{Diam}(\Theta_d) \in \text{Poly}(d)$, and $\forall x \in \mathcal{X}, \theta \in \Theta_d, h(x; \theta)$ can be evaluated in $\text{Poly}(d)$ time, and $\hat{\theta} \in \mathbb{R}^{\text{Poly}(d)}$ can be Euclidean-projected onto Θ_d in $\text{Poly}(d)$ time. Suppose also ℓ such that $\forall x \in \mathcal{X}, y \in \mathcal{Y} : \theta \mapsto \ell(y, h(x; \theta))$ is a *convex function*, and suppose Lipschitz constants $\lambda_\ell, \lambda_{\mathcal{H}} \in \text{Poly}(d)$ and some norm $\|\cdot\|_{\mathcal{Y}}$ over \mathcal{Y} s.t. ℓ is $\lambda_\ell \|\cdot\|_{\mathcal{Y}}$ -Lipschitz in \hat{y} , i.e.,

$$\forall y, \hat{y}, \hat{y}' \in \mathcal{Y} : |\ell(y, \hat{y}) - \ell(y, \hat{y}')| \leq \lambda_\ell \|\hat{y} - \hat{y}'\|_{\mathcal{Y}} ,$$

and also that each \mathcal{H}_d is $\lambda_{\mathcal{H}} \|\cdot\|_2 \|\cdot\|_{\mathcal{Y}}$ -Lipschitz in θ , i.e.,

$$\forall x \in \mathcal{X}, \theta, \theta' \in \Theta_d : \|h(x; \theta) - h(x; \theta')\|_{\mathcal{Y}} \leq \lambda_{\mathcal{H}} \|\theta - \theta'\|_2 .$$

Finally, assume $\ell \circ \mathcal{H}_d$ exhibits ε - δ *uniform convergence* with sample complexity $m_{\text{UC}}(\varepsilon, \delta, d) \in \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d)$.

It then holds that, for arbitrary initial guess $\theta_0 \in \Theta_d$, for any group distributions $\mathcal{D}_{1:g}$, group weights \mathbf{w} , and fair malfare function $\mathbb{M}(\cdot; \cdot)$, the algorithm (see algorithm 2)

$$\mathcal{A}(\mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}(\cdot; \cdot), \varepsilon, \delta, d) \doteq \mathcal{A}_{\text{PSG}}(\ell, \mathcal{H}_d, \theta_0, m_{\text{UC}}(\cdot, \cdot, d), \mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}(\cdot; \cdot), \varepsilon, \delta)$$

Algorithm 2 Approximate Empirical Malfare Minimization via the Subgradient Method

- 1: **procedure** $\mathcal{A}_{\text{PSG}}(\ell, \mathcal{H}, \theta_0, \text{m}_{\text{UC}}(\cdot, \cdot), \mathcal{D}_{1:g}, \mathbf{w}, \Lambda(\cdot; \cdot), \varepsilon, \delta)$
 - 2: **Input:** λ_ℓ -Lipschitz loss function ℓ , $\lambda_{\mathcal{H}}$ -Lipschitz hypothesis class \mathcal{H} with parameter space Θ , initial guess $\theta_0 \in \Theta$, uniform-convergence sample-complexity bound $\text{m}_{\text{UC}}(\cdot, \cdot)$, group distributions $\mathcal{D}_{1:g}$, group weights \mathbf{w} , malfare function $\Lambda(\cdot; \cdot)$, solution optimality guarantee ε - δ .
 - 3: **Output:** ε - δ - $\Lambda(\cdot; \cdot)$ -optimal $h \in \mathcal{H}$ (under the conditions of theorem 5.3.1).
 - 4: $m_{\mathcal{A}} \leftarrow \text{m}_{\text{UC}}(\frac{\varepsilon}{3}, \frac{\delta}{g})$ ▷ Determine sufficient sample size
 - 5: $\mathbf{z}_{1:g, 1:m_{\mathcal{A}}} \sim \mathcal{D}_1^{m_{\mathcal{A}}} \times \dots \times \mathcal{D}_g^{m_{\mathcal{A}}}$ ▷ Draw training sample for each group
 - 6: $n \leftarrow \left\lceil \left(\frac{3 \text{Diam}(\Theta) \lambda_\ell \lambda_{\mathcal{H}}}{\varepsilon} \right)^2 \right\rceil$ ▷ Iteration count
 - 7: $\alpha \leftarrow \frac{\text{Diam}(\Theta)}{\lambda_\ell \lambda_{\mathcal{H}} \sqrt{n}}$ ▷ Learning rate ($\approx \frac{\varepsilon}{3 \lambda_\ell^2 \lambda_{\mathcal{H}}^2}$)
 - 8: $f(\theta) : \Theta \mapsto \mathbb{R}_{0+} \doteq \Lambda(i \mapsto \hat{\text{R}}(h(\cdot; \theta); \ell, \mathbf{z}_i); \mathbf{w})$ ▷ Define empirical malfare objective
 - 9: $\hat{\theta} \leftarrow \text{PROJECTEDSUBGRADIENT}(f, \Theta, \theta_0, n, \alpha)$ ▷ Run PSG algorithm on empirical malfare
 - 10: **return** $h(\cdot; \hat{\theta})$ ▷ Return ε - δ optimal model
 - 11: **end procedure**
-

fair-PAC-learns (\mathcal{H}, ℓ) with sample complexity $\text{m}(\varepsilon, \delta, d, g) = g \cdot \text{m}_{\text{UC}}(\frac{\varepsilon}{3}, \frac{\delta}{g}, d)$, and (training) time-complexity $\in \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$, thus $(\mathcal{H}, \ell) \in \text{FPAC}_{\text{Poly}}^{\text{Agn}}$.

It is of course possible to show similar guarantees under relaxed conditions. and with sharper sample complexity and time complexity bounds. The above result merely characterizes a simple and standard convex optimization setting under which standard convex optimization risk minimization guarantees readily translate to malfare minimization.

5.3.2 Uniform Convergence and Efficient Covering

As we have seen in section 5.2, uniform convergence implies, and is often equivalent to, (fair) PAC-learnability. However, these results all consider only *statistical learning*, and to analyze *computational learning* questions, we introduce a strengthening of uniform convergence that considers *computation*. We now show sufficient conditions for polynomial-time fair PAC-learnability via *covering numbers*, which we use both to show uniform convergence and to construct an efficient training algorithm.

In particular, we show that if a *polynomially-large cover* of each $\ell \circ \mathcal{H}_d$ exists and can be efficiently enumerated, then $(\mathcal{H}, \ell) \in \text{FPAC}_{\text{Poly}}$. In what follows, an ℓ_2 - γ -empirical-cover of loss family $(\ell \circ \mathcal{H}_d) \subseteq \mathcal{X} \mapsto \mathbb{R}_{0+}$ on a sample $\mathbf{z} \in (\mathcal{X} \times \mathcal{Y})^m$ is any $\mathcal{H}_{d, \gamma}$ such that

$$\forall h \in \mathcal{H}_d : \min_{h_\gamma \in \mathcal{H}_{d, \gamma}} \sqrt{\frac{1}{m} \sum_{i=1}^m ((\ell \circ h)(\mathbf{z}_i) - (\ell \circ h_\gamma)(\mathbf{z}_i))^2} \leq \gamma .$$

We take $\mathcal{C}(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$ to denote such a cover, and $\mathcal{C}^*(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$ to denote such a cover of *minimum cardinality*. Finally, we define the *uniform covering numbers*

$$\mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma) \doteq \sup_{\mathbf{z} \in (\mathcal{X} \times \mathcal{Y})^m} |\mathcal{C}^*(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)| \quad \& \quad \mathcal{N}(\ell \circ \mathcal{H}_d, \gamma) \doteq \sup_{m \in \mathbb{N}} \mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma) .$$

This concept is crucial to both our *uniform convergence* and *optimization efficiency* guarantees. In particular, our construction ensures that $\mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma)$ is sufficiently small so as to ensure *polynomial training time* on a *polynomially-large* training sample is sufficient to FPAC-learn (ℓ, \mathcal{H}) .

With this exposition complete, the FPAC-learning algorithm we present is simply EMM on a cover $\hat{\mathcal{C}}(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$ that is not superpolynomially larger than $\mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma)$, which our construction ensures exists and may be

efficiently enumerated, and upon which we guarantee that a polynomially-large training sample is sufficient. We now state the result, noting that full derivation with exact constants appears in the appendix.

Theorem 5.3.2 (Efficient FPAC-Learning by Covering). Suppose loss function ℓ of bounded codomain (i.e., $\|\ell\|_\infty$ is bounded), and hypothesis class sequence \mathcal{H} , s.t. $\forall m, d \in \mathbb{N}, \mathbf{z} \in (\mathcal{X} \times \mathcal{Y})^m$, there exist

1. a γ - ℓ_2 cover $\mathcal{C}^*(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$, where $|\mathcal{C}^*(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)| \leq \mathcal{N}(\ell \circ \mathcal{H}_d, \gamma) \in \text{Poly}(\frac{1}{\gamma}, d)$; and
2. an algorithm to enumerate a γ - ℓ_2 cover $\hat{\mathcal{C}}(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$ of size $\text{Poly} \mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma)$ in $\text{Poly}(m, \frac{1}{\gamma}, d)$ time.

Then $(\ell, \mathcal{H}) \in \text{FPAC}_{\text{Poly}}$, where (1) is sufficient to show that (ℓ, \mathcal{H}) is fair-PAC learnable with polynomial *sample complexity*, and (2) is required only to show polynomial *training time complexity*.

This immediately implies that fixed \mathcal{H} that are *finite*, or of bounded *pseudodimension* or γ -*fat-shattering dimension*² are fair-PAC-learnable. For instance, this includes *classifiers* such as *all possible languages of Boolean formulae* over (constant) d variables, or *halfspaces* (i.e., linear hard classifiers $\mathcal{H} \doteq \{\mathbf{x} \mapsto \text{sgn}(\mathbf{x} \cdot \mathbf{w}) \mid \mathbf{w} \in \mathbb{R}^d\}$), as well as GLM, subject to regularity constraints to appropriately control the loss function. However, it is perhaps not as powerful as it appears; it applies to *fixed hypothesis classes*, thus each of the above linear models over \mathbb{R}^d is efficiently fair-PAC learnable, but it says nothing about their performance as $d \rightarrow \infty$.

This is essentially because the *statistical analysis* to show polynomial *sample complexity* requires only that $\ln \mathcal{N}(\ell \circ \mathcal{H}_d, \gamma) \in \text{Poly}(\frac{1}{\gamma}, d)$, whereas our *training algorithm* must actually *enumerate* an empirical cover, which yields the (exponentially) stronger requirement that $\mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma) \in \text{Poly}(\frac{1}{\gamma}, d)$ for polynomial *time complexity*. Indeed, we see that while the covering numbers we assume imply uniform convergence with *sample complexity* polynomial in d , when covering numbers grow exponentially in d , then our algorithm yields only *exponential* time complexity in d . Consequently, the result only implies polynomial-time algorithms w.r.t. sequences that grow slowly in complexity; e.g., sequences of linear classifiers that grow only *logarithmically* in dimension $\mathcal{H}_d \doteq \{\mathbf{x} \mapsto \text{sgn}(\mathbf{x} \cdot \mathbf{w}) \mid \mathbf{w} \in \mathbb{R}^{\lfloor \ln d \rfloor}\}$.

Note also that theorem 5.3.2 leverages *covering arguments* in both their *statistical* and *computational* capacities. Statistical bounds based on covering are generally well-regarded, particularly when strong analytical bounds on covering numbers are available, although sharper results are possible (e.g., through the *entropy integral* or *majorizing measures*). Furthermore, while we do construct a *polynomial time* training algorithm, in many cases, specific optimization methods (e.g., stochastic gradient descent or Newton’s method) exist to perform EMM *more efficiently* and with *higher accuracy*. Worse yet, *efficient enumerability* of a cover may be non-trivial in some cases; while most covering arguments in the wild are either constructive, or compositional to the point where each component can easily be constructed, it may hold for some problems that computing or enumerating a cover is computationally prohibitive.

On Compositionality and Coverability Conditions The covers and covering numbers discussed above are of course properties of each $\ell \circ \mathcal{H}_d$, rather than ℓ and each \mathcal{H}_d individually. This creates proof obligation for each loss function of interest, in contrast to theorem 5.3.1, wherein only Lipschitz continuity of ℓ is assumed, and the remaining analysis is on \mathcal{H} . Fortunately, in many cases it is still possible to analyze covers of each \mathcal{H}_d in isolation, and then draw conclusions across a broad family of ℓ composed with each \mathcal{H}_d . In particular, via standard properties of covering numbers, if ℓ is Lipschitz-continuous w.r.t. some pseudonorm $\|\cdot\|_{\mathcal{Y}}$ over \mathcal{Y} , and γ - ℓ_2 covering numbers of each \mathcal{H}_d w.r.t. $\|\cdot\|_{\mathcal{Y}}$ are well-behaved, it can be shown that the conditions of theorem 5.3.2 are met. This is useful as, for example, regression losses like *square error*, *absolute error*, and *Huber loss* are all Lipschitz continuous on bounded domains, and thus analysis on each \mathcal{H}_d alone is sufficient to apply theorem 5.3.2 with each such loss function.

²The reader is invited to consult [Anthony and Bartlett, 2009] for an encyclopedic overview of various combinatorial dimensions, associated covering-number and shattering-coefficient concepts, and their applications to statistical learning theory.

5.4 Conclusion

The central thrust of this paper introduces malfare minimization as a fair learning task, and shows relationships between the statistical and computational issues of malfare and risk minimization. In particular, we argue that our method is more in line with welfare-centric machine learning theory than demographic-parity theory, however malfare is better aligned to address machine learning tasks cast as loss minimization problems than is welfare. As such, the first half of this manuscript is dedicated to deriving and motivating malfare minimization, while the latter studies the problem from statistical and computational learning theoretic perspectives.

5.4.1 Contrasting Malfare and Welfare

With our framework now fully laid out and initial results presented, we now contrast our malfare-minimization framework with traditional welfare-maximization approaches in greater detail. We do not claim that malfare is a better or more useful concept than welfare; but rather we argue only that it is *significantly different* (with surprising non-equivalence results between power-mean welfare and malfare functions), stands on an equal axiomatic footing, and it stands to reason that the right tool (malfare) should be used for the right job (fair loss minimization). With this said, we acknowledge that some learning tasks, e.g., bandit problems and reinforcement learning tasks, are more naturally phrased as *maximizing* utility or (discounted) reward. However, with a few exceptions, e.g., the *spherical scoring rule* from decision theory, most supervised learning problems are naturally cast as minimizing nonnegative *loss functions* (arguably via cross-entropy or KL-divergence minimization through maximum-likelihood, either as explicitly intended [Nelder and Wedderburn, 1972], or *ex-post-facto* through subsequent analysis [Cousins and Riondato, 2019]).

We are highly interested in exploring a parallel theory of fair welfare optimization, however some key malfare properties do not hold for welfare. In particular, fair welfare functions $W_p(\cdot; \cdot)$ for $p \in [0, 1)$ are not Lipschitz continuous; for example, the *Nash social welfare* (a.k.a. *unweighted geometric welfare*) $W_0(\mathcal{S}; \omega \mapsto \frac{1}{g}) = \sqrt[g]{\prod_{i=1}^g \mathcal{S}_i}$ is unstable to perturbations of each \mathcal{S}_i around 0, which causes difficulty in both the *statistical* and *computational* aspects of learning. Leveraging this fact, it is trivial to construct welfare-maximization learning problems for which sample complexity is unbounded, making straightforward translation of our FPAC framework into a welfare setting rather vacuous, except in contrived, trivial, or degenerate cases. Essentially, this is because while corollary 3.4.2 holds for both welfare and malfare, it does not imply *uniform sample-complexity* bounds, whereas, such bounds are trivial for fair malfare (see lemma 3.4.1), due to the *contraction property* (theorem 3.3.7 item 3). It thus seems such a theory would need either to either impose additional assumptions to avoid non-Lipschitz behavior (e.g., artificially limit the permitted range of p), or otherwise provide weaker (non-uniform) learning guarantees.

5.4.2 FPAC Learning: Contributions and Open Questions

After motivating the malfare-minimization machine learning setting, we introduce fair PAC-learning to study the statistical and computational difficulty of malfare minimization. As a generalization of PAC-learning, known hardness results (e.g., lower-bounds on computational and sample complexity of loss minimization) immediately apply, thus, coarsely speaking, the interesting question is whether, for some tasks, malfare minimization is harder than risk minimization. Theorem 5.2.4 answers this question in the negative for *sample complexity*, under appropriate conditions on the loss function, as does theorem 5.1.6 under realizability. The remaining cases are left open, though section 5.3 at least shows that many conditions sufficient for PAC-learnability are also sufficient for FPAC-learnability.

We are optimistic that our FPAC-learning definitions will motivate the community to further pursue the deep connections between various PAC and FPAC learning settings, as well as promote cross-pollination between

computational learning theory and fair machine learning research. We believe that deeper inquiry into these questions will lead to both a better understanding of what is and is not FPAC-learnable, as well as more practical and efficient reductions and FPAC-learning algorithms.

Part III

Sample-Efficient Mean-Estimation with Dependent Sequential Data

In this part, I examine sampling and data science problems where we do not assume our setting can be reduced to an independently and identically distributed (i.i.d.) statistical estimation task. As the analysis of non-i.i.d. settings requires far more sophisticated techniques than i.i.d. settings, I focus on the simpler statistical problem of *mean estimation*, rather than uniform convergence bounds.³ I am confident that our methods can be generalized to handle more sophisticated concentration-of-measure techniques, such as Rademacher averages, though with significantly more analytical complexity than in the i.i.d. case.

In particular, I consider two settings where the data are somehow temporally or sequentially dependent (e.g., *correlated*), and we still wish to recover guarantees similar to those in the i.i.d. case. While the examples I give are for *sampling* and *data science* problems, there is no reason that similar methods can not be applied to machine learning tasks and other application domains.

As in the machine learning settings of parts I and II, where data dependent bounds are desirable to leverage the structure of the data-distribution, without unreasonably assuming *a priori* knowledge of properties like variances, I wish to show similar guarantees (w.h.p.) while assuming as little as possible *a priori*, and providing efficient algorithms, with sample complexity bounds comparable to those attainable assuming additional information were available. This is very important in practical sampling problems, as the reason we are sampling in the first place is generally that we know very little about a distribution, and it is infeasible to compute such properties exactly, so assuming known properties greatly hinders the practical usability of such methods.

Two Settings of Non-I.I.D. Mean Estimation I now describe our two settings, while aggressively focusing on their similarities, perhaps to the point of extreme reductionism. Note that despite their intuitive similarities, and the great similarity between the algorithms and sample complexity guarantees they admit, the subtle technical differences between these results are quite substantial. In particular, results in the setting are essentially self-contained, depending primarily on known bounds in the i.i.d. setting, while the second has deep roots in recent developments in the theory of concentration inequalities for Markov chains.

I first consider an important setting in databases and data analysis, where data are stored in such a way that *random access* is expensive, but access to contiguous blocks of data is cheap, e.g., in *sequentially accessible* media. This includes many physical media, such as *magnetic tapes*, *LaserDiscs* (and other optical media), *spinning magnetic disk drives*, and *wax cylinders*, but also more generally, in modern digital systems, including in blocks stored across multiple *files*, *drives*, or *servers* across a network. In each of these settings, there is a simple and efficient access pattern for the data, however standard sampling methods require *random access*, which may be orders of magnitude slower.

I give an algorithm that is able to leverage this database structure by sampling large (but dependent) blocks, at cost comparable to drawing the same quantity of independent samples. In the worst case, samples within blocks are *highly correlated*, and the method is no better than if one sample were drawn from each block. However, often, in particular when samples within each block are actually independent, the sample complexity of my method nearly matches that of methods that assume each sample in each block were drawn independently.

In the second setting, I adopt a widely-used theoretical model of dependent processes, namely, I assume we wish to sample to estimate the expectation of a *function* f over the *stationary distribution* of a *Markov chain* \mathcal{M} . Here I continue the theme of *approximate independence*, where now rather than sampling blocks with independence between them and dependence within them, we draw a long chain of samples, where nearby samples may be dependent, but the chain is guaranteed to *mix*, in the sense that sufficiently-distant samples are *nearly independent*. Precise definitions of *mixing* and *mixing time*, as well as the closely-related *relaxation*

³Do note that elementary machine learning guarantees for finite families, early VC theoretic bounds, and to some extent covering number analysis methods are all based on simple mean-estimation, making these results applicable in these settings as well.

time, are given in chapter 7, but for now it suffices to understand that there exists some T such that samples at least T steps apart in the chain behave near-independently for the purposes of mean-estimation.

Preview of the Method At a high level, both settings are the same, in the sense that we have an *a priori* guarantee of *approximate independence*, i.e., between blocks, or across large gaps in Markov chains. In fact, both have a parameter T describing the degree of *known independence*, i.e., the *size of each block*, or the *relaxation time* (which, roughly describes the length between samples in the chain for them to behave *near-independently* for *mean estimation*). They also may exploit structure in the potentially-dependent space between these near-independent milestones, which yields stronger guarantees when more independence exists in the data (distribution) than would be implied by the *a priori* guarantees.

Upon closer examination, one may note that in the block sampling setting, we have *complete independence* across spans of T records, whereas in the Markov chain setting, we have only *approximate independence*. Fortunately, through modern concentration inequalities, this only results in small constant-factor changes to tail bounds and sample complexity, thus despite the significantly more challenging *partial independence* setting of Markov chains, I recover the same category of guarantee, with comparable constants in all bounds. It is also worth noting that in both cases, the guarantees and bounds are shown by analyzing novel *variance concepts*. In particular, in block sampling I use the *inter-block variance*, defined as the variance *over blocks* of the *average of f* (rather than the variance of f *over all records*), and in the Markov chain setting, I use the *inter-trace variance*, defined as the *stationary variance over length- T traces* of \mathcal{M} of the *average of f* (rather than the standard *stationary variance*, i.e., the variance of f over the *stationary distribution* of \mathcal{M}). In both cases, I show that when more independence exists than an *a priori* assumption on T indicates, these variance concepts are commensurately smaller. Such independence structure is sufficient but not necessary for our methods to yield improvements, and this analysis leads to intuitive and simple conditions under which the method is guaranteed to outperform classical alternatives.

As for the actual mechanism of these methods, both work by taking averages over all available samples (regardless of their potential dependence). This doesn't impact expectation, via linearity of expectation, and can only reduce the variance of the estimate. Because they don't require *a priori* variance bounds, the algorithms employ *progressive sampling*, meaning they operate on an iteratively-doubling sample of blocks or traces, and at each iteration, upper-bound the appropriate variance-concept, and terminate when the sample size is sufficiently small given said variance bound to ensure that the user-supplied ε - δ mean-estimation guarantee is met. Due to this dynamic sampling schedule, these algorithms dynamically adapt their sample-consumption to structure found within the data, and consequently, the dominant terms of their sample complexities match known lower-bounds for (variance-aware) mean-estimation to within logarithmic factors.

The Form of the Bounds I now briefly preview the sample complexity guarantees derived in this part. Let T denote the *block size* or a *relaxation time* bound, r the range of f , and v_T the inter-block or inter-trace variance of f . Then both algorithms consume no more than

$$\mathbf{O} \left(\log \left(\frac{\log r/\varepsilon}{\delta} \right) \left(\frac{Tr}{\varepsilon} + \frac{Tv_T}{\varepsilon^2} \right) \right) \quad (39)$$

records or Markov-chain steps, with probability at least $1 - \delta$. First note that the $\log \log r/\varepsilon$ term is incurred by correcting for *multiple comparisons* (early stopping) made over iterations of *progressive sampling*. Ignoring this term, if we then consider that this bounds the number of *dependent samples*, and we have a factor- T fewer of *near-independent* samples, the result then matches *Bennett's inequality*, with v_T playing the role of v .

In the worst case, $v_T \in \Theta(v_1)$, where v_1 is the variance or stationary variance of f , and we match i.i.d. guarantees that take a single record from each length- T block or sequence. However, when more independence exists in the data, $v_T \ll v_1$, and we get a more interesting sample complexity bound. Note that we assume

no *a priori* knowledge of v_T , v_1 , or their relationship, yet nevertheless, these quantities appear in sample complexity bounds.

In particular, in the block sampling setting, if we assume $v_T \leq \frac{\alpha v_1}{T}$, which occurs with $\alpha = 1$ assuming perfect independence within blocks, then the number of *blocks* consumed by the block sampling algorithm, via equation 39 is

$$\mathbf{O} \left(\log \left(\frac{\log r/\varepsilon}{\delta} \right) \left(\frac{r}{\varepsilon} + \frac{\alpha v_1}{T\varepsilon^2} \right) \right) . \quad (40)$$

Similarly, in the Markov-chain setting, I show that $Tv_T \in \mathbf{O}(\tau_{\text{rel}}v_1)$, where τ_{rel} is the *relaxation time* of chain \mathcal{M} , thus equation 39 yields a sample complexity bound of

$$\mathbf{O} \left(\log \left(\frac{\log r/\varepsilon}{\delta} \right) \left(\frac{Tr}{\varepsilon} + \frac{\tau_{\text{rel}}v_1}{\varepsilon^2} \right) \right) . \quad (41)$$

The fascinating insight here is that, although *estimating the relaxation time* τ_{rel} is generally impossible without strong assumptions on \mathcal{M} , τ_{rel} still appears in the sample complexity bound. In particular, the loose bound T appears in equation 41, but is asymptotically dominated by a term involving τ_{rel} as $\varepsilon \rightarrow 0$. Again these bounds are presented in loose simplified forms to emphasize their similarity; in the subsequent two chapters, sharper forms, non-asymptotic guarantees, and corresponding insights are derived.

Chapter 6

Mean Estimation in Block Databases

I formulate and analyze *BlockSample*, a suite of sampling algorithms that use the entire information in each sampled block to estimate the mean values of arbitrary functions of record values. *BlockSample* is tailored to modern data storage where the cost of retrieving a random record is essentially the same as the cost of reading a block of records. The algorithms *progressively sample* blocks of records, rather than individual records, and estimate the variance of the mean functions between blocks. By applying a specially tailored version of Bernstein’s large deviation bound, the algorithm adaptively computes the required number of blocks to yield provable confidence intervals (or ε - δ -approximations) of function means. Thus, the sampling cost of *BlockSample* adapts to the *amount of variance* and *degree of independence* within each block, despite assuming no *a priori* information about either. In experiments on real data, a 40 \times speedup over i.i.d. sampling was achieved for a block size of 128 records. This chapter is based on joint work with Larry Rudolph and Eli Upfal.

6.1 Preliminaries

We assume that our database \mathbf{D} is organized into M *fixed-size blocks*, each containing exactly B *records* in space \mathcal{X} , thus \mathbf{D} may be thought of as a matrix in $\mathcal{X}^{M \times B}$, where $D_{i,j}$ is the j record in block i . Records may contain several *fields*, e.g., an `Employee` record $x \in \mathcal{X}$ may be (`Name` \in `STRING`, `ID` \in \mathbb{N} , `Salary` \in \mathbb{R}). Let \mathcal{D} be uniform distribution over the $M \times B$ records of \mathbf{D} . Given a function $f : \mathcal{X} \rightarrow \mathbb{R}$, or a function family $\mathcal{F} \subseteq \mathcal{X} \rightarrow \mathbb{R}$, where $\mathcal{F} = \{f_1, \dots, f_n\}$, we want to estimate the expected value, i.e., the scalar $\mathbb{E}_{\mathcal{D}}[f]$ or the vector $\mathbb{E}_{\mathcal{D}}[\mathcal{F}] = (\mathbb{E}_{\mathcal{D}}[f_1], \dots, \mathbb{E}_{\mathcal{D}}[f_n])$.

Our algorithms do not have access to the distribution \mathcal{D} over the individual records. Instead the algorithms sample from a uniform distribution over the M blocks. To estimate $\mathbb{E}_{\mathcal{D}}[f]$ the algorithm samples *blocks* $\mathbf{x}_{r,1:B} = \mathbf{D}_{i,1:B}$, for $i \sim \mathcal{U}(1 : M)$, and computes $\hat{\mathbb{E}}_{\mathbf{x}_{r,1:B}}[f] = \frac{1}{B} \sum_{j=1}^B f(\mathbf{x}_{r,j})$, i.e., the mean value of the block’s B records, for each block. Since by linearity,

$$\mathbb{E}_{\mathcal{D}}[f] = \mathbb{E}_{i \sim \mathcal{U}(1:M)} \left[\hat{\mathbb{E}}_{\mathbf{D}_{i,1:B}}[f] \right] .$$

we can estimate $\mathbb{E}_{\mathcal{D}}[f]$ by $\frac{1}{m_B} \sum_{i=1}^{m_B} \hat{\mathbb{E}}_{\mathbf{x}_{i,1:B}}[f]$ for m_B blocks of size B , chosen uniformly at random.

We are interested in the minimum number of block samples needed to obtain an (ε, δ) -estimate of $\mathbb{E}_{\mathcal{D}}[f]$, i.e.,

$$m_B^*(\varepsilon, \delta) = \inf_{m_B^* \in \mathbb{N}} : \sup_{m_B > m_B^*} \mathbb{P} \left(\left| \frac{1}{m_B} \sum_{i=1}^{m_B} \hat{\mathbb{E}}_{\mathbf{x}_{i,1:B}}[f] - \mathbb{E}_{\mathcal{D}}[f] \right| \geq \varepsilon \right) \leq \delta .$$

Thus, our algorithms and analysis focus on the random variable $\hat{\mathbb{E}}_{\mathbf{x}_{i,1:B}}[f]$ for block i chosen uniformly at random. The mean and range of $\hat{\mathbb{E}}_{\mathbf{x}_{i,1:B}}[f]$ is equal to that of f , its variance, the *inter-block-variance* $V_B \doteq \mathbb{V}_{i \sim \mathcal{U}(1:M)}[\hat{\mathbb{E}}_{\mathbf{D}_{i,1:B}}[f]]$, is typically smaller than $v \doteq \mathbb{V}_{\mathcal{D}}[f]$, the variance of f .

To show this, we apply the *law of total variance* to decompose the *total variance* $v = \mathbb{V}_{\mathcal{D}}[f]$ into *intra-block* variance V_B° and *inter-block* variance V_B as

$$\begin{aligned} \mathbb{V}_{\mathcal{D}}[f] &= \frac{1}{MB} \sum_{i=1}^M \sum_{j=1}^B \left(f(\mathbf{D}_{i,j}) - \hat{\mathbb{E}}_{\mathbf{D}_{1:M,1:B}}[f] \right)^2 = \\ &= \underbrace{\frac{1}{M} \sum_{i=1}^M \frac{1}{B} \sum_{j=1}^B \left(f(\mathbf{D}_{i,j}) - \hat{\mathbb{E}}_{\mathbf{D}_{i,1:B}}[f] \right)^2}_{\text{Intra-Block Variance } V_B^\circ} + \underbrace{\frac{1}{M} \sum_{i=1}^M \left(\hat{\mathbb{E}}_{\mathbf{D}_{i,1:B}}[f] - \hat{\mathbb{E}}_{\mathbf{D}_{1:M,1:B}}[f] \right)^2}_{\text{Inter-Block Variance } V_B}. \end{aligned}$$

Here we are primarily interested in the inter-block variance V_B , and its impact on the required block-sample size. The above shows us that $0 \leq V_B \leq v$; the minimal $V_B = 0$ is obtained when all blocks are identical, while the maximum value $V_B = v$ is obtained when the records within each block are identical, but different between blocks. A third case of interest is $V_B = v/B$ which corresponds to a uniform random partition of record to blocks. In general, V_B corresponds to the dependence between the allocations of items to blocks and the value of f . For example, if the data represents some sequential process such as time series, positive sequential dependence leads to higher V_B , while negative sequential dependence gives smaller V_B . Note that $B = 1$ corresponds to standard uniform random sample of individual items, with $V_1^\circ = 0$, thus $V_1 = v$.

In addition to *variance* our algorithms are often concerned with the *range* of the function f . In particular, we assume knowledge of an interval r such that $\forall x \in \mathcal{X} : f(x) \in r$, and we denote the lower and upper limits of r as $\lfloor r \rfloor$ and $\lceil r \rceil$, and the *length* as $\lceil r \rceil - \lfloor r \rfloor$.

6.1.1 Variance and Concentration Inequalities

We saw in the above discussion that averaging the value of a function over all records in a randomly chosen block gives an unbiased estimate of $\mathbb{E}_{\mathcal{D}}[f]$ with possibly smaller variance than the value of the function on one random record. To fully utilise the possible reduction in variance we need:

1. A tight bound on the minimum sample size $m_B^*(\varepsilon, \delta)$ as a function of the sample variance; and
2. Adapting the bound in (1) to use the empirical variance instead of the true variance of the sample.

In the asymptotic regime, the *Central Limit Theorem* (CLT) gives an optimal solution to (1) and the t -distribution properties resolve (2). In the finite sample regime these tasks are harder and do not have optimal solution. Nevertheless, rigorous finite-sample guarantees are needed in algorithmic applications because asymptotic bounds for samples from common distributions such as heavy-tailed distributions can be extremely inaccurate on finite samples.

The standard CS tool for *finite-sample* bounds is *Hoeffding's inequality* [Hoeffding, 1963], which generalizes Chernoff's inequality (see, e.g., [Mitzenmacher and Upfal, 2017]) to provide *sub-Gaussian* guarantees that depend on the range of the random variable.

Theorem 6.1.1 (Hoeffding's Inequality). Suppose distribution \mathcal{D} over $r \subseteq \mathbb{R}$ and mean μ , and sample $\mathbf{x} \sim \mathcal{D}^m$. We then have

$$\begin{aligned} \mathbb{P}_{\mathbf{x}_{1:m}} \left(\left| \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i - \mu \right| \geq \sqrt{\frac{2 \frac{1}{4} \lceil r \rceil^2 \ln(\frac{2}{\delta})}{m}} \right) &\leq \delta \quad \& \\ m^*(\varepsilon, \delta, \lceil r \rceil) &\leq \frac{2 \frac{1}{4} \lceil r \rceil^2 \ln(\frac{2}{\delta})}{\varepsilon^2}. \end{aligned} \tag{42}$$

Hoeffding’s inequality is tight to within constant factors when variance is maximal (i.e., $v = \frac{1}{4}[r]^2$), but is quite loose as $v \rightarrow 0$.

In our work, we seek bounds that merge the *simple assumptions* and *rigorous finite-sample guarantees* of Hoeffding’s inequality [Hoeffding, 1963] with the *sharpness* and *variance-dependence* of the asymptotic CLT-Gaussian bounds. To this end, we build on *Bennett’s (1962) inequality* (which improves an earlier result of Sergei Bernstein [1924] under identical conditions), by replacing dependence on *known a priori variance* with *empirical variance*.

Theorem 6.1.2 (Bennett’s inequality (sub-gamma form)). Suppose \mathcal{D} a distribution over $r \subseteq \mathbb{R}$, where $\lfloor r \rfloor$ and variance v . Then

$$\mathbb{P}_{\mathbf{x}_{1:m}} \left(\left| \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i - \mu \right| \geq \underbrace{\frac{\lfloor r \rfloor \ln(\frac{2}{\delta})}{3m}}_{\text{SCALE}} + \underbrace{\sqrt{\frac{2v \ln(\frac{2}{\delta})}{m}}}_{\text{SCALE}} \right) \leq \delta \quad \& \quad (43)$$

$$m^*(\varepsilon, \delta, \lfloor r \rfloor, v) \leq \underbrace{\frac{2\lfloor r \rfloor \ln(\frac{2}{\delta})}{3\varepsilon}}_{\text{SCALE}} + \underbrace{\frac{2v \ln(\frac{2}{\delta})}{\varepsilon^2}}_{\text{VARIANCE}} .$$

This bound has *sub-gamma form*, i.e., resembles the Chernoff bound of a *gamma random variable* (see [Boucheron et al., 2013]). The form of ε is a convenient lens through which to view the result; we decompose the bound into a fast-decaying (in m) $\lfloor r \rfloor$ -dependent *scale term* and a slow-decaying (in m) v -dependent *variance-term*. If we ignore the scale term, Bennett’s inequality matches the asymptotic CLT-Gaussian bound, *despite holding non-asymptotically*, and improves upon Hoeffding’s inequality, by replacing the worst-case $\frac{1}{4}\lfloor r \rfloor^2$ sub-Gaussian variance-proxy with a sub-gamma variance-proxy of v . The fast-decaying $\lfloor r \rfloor$ -term is usually quite small, and it acts as a finite-sample correction for an asymptotic sub-Gaussian bound.

Our analysis relies of two large deviation bounds. We apply a sub-gamma bound to the error in estimating the variance by the empirical variance, and then use that estimate of the variance to compute a bound on the required sample size using a Bennett’s type bound. Together, this results can be viewed as a sub-gamma finite sample approximation for the t-statistic.

6.1.2 Prior Work

Sampling techniques and their analysis have been studied extensively, mostly in the asymptotic regime. Here we are interested in finite-sample analysis, as required for algorithmic analysis. We are particularly focused on large deviation bounds that are sensitive to the sample variance, and on tight substituting the variance with empirical one. Similar bounds exist in the literature [Audibert et al., 2007, Mnih et al., 2008, Audibert et al., 2009, Maurer and Pontil, 2009], however our analysis is sharper, and we show a complementary lower-tail bound, which we required for high-probability sample-complexity guarantees.

With *known variance*, it is easy to determine a sufficient sample size m to satisfy an ε - δ error guarantee, but with *empirical variance*, we can’t determine if a sample-size is sufficient until we’ve already drawn it, necessitating a sort of *guess and check* approach. We address this issue with *progressive sampling*, which has been used in data-science settings where sufficient sample sizes are not known *a priori*, e.g., with Rademacher averages in *pattern mining* problems, [Pietracaprina et al., 2010, Riondato and Upfal, 2015] and with *variance* in *relative confidence interval estimation* [Mnih et al., 2008]. The basic idea is that we pick a *sampling schedule (a priori)* that tells us how many samples to *progressively draw* at each timestep, and we follow the schedule until our ε - δ guarantee is met, after which we can stop sampling.

6.2 Algorithms

In this section, we present our mean-query sampling algorithms for block data, alongside intuitive explanation, and rigorous finite-sample correctness and efficiency guarantees. We first consider in section 6.2.1 the simple case of *single-function mean estimation* with *additive guarantees*, emphasizing simplicity and intuitive presentation over sampling efficiency and flexibility. We then present an optimized algorithm for *multi-function mean estimation* in section 6.2.2. Our first algorithm assumes that the function f being estimated has range $[0, 1]$, but we lift this restriction in the generalized algorithm. We close by sketching how these algorithms may be modified to support *arbitrary error guarantees* and address other related tasks.

6.2.1 A Simple Block Sampling Algorithm

Algorithm 3 Simple Block Sampling

```

1: procedure UBMS( $\mathbf{D}, f, \varepsilon, \delta$ )  $\rightarrow \hat{\boldsymbol{\mu}}$ 
2:   Input: Database  $\mathbf{D}$  of  $M$  blocks of size  $B$ , function  $f : \mathcal{X} \rightarrow [0, 1]$ , confidence interval width  $\varepsilon$ , and failure probability  $\delta \in (0, 1)$ 
3:   Output: Additive  $\varepsilon$ - $\delta$  approximation  $\hat{\boldsymbol{\mu}}$  of  $\mathbb{E}[f]$ 
4:    $N \leftarrow \left\lceil \log_2 \left( \frac{3}{28\varepsilon} \right) \right\rceil$  ▷ Maximum iteration index
5:    $\alpha \leftarrow \frac{14 \ln(\frac{3N}{\delta})}{3\varepsilon}$  ▷ Minimal Sufficient Sample Size
6:    $\beta \leftarrow 2$  ▷ Geometric Schedule Ratio
7:    $\mathbf{U}_0 \leftarrow 0; \mathbf{V}_0 \leftarrow 0$  ▷ Avg & avg2 accumulators
8:    $\mathbf{m}_0 \leftarrow 0$  ▷ Sample count begins at 0
9:   for  $t \in 1, 2, \dots, N$  do
10:     $\mathbf{m}_t \leftarrow \lceil \alpha \beta^t \rceil$  ▷ Geometric sampling schedule
11:     $\mathbf{m}'_t \leftarrow \mathbf{m}_t - \mathbf{m}_{t-1}$  ▷ Marginal sample size
12:     $\mathcal{I} \sim \mathcal{U}^{\mathbf{m}'_t}(1, \dots, M)$  ▷ Draw uniform block indices
13:     $\mathbf{x}_{t,\cdot} \leftarrow \text{READ}(\mathbf{D}, \mathcal{I})$  ▷ Read of  $\mathbf{x}_{t,\cdot} \in \mathcal{X}^{\mathbf{m}'_t \times B}$ 
14:     $\mathbf{U}_t \leftarrow \mathbf{U}_{t-1} + \sum_{j=1}^{\mathbf{m}'_t} \hat{\mathbb{E}}[\mathbf{x}_{t,j,1:B}]$  ▷ Sum block avgs
15:     $\hat{\boldsymbol{\mu}}_t \leftarrow \frac{1}{\mathbf{m}_t} \mathbf{U}_t$  ▷ Grand mean
16:     $\mathbf{V}_t \leftarrow \mathbf{V}_{t-1} + \sum_{j=1}^{\mathbf{m}'_t} \hat{\mathbb{E}}^2[\mathbf{x}_{t,j,1:B}]$  ▷ Sum square block avgs
17:     $\hat{\mathbf{v}}_t \leftarrow \frac{\mathbf{m}_t}{\mathbf{m}_t - 1} \left( \frac{1}{\mathbf{m}_t} \mathbf{V}_t - \hat{\boldsymbol{\mu}}_t^2 \right)$  ▷ Inter-block variance
18:    if  $(t = N) \vee \left( \frac{7 \ln(\frac{3N}{\delta})}{3\mathbf{m}_t} + \sqrt{\frac{2 \ln(\frac{3N}{\delta}) \hat{\mathbf{v}}_t}{\mathbf{m}_t}} \leq \varepsilon \right)$  then
19:      return  $\hat{\boldsymbol{\mu}}_t$ 
20:    end if
21:  end for
22: end procedure

```

We first present a simple technique for single-function block-mean estimation, pseudocode given as algorithm 3. Algorithm 3 takes as input *query parameters*, which are a database \mathbf{D} of M blocks of size B and a function $f : \mathcal{X} \rightarrow [0, 1]$ upon which to perform mean estimation, and the *confidence parameters* ε and δ , which determine the width and confidence of the resulting additive- ε guarantee. In each iteration, algorithm 3 draws a geometrically-increasing number of blocks to sample uniformly from the database. It then averages f over all the records in each block, and then applies an empirical-variance sensitive variant of Bennett's inequality, repeating until it produces a 2ε -diameter CI.

Intuitively, algorithm 3 enforces its ε - δ additive guarantees by establishing a *fixed schedule* of geometrically

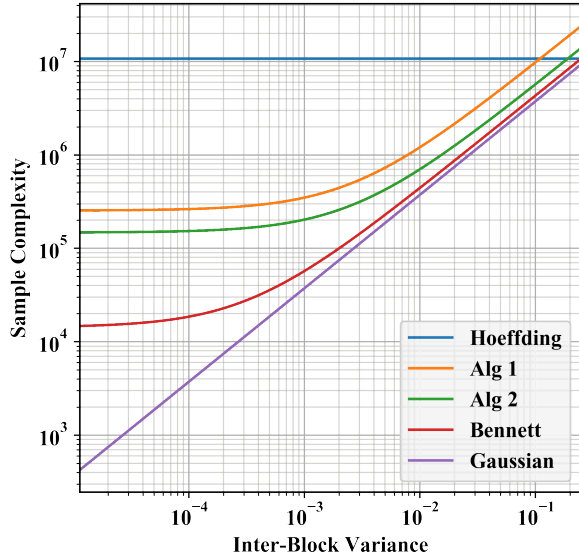


Figure 6.1: Sample complexity (y -axis) of algorithm 3 (theorem 6.2.1.1), and algorithm 4 (theorem 6.2.3) with optimized β , compared to Hoeffding, Bennett, and asymptotic Gaussian bounds, as a function of variance (x -axis). Note that variance is not known *a priori* to algorithm 3, but is used in the other bounds. Here $r = 1$, $\varepsilon = 10^{-3}$ and $\delta = 10^{-9}$.

increasing sample-sizes and bound-applications that are all valid simultaneously (via union bound). Essentially, algorithm 3 initially draws a small sample, after which it may terminate immediately, if the *empirical variance* (and thus w.h.p. the true variance) is ≈ 0 , via application of theorem 6.1.2; otherwise, it samples additional blocks to double the total sample size at each iteration, until *either* the Bennett confidence interval satisfies the requested ε - δ guarantee, *or* the *maximum iteration count* N is reached, after which a Hoeffding confidence interval implies the guarantee. In both cases, the bounds depend on $\ln(N)$ to account for the multiple-comparisons made by the algorithm in iterations before termination (i.e., the possibility that a small sample is drawn with *empirical variance* $\hat{v} \ll v$, leading to early termination). Formally, we characterize the guarantees of algorithm 3 below, with proof given in the appendix.

Theorem 6.2.1 (Algorithm 3 Guarantees). Suppose $\varepsilon > 0$ and $\delta \in (0, 1)$, then with respect to the distribution \mathcal{M} of sampling blocks uniformly at random,

1. $\mathbb{P}\left(|\hat{\mu} - \mathbb{E}_{\mathcal{D}}[f]| \leq \varepsilon\right) \geq 1 - \delta$; &
2. If $\delta \leq 0.01$, then $\mathbb{P}\left(\hat{m} \leq 2 + \frac{4 \ln(\frac{3N}{\delta})}{\varepsilon} \left(\frac{v}{\varepsilon} + \frac{8}{3}\right)\right) \leq 1 - \delta$.

Item 1 simply tells us that algorithm 3 works “as advertised,” and produces a confidence interval of the user-requested diameter (tolerance) with the user-supplied certainty (failure probability). Item 2 is rather more informative; it tells us that, at least with high probability, the *sample consumption* of algorithm 3 is small, and *mostly* dependent on the *inter-block variance* of f . It is interesting to note that although this variance is *not known a priori*, it still characterizes the sample consumption of algorithm 3.

In figure 6.1, we contrast the high-probability sample-complexity bound of algorithm 3 with the corresponding sample complexities of Bennett and asymptotic Gaussian bounds. We see that, although the latter two methods are not usable in practice, as they require *known variance*, algorithm 3 is nearly as sample-efficient, despite only using *empirical variance*. Note also that the figure is not entirely fair, as it compares a rather loose *high probability bound* on algorithm 3 sample consumption to the exact required sample-size of the Bennett and Hoeffding bounds.

Algorithm 4 Block Mean Sampling

```

1: procedure BMS( $\mathbf{D}$ , ( $\mathcal{F}$ ,  $\mathbf{r}_{1:n}$ ,  $\boldsymbol{\mu}_{0,1:n}$ ,  $\mathbf{v}_{0,1:n}$ ),  $(\varepsilon, \delta)$ ,  $\beta$ )  $\rightarrow \boldsymbol{\mu}_{1:n}$ 
2:   Input: Database  $\mathbf{D}$  of  $M$  size  $B$  blocks, function family  $\mathcal{F} \in \prod_{i=1}^n (\mathcal{X} \rightarrow \mathbf{r}_i)$ , (range, mean, variance)
   interval bounds  $(\mathbf{r}_i, \boldsymbol{\mu}_{0,i}, \mathbf{v}_{0,i})$  for each  $f_i \in \mathcal{F}$ , target CI (radius, error probability) bounds  $(\varepsilon, \delta)$ , geometric
   ratio  $\beta$ 
3:   Output: Confidence intervals  $\boldsymbol{\mu}_{i \in 1:n}$  of each  $\mathbb{E}_{\mathbf{D}}[f_i]$ 
4:    $N \leftarrow 1 \vee \left\lceil \log_{\beta} \left( \frac{\max_{i \in 1:n} 4(\lceil \mathbf{v}_i \rceil + \frac{7\lceil \mathbf{r}_i \rceil}{3}\varepsilon) \wedge \lfloor \mathbf{r}_i \rfloor^2}{\max_{i \in 1:n} 4(\lfloor \mathbf{v}_i \rfloor + \frac{7\lfloor \mathbf{r}_i \rfloor}{3}\varepsilon) \wedge \lceil \mathbf{r}_i \rceil^2} \right) \right\rceil$   $\triangleright$  Compute maximum round  $N$ 
5:    $\alpha \leftarrow \frac{1}{2\varepsilon^2} \ln\left(\frac{3nN}{\delta}\right) \left( \max_{i \in 1:n} 4(\lceil \mathbf{v}_i \rceil + \frac{7\lceil \mathbf{r}_i \rceil}{3}\varepsilon) \wedge \lfloor \mathbf{r}_i \rfloor^2 \right)$   $\triangleright$  Compute minimum sufficient sample size  $\lceil \alpha \rceil$ 
6:    $\mathbf{U}_{0,1:n} \leftarrow \mathbf{0}$ ;  $\mathbf{V}_{0,1:n} \leftarrow \mathbf{0}$   $\triangleright$  Avg & avg2 accumulators
7:    $\mathbf{m}_0 \leftarrow 0$   $\triangleright$  Initialize sample count
8:    $T \leftarrow 3nN$   $\triangleright$  Compute total # of tail bounds
9:   for  $t \in 1, 2, N$  do  $\triangleright$  Main loop
10:     $\mathbf{m}_t \leftarrow \lceil \alpha \beta^t \rceil$   $\triangleright$  Geometric sampling schedule (cumulative)
11:     $\mathbf{m}'_t \leftarrow \mathbf{m}_t - \mathbf{m}_{t-1}$   $\triangleright$  # of blocks to process this round
12:     $\mathcal{I} \sim \mathcal{U}^{\mathbf{m}'_t}(1, \dots, M)$   $\triangleright$  Draw  $\mathbf{m}'_t$  uniform block indices
13:     $\mathbf{x}_{t,\cdot} \leftarrow \text{READ}(\mathbf{D}, \mathcal{I})$   $\triangleright$  Read of  $\mathbf{x}_{t,\cdot} \in \mathcal{X}^{\mathbf{m}'_t \times B}$ 
14:    for  $f_i \in \mathcal{F}$  do
15:       $\mathbf{U}_{t,i} \leftarrow \mathbf{U}_{t-1,i} + \sum_{j=1}^{\mathbf{m}'_t} \mathbb{E}_{\mathbf{x}_{t,j,\cdot}}[f_i]$   $\triangleright$  Sum of block avgs
16:       $\hat{\boldsymbol{\mu}}_{t,i} \leftarrow \frac{1}{\mathbf{m}_t} \mathbf{U}_{t,i}$   $\triangleright$  Grand mean
17:       $\mathbf{V}_{t,i} \leftarrow \mathbf{V}_{t-1,i} + \sum_{j=1}^{\mathbf{m}'_t} \mathbb{E}[f_i^2]$   $\triangleright$  Sum of square block means
18:       $\hat{\mathbf{v}}_{t,i} \leftarrow \frac{\mathbf{m}_t}{\mathbf{m}_t - 1} \left( \frac{1}{\mathbf{m}_t} \mathbf{V}_{t,i} - \hat{\boldsymbol{\mu}}_{t,i}^2 \right)$   $\triangleright$  Inter-block variance
19:       $\varepsilon_v \leftarrow \frac{2\lceil \mathbf{r}_i \rceil^2 \ln(\frac{T}{\delta})}{3\mathbf{m}_t} + \sqrt{\left( \frac{\lceil \mathbf{r}_i \rceil^2 \ln(\frac{T}{\delta})}{\sqrt{3}\mathbf{m}_t} \right)^2 + \frac{2\lceil \mathbf{r}_i \rceil^2 \ln(\frac{T}{\delta}) \left( \hat{v} + \frac{\lceil \mathbf{r}_i \rceil^2 \mathbf{m}_t}{4(\mathbf{m}_t - 1)^2} \right)}{\mathbf{m}_t}}$ 
20:       $\mathbf{v}_{t,i} \leftarrow \mathbf{v}_{t-1,i} \cap [0, \hat{\mathbf{v}}_{t,i} + \varepsilon_v]$   $\triangleright$  Refine variance bound
21:       $\varepsilon_- \leftarrow \min \left( \frac{(\lceil \mathbf{r}_i \rceil - \lfloor \boldsymbol{\mu}_{t-1,i} \rfloor) \ln(\frac{T}{\delta})}{3\mathbf{m}_t} + \sqrt{\frac{2\lceil \mathbf{v}_{t,i} \rceil \ln(\frac{T}{\delta})}{\mathbf{m}_t}}, \lfloor \mathbf{r}_i \rfloor \sqrt{\frac{\ln(\frac{T}{\delta})}{2\mathbf{m}_t}} \right)$ 
22:       $\varepsilon_+ \leftarrow \min \left( \frac{(\lfloor \boldsymbol{\mu}_{t-1,i} \rfloor - \lfloor \mathbf{r}_i \rfloor) \ln(\frac{T}{\delta})}{3\mathbf{m}_t} + \sqrt{\frac{2\lceil \mathbf{v}_{t,i} \rceil \ln(\frac{T}{\delta})}{\mathbf{m}_t}}, \lfloor \mathbf{r}_i \rfloor \sqrt{\frac{\ln(\frac{T}{\delta})}{2\mathbf{m}_t}} \right)$ 
23:       $\boldsymbol{\mu}_{t,i} \leftarrow \boldsymbol{\mu}_{t-1,i} \cap [\hat{\boldsymbol{\mu}}_{t,i} - \varepsilon_-, \hat{\boldsymbol{\mu}}_{t,i} + \varepsilon_+]$   $\triangleright$  Refine mean
24:    end for
25:    if  $\max_{i \in 1:n} \lfloor \boldsymbol{\mu}_{t,i} \rfloor \leq 2\varepsilon$  then  $\triangleright$  Check additive  $\varepsilon$ -error
26:      return  $\boldsymbol{\mu}_{t,1:n}$ 
27:    end if
28:  end for
29: end procedure

```

6.2.2 A Refined Block Sampling Algorithm

We now present algorithm 4, which supports compound queries (multiple functions), can use *a priori* variance and expectation knowledge, and is generally more sample-efficient than algorithm 3. Briefly put, the sample complexity improvements derive from using tighter concentration inequalities, propagating bounds between timesteps, though the overall structure is largely unchanged.

Here we walk through algorithm 4, in particular emphasizing where and how it differs from algorithm 3. The first major difference is in its input; rather than taking a single function f , algorithm 4 operates on a *function family* \mathcal{F} . Additionally, algorithm 4 relaxes the *unit range* assumption of algorithm 3, and thus

needs the *range* r_i of each $f_i \in \mathcal{F}$, and also takes (possibly vacuous) *a priori mean* and *variance* bounds $\mu_{0,\cdot}$ and $v_{0,\cdot}$ of each $f_i \in \mathcal{F}$. The algorithm also adds a β parameter, which determines the *geometric ratio* between successive sample sizes (discussed further in section 6.2.3), whereas algorithm 3 simply fixes $\beta = 2$.

Another major difference is in the *sampling schedule* of algorithm 4. Because the algorithm must operate on a *family* of functions, each with their own range, and possibly variance bounds, determining the *range* of sample sizes to draw, and thus the *number* N of doubling geometric samples to draw, is more involved. We use a union bound over a variance and 2-tailed mean bound for each function at each timestep, for a total of $3nN$ bounds. Because algorithm 4 can not terminate until *each* $f_i \in \mathcal{F}$ has been ε -additively estimated, the smallest and largest sample sizes are both worst-case over \mathcal{F} , and depend on the known range r_i , which increases both, and *variance bounds* v_i , where *lower bounds* can *increase* minimal sample sizes, and *upper-bounds* can *decrease* maximal sample sizes. The details are rather complicated (see proof of theorem 6.2.2, but intuitively, a sampling schedule is selected using *a priori* knowledge to avoid statistical inefficiency by ignoring infeasible sample sizes. However, if we assume $\mathcal{F} = \{f_1\}$ & $r_1 = [0, 1]$, (as in algorithm 3), the schedule of algorithm 4 is more efficient, using fewer iterations, ($N \leq \lceil \log_2(3/28\varepsilon) \rceil$) and a finer geometric sample size grid ($\beta \leq 2$).

Each step of the main loop progressively samples blocks (line 13), and for each $f \in \mathcal{F}$ computes the empirical variances and means of block-averages (lines 16 & 18). The algorithm then refines its existing bounds on variances and means (lines 20-23), which synthesizes the information gained this iteration with the *a priori* bounds and previous iterations. Algorithm 4 may require fewer samples than algorithm 3, as it also considers upper and lower bounds from previous timesteps, which due to random fluctuations in empirical means and variances, may be sharper than the current bound.

While the exact form of the tail bounds in algorithm 4 is rather subtle, they are defined and discussed further in section 6.2.2, and they always improve upon the simple empirical Bennett inequality (line 18) of algorithm 3. The last step of each iteration is then to check the additive ε -error *termination condition*, which is satisfied with certainty by the end of iteration N .

Concentration Inequalities Algorithm 3 uses a generic sub-gamma empirical Bennett bound, whereas algorithm 4 uses an asymmetric refined bound with weaker dependence on *function ranges*. This is significant in the block-sampling setting, as we generally expect very small inter-block variances ($V_B \ll V_1 \leq \frac{1}{4} \lfloor r \rfloor^2$), in which case the fast-decaying $\Theta(\lfloor r \rfloor \ln(\frac{1}{\delta})/m)$ *scale term* can dominate the slow-decaying $\Theta(\sqrt{\ln(\frac{1}{\delta})}v/m)$ *variance term*, thus moving closer to purely sub-Gaussian tails greatly reduces sample complexity.

Algorithm 4 bounds the *variance* v in terms of the unbiased empirical variance \hat{v} with probability at least $1 - \delta$ as

$$v \leq \hat{v} + \underbrace{\frac{2\lfloor r \rfloor \ln(\frac{1}{\delta})}{3m}}_{\text{SCALE}} + \underbrace{\sqrt{\left(\frac{\lfloor r \rfloor \ln(\frac{1}{\delta})}{\sqrt{3m}}\right)^2 + \frac{2\lfloor r \rfloor^2 \ln(\frac{1}{\delta})(\hat{v} + \frac{m}{4(m-1)^2})}{m}}}_{\text{VARIANCE}}, \quad (44)$$

(used on line 19) and similarly, our efficiency guarantees require a bound in the opposite direction, namely, with probability at least $1 - \delta$, we have

$$\hat{v} \leq v + \underbrace{\frac{\lfloor r \rfloor \ln(\frac{1}{\delta})}{3m}}_{\text{SCALE}} + \underbrace{\sqrt{\frac{2\lfloor r \rfloor^2 \ln(\frac{1}{\delta})(v + \frac{m}{4(m-1)^2})}{m}}}_{\text{VARIANCE}}, \quad (45)$$

both discussed further in the appendix. Additionally we apply Bennett's inequality more carefully, noting that the upper and lower tails can use scale terms $\frac{\lfloor r_i \rfloor - \mathbb{E}[f_i]}{3}$ and $\frac{\mathbb{E}[f_i] - \lfloor r_i \rfloor}{3}$, respectively (see lines 21 & 22), which gives asymmetric tail bounds up to a factor-2 sharper than the scale term $\frac{\lfloor r_i \rfloor}{3}$ used by algorithm 3.

6.2.3 Setting Parameters

Schedule Selection With the introduction of the β parameter, algorithm 4 allows the user to tune the geometric sampling schedule to tradeoff between a large number of closely-spaced sample sizes (and thus a large union bound), or a small number of distant sample sizes (and thus potentially overshooting the sufficient sample size).

We note that while algorithm 3 simply fixes $\beta = 2$, as this intuitively simple choice leads to a straightforward *doubling* of sample size at every step, this choice is often quite sub-optimal, by up to a 2-factor (hence the ≈ 2 -factor gap between UBMS and Bennett’s inequality in figure 6.1). Decreasing β can remove of this 2-factor, though only to a point, as taking β too close to 1 increases N , which counteracts this improvement.

The other difficulty of using algorithm 4 is in setting the *range*, *mean*, and *variance* parameters \mathbf{r} , $\boldsymbol{\mu}$, and \mathbf{v} . First note that $f : \mathcal{X} \rightarrow r$ implies $\mathbb{E}_{\mathcal{D}}[f] \in r$ and $\mathbb{V}_{\mathcal{D}}[f] \leq \frac{1}{4}[r]^2$, thus bounding \mathbf{r} is *sufficient* to invoke algorithm 4, and sharper bounds on $\boldsymbol{\mu}$ & \mathbf{v} are only used to improve performance.

Often f has bounded range (e.g., indicator functions, percentages, etc.), but in general, we need to assume the database and its indices maintain enough information to compute appropriate bounds. In particular, we assume they maintain a few elementary statistics on each field. If f is defined on a *single field* (similar methods apply for multivariate f) with range $\mathbf{r}_{\mathcal{X}}$, then

$$r \subseteq \left[\min_{x \in \mathbf{r}_{\mathcal{X}}} f(x), \max_{x \in \mathbf{r}_{\mathcal{X}}} f(x) \right] ,$$

which simplifies to $[f(\lfloor \mathbf{r}_{\mathcal{X}} \rfloor), f(\lceil \mathbf{r}_{\mathcal{X}} \rceil)]$ for increasing f .

Expectation and variance bounds are trickier, and depend on exactly which statistics we maintain and the properties of f . To illustrate the principle, we consider the case of *contraction functions*, i.e., $f : \mathbb{R} \rightarrow \mathbb{R}$ s.t. $\forall x \in \mathbb{R} : \left| \frac{d}{dx} \right| \leq 1$. If f is a contraction, it holds that $V_B \leq \mathbb{V}_{\mathcal{D}}[f] \leq \mathbb{V}[\mathcal{D}]$, thus this is another case where field-statistics maintained by the database yield *a priori* knowledge exploitable by algorithm 4. Similar properties requiring various assumptions on f and database statistics can be shown to bound \mathbf{v} and $\boldsymbol{\mu}$, which in all cases can be given to algorithm 4 to improve sampling efficiency.

Algorithm 4 Guarantees The correctness and efficiency guarantees of algorithm 4 are much like in algorithm 3, although they must depend on $|\mathcal{F}|$ to account for multiple comparisons, and are also sensitive to *a priori* range and variance knowledge.

Theorem 6.2.2 (Algorithm 4 Correctness Guarantees). Suppose $\varepsilon > 0$ and $\delta \in (0, 1)$, and take $\hat{\boldsymbol{\mu}} \leftarrow \text{BMS}(\mathcal{D}, (\mathcal{F}, \mathbf{r}, \boldsymbol{\mu}, \mathbf{v}), (\varepsilon, \delta), \beta)$. Let $n = |\mathcal{F}|$ and take N as computed on line 4. Then with probability at least $1 - \delta$ over randomness of BMS, it holds that

$$\forall i \in 1, \dots, n : \left(\mathbb{E}_{\mathcal{D}}[f_i] \in \hat{\boldsymbol{\mu}} \wedge \lfloor \hat{\boldsymbol{\mu}}_i \rfloor \leq 2\varepsilon \right) ,$$

and therefore $\max_{i \in 1, \dots, n} \left| \mathbb{E}_{\mathcal{D}}[f_i] - \frac{\lceil \boldsymbol{\mu}_i \rceil - \lfloor \boldsymbol{\mu}_i \rfloor}{2} \right| \leq \varepsilon$.

Note that algorithm 4 produces *interval estimates*, which are more informative, not necessarily centered on the empirical mean, and often sharper than requested, particularly in multi-function estimation, when some function means are more easily estimated than others, and this information may be useful to the user downstream. The first statement in theorem 6.2.2 is a guarantee for these interval estimates, and the second converts their centers to *point estimates*, and more closely resembles theorem 6.2.1.1.

Theorem 6.2.3. Suppose $\varepsilon > 0$ and $\delta \leq 0.01n$, and take

$$\hat{\boldsymbol{\mu}} \leftarrow \text{BLOCKMEANSAMPLING}(\mathcal{D}, (\mathcal{F}, \mathbf{r}, \boldsymbol{\mu}, \mathbf{v}), (\varepsilon, \delta), \beta) .$$

Let $n = |\mathcal{F}|$ and take N as computed on line 4. Then, taking V_{\max} and r_{\max} to be the largest inter-block variance and range among \mathcal{F} , with probability at least $1 - \frac{\delta}{3N}$ over randomness of BLOCKMEANSAMPLING, it holds that

$$\hat{m}(\varepsilon, \delta, v) \leq \left\lceil \frac{2\beta \ln(\frac{3nN}{\delta})}{\varepsilon} \left(\frac{V_{\max}}{\varepsilon} + \frac{8r_{\max}}{3} \right) \right\rceil .$$

Proof. To see this result, note that together, eqs. 44 & 45 and Bennett’s inequality are sufficient to show that if algorithm 4 reaches a sample of size at least

$$\frac{2 \ln(\frac{3nN}{\delta})}{\varepsilon} \left(\frac{V_{\max}}{\varepsilon} + \frac{8r_{\max}}{3} \right) ,$$

then with probability at least $1 - \frac{\delta}{3N}$, the *empirical variance* of each f_i will be small enough to satisfy the termination condition. This exact sample size is not necessarily considered by algorithm 4, but because of its geometric structure, we can guarantee that a sample size a factor $\leq \beta$ in excess of the above (rounded up) is present in the sampling schedule, and termination w.h.p. will occur there. Note that here equation 44 is required, as we need to show that the *empirical variance* is not much greater than the *true variance*, whereas in the correctness proof, we only cared about the reverse bound. Note also that the statement holds with probability independent of N , as it only considers the largest single timestep where the sample size is no greater than the stated bound: terminating earlier only decreases \hat{m} , and we show that terminating later occurs with probability $\leq \frac{\delta}{3N}$. The $\delta \leq 0.01n$ requirement is necessary only to get the $\frac{8}{3}$ constant from eqs. 44 & 45; see appendix for details. \square

This result mirrors that of theorem 6.2.1.1, although a more careful analysis gives a higher-probability guarantee, removes a nuisance constant term, and is now sensitive to β and the other parameters of algorithm 4. This bound with optimal selection of β is plotted against the corresponding result for theorem 6.2.1 and various tail bounds in figure 6.1.

Comparative Sample Complexity Analysis In some sense, the strongest possible way to show that algorithm 4 is near optimal is to compare its sample complexity to that of the *minimax-optimal mean estimation algorithm* under identical conditions. In particular, assuming variance v and range $\geq \sqrt{4v}$, suppose that \mathcal{A} is a *non-adaptive* (fixed-sample size) ε - δ single-function mean-estimation algorithm, and $m^*(\varepsilon, \delta, v)$ its sample complexity. [Devroye et al., 2016] show that under mild-conditions on δ ,

$$m^*(\varepsilon, \delta, v) \geq \frac{\ln(\frac{1}{\delta})v}{\varepsilon^2} ,$$

i.e., the *minimax optimal mean estimator* requires at least this many samples. The best-known upper-bounds nearly match up to constant factors [Lugosi and Mendelson, 2019], but induce additional dependence on other parameters. In our case, we have *additive dependence* on *range*, and thus can’t get a purely competitive ratio, but the following theorem gives an affine-characterization of our bound, which is nearly multiplicative when r_{\max}/ε is small.

Theorem 6.2.4. Suppose as in theorem 6.2.3, and also $\ln(\frac{1}{\delta})v/\varepsilon^2 > 5$ & $\delta < \min(\frac{1}{2}, 0.01n, 2^{4\varepsilon^2/v})$. With probability at least $1 - \delta$, it then holds that

$$\hat{m}(\varepsilon, \delta, v) \leq 2\beta \left(1 + \frac{\ln(3nN)}{\ln(\frac{1}{\delta})} \right) m^*(\varepsilon, \delta, v) + \beta \frac{16 \ln(\frac{3nN}{\delta}) r_{\max}}{3\varepsilon} .$$

Proof. This result is a consequence of theorem 6.2.1 item 2 and the sample-complexity lower-bound of [Devroye et al., 2016]. Note that the sub-gamma variance term is purely multiplicative term, but the sub-gamma scale term is additive, as the lower-bound has no range-dependence. The assumptions of δ are required by the sample-complexity lower-bound and theorem 6.2.3. \square

This result is significant, as it tells us that despite its relative simplicity, no amount of cleverness can yield a general algorithm that performs significantly better than algorithm 4.

6.3 Experiments

This section compares the behavior of the two algorithms for various mean queries on a very large industrial dataset containing public disk block traces on block accesses in a data-center¹. There is no need to compare the results of the sample to the full dataset, rather what is important is to see how the results of the i.i.d. sample compare to the block samples. The experiments vary block sizes $B \in \{1, \dots, 128\}$, with 1 being i.i.d. sampling, and 128 being blocks of 128 records. We require more samples for higher confidence (small δ) and tighter intervals (small ϵ), and similarly, when inter-block variance V_B (which generally decreases with B) is large. Here we evaluate performance for various queries, parameter settings, and block sizes.

In each experiment, we estimate means to within target confidence interval diameter 2ϵ . The first two experiments show how various bounds behave *as a function of sample size m* , and the final experiment examines how *sample size \hat{m}* changes as a function of *block size B* . Our primary goal in these experiments is to show that the sample complexity of block sampling is much less than that of independent sampling, which we see by comparing performance across different block sizes.

The Hoeffding bound is our baseline as a non-asymptotic sufficient sample size is computed *a priori* given only range information, thus is invariant under *block size* and *variance*. The Bennett and asymptotic Gaussian bounds, on the other hand, depend on the (unknown) true variance V_B of the data, so are not usable in practice. They are presented only for comparison purposes. A secondary goal of these experiments is to show that our algorithms, despite having less *a priori* information, are still competitive with the Bennett and Gaussian bounds.

All plots are presented on log-log axes, so as to distinguish many similar curves across a broad range of sample sizes and confidence interval diameters. This is quite helpful as it visually contrasts the *linear* $\Theta(m^{-1/2})$ *convergence rates* of the sub-Gaussian Hoeffding and Gaussian bounds to the *curved mixed-rate* sub-Gamma Bennett and algorithm 4 bounds.

6.3.1 Frequency Estimation Tasks

We first address several *frequency estimation tasks* on each categorical fields of the dataset (**Job**, **Rack**, **Dataset**, **Machine**, and **Object** IDs), as well as the joint task of frequency estimation on *all five categorical fields*, and on categorical estimation of (**Machine**, **Dataset**) tuples. In each task, we require a single indicator function for *each category* of the relevant (combinations of) *categorical field(s)*, the expectation of which is the *category frequency*, thus all $f \in \mathcal{F}$ have range $[0, 1]$, and variance $\leq 1/4$. Each query is performed with algorithm 4 using $\beta = 1.25$, and results are presented in figure 6.2.

The queries are arranged by increasing *function family size* $|\mathcal{F}|$, thus Hoeffding sample complexity is monotonic in \mathcal{F} . However, the remaining bounds, which are *variance-sensitive*, often attain *improved sample complexity* for larger \mathcal{F} , as *maximum variance* more strongly impacts performance than *function family size*.

Except for the **Job**, **Dataset**, and **All** experiments, which each contain a single high-variance function, we see that the variance-dependent bounds outperform Hoeffding’s inequality, even for independent sampling. Furthermore, using $B = 10$ or $B = 100$ reduces the variance substantially, and algorithm 4 beats Hoeffding’s inequality in all datasets.

In all experiments, we see sample consumption decrease as a function of block size, empirically validating the thesis that block-information yields superior convergence rates over independent sampling. Indeed

¹The Two Sigma data object trace set will be made generally available soon; it has been used in several other publications.

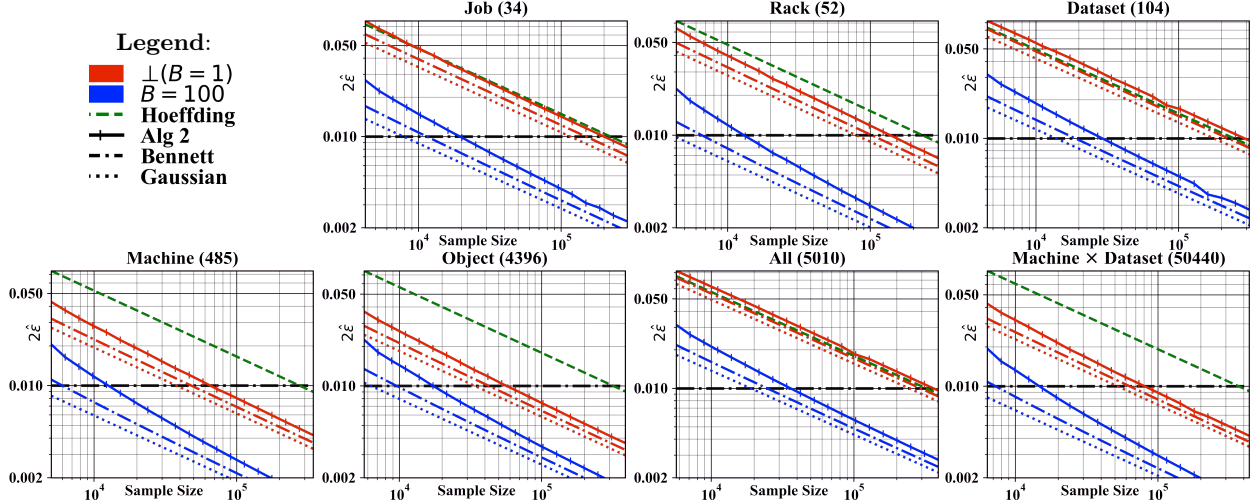


Figure 6.2: Confidence interval width 2ϵ vs. sample size for frequency estimation tasks on the *Two Sigma disk block traces* dataset, with base-10 log-log axes. In each experiment, target $\epsilon = \frac{1}{100}$ and $\delta = \frac{1}{1000}$, and the column field and $|\mathcal{F}|$ are given in the title. *Independent sampling* shown in red, and *block sampling* with block size $B = 100$ in blue. We show algorithm 4 (defined only at specific sample sizes, denoted by \dagger marks) comparing to Hoeffding, Bennett, and (asymptotic) Gaussian tail bounds. Termination occurs when a bound crosses the $\epsilon = \frac{1}{100}$ line, so lower curves are better.

we generally see sample complexity behave roughly $\propto \frac{1}{B}$, though due to *dependence within blocks*, the constant-of-proportionality is ≥ 1 . Nevertheless, the improvement with increasing block size is still quite substantial. This also illustrates why standard i.i.d. bounds can't be applied to the mB samples at each timestep directly, and so our variance-based approach is required.

Also of note is the clear curvature of the *Bennett* and algorithm 4 bounds in log-log space. Both the *Hoeffding* and *Gaussian* plots are linear (slope $-1/2$), as confidence interval radius ϵ obeys $\epsilon \propto c_{\max}/\sqrt{m}$ and $\epsilon \propto \sqrt{V_{\max}/m}$, respectively, whereas the *Bennett* and empirical *Bennett* (algorithm 4) inequalities yield curves with initial c_{\max}/m -rate decay (slope -1), followed by subsequent $\sqrt{V_{\max}/m}$ -rate decay (slope $-1/2$). This curvature is more visible when V_{\max} is very close to 0 (e.g., when $B = 100$), as with the chosen parameters, the algorithm 4 generally starts with a large enough sample size that the $\sqrt{V_{\max}/m}$ term dominates. This is significant in block sampling, as for a sufficiently large block size (thus small variance), we can expect to remain in the fast-decaying (slope -1) regime for longer, and thus quickly reach a sufficient sample size.

A final point of interest is that algorithm 4 bounds exhibit slight deviations from the ideal sub-Gaussian and sub-Gamma curves of the other bounds, which generally appear as dips in the algorithm 4 curves. These deviations are most visible in the *Job* and *Dataset* field experiments, and are mainly due to using empirical variance instead of variance. Correcting for these fluctuations is also the reason why the algorithm 4 bounds are necessarily looser than the *Bennett* bounds.

6.3.2 Sigmoid Functions

To understand how the algorithms behave for particular functions this subsection determines averages of *sigmoid functions* of *read sizes* (in mb) on the *Two Sigma block traces* dataset. We consider 15 sigmoid functions (see appendix), each a contraction with range $[-10^3, 10^3]$, and want simultaneous $\epsilon = 0.5$, $\delta = 10^{-6}$ confidence intervals on each. Unlike the frequency experiments, these functions have large range, so are expected to have much higher variance. Additionally, we assume *a priori* knowledge of the *variance* v and *range* r of the *read size* field. Because each sigmoid is a monotonic contraction function, we use v and r to provide algorithm 4 with *a priori* r_i and $v_{0,i}$ bounds (see section 6.2.3), to improve sampling efficiency.

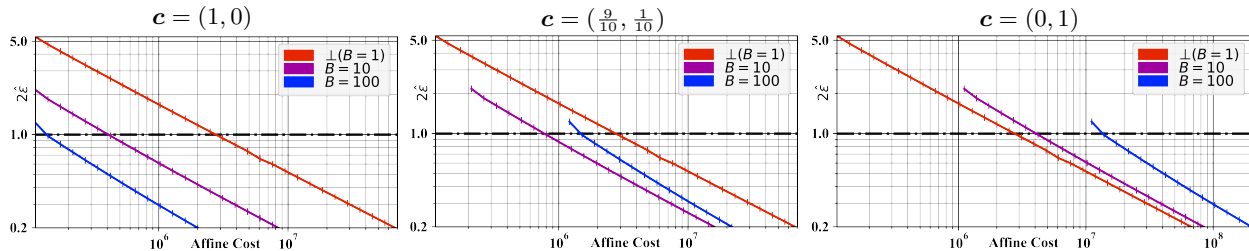


Figure 6.3: Sigmoid read size experiments, with log-log axes and various affine cost models. Confidence interval width $2\hat{\varepsilon}$ is plotted as a function of *affine cost* $\mathbf{c}_0m + \mathbf{c}_1Bm$ for $B \in \{1, 10, 100\}$. The left plot only pays per-block, and the right plot pays per record.

We again compare Hoeffding, Bennett, and Gaussian bounds to algorithm 4 with $\beta = 1.25$, but now plot, figure 6.3, confidence interval width against various measures of *cost*. In particular, we consider *affine costs* of *block size*, measured as $m\mathbf{c}_1 + mB\mathbf{c}_2$ where \mathbf{c}_1 is the cost of reading and processing each *block*, and \mathbf{c}_2 is the cost of reading and processing each *record*. Intuitively, \mathbf{c}_1 represents the networking and communication costs of *accessing* a block, and \mathbf{c}_2 represents the cost of *processing* each element of a block.

In the left plot, where $\mathbf{c} = (1, 0)$ (pay-per-block), we see similar trends as in the frequency estimation tasks; again algorithm 4 performance is quite near that of the Gaussian bound, and increasing block size greatly improves all variance-sensitive bounds. In the middle plot, where $\mathbf{c} = (\frac{9}{10}, \frac{1}{10})$, we see great improvement moving from $B = 1$ to $B = 10$, but diminishing returns thereafter, as large blocks, with their high-cost and low variance, can't compete with a larger number of low-cost small blocks. In the right plot, we have $\mathbf{c} = (0, 1)$ (pay-per-sample), where independent sampling outperforms the larger block sizes. This is unsurprising, as here dependent and independent data have the same cost, and due to autocorrelation of nearby read-sizes, block-averages have higher variance than averages of B independent samples (thus Gaussian performance is worse), and Bennett and algorithm 4 performance is significantly worse, as they also pay the price of having fewer independent samples in their sub-gamma scale terms.

These experiments nicely illustrate the limitations of block sampling; the improvement one gets depends entirely upon how one measures, and how one measures should depend on a specific application or architecture. In the convenient model where one pays only for blocks, one can generally expect large improvement moving from independent sampling to block sampling. With a hybrid affine cost, which may be reasonable for balancing IO and computation costs the $\mathbf{c} = (\frac{9}{10}, \frac{1}{10})$ plot is illustrative, perhaps with even greater gains, when compute costs are much less than $1/9$ of IO costs. Finally, the takeaway of the $\mathbf{c} = (0, 1)$ is almost tautological: if there is no cost benefit to operating on blocks, then we generally can't expect to benefit from block sampling.

6.3.3 Block Size and Sample Complexity

In this experiment, we estimate the *mean read offset*, truncated to 4GiB (i.e., reads above 4GiB were considered to be 4GiB), in the block-traces dataset. Like the sigmoid experiments, this is a continuous task, but here we provide no *a priori* variance bounds, and only the range bound $r = [0, 2^{32}]B$. Our goal is to see how efficient block-sampling schedules are, and how much they benefit from large block sizes.

In section 6.3.3, we plot ($\varepsilon = 0.1, \delta = 10^{-6}$) *sample consumption* \hat{m} , as well as the sample complexity bounds m^* (from theorems 6.2.1 and 6.2.3) for algorithm 3 (scaled as necessary to establish unit range), and algorithm 4 with $\beta = 1.05$. With $\beta = 1.05$, taking $B = 100$, we see that block-sampling yields a greater than 40-fold improvement in sample complexity over independent sampling. We also see that sample consumption decreases roughly hyperbolically as a function of block size (as under independence), which explains the

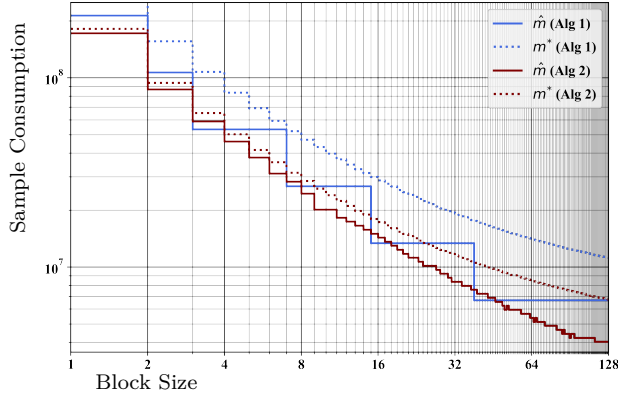


Figure 6.4: Sample Complexity versus Block Size on Mean Read Offset Task. For all $B \in 1, \dots, 100$, we plot empirical sample consumption \hat{m} , and sample consumption bound m^* for both algorithm 3 and algorithm 4 with $\beta = 1.05$.

strong performance of block sampling on this dataset.

These experiments are also quite revealing as to the differences between algorithms 3 and 4 and the importance of the geometric schedule constant β . In algorithm 3, where effectively $\beta = 2$, we see large and infrequent *discrete jumps* in sample consumption, due to the *coarse sampling schedule*. We can see how inefficient this is by comparing it to the much finer $\beta \approx 1$ sampling schedule of algorithm 4, where a smaller sample size is almost always sufficient.

6.4 Conclusion

Motivated by the fact that in modern data storage the cost of accessing a random record is essentially the same as the cost of reading a block of records, we present a rigorous technique for taking full advantage of all the record in the sampled block instead of just a random one, as done in standard sampling methods. The algorithms are progressive, continually fetching random blocks until the stopping conditions are met. This is well suited to progressively refined results when a coarse-grain result is shown almost immediately and then improved over time.

While we focused here on mean-estimation algorithms, the concept of block sampling is far more general. Our methods extend to any estimation task with a *termination condition* that can be progressively estimated at each step. In particular, with trivial modifications, our can estimate *relative confidence intervals* and *conditional means*. Similarly, in the multi-function case, (e.g., estimating averages over multiple fields), the uniform bound (FWER) over \mathcal{F} is rather stringent, and it is a simple matter to instead bound the ℓ_2 -distance between the point-estimate-means and true means, or provide a *false discovery rate* (FDR) guarantee, both of which are less sensitive to bad behavior of small subsets of \mathcal{F} .

The disaggregation of storage from compute, especially in the cloud, means that there no single model for the cost of accessing data. Large datasets are distributed across multiple storage servers, which is appropriate for parallel processing over large data sets. Parallel or distributed processing often involves higher compute costs, non-uniform startup times, and delays due to stragglers. As a result, it is sub-second response time for simple operations are difficult to achieve without block sampling.

Chapter 7

Mean Estimation and the Markov-Chain Monte-Carlo Method

I introduce a novel method for MCMC-mean estimation, whose sample complexity depends on a novel measure: the *inter-trace variance*, which characterizes the sample complexity of MCMC-mean estimation better than stationary variance (used widely in classical mean estimators). For a range R function f with stationary variance v_π , standard *variance-agnostic* MCMC mean-estimators run the chain for $\tilde{\mathcal{O}}(\frac{TR^2}{\varepsilon^2})$ steps, when given as input an (often loose) upper-bound T on the *relaxation time* τ_{rel} (\leq mixing time τ_{mix}), and when an upper-bound V on v_π is known, $\tilde{\mathcal{O}}(\frac{TR}{\varepsilon} + \frac{TV}{\varepsilon^2})$ steps suffice.

For an arbitrary trace length τ , I define the *inter-trace variance* $\text{trv}^{(\tau)}$, and using it I introduce the DYNAmic MCMC Inter-Trace variance Estimation (DYNAMITE) algorithm for mean-estimation. DYNAMITE estimates w.h.p. the mean to within ε additive error in $\tilde{\mathcal{O}}(\frac{TR}{\varepsilon} + \frac{\tau_{\text{rel}} \cdot \text{trv}^{(\tau_{\text{rel}})}}{\varepsilon^2})$ steps, without *a priori* variance bounds (w.r.t. f) on the *stationary variance* v_π , or the *trace variance* $\text{trv}^{(\tau_{\text{rel}})}$. When ε is small, the dominating complexity term is $\tau_{\text{rel}} \cdot \text{trv}^{(\tau_{\text{rel}})}$, thus the complexity of DYNAMITE principally depends on the *a priori unknown* τ_{rel} and $\text{trv}^{(\tau_{\text{rel}})}$. I show that often $\text{trv}^{(T)} = o(v_\pi)$, and furthermore, it always holds that $\text{trv}^{(\tau_{\text{rel}})} \leq 2v_\pi$, thus the worst-case complexity of DYNAMITE is $\tilde{\mathcal{O}}(\frac{TR}{\varepsilon} + \frac{\tau_{\text{rel}} \cdot v_\pi}{\varepsilon^2})$, improving the dependence of classical methods on the loose bounds T and V . This chapter is adapted from Cousins, Haddadan, and Upfal [2020].

7.1 Introduction and Related Work

Assume we have a bounded, real-valued function $f : \Omega \rightarrow [a, b]$, and a distribution π defined on the domain Ω , and our goal is to estimate the average of this function, $\mu \doteq \mathbb{E}_\pi[f(x)]$. Having enough samples from π , one can estimate μ by taking the average of such samples, usually referred to as the *empirical mean* $\hat{\mu}$. The accuracy of this estimate depends on the variance of f and the number of samples. Markov chains are often used to *approximately sample* from π efficiently Aldous et al. [1997], Brightwell and Winkler [1991], Wilson [2004], Randall [2006], Levin and Peres [2017], which can be used to construct a FPRAS for many #P-hard problems which aim to find the size of Ω .

Classical *Markov Chain Monte Carlo* (MCMC) approaches work by analytically determining a sufficient quantity m of approximately-independent samples to estimate μ , and then generating such samples by running a Markov chain. One approach is to run m copies of a Markov chain \mathcal{M} converging to a stationary distribution π for T steps, where T is chosen to exceed the *mixing time* τ_{mix} , which ensures each sample is approximately π -distributed. Another approach is to run a *single instance* of a Markov chain for a larger

number of steps, and estimate μ from these samples (rather than restarting the chain every T steps to collect a single sample). In the latter approach, m is taken proportional to an upper-bound T on the *relaxation time* $\tau_{\text{rel}} (\leq \tau_{\text{mix}})$; both τ_{mix} and τ_{rel} are formally defined in section 7.2. Unfortunately, because MCMC is generally used with distributions too complicated to analyze precisely, we lack sharp bounds on $v_\pi \doteq \mathbb{V}_\pi[f]$, and furthermore, it is often difficult to sharply bound τ_{rel} (or τ_{mix}) analytically Guruswami [2016], Levin and Peres [2017], Lovász and Winkler [1998].

A series of papers analyze the sample complexity of such approaches Lezaud [1998], Miasojedow [2012], Paulin [2015], Jiang et al. [2018], Mitzenmacher and Upfal [2017], Chung et al. [2012], Fan et al. [2018], Leon and Perron [2004], Peskun [1973], finding it to be $\tilde{O}(T(V/\varepsilon^2 + R/\varepsilon))$ for variance-aware bounds, where $V \geq v_\pi$, $r \doteq b - a$, and \tilde{O} hides log factors. For variance-agnostic bounds, which assume no prior knowledge of v_π , the bound becomes $\tilde{O}(TR^2/\varepsilon^2)$ (see theorems 7.2.5 and 7.2.6 for more details). The drawbacks of these approaches are twofold: (1) among these algorithms, the more-efficient variance-aware methods need an *a priori* bound $V \geq v_\pi$, and (2) the *complexity* of these methods depend on the upper-bounds they receive on *relaxation time* and *the stationary variance* as input, and are thus inefficient when given loose upper-bounds.

Mean estimation is an interesting problem, well-studied in theoretical computer science, and has also borne fruit as a tool for solving other difficult theoretical computer science problems. For example, Jerrum, Valiant, and Vazirani reduced estimating the *size* of a *self-reducible set* to solving a series of *mean-estimation* tasks, and ongoing work in the area seeks to reduce the complexity of such reductions Stefankovic et al. [2007], Huber [2015], David G. Harris [2007], Kolmogorov [2018]. Furthermore, there has been an emerging line of research to analyze *sublinear algorithms* having query access to input graphs using random walks Chierichetti and Haddadan [2018], Ben-Hamou et al. [2018], Bera and Seshadhri [2020], for example to estimate the mean of a real-valued function on vertices Chierichetti and Haddadan [2018]. In addition, mean estimation using MCMC has enormous applications in other fields of science such as *computational biology* Vandin et al. [2016, 2012a,b, 2011], Enright et al. [2002], Azad et al. [2018], *chemistry* Gillespie [2007], *statistical physics* Del Moral and Miclo [2007], Mignani and Rodolfo [2001], Diaconis [2009], Navarro et al. [2016], and various subfields of *computer science* e.g. *statistical learning*, *image processing*, etc Fan et al. [2013], Salatino et al. [2019], Yuan and Kendzioriski [2006], Tu and Zhu [2002].

In this work, we show that the *inter-trace variance* v better characterizes MCMC sample-complexity than does the *stationary variance* v_π (we formally define v_τ for arbitrary τ in section 7.2.1). In particular, we present and analyze the DYNAMITE algorithm, which begins by running a Markov chain for $T \geq \tau_{\text{rel}}$ steps, and *averaging* f along this trace. We call any τ consecutive steps of a chain a τ -*trace* (dropping τ when it is clear from context). Using a *progressive sampling schedule*, DYNAMITE repeatedly doubles the chain length, taking more and more samples $f_{\text{avg}}^1, f_{\text{avg}}^2, \dots$, each f_{avg}^i being the average of the i th T -trace. At each iteration, DYNAMITE approximates the inter-trace variance v_T by empirically estimating $\mathbb{V}[f_{\text{avg}}]$, and stops when the termination condition is satisfied.

We show that w.h.p., the sample consumption of DYNAMITE is $\tilde{O}\left(\frac{Tr}{\varepsilon} + \frac{T \cdot v_T}{\varepsilon^2}\right) = \tilde{O}\left(\frac{Tr}{\varepsilon} + \frac{\tau_{\text{rel}} \cdot v_{\tau_{\text{rel}}}}{\varepsilon^2}\right)$, where $v_{\tau_{\text{rel}}}$ is the trace variance of a τ_{rel} -trace. Note that this improves the sample complexity of traditional variance-agnostic approaches, and also dominates variance-aware approaches when ε is small, *despite assuming no a priori knowledge of v_T* . Moreover, the complexity of DYNAMITE is primarily dependent on the *a priori unknown* relaxation time and inter-trace variance, rather than their loose upper bounds, which limits the efficiency of standard approaches. Note that while T appears in DYNAMITE and its sample complexity bound, its contribution is negligible for small ε , and a loose T bound is compensated for by earlier termination.

Our goal in this paper is to demonstrate the significance of using the *inter trace variance* v_τ instead of the stationary variance v_π . We develop techniques to theoretically guarantee $v_{\tau_{\text{rel}}} = o(v_\pi)$ in specific scenarios, and then prove this phenomenon occurs in example 7.2.2 (see lemma 7.2.3).

Road map The paper is organized as follows: Section 7.2 contains all the preliminaries, important definitions,

and classical results that we use in our proofs. In section 7.3, we introduce our algorithm, and provide a proof sketch for its correctness and efficiency. Finally, proofs are presented in full detail in appendix F.1.

Our Contributions:

1. We introduce the *inter-trace variance* v , and show that $v_{\tau_{\text{rel}}}$ characterizes the sample complexity of MCMC-mean estimators better than v_{π} . It always holds that $v_{\tau_{\text{rel}}} \leq 2v_{\pi}$, furthermore, we identify conditions under which $v_{\tau_{\text{rel}}} = o(v_{\pi})$. We conjecture that this is common, and show a technique for bounding $v_{\tau_{\text{rel}}}$, when a simpler *projection chain* approximates f 's image on MCMC-traces.

2. Leveraging v , we design DYNAMITE. Under the standard MCMC assumptions of known relaxation time bound $T \geq \tau_{\text{rel}}$ and bounded range r , DYNAMITE requires $\tilde{O}\left(\frac{Tr}{\varepsilon} + \frac{\tau_{\text{rel}} \cdot v_{\tau_{\text{rel}}}}{\varepsilon^2}\right)$ samples, *without assuming* any *a priori* knowledge of τ_{rel} , $v_{\tau_{\text{rel}}}$ or v_T . For small ε , note that while DYNAMITE takes a loose upper bound on τ_{rel} as input, its complexity depends on τ_{rel} and $v_{\tau_{\text{rel}}}$; this is particularly impactful, as practical MCMC efficacy is often limited by loose bounds on τ_{rel} , and sharply upper bounding $v_{\tau_{\text{rel}}}$ is seldom an easy task.

3. DYNAMITE's sample complexity is always better than classic MCMC mean-estimators, since it reduces the dependency on loose upper bounds T and V . Furthermore, if $v_T = o(v_{\pi})$, then DYNAMITE is asymptotically better than classic algorithms. We note that complexity of DYNAMITE depends on $v_{\tau_{\text{rel}}}$ oblivious to any prior knowledge of $v_{\tau_{\text{rel}}}$.

7.2 Preliminaries

In this section, we first introduce notation that we use throughout the text and relevant theorems from the literature. Throughout this work, we take $f : \mathcal{X} \rightarrow [a, b]$ to be a function, and π a probability distribution on \mathcal{X} . We define the *mean* and *stationary variance* of f as $\mu \doteq \mathbb{E}_{x \sim \pi}[f] \doteq \int_{\mathcal{X}} f(x) d\pi(x)$ and $v_{\pi} \doteq \mathbb{V}_{\pi}(f) = \mathbb{E}_{x \sim \pi}[(f(x) - \mu)^2]$, respectively. An *MCMC mean-estimator* estimates μ by collecting samples running a Markov chain.

7.2.1 The Inter-Trace Variance

MCMC Terminology A Markov chain \mathcal{M} on sample space \mathcal{X} is a sequence of random variables X_1, X_2, \dots that is *memoryless*. The evolution of \mathcal{M} is characterized by its *transition matrix*; here we abuse notation and also use \mathcal{M} to represent its transition matrix, i.e., for any $x, y \in \mathcal{X}$, $\mathcal{M}(x, y) \doteq \mathbb{P}(X_i = y | X_{i-1} = x)$, and thus $\mathbb{P}(X_{i+k} = y | X_i = x) = \mathcal{M}^k(x, y)$, where \mathcal{M}^k is the standard matrix powering notation. When ν is a probability distribution on \mathcal{X} , by $\mathcal{M}(\nu)$ we mean the distribution of X_i when $X_{i-1} \sim \nu$. The distribution of the chain after k steps of \mathcal{M} is denoted by $X_{i+k} = \mathcal{M}^k(X_i)$. For arbitrary τ , we call a sequence of random variables generated by \mathcal{M} a τ -*trace* of Markov chain \mathcal{M} . We use *vector notation* to denote traces of Markov chains, i.e., $\mathbf{X}_{1:\tau} : X_1, X_2, \dots, X_{\tau}$ is a τ -trace of \mathcal{M} (we drop the subscript $1 : \tau$ when the length is clear from context). We assume all the Markov chains we discuss are ergodic and lazy¹ (see Levin and Peres [2017] if not familiar with definitions), and denote the *stationary distribution* by π . For ergodic chains, and w.r.t. a precision parameter ε , the *mixing time*, denoted by $\tau_{\text{mix}}(\varepsilon)$, is defined as the minimum τ satisfying $\text{TVD}(\mathcal{M}^{\tau}(\nu), \pi) \leq \varepsilon$ for any arbitrary starting distribution ν , and we take $\tau_{\text{mix}} \doteq \tau_{\text{mix}}(1/4)$. It is known that $\tau_{\text{mix}}(\varepsilon) \leq \tau_{\text{mix}} \log(\varepsilon^{-1})$. We denote the *second largest absolute eigenvalue* \mathcal{M} by λ , and the *relaxation time* by τ_{rel} . Note that $\tau_{\text{rel}} \doteq (1 - \lambda)^{-1}$ and it is closely related to τ_{mix} , as $(\tau_{\text{rel}}(\mathcal{M}) - 1) \ln(2) \leq \tau_{\text{mix}}(\mathcal{M}) \leq \tau_{\text{rel}}(\mathcal{M}) \ln(2/\sqrt{\pi_{\text{min}}})$, see Levin and Peres [2017]. We use T to denote a known upper-bound on the relaxation time.

¹We explicitly note when reversibility is needed.

We are interested in designing algorithms to estimate μ using samples from π collected by running \mathcal{M} . Assume $\mathbf{X} : X_1, X_2, \dots, X_\tau$ is an τ -trace of \mathcal{M} , we define the *empirical mean* $\hat{\mu}(\mathbf{X}) \doteq \frac{1}{\tau} \sum_{i=1}^{\tau} f(X_i)$. Note that $\hat{\mu}$ is an *unbiased estimator* for μ , i.e., $\mathbb{E}[\hat{\mu}(\mathbf{X})] = \mu$. Since \mathbf{X} are not i.i.d. samples, we can not use the standard Bessel's correction as an unbiased variance estimator. Therefore, we devise a novel MCMC *unbiased estimator for variance*, which is based on running two independent Markov chains in parallel, and averaging the square difference of f between them (see definition 7.3.4).

We now introduce the *inter-trace variance* v , which is the backbone of our proposed algorithm DYNAMITE. Letting T be an upper-bound on τ_{rel} , DYNAMITE uses progressive sampling to estimate v_T , thus it requires no *a priori* knowledge of v_T , $v_{\tau_{\text{rel}}}$ or v_π .

Inter-Trace Variance Let τ be an arbitrary length. On the set of all τ -traces, we define the probability distribution $\pi^{(\tau)}$ as $\pi^{(\tau)}(\mathbf{X}_{1:\tau}) \doteq \pi(X_1) \prod_{i=1}^{\tau-1} \mathcal{M}(X_i, X_{i+1})$. For an arbitrary F defined on Ω^τ , by $\mathbb{E}_{\mathbf{X} \sim \pi^{(\tau)}}[F(\mathbf{X})]$, we mean the expectation of F on any τ -trace of the Markov chain when $X_1 \sim \pi$, and other consecutive X_i s follow the transition of the Markov chain. For arbitrary τ and $\mathbf{X}_{1:\tau}$, we define $f_{\text{avg}}(\mathbf{X}_{1:\tau}) \doteq \frac{1}{\tau} \sum_{i=1}^{\tau} f(X_i)$. Note that by linearity of expectation, $\mathbb{E}_{\mathbf{X} \sim \pi^{(\tau)}}[f_{\text{avg}}(\mathbf{X})] = \mu$. We denote the variance of f_{avg} by v_τ , i.e., $v_\tau \doteq \mathbb{E}_{\mathbf{X} \sim \pi^{(\tau)}}[(f_{\text{avg}}(\mathbf{X}) - \mu)^2]$. Note that under independence, we have $v_\tau = \frac{v_\pi}{\tau}$, and for mixing processes (see Thm. 3.1 of Paulin [2015]), we get

$$\frac{1}{\tau} v_\pi \leq v_\tau \leq \frac{2\tau_{\text{rel}}}{\tau} v_\pi . \quad (46)$$

Receiving as input T , an upper-bound on relation time, DYNAMITE uses averages along T -traces of \mathcal{M} and it estimates v_T . We show that DYNAMITE's sample complexity is dominated by $v_{\tau_{\text{rel}}}$, thus when $v_{\tau_{\text{rel}}}$ is significantly smaller than the stationary variance (i.e., $v_{\tau_{\text{rel}}} = o(v_\pi)$), DYNAMITE outperforms classical methods. We show this occurs, when f partitions the state space of \mathcal{M} in such a way that the *projection chain* (defined below in definition 7.2.1) has a much smaller relaxation time ($v_{\tau_{\text{rel}}} = o(v_\pi)$).

Definition 7.2.1 (Projection chain Levin and Peres [2017]). Having an equivalence relationship \simeq on \mathcal{X} and classes $\tilde{\mathcal{X}} = \{[x]; x \in \mathcal{X}\}$ such that if $x \simeq x'$, then $\mathcal{M}(x, [y]) = \mathcal{M}(x', [y])$, we call the Markov chain $\tilde{\mathcal{M}}$ with state space $\tilde{\mathcal{X}}$ transition probabilities $\tilde{\mathcal{M}}([x], [y]) = \mathcal{M}(x, [y])$ a *projection chain* (see section 2.3 of Levin and Peres [2017] for full discussion).

In example 7.2.2, we introduce a class of functions having equal *means* and *stationary variances*, but projecting differently on traces of a fixed chain. Thus we can prove that while they all have equal v_π , under appropriate parameterization, $v_{\tau_{\text{rel}}}$ can take any arbitrary value in the range of equation 46.

Example 7.2.2. Consider the Markov chain \mathcal{C} , known as the cycle, defined on $[n] = \{1, 2, \dots, n\}$ having transition probabilities: $\mathcal{C}(i, i) = 1/2$, $\mathcal{C}(i, i+1) = 1/4$ and $\mathcal{C}(i, i-1) = 1/4$, where $i-1$ and $i+1$ are taken mod n . Clearly the stationary distribution on this chain is uniform, and it is well known that the mixing time and relaxation time are both $\Theta(n^2)$.

The following class of functions defined on $[n]$ all satisfy $\mathbb{E}_\pi(f_i) = 1/2$ and $v_\pi(f_i) = 1/4$. However, as shown in lemma 7.2.3, each has a different value for $v_{\tau_{\text{rel}}}$, covering the entire spectrum in equation 46:

For $1 \leq i \leq \frac{n}{2}$, let $f_i : [n] \rightarrow \{0, 1\}$ be $f_i(x) = 0$ if and only if $x \bmod 2i < i$, so f_i 's image on the cycle is consecutive length- i runs of 0s and 1s (see figure 7.1). Note that in the two extreme cases, we have, (1) $f_1 : [n] \rightarrow \{0, 1\}$, $f_1(x) = 0$ if and only if $x \bmod 2 = 0$, and (2) $f_{n/2} : [n] \rightarrow \{0, 1\}$, $f_{n/2}(x) = 0$ if and only if $x \leq \frac{n}{2}$.

On equivalence classes of mod 2, we find the corresponding projection chain, denoted by \mathcal{C}_{f_1} , which has transition probabilities: $\mathcal{C}_{f_1}(0, 1) = \mathcal{C}_{f_1}(1, 0) = 1/4$ and $\mathcal{C}_{f_1}(0, 0) = \mathcal{C}_{f_1}(1, 1) = 1/2$. Similarly, for arbitrary i , we denote the projection chain on equivalence classes of mod $2i$ by \mathcal{C}_{f_i} . Clearly the trace variance of f_i on \mathcal{C} and on \mathcal{C}_{f_i} are the same. Figure 7.1 illustrates these projections.

We now present the following lemma with a proof sketch. The complete proof is presented in the appendix.

Lemma 7.2.3. For $1 \leq i \leq \frac{n}{2}$, let f_i be defined as above, for arbitrary τ we have $v_\tau(f_i) = \Theta(i^2/\tau)$. In particular, for $\tau = \Theta(\tau_{\text{rel}}(\mathcal{M})) = \Theta(n^2)$, we have $v_\tau(f_{n/2}) = \Theta(1)$, and $v_\tau(f_1) = \Theta(1/n^2)$.

Proof sketch. We first prove that $v_\tau(f_i) \leq \Theta(i^2/\tau)$. Note that $x \mapsto x \bmod 2i$ partitions the set $[n]$ into $2i$ partitions, and the relaxation time of the projection chain \mathcal{C}_{f_i} is $\Theta(i^2)$. Since the trace variance of f_i on \mathcal{C} and \mathcal{C}_{f_i} are the same, by applying inequality 46, we get $v_\tau(f_i) \leq \Theta(i^2/\tau)$.

We now proceed to prove that $v_\tau(f_i) \geq \Theta(i^2/\tau)$. Let $S_0 = \{i/3 + 1, i/3 + 2, \dots, 2i/3\} \subseteq [n]$. In the appendix by an application of Lévy's inequality and Hoeffding's inequality, we show that the trace variance of any trace of length at most i^2 conditioned starting at S_0 is $\Theta(1)$. Similarly, we define S_k to be $S_k \doteq \{k(2i) + (i/3 + 1), k(2i) + (i/3 + 2), \dots, k(2i) + 2i/3\}$, and using a similar argument we can show that any walk starting at $S = \bigcup_{k=0}^{n/2i} S_k$ has trace variance at least constant. Since the stationary probability of S is $\frac{1}{3}$, we conclude $v_\tau \geq \Theta(1)$. See Figure 7.2 for visualization of f_i on \mathcal{C} , and the region S . \square

7.2.2 Classic MCMC-Mean Estimators

In contrast to our proposed algorithm DYNAMITE, which increases the sample size dynamically, the classic methods are *static*; they receive a fixed m as input and run the chain for m steps to generate the trace $\mathbf{X}_{1:m} : X_1, X_2, \dots, X_m$. These algorithms then returns the empirical mean $\hat{\mu}$ as the output.

In the following theorems $\mathbf{X}_{1:m} : X_1, X_2, \dots, X_m$, is a length m stationary trace of \mathcal{M} , (i.e., $\mathbf{X}_{1:m} \sim \pi^{(T)}$), with mixing time τ_{mix} , relaxation time τ_{rel} , and second largest eigenvalue λ . Based on each of these theorems we obtain the sample complexity (sufficient m) of a variance-aware or variance-agnostic algorithm. Remember that f is a bounded function with range $[a, b]$, we define $r \doteq b - a$ as the range of f . We use the following definition for *sample complexity*.

Definition 7.2.4 (Sample complexity). Assume algorithm A runs a Markov chain for m steps and generates a trace X_1, X_2, \dots, X_m . Algorithm A uses this trace to find an estimation for μ , namely $\hat{\mu}$. If there exists m_A such that for any $m \geq m_A$, we have $\mathbb{P}(|\mu - \hat{\mu}| \geq \varepsilon) \leq \delta$, we call m_A the *sample complexity* of algorithm A .

Theorem 7.2.5 (Hoeffding-Type Bounds for Mixing Processes, (see Thm. 2.1 of [Fan et al., 2018])). For any $\delta \in (0, 1)$, we have

$$\mathbb{P}\left(|\hat{\mu} - \mu| \geq \sqrt{\frac{2(1 + \lambda)(\frac{r^2}{4}) \ln(\frac{2}{\delta})}{(1 - \lambda)m}}\right) \leq \delta. \quad (47)$$

This implies sample complexity

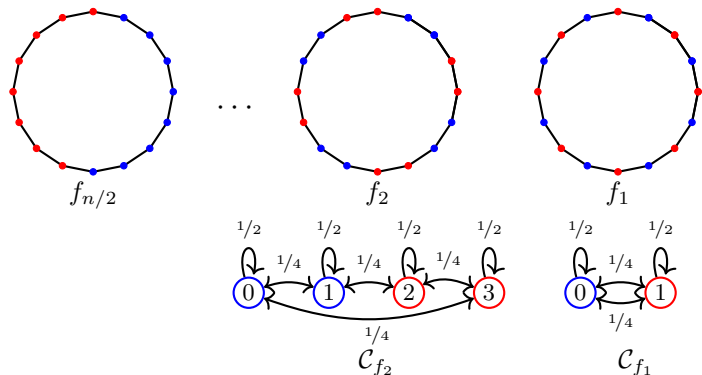
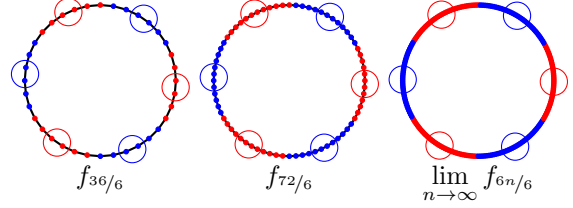


Figure 7.1: The image of f_i s on a cycle of length $n = 16$, and corresponding *projection chains* with the same trace variance.

Figure 7.2: Image of $f_{n/6}$ on a cycle for various values of n , zero values are colored in red, and one values in blue. The trace variance for trace length than $(n/52)^2$, conditioned starting at $S = \bigcup_{k=1}^6 S_k$ (circled regions), is constant.



$$m_H(\lambda, r, \varepsilon, \delta) = \frac{1 + \lambda}{1 - \lambda} \ln\left(\frac{2}{\delta}\right) \frac{r^2}{2\varepsilon^2} \in \Theta\left(\tau_{\text{rel}} \ln\left(\frac{1}{\delta}\right) \frac{r^2}{\varepsilon^2}\right) .$$

Largely due to their simplicity and convenience, Hoeffding-type bounds enjoy ubiquity in computer science. Unfortunately, they are generally much looser than variance-aware bounds, as the dependence on variance v is replaced by *variance proxy* $(\frac{r}{2})^2$ (maximal assuming range r). While it is not generally possible to remove dependence on r entirely, Bernstein-type inequalities greatly reduce this dependence, often yielding *asymptotic improvement* to sample complexity Audibert et al. [2007], Mnih et al. [2008], Audibert et al. [2009], Maurer and Pontil [2009].

Theorem 7.2.6 (Bernstein-Type Bound for Mixing Process [Jiang et al., 2018, Thm. 1.2]). ² For any $\delta \in (0, 1)$, we have

$$\mathbb{P}\left(|\hat{\mu} - \mu| \geq \frac{10r \ln(\frac{2}{\delta})}{(1 - \lambda)m} + \sqrt{\frac{2(1 + \lambda)v_{\pi} \ln(\frac{2}{\delta})}{(1 - \lambda)m}}\right) \leq \delta . \quad (48)$$

Sample complexity of the static variance-aware algorithm: This implies sample complexity

$$m_B(\lambda, r, v, \varepsilon, \delta) = \frac{2}{1 - \lambda} \ln\left(\frac{2}{\delta}\right) \left(\frac{5r}{\varepsilon} + \frac{(1 + \lambda)v_{\pi}}{\varepsilon^2}\right) \in \Theta\left(\tau_{\text{rel}} \ln\left(\frac{1}{\delta}\right) \left(\frac{r}{\varepsilon} + \frac{v_{\pi}}{\varepsilon^2}\right)\right) .$$

7.3 DynaMITE

In this section we present the DYNAMITE: DYNAMIC MCMC INTER-TRACE VARIANCE ESTIMATION method. We show that its sample complexity is dependent on the *a priori unknown* trace variance v_T . As discussed previously, in various scenarios, v_T is significantly smaller than the stationary variance v_{π} .

Our algorithm consists of three subroutines: (1) MCMCPRO employs *progressive sampling*, beginning with a small sample obtained by running the chain, and we *progressively* increase the sample size until a *stopping condition* is met, namely the mean is ε - δ well-estimated. In each round of MCMCPRO, we bound v_T using an unbiased estimator of it, which we obtain from running two *independent copies* of \mathcal{M} in parallel (see definition 7.3.4). In order to find a suitable stopping conditioning guaranteeing that the empirical mean is sufficiently accurate, we employ the Bernstein bound (theorem 7.2.6). (2) Our main procedure DYNAMITE runs the above progressive schedule on a *trace chain* (defined in definition 7.3.3), using the *trace averaging technique* explained in section 7.2. (3) Finally, in procedure WARMSTARTDYNAMITE, we begin at an arbitrary (nonstationary) state, but run it chain for at least the *uniform mixing time* τ_{unif} steps before starting sampling, and apply an appropriate *nonstationarity correction*. Note that this is performed only once, and does not significantly impact sample complexity, which thus depends on τ_{rel} rather than τ_{unif} .

The following theorems, proved in appendix F.1, guarantee correctness and efficiency of DYNAMITE.

²Note that thms. 7.2.5 and 7.2.6 assume *stationarity*, i.e., they consider traces $\mathbf{X}_{1:m} : X_1, X_2, \dots, X_m$ assuming $X_1 \sim \pi$. This can be done using what is known as a *warm start*. For more details see subsection F.1.1 in the appendix.

Algorithm 5 MCMC**PRO**, DYNAMITE, and WARMSTARTDYNAMITE routines

- 1: **procedure** MCMC**PRO** $((X_0, X_1), \mathcal{M}, \Lambda, f, \varepsilon, \delta) \mapsto \hat{\mu}$
 - 2: **Input:** Initial state $(X_0, X_1) \in \mathcal{X}^2$ of tensor product chain, *Markov chain* \mathcal{M} over \mathcal{X} , *second absolute eigenvalue upper-bound* Λ , *function* $f : \mathcal{X} \rightarrow [a, b]$ with $r \doteq b - a$, *confidence interval radius* ε , and *failure probability* $\delta \in (0, 1)$.
 - 3: **Output:** Additive ε, δ approximation $\hat{\mu}$ of $\mu = \mathbb{E}_\pi[f]$.
-
- 4: $N \leftarrow \max\left(1, \lceil \log_2\left(\frac{r}{2\varepsilon}\right) \rceil\right)$; $\alpha \leftarrow \frac{(1 + \Lambda)r \ln \frac{3N}{\delta}}{(1 - \Lambda)\varepsilon}$; $m_0 \leftarrow 0$ ▷ Initialize *sampling schedule*
 - 5: **for** $i \in 1, 2, \dots, N$ **do**
 - 6: $m_i \leftarrow \lceil \alpha 2^i \rceil$ ▷ Total sample count at iteration i
 - 7: **for** $j \in (m_{i-1} + 1), \dots, m_i$ **do**
 - 8: $(X_{j,1}, X_{j,2}) \sim (\mathcal{M} \otimes \mathcal{M})(X_{j-1,1}, X_{j-1,2})$ ▷ Run both chains up to step m_i
 - 9: **end for**
 - 10: $\hat{\mu}_i \leftarrow \frac{1}{2m_i} \sum_{j=1}^{m_i} (f \oplus f)(X_{j,0}, X_{j,1})$ ▷ Empirical mean
 - 11: $\hat{v}_i \leftarrow \frac{1}{2m_i} \sum_{j=1}^{m_i} (f \ominus f)^2(X_{j,0}, X_{j,1})$ ▷ Empirical variance
 - 12: $\mathbf{u}_i \leftarrow \hat{v}_i + \frac{(11 + \sqrt{21})(1 + \Lambda/\sqrt{21})r^2 \ln \frac{3N}{\delta}}{(1 - \Lambda)m_i} + \sqrt{\frac{(1 + \Lambda)r^2 \hat{v}_i \ln \frac{3N}{\delta}}{(1 - \Lambda)m_i}}$ ▷ Variance upper bound
 - 13: $\hat{\varepsilon}_i \leftarrow \frac{10r \ln \frac{3N}{\delta}}{(1 - \Lambda)m_i} + \sqrt{\frac{(1 + \Lambda)\mathbf{u}_i \ln \frac{3N}{\delta}}{(1 - \Lambda)m_i}}$ ▷ Apply Bernstein bound
 - 14: **if** $(i = N) \vee (\hat{\varepsilon}_i \leq \varepsilon)$ **then** ▷ Terminate if accuracy guarantee is met
 - 15: **return** $\hat{\mu}_i$
 - 16: **end if**
 - 17: **end for**
 - 18: **end procedure**
-
- 19: **procedure** DYNAMITE $((x_0, x_1), \mathcal{M}, \Lambda, f, \varepsilon, \delta) \mapsto \hat{\mu}$
 - 20: **Input:** Initial state $(x_0, x_1) \in \mathcal{X} \times \mathcal{X}$, *lazy Markov chain* \mathcal{M} over \mathcal{X} , *second absolute eigenvalue upper-bound* Λ , *function* $f : \mathcal{X} \rightarrow [a, b]$, *confidence interval radius* ε , and *failure probability* $\delta \in (0, 1)$.
 - 21: **Output:** Additive ε, δ approximation $\hat{\mu}$ of $\mu = \mathbb{E}_\pi[f]$.
-
- 22: $T \leftarrow \lceil \frac{1+\Lambda}{1-\Lambda} \ln \sqrt{2} \rceil$ ▷ Select T s.t. $\tau_{\text{rel}}(\mathcal{M}^{(T)}) \leq 2$
 - 23: $(X_0, X_1) \leftarrow \left(\underbrace{(\omega_1, \omega_2, \omega_3, \dots, \omega_{T-1})}_{\text{ARBITRARY } \omega_{1:T-1} \in \Omega^{T-1}}, \underbrace{(x_0, (\omega_1, \omega_2, \omega_3, \dots, \omega_{T-1}), x_1)}_{\text{ARBITRARY } \omega_{1:T-1} \in \Omega^{T-1}} \right)$ ▷ Initialize trace chain state
 - 24: **return** MCMC**PRO** $((X_0, X_1), \mathcal{M}^{(T)}, \Lambda^T, f_{\text{avg}}, \varepsilon, \delta)$ ▷ Run MCMC**PRO** on trace chain
 - 25: **end procedure**
-
- 26: **procedure** WARMSTARTDYNAMITE $(\omega_0, \mathcal{M}, \Lambda, \pi_{\min}, f, \varepsilon, \delta) \mapsto \hat{\mu}$
 - 27: **Input:** Initial state $\omega_0 \in \text{Support}(\pi(\mathcal{M}))$, *lazy reversible Markov chain* \mathcal{M} over \mathcal{X} , *second absolute eigenvalue upper-bound* Λ , *minimum probability lower-bound* π_{\min} , *function* $f : \mathcal{X} \rightarrow [a, b]$, *confidence interval radius* ε , and *failure probability* $\delta \in (0, 1)$.
 - 28: **Output:** Additive ε, δ approximation $\hat{\mu}$ of $\mu = \mathbb{E}_\pi[f]$.
-
- 29: $\tau_{\text{unif}} \leftarrow \left\lceil \frac{\ln \frac{1}{\pi_{\min}}}{\ln \frac{1}{\Lambda}} \right\rceil$ ▷ *Uniform mixing time* bound
 - 30: $(x_0, x_1) \sim (\mathcal{M} \otimes \mathcal{M})^{\tau_{\text{unif}}}(x_0, x_0)$ ▷ Run product chain until *uniformly mixed*
 - 31: **return** DYNAMITE $((x_0, x_1), \mathcal{M}, \Lambda, f, \varepsilon, \frac{\delta}{4})$ ▷ Run DYNAMITE with nonstationarity correction
 - 32: **end procedure**
-

Theorem 7.3.1 (Correctness of DYNAMITE). Consider a Markov chain \mathcal{M} over Ω , and assume its second absolute eigenvalue is less than Λ . Assume we have function $f : \Omega \rightarrow [a, b]$ with $r \doteq b - a$, and we want to estimate *true mean*, $\mu \doteq \mathbb{E}_\pi[f]$ with precision $\varepsilon > 0$, and confidence $(1 - \delta) \in (0, 1)$.

If we start at stationarity, i.e., $(x_0, x_1) \sim \pi(\mathcal{M} \otimes \mathcal{M})$, we take *mean estimate* $\hat{\mu}$ as either

1. $\hat{\mu} \leftarrow \text{MCMCPRo}((x_0, x_1), \mathcal{M}, \Lambda, f, \varepsilon, \delta)$; or
2. $\hat{\mu} \leftarrow \text{DYNAMITE}((x_0, x_1), \mathcal{M}, \Lambda, f, \varepsilon, \delta)$ for lazy \mathcal{M} .

More generally, for a nonstationary start from any $\omega \in \text{Support}(\pi(\mathcal{M}))$, given minimum supported stationary probability at least π_{\min} , we may take

3. $\hat{\mu} \leftarrow \text{WARMSTARTDYNAMITE}(\omega, \mathcal{M}, \Lambda, \pi_{\min}, f, \varepsilon, \delta)$ for lazy, reversible \mathcal{M} .

In all cases, with probability at least $1 - \delta$, we have $|\hat{\mu} - \mu| \leq \varepsilon$.

Theorem 7.3.2 (Efficiency of DYNAMITE). Suppose as in theorem 7.3.1. With probability at least $1 - \delta$, it holds that total sample consumption \hat{m} of DYNAMITE obeys

$$\hat{m} \in \mathcal{O} \left(\log \left(\frac{\log(r/\varepsilon)}{\delta} \right) \left(\frac{r}{(1-\Lambda)\varepsilon} + \frac{\tau_{\text{rel}} v_{\text{rel}}}{\varepsilon^2} \right) \right)$$

Note on input parameters and their interplay with efficiency The improvement of DYNAMITE over standard methods discussed in section 7.2 is threefold: (1) First, in MCMCPRo, we develop a *dynamic sampling schedule*, which is called in line 24 of DYNAMITE. Note that MCMCPRo also works if we simply run it on the original chain. In this simpler version, the sample complexity depends on the *a priori unknown* v_π . Thus we get an improvement over standard methods, whose complexity depend on loose upper bound V . (2) In DYNAMITE, we employ the trace averaging technique to reduce the sample complexity from Tv_π to Tv_T , thus, if $v_T = o(v_\pi)$ we obtain asymptotic improvement over standard methods (again we require no *a priori* knowledge of v_T). (3) Since it always holds $Tv_T \leq 2\tau_{\text{rel}}v_{\text{rel}}$, DYNAMITE’s complexity depends on $\tau_{\text{rel}}v_{\text{rel}}$, i.e., the *a priori unknown* τ_{rel} and v_{rel} . Thus, even in the worst case when $v_{\text{rel}} \approx v_\pi$, our algorithm outperforms static MCMC-mean estimators whose complexity is TV (see theorem 7.2.6 in section 7.2).

Starting at Nonstationarity Often we can’t assume even a single (perfectly) stationary sample may be efficiently drawn, and realistically can only start with an *arbitrary* $\omega \in \text{Support}(\pi)$. WARMSTARTDYNAMITE uses a standard *warm start technique* (widely used in MCMC algorithms), where the chain is run from an arbitrary point for its *uniform mixing time*, and then DYNAMITE is run, applying a standard *nonstationarity correction* (see appendix F.1.1) to account for the nonstationary start. Note that this correction applies to both theorems 7.2.5 and 7.2.6 of section 7.2.

Prelude to a Proof We show theorem 3.1 in several steps. First, given \mathcal{M} , we define a *trace chain* $\mathcal{M}^{(T)}$, whose state space is all possible traces of length T (definition 7.3.3), and we show that for $T \geq \tau_{\text{mix}}(\mathcal{M})$, mixing and relaxation times of $\mathcal{M}^{(T)}$ are at most constant (lemma F.1.5). Second, in definition 7.3.4 we introduce a novel MCMC *unbiased variance estimator*, which runs two copies of \mathcal{M} in parallel, denoted $\mathcal{M} \otimes \mathcal{M}$, and averages the square difference of f in them, i.e., we note that $v_\pi = \frac{1}{2} \mathbb{E}_{\pi(\mathcal{M} \otimes \mathcal{M})}[(f \ominus f)^2]$. We prove the correctness of our variance estimator in lemma 7.3.5, and in line 11 of DYNAMITE we employ it to estimate v_T . Third, we use variance-sensitive concentration bounds for mixing processes to bound the *true mean* μ in terms of the *empirical mean* $\hat{\mu}$ and *empirical variance* \hat{v} . Note that since we apply the variance estimation algorithm to the trace chain $\mathcal{M}^{(T)}$, \hat{v} is an estimator for v_T . Since our algorithm depends on the *empirical variance*, we don’t know *a priori* how many samples will be required to estimate the mean, so we run with a *doubling sampling schedule* until the desired guarantee is met.

We now define the T -trace Markov chain of \mathcal{M} , which we denote $\mathcal{M}^{(T)}$.³ We use the trace chain knowing (for proof, see lemma F.1.5) that $\mathcal{M}^{(T)}$ has constant mixing and relaxation time for $T \geq \tau_{\text{rel}}$.

Definition 7.3.3 (Trace chain). For a Markov chain \mathcal{M} on state space Ω , we define $\mathcal{M}^{(T)}$ on state space Ω^T as follows: given $\mathbf{a} = (a_1, a_2, \dots, a_T)$ and $\mathbf{b} = (b_1, b_2, \dots, b_T)$ in Ω^T , the probability of going from \mathbf{a} to \mathbf{b} is $\mathcal{M}^{(T)}(\mathbf{a}, \mathbf{b}) \doteq \mathcal{M}(a_T, b_1) \prod_{i=1}^{T-1} \mathcal{M}(b_i, b_{i+1})$.

The following definition provides us with an *unbiased MCMC estimator for variance*.

Definition 7.3.4 (Unbiased MCMC variance estimator). Suppose function $f : \mathcal{X} \mapsto \mathbb{R}$, chain \mathcal{M} over \mathcal{X} , and assume \mathbf{X}_1 and \mathbf{X}_2 are two length m independent traces of \mathcal{M} . The following definition for \hat{v} satisfies $\mathbb{E}[\hat{v}] = v_\pi$, thus \hat{v} is an unbiased estimator for v_π .

$$\hat{v} \doteq \frac{1}{2m} \sum_{i=1}^m (f(X_{1,i}) - f(X_{2,i}))^2 .$$

In other words, we can think \mathbf{X}_1 and \mathbf{X}_2 as the *two factors* of a single trace \mathbf{X} of the *tensor product chain* $\mathcal{M} \otimes \mathcal{M}$, which steps forward by simultaneously running both factor chains for one step. Given a sample $\mathbf{X} \sim (\mathcal{M} \otimes \mathcal{M})^{(m)}$, this variance estimate \hat{v} can then be represented as the *empirical mean* of half the *square difference function*, i.e., the function $\frac{1}{2}(f \ominus f)^2$.

Note that applying the above estimator to \mathcal{M} yields an unbiased estimate of v_π , but applying it to $\mathcal{M}^{(T)}$ yields an unbiased estimator of v_T , thus it is sufficient to estimate both the stationary and trace variances. The following lemma, proved in the appendix, bounds the error of this estimator.

Lemma 7.3.5 (Trace Variance Estimation). Let \hat{v} be defined as in definition 7.3.4, it then holds that, for any $\delta \in (0, 1)$, we have

$$\mathbb{P} \left(v_\pi \geq \hat{v} + \frac{(11 + \sqrt{21})(1 + \frac{\lambda}{\sqrt{21}})r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m} + \sqrt{\frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} \right) \leq \delta .$$

Sampling Schedule and Proof Sketch We first note that correctness of DYNAMITE and WARM-STARTDYNAMITE follow from that of MCMCPRO, as they simply apply MCMCPRO to a trace chain and/or apply a *nonstationarity correction*.

The key to MCMCPRO is an *a priori* fixed *sampling schedule*, which determines the sizes of *progressively larger samples*. To ensure correctness, a sequence of tail bounds must all hold simultaneously w.h.p. by union bound. The main difficulty is that the schedule length N and the probability concentration bounds are codependent. A shorter schedule is more *statistically efficient*, but can *overshoot* the sufficient sample size, and the opposite holds for longer schedules. We explain how to resolve this cyclic dependence in the next paragraph.

Over a run of DYNAMITE, we take (up to) $3N$ probability concentration bounds (N bounds for variance, line 12, and N bounds each for upper and lower mean bounds, line 13). We first establish the *worst-case* Hoeffding sample complexity m_N (theorem 7.2.5) and *best-case* Bernstein sample complexity α (theorem 7.2.6) by taking $v = 0$,

$$m_I \doteq m_H(\Lambda, r, \varepsilon, \frac{2\delta}{3N}) = \frac{(1 + \Lambda)r^2 \ln \frac{3N}{\delta}}{2(1 - \Lambda)\varepsilon^2} , \quad \& \quad \alpha \doteq \frac{(1 + \Lambda)r \ln \frac{3N}{\delta}}{(1 - \Lambda)\varepsilon} \approx m_B(\Lambda, r, 0, \varepsilon, \frac{2\delta}{3N}) .$$

We know from these bounds that once a sample of size m_N is drawn, the desired guarantee has been met, and similarly, before a sample of size $\approx \alpha$ is drawn, the desired Bernstein bound *can not* be met. We now select the minimal N such that each sample size m_i obeys $m_i \leq 2m_{i-1}$. Observe that in the ratio $\frac{m_N}{\alpha}$,

³We use parentheses in the exponent to preserve standard matrix powering notation.

all dependence on N is divided out ($\ln \frac{3N}{\delta}$ terms cancel). We initially run the chain for m_1 to be $\lceil 2\alpha \rceil$, and in iteration i , we run up to $\lceil 2\alpha^i \rceil$ steps, thus we conclude $N \doteq \lceil \log_2(\frac{r}{2\varepsilon}) \rceil \geq \lceil \log_2(\frac{r}{2\varepsilon}) \rceil - 1$ (doubling) iterations are sufficient. These computations are repeated verbatim by MCMCPRO (lines 4 & 6) to compute the sampling schedule.

Conclusion and Discussion

Part I introduces new statistical techniques for variance-sensitive uniform convergence bounds and generalization guarantees in machine learning. The centralization strategy of chapter 1 achieves asymptotically-optimal bound convergence rates, and our Monte-Carlo estimation procedure yields sharp bounds with both centralized (theorem 1.3.7) and non-centralized (corollary 1.3.8) Rademacher averages. This has broad implications in machine learning, data science, and statistical settings, and can also be used as a component statistical method, *mutatis mutandis*, in all subsequent chapters. Chapter 2 then continues this theme of efficient use of data by introducing an algorithm that augments a small amount of labeled data with the output of *weak labelers*, which are assumed to be *machine learning models* associated with *correlated tasks*. I show both *computational guarantees* on efficient learnability and statistical guarantees on generalization bounds involving both the number of labeled and unlabeled samples.

Overall, the work of part I may be interpreted as novel methods to reduce the burden of acquiring large amounts of labeled data to train sophisticated machine learning models. Indeed, with the ever-expanding available computational power available (i.e., as characterized by Moore’s law, and more recently with learning on GPUs, FPGAs, and specialized hardware) we have trained more complicated models (in particular, deeper neural networks), but datasets have had to grow commensurately to support training such models without overfitting. However, as computation becomes cheaper, the economic costs of collecting larger datasets do not always scale similarly, leaving learning from small datasets a problem that is both practically important and theoretically challenging. In fact, these methods are particularly impactful, as poor performance and analysis in the small sample setting can be a fairness issue, in the sense that by nature of their size and visibility, more data may be available on majority groups than minority or understudied groups, which contributes to marginalized groups being often poorly-served by machine learning systems.

Following the pure statistical methods of part I, in part II, I define *malfare* parallel to *welfare* (definition 3.3.1), and show that subject to several intuitive and basic axioms (section 3.3.1), all welfare and malfare functions are *power means* (definition 3.3.6, see theorem 3.3.8). I then argue that fair ML should seek to *minimize malfare*, as this naturally generalizes *loss minimization* to multiple groups. I then combine the uniform convergence guarantees of part I with the fairness setting of part II in chapter 4, where I introduce the *fair codec selection problem*. Finally, chapter 5 defines a concept of *fair-PAC-learnability* in terms of malfare minimization, and characterizes both necessary and sufficient conditions for various flavors of fair learnability. Following Valiant’s introduction of classical PAC-learning in order to rigorously characterize machine learning problems in the lexicon of computer scientists, I define malfare minimization as the target of fair-PAC learning, with axiomatically justified fairness characteristics, the precise details of which also lead to interesting theoretical computer science and statistical estimation questions. In particular, I show a hierarchy of PAC-learning and fair-PAC-learning settings, and am excited to see additional connections drawn in future work.

Finally, in part III I leave the i.i.d. setting behind, with two dependent *mean-estimation* settings of great practical and theoretical interest. Chapter 6 describes a common real-world problem in databases, where a user wants to estimate the *mean of a function*, without waiting for the database to iterate over all records. Motivated by the real-world performance characteristics of modern database systems, I show that, while in

the worst case, sampling entire blocks may not improve over sampling individual records, my algorithm is able to detect independence within blocks automatically, and adapts accordingly. Chapter 7 adapts this idea of *algorithmically adapting* to independence detected within the data, and applies it to sample from Markov chains. Such settings are both of great interest to the theoretical computer science community, as increasingly attention is placed on *randomized algorithms*, and also to practitioners in diverse fields, from *significance testing* in the sciences to *machine learning* and *artificial intelligence*, wherein such methods are vital to efficiently reason under partial knowledge in an uncertain world.

I have additionally published works in several related areas that have not made it into this thesis, namely *anomaly detection* [Cousins et al., 2017], *automated machine learning* [Binnig et al., 2015, 2018], *empirical game theory* and *mechanism design* [Viqueira et al., 2019, 2020, 2021], and *statistical data science* [Pellegrina et al., 2020]. Of special note are two somewhat more theoretical papers in pure machine learning. In Cousins and Upfal [2017], I show generalization bounds for a distance-based classifier, including what is to my knowledge the only known *exact (closed form) expression* for the *empirical Rademacher average* of a nontrivial classification model. In Cousins and Riondato [2019], I theoretically unify *classification trees* with the *entropy* impurity criterion and *regression trees* with the *square-error* impurity criterion, as well as an infinite family of parametric conditional-density estimation trees, under the *cross-entropy* information criterion, and show all to be instances of the *minimax-entropy principle*, with consequent efficient training heuristics.

The theme running through all of my work is the importance placed on rigorous guarantees for domain-appropriate properties in various probabilistic processes, most notably in *machine learning* (generalization bounds), *sampling* (mean estimation guarantees), and *data science* (frequent itemsets, equilibria in empirical game theory, etc.) settings. In particular, almost all of my papers rigorously show *finite-sample probabilistic guarantees* in some probabilistic setting, and overall my work illustrates that while such guarantees may induce significant proof complexity, often they asymptotically match or nearly match known lower bounds for sample complexity and approximate central-limit-theorem bounds. I argue that this additional analytical overhead is a price worth paying, as finite-sample guarantees are a necessary component in understanding the behavior and failure modes of machine learning and data science systems. As we see machine learning increasingly employed in our day-to-day lives, and witness the catastrophic consequences of failures of such systems, the importance of rigorous analysis of these methods to avoid negative outcomes becomes increasingly apparent.

Bibliography

- D. Aldous, L. Lovasz, and P. Winkler. Mixing times for uniformly ergodic Markov chains. *Stochastic Processes and their Applications*, 71:165–185, 1997.
- Noga Alon, Shai Ben-David, Nicolo Cesa-Bianchi, and David Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of the ACM (JACM)*, 44(4):615–631, 1997.
- Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. Cambridge University Press, 2009.
- Chidubem Arachie and Bert Huang. Adversarial label learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- Enrique Areyan Viqueira, Amy Greenwald, Cyrus Cousins, and Eli Upfal. Learning simulation-based games from data. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1778–1780. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. Improved algorithms for learning equilibria in simulation-based games. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 79–87, 2020.
- Ahmed Ashraf, Shehroz Khan, Nikhil Bhagwat, Mallar Chakravarty, and Babak Taati. Learning to unlearn: Building immunity to dataset bias in medical imaging studies. *arXiv preprint arXiv:1812.01716*, 2018.
- Anthony B Atkinson et al. On the measurement of inequality. *Journal of Economic Theory*, 2(3):244–263, 1970.
- Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Tuning bandit algorithms in stochastic environments. In *International conference on algorithmic learning theory*, pages 150–165. Springer, 2007.
- Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- Ariful Azad, Georgios Pavlopoulos, Christos Ouzounis, Nikos Kyrpides, and Aydin Buluç. HipMCL: a high-performance parallel implementation of the Markov clustering algorithm for large-scale networks. *Nucleic Acids Research*, 46: 1–11, 01 2018. doi: 10.1093/nar/gkx1313.
- Stephen H. Bach, Bryan He, Alexander Ratner, and Christopher Ré. Learning the structure of generative models without labeled data. In *International Conference on Machine Learning (ICML)*, 2017.
- Stephen H. Bach, Daniel Rodriguez, Yintao Liu, Chong Luo, Haidong Shao, Cassandra Xia, Souvik Sen, Alexander Ratner, Braden Hancock, Houman Alborzi, Rahul Kuchhal, Christopher Ré, and Rob Malkin. Snorkel DryBell: A case study in deploying weak supervision at industrial scale. 2019.
- Peter L Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Anna Ben-Hamou, Roberto I. Oliveira, and Yuval Peres. Estimating graph parameters via random walks with restarts. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, page 1702–1714, USA, 2018. Society for Industrial and Applied Mathematics. ISBN 9781611975031.
- George Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962.
- Suman K. Bera and C. Seshadhri. How to count triangles, without seeing the whole graph. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 306–316, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. doi: 10.1145/3394486.3403073. URL

<https://doi.org/10.1145/3394486.3403073>.

- Sergei Bernstein. On a modification of Chebyshev’s inequality and of the error formula of Laplace. *Annals of the Science Institute SAV. Ukraine, Sect. Math*, 1(4):38–49, 1924.
- D.P. Bertsekas. *Convex Optimization Algorithms*. Athena Scientific, 2015.
- Carsten Binnig, Fuat Basik, Benedetto Buratti, Ugur Cetintemel, Yeounoh Chung, Andrew Crotty, Cyrus Cousins, Dylan Ebert, Philipp Eichmann, Alex Galakatos, et al. Towards interactive data exploration. In *Real-Time Business Intelligence and Analytics*, pages 177–190. Springer, 2015.
- Carsten Binnig, Benedetto Buratti, Yeounoh Chung, Cyrus Cousins, Tim Kraska, Zeyuan Shang, Eli Upfal, Robert Zeleznik, and Emanuel Zraggen. Towards interactive curation & automatic tuning of ML pipelines. In *Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*, pages 1–4, 2018.
- Dankmar Böhning. Multinomial logistic regression algorithm. *Annals of the institute of Statistical Mathematics*, 44(1): 197–200, 1992.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. A sharp concentration inequality with applications. *Random Structures & Algorithms*, 16(3):277–292, 2000.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. Concentration inequalities using the entropy method. *The Annals of Probability*, 31(3):1583–1614, 2003.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. On concentration of self-bounding functions. *Electronic Journal of Probability*, 14:1884–1899, 2009.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Olivier Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *Comptes Rendus Mathématique*, 334(6):495–500, 2002.
- Joseph Bradley and Carlos Guestrin. Sample complexity of composite likelihood. In *Artificial Intelligence and Statistics*, pages 136–160, 2012.
- Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- Andreas Brieden, Peter Gritzmann, Ravindran Kannan, Victor Klee, László Lovász, and Miklós Simonovits. Deterministic and randomized polynomial-time approximation of radii. *Mathematika*, 48(1-2):63–105, 2001.
- G. Brightwell and P. Winkler. Counting linear extensions. *Order*, 8:225–242, 1991.
- Peter S Bullen. *Handbook of means and their inequalities*, volume 560. Springer Science & Business Media, 2013.
- Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.
- Jacqueline G Cavazos, P Jonathon Phillips, Carlos D Castillo, and Alice J O’Toole. Accuracy comparison across face recognition algorithms: Where are we on measuring race bias? *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2020.
- Vincent S Chen, Paroma Varma, Ranjay Krishna, Michael Bernstein, Christopher Ré, and Li Fei-Fei. Scene graph prediction with limited labels. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- Flavio Chierichetti and Shahrzad Haddadan. On the complexity of sampling vertices uniformly from a graph. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 149:1–149:13. Schloß Dagstuhl–Leibniz-Zentrum für Informatik, 2018. ISBN 978-3-95977-076-7.
- Kai-Min Chung, Henry Lam, Zhenming Liu, and Michael Mitzenmacher. Chernoff-Hoeffding bounds for Markov chains: Generalized and simplified. *arXiv preprint arXiv:1201.0559*, 2012.
- Cynthia M Cook, John J Howard, Yevgeniy B Sirotin, Jerry L Tipton, and Arun R Vemury. Demographic effects in facial recognition and their dependence on image acquisition: An evaluation of eleven commercial systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 1(1):32–41, 2019.
- Cyrus Cousins. An axiomatic theory of provably-fair welfare-centric machine learning. *arXiv preprint arXiv:2104.14504*, 2021.
- Cyrus Cousins and Matteo Riondato. CaDET: interpretable parametric conditional density estimation with decision trees and forests. *Machine Learning*, 108(8-9):1613–1634, 2019.

- Cyrus Cousins and Matteo Riondato. Sharp uniform convergence bounds through empirical centralization. In *Advances in Neural Information Processing Systems*, 2020.
- Cyrus Cousins and Eli Upfal. The k -nearest representatives classifier: A distance-based classifier with strong generalization bounds. In *2017 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10. IEEE, 2017.
- Cyrus Cousins, Christopher M Pietras, and Donna K Slonim. Scalable FRaC variants: Anomaly detection for precision medicine. In *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 253–262. IEEE, 2017.
- Cyrus Cousins, Shahrzad Haddadan, and Eli Upfal. Making mean-estimation more efficient using an MCMC trace variance approach: DynaMITE. *arXiv preprint arXiv:2011.11129*, 2020.
- Hugh Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30(119):348–361, 1920.
- Nilesh Dalvi, Anirban Dasgupta, Ravi Kumar, and Vibhor Rastogi. Aggregating crowdsourced binary ratings. In *Proceedings of the 22nd international conference on World Wide Web*, pages 285–294, 2013.
- Vladimir Kolmogorov David G. Harris. Parameter estimation for Gibbs distributions. 2007. URL <https://arxiv.org/abs/2007.10824>.
- A. P. Dawid and A. M. Skene. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Journal of the Royal Statistical Society C*, 28(1):20–28, 1979.
- Gerard Debreu. Topological methods in cardinal utility theory. Technical report, Cowles Foundation for Research in Economics, Yale University, 1959.
- Pierre Del Moral and L. Miclo. Branching and interacting particle systems approximations of Feynman-Kac formulae with applications to non-linear filtering. *Sem. Probab. Stras.*, 34:1–145, 05 2007. doi: 10.1007/BFb0103798.
- Luc Devroye, Matthieu Lerasle, Gábor Lugosi, and Roberto I. Oliveira. Sub-Gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016.
- Persi Diaconis. The Markov chain Monte Carlo revolution. *Bulletin of the American Mathematical Society*, 46:179–205, 04 2009. doi: 10.1090/S0273-0979-08-01238-X.
- Khaled Elbassioni and Hans Raj Tiwary. On computing the vertex centroid of a polyhedron. *arXiv preprint arXiv:0806.3456*, 2008.
- Yonina C Eldar, Amir Beck, and Marc Teboulle. A minimax Chebyshev estimator for bounded error estimation. *IEEE Transactions on Signal Processing*, 56(4):1388–1397, 2008.
- Anton Enright, S Dongen, and C.A. Ouzounis. An efficient algorithm for large-scale detection of protein families. *Nucleic acids research*, 30:1575–84, 05 2002. doi: 10.1093/nar/30.7.1575.
- J. Erfurt, C. R. Helmrich, S. Bosse, H. Schwarz, D. Marpe, and T. Wiegand. A study of the perceptually weighted peak signal-to-noise ratio (WPSNR) for image compression. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2339–2343, Sep. 2019. doi: 10.1109/ICIP.2019.8803307.
- Jianqing Fan, Yuan Liao, and Martina Mincheva. Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society. Series B, Statistical methodology*, 75, 09 2013. doi: 10.1111/rssb.12016.
- Jianqing Fan, Bai Jiang, and Qiang Sun. Hoeffding’s lemma for Markov chains and its applications to statistical learning. *arXiv preprint arXiv:1802.00211*, 2018.
- Yoav Freund. Boosting a weak learning algorithm by majority. *Information and Computation*, 121(2):256–285, 1995.
- Bolin Gao and Laca Pavel. On the properties of the softmax function with application in game theory and reinforcement learning, 2018.
- Chao Gao and Dengyong Zhou. Minimax optimal convergence rates for estimating ground truth from crowdsourced labels. *CoRR*, abs/1207.0016, 2013.
- Arpita Ghosh, Satyen Kale, and Preston McAfee. Who moderates the moderators? Crowdsourcing abuse detection in user-generated content. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 167–176, 2011.
- Robert Gibbons. *Game theory for applied economists*. Princeton University Press, 1992.
- Daniel T. Gillespie. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58(1):35–55, 2007. doi: 10.1146/annurev.physchem.58.032806.104637.

- Evarist Giné and Vladimir Koltchinskii. Concentration inequalities and asymptotic results for ratio type empirical processes. *The Annals of Probability*, 34(3):1143–1216, 2006.
- Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Addison-Wesley Longman Publishing Co., Inc., USA, 2nd edition, 1992. ISBN 0201508036.
- William M Gorman. The structure of utility functions. *The Review of Economic Studies*, 35(4):367–390, 1968.
- V. Guruswami. Rapidly mixing Markov chains: A comparison of techniques (a survey). *ArXiv*, abs/1603.01512, 2016.
- Uffe Haagerup. The best constants in the Khintchine inequality. *Studia Mathematica*, 70(3):231–283, 1982.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Robert Hegemann, Alexander Leiding, and Rogério Brito. LAME MP3 encoder, 2017. URL <http://lame.sourceforge.net/>.
- Hoda Heidari, Claudio Ferrari, Krishna Gummadi, and Andreas Krause. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Advances in Neural Information Processing Systems*, pages 1265–1276, 2018.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- Darrell Hoy, Denis Nekipelov, and Vasilis Syrgkanis. Welfare guarantees from data. In *Advances in Neural Information Processing Systems*, pages 3768–3777, 2017.
- Lily Hu and Yiling Chen. Fair classification and social welfare. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 535–545, 2020.
- Mark Huber. Approximation algorithms for the normalizing constant of Gibbs distributions. *The Annals of Applied Probability*, 25(2):974–985, 2015. ISSN 10505164. URL <http://www.jstor.org/stable/24519939>.
- Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, pages 507–523. Springer, 2011.
- Bai Jiang, Qiang Sun, and Jianqing Fan. Bernstein’s inequality for general Markov chains. *arXiv preprint arXiv:1805.10721*, 2018.
- David R. Karger, Sewoong Oh, and Devavrat Shah. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 62(1):1–24, 2014.
- Maximilian Kasy and Rediet Abebe. Fairness, equality, and power in algorithmic decision-making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 576–586, 2021.
- Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67, page 43. Schloß Dagstuhl–Leibniz-Zentrum für Informatik, 2017.
- V. Kolmogorov. A faster approximation algorithm for the Gibbs partition function. *ArXiv*, abs/1608.04223, 2018.
- Vladimir Koltchinskii. Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47(5):1902–1914, 2001.
- Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, pages 489–504. ACM, 2018.
- Carlos Leon and François Perron. Optimal Hoeffding bounds for discrete reversible Markov chains. *The Annals of Applied Probability*, 14, 05 2004. doi: 10.1214/105051604000000170.
- David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Society, 2017.
- Pascal Lezaud. Chernoff-type bound for finite Markov chains. *The Annals of Applied Probability*, 8, 08 1998. doi: 10.1214/aoap/1028903453.
- László Lovász and Peter Winkler. Mixing times. *AMS DIMACS Series*, 41:189–204, 1998.
- Gábor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190, 2019.

- Brian Mac Namee, Padraig Cunningham, Stephen Byrne, and Owen I Corrigan. The problem of bias in training data in regression problems in medical decision support. *Artificial intelligence in medicine*, 24(1):51–70, 2002.
- Pascal Massart. Some applications of concentration inequalities to statistics. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 9, pages 245–303, 2000.
- Andreas Maurer. Concentration inequalities for functions of independent variables. *Random Structures & Algorithms*, 29(2):121–138, 2006.
- Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- A. Mazzetto, D. Sam, A. Park, E. Upfal, and S. H. Bach. Semi-supervised aggregation of dependent weak supervision sources with performance guarantees. In *Artificial Intelligence and Statistics (AISTATS)*, 2021.
- Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.
- Colin McDiarmid and Bruce Reed. Concentration for self-bounding functions and an inequality of Talagrand. *Random Structures & Algorithms*, 29(4):549–557, 2006.
- Hugh McGuire. LibriVox: Free public domain audiobooks. URL <https://librivox.org/>.
- Blazej Miasojedow. Hoeffding’s inequalities for geometrically ergodic Markov chains on general state space. *Statistics & Probability Letters*, 87, 01 2012. doi: 10.1016/j.spl.2014.01.013.
- Stefania Mignani and Rosa Rodolfo. Markov chain Monte Carlo in statistical mechanics: The problem of accuracy. *Technometrics*, 43:347–55, 2001.
- Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, second edition, 2017.
- Volodymyr Mnih, Csaba Szepesvári, and Jean-Yves Audibert. Empirical Bernstein stopping. In *Proceedings of the 25th international conference on Machine learning*, pages 672–679. ACM, 2008.
- Hervé Moulin. *Fair division and collective welfare*. MIT Press, 2004.
- Cristóbal A. Navarro, Wei Huang, and Youjin Deng. Adaptive multi-GPU exchange Monte Carlo for the 3D random field Ising model. *Computer Physics Communications*, 205:48–60, 2016.
- John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.
- Luca Oneto, Alessandro Ghio, Davide Anguita, and Sandro Ridella. An improved analysis of the Rademacher data-dependent bound using its self bounding property. *Neural Networks*, 44:107–111, 2013.
- Daniel Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.
- Leonardo Pellegrina, Matteo Riondato, and Fabio Vandin. SPuManTE: Significant pattern mining with unconditional testing. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1528–1538. ACM, 2019.
- Leonardo Pellegrina, Cyrus Cousins, Fabio Vandin, and Matteo Riondato. MCRapper: Monte-Carlo Rademacher averages for POSET families and approximate pattern mining. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2165–2174, 2020.
- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1406–1415, 2019.
- P. H. Peskun. Optimum Monte-Carlo sampling using Markov chain. *Biometrika*, 60(3):607–612, 1973. URL [JSTOR, www.jstor.org/stable/2335011](https://www.jstor.org/stable/2335011).
- Andrea Pietracaprina, Matteo Riondato, Eli Upfal, and Fabio Vandin. Mining top- k frequent itemsets through progressive sampling. *Data Mining and Knowledge Discovery*, 21(2):310–326, 2010.
- Arthur Cecil Pigou. *Wealth and welfare*. Macmillan and Company, limited, 1912.
- David Pollard. *Convergence of Stochastic Processes*. Springer, New York, 1984.
- Dana Randall. Rapidly mixing Markov chains with applications in computer science and physics. 8(2):30–41, March 2006. ISSN 1521-9615. doi: 10.1109/MCSE.2006.30. URL <https://doi.org/10.1109/MCSE.2006.30>.
- Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid

- training data creation with weak supervision. *Proceedings of the VLDB Endowment*, 11(3):269–282, 2017.
- Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. Data programming: Creating large training sets, quickly. In *Neural Information Processing Systems (NeurIPS)*, 2016.
- Alexander J Ratner, Braden Hancock, Jared Dunnmon, Frederic Sala, Shreyash Pandey, and Christopher Ré. Training complex models with multi-task weak supervision. In *AAAI*, 2019.
- Matteo Riondato and Eli Upfal. Mining frequent itemsets through progressive sampling with Rademacher averages. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1005–1014. ACM, 2015.
- Matteo Riondato and Eli Upfal. ABRA: Approximating betweenness centrality in static and dynamic graphs with Rademacher averages. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1145–1154. ACM, 2016.
- Matteo Riondato and Eli Upfal. ABRA: Approximating betweenness centrality in static and dynamic graphs with Rademacher averages. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 12(5):61, 2018.
- Kevin WS Roberts. Interpersonal comparability and social choice theory. *The Review of Economic Studies*, pages 421–439, 1980.
- Esther Rolf, Max Simchowitz, Sarah Dean, Lydia T Liu, Daniel Björkegren, Moritz Hardt, and Joshua Blumenstock. Balancing competing objectives with noisy data: Score-based classifiers for welfare-aware machine learning. *arXiv preprint arXiv:2003.06740*, 2020.
- Esteban Safranchik, Shiyang Luo, and Stephen H. Bach. Weakly supervised sequence tagging from noisy rules. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- M. Salatino, J. Austermann, J. A. Beall, S. Choi, K. T. Crowley, S. Duff, S. W. Henderson, G. Hilton, S. Ho, J. Hubmayr, Y. Li, M. D. Niemack, S. M. Simon, S. T. Staggs, and E. J. Wollack. Machine learning, Markov chain Monte Carlo, and optimal algorithms to characterize the AdvACT kilopixel transition-edge sensor arrays. *IEEE Transactions on Applied Superconductivity*, 29(5):1–5, 2019.
- Paul-Marie Samson. Infimum-convolution description of concentration properties of product probability measures, with applications. In *Annales de l’IHP Probabilités et statistiques*, volume 43, pages 321–338, 2007.
- Sandvine. The global internet phenomena report October 2018. <https://www.sandvine.com/hubfs/downloads/phenomena/2018-phenomena-report.pdf>, October 2018.
- Clayton Sanford, Cyrus Cousins, and Eli Upfal. Uniform convergence bounds for codec selection. *arXiv preprint arXiv:1812.07568*, 2018.
- Robert E Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.
- Amartya Sen. On weights and measures: Informational constraints in social welfare analysis. *Econometrica: Journal of the Econometric Society*, pages 1539–1572, 1977.
- Amartya Sen, Master Amartya Sen, Sen Amartya, James E Foster, James E Foster, et al. *On economic inequality*. Oxford university press, 1997.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- Naum Zuselevich Shor. *Minimization methods for non-differentiable functions*, volume 3. Springer Science & Business Media, 2012.
- Till Speicher, Hoda Heidari, Nina Grgic-Hlaca, Krishna P Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. A unified approach to quantifying algorithmic unfairness: Measuring individual & group unfairness via inequality indices. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2239–2248, 2018.
- D. Stefankovic, S. Vempala, and E. Vigoda. Adaptive simulated annealing: A near-optimal connection between sampling and counting. In *48th Annual IEEE Symposium on Foundations of Computer Science*, pages 183–193, 2007. doi: 10.1109/FOCS.2007.67.
- Henri Theil. Economics and information theory. Technical report, Econometric Institute, Netherlands School of Economics, 1967.
- Thilo Thiede, William C. Treurniet, Roland Bitto, Christian Schmidmer, Thomas Sporer, John G. Beerends, Catherine Colomes, Michael Keyhl, Gerhard Stoll, Karlheinz Brandenburg, and Bernhard Feiten. PEAQ—the ITU standard

- for objective measurement of perceived audio quality. *Journal of the Audio Engineering Society*, 48(1):3–29, 2000.
- Philip S Thomas, Bruno Castro da Silva, Andrew G Barto, Stephen Giguere, Yuriy Brun, and Emma Brunskill. Preventing undesirable behavior of intelligent machines. *Science*, 366(6468):999–1004, 2019.
- Z. Tu and Song Zhu. Image segmentation by data-driven Markov chain Monte Carlo. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24:657–673, 06 2002. doi: 10.1109/34.1000239.
- Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- AW Van der Vaart and JA Wellner. *Weak Convergence and Empirical Processes*. Springer, New York, 1996.
- Fabio Vandin, Eli Upfal, and Ben Raphael. Algorithms for detecting significantly mutated pathways in cancer. *Journal of computational biology : a journal of computational molecular cell biology*, 18:507–22, 03 2011. doi: 10.1089/cmb.2010.0265.
- Fabio Vandin, Patrick Clay, Eli Upfal, and Benjamin J Raphael. Discovery of mutated subnetworks associated with clinical data in cancer. In *Biocomputing 2012*, pages 55–66. World Scientific, 2012a.
- Fabio Vandin, Eli Upfal, and Benjamin J Raphael. De novo discovery of mutated driver pathways in cancer. *Genome research*, 22(2):375–385, 2012b.
- Fabio Vandin, Benjamin J Raphael, and Eli Upfal. On the sample complexity of cancer pathways identification. *Journal of Computational Biology*, 23(1):30–41, 2016.
- Vladimir Vapnik. Principles of risk minimization for learning theory. In *Advances in neural information processing systems*, pages 831–838, 1992.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- Paroma Varma, Bryan He, Payal Bajaj, Imon Banerjee, Nishith Khandwala, Daniel L Rubin, and Christopher Ré. Inferring generative model structure with static analysis. *Advances in neural information processing systems*, 30: 239, 2017.
- Paroma Varma, Frederic Sala, Ann He, Alexander Ratner, and Christopher Ré. Learning dependency structures for weak supervision models. In *International Conference on Machine Learning (ICML)*, 2019.
- Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. Learning simulation-based games from data. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019.
- Enrique Areyan Viqueira, Cyrus Cousins, Yasser Mohammad, and Amy Greenwald. Empirical mechanism design: Designing mechanisms from data. In *Uncertainty in Artificial Intelligence*, pages 1094–1104. PMLR, 2020.
- Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. Learning competitive equilibria in noisy combinatorial markets. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021.
- Weiran Wang and Miguel A Carreira-Perpinán. Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. *arXiv preprint arXiv:1309.1541*, 2013.
- David Bruce Wilson. Mixing times of lozenge tiling and card shuffling Markov chains. *The Annals of Applied Probability*, (1):274–325, 2004.
- Sen Wu, Luke Hsiao, Xiao Cheng, Braden Hancock, Theodoros Rekatsinas, Philip Levis, and Christopher Ré. Fonduer: Knowledge base construction from richly formatted data. In *International Conference on Management of Data*, 2018.
- Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2018.
- Ming Yuan and Christina Kendzierski. Hidden Markov models for microarray time course data in multiple biological conditions. *Journal of the American Statistical Association*, 101:1323–1332, 02 2006. doi: 10.1198/016214505000000394.
- Cha Zhang and Yunqian Ma. *Ensemble Machine Learning: Methods and Applications*. Springer, 2012.
- Yuchen Zhang, Xi Chen, Dengyong Zhou, and Michael I Jordan. Spectral methods meet EM: A provably optimal algorithm for crowdsourcing. *The Journal of Machine Learning Research*, 17(1):3537–3580, 2016.

Appendices and Supplementary Material

Appendix A

Supplementary Material for Chapter 1

A.1 Proofs

Lemma 1.2.1. Suppose $m \geq 4$. Then

$$\frac{\mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right]}{1 + 2b(m)} \leq \mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \frac{\mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right]}{1 - 2b(m)} .$$

Proof. We first show the rightmost inequality. Starting from the definition of the RA of the distributional centralization, and then subtracting and adding $\hat{\mathbb{E}}_{\mathbf{x}}[f]$, it holds

$$\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) = \mathbb{E}_{\sigma, \mathbf{x}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \left((f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f]) + (\hat{\mathbb{E}}_{\mathbf{x}}[f] - \mathbb{E}_{\mathcal{D}}[f]) \right) \right| \right] .$$

The subadditivity of the supremum and of the absolute value, and the linearity of the expectation allow us to split the r.h.s. into two summands and obtain

$$\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \mathbb{E}_{\sigma, \mathbf{x}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f]) \right| \right] + \mathbb{E}_{\sigma, \mathbf{x}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i (\hat{\mathbb{E}}_{\mathbf{x}}[f] - \mathbb{E}_{\mathcal{D}}[f]) \right| \right] .$$

Both terms on the r.h.s. can be seen as expectations w.r.t. \mathbf{x} of the ERAs on \mathbf{x} of two sample-dependent families: the empirical centralization of \mathcal{F} , and the family

$$\mathcal{K}_{\mathbf{x}} \doteq \{y \mapsto \hat{\mathbb{E}}_{\mathbf{x}}[f] - \mathbb{E}_{\mathcal{D}}[f], f \in \mathcal{F}\} .$$

Each function in $\mathcal{K}_{\mathbf{x}}$ is *constant*. Thus, we can write

$$\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right] + \mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\mathcal{K}_{\mathbf{x}}, \mathbf{x}) \right] . \quad (49)$$

Using (7) and the linearity of expectation we have that, for each $\mathbf{x} \in \mathcal{X}^m$, it holds

$$\hat{\mathbf{R}}_m(\mathcal{K}_{\mathbf{x}}, \mathbf{x}) = \sup_{f \in \mathcal{F}} |\hat{\mathbb{E}}_{\mathbf{x}}[f] - \mathbb{E}_{\mathcal{D}}[f]| b(m) = \text{SD}(\mathcal{F}, \mathbf{x}) b(m) = \text{SD}(C_{\mathcal{D}}(\mathcal{F}), \mathbf{x}) b(m), \quad (50)$$

where in the last step we use the fact that the SD is invariant to shifting of functions. Continuing from (49) and using (50) and the rightmost inequality of (6), we obtain

$$\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \mathbb{E}_{\mathbf{x}} \left[\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right] + 2\mathbf{R}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) b(m) .$$

The hypothesis $m \geq 4$ implies $1 - 2b(m) > 0$ (see (7)), so we can rewrite the above as

$$\mathfrak{R}_m(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \leq \frac{1}{1 - 2b(m)} \mathbb{E}_{\mathbf{x}} \left[\hat{\mathfrak{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right],$$

which completes the proof of the upper bound.

We next show the lower bound. Starting from the definition of $\hat{\mathfrak{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$ and subtracting and adding $\mathbb{E}_{\mathcal{D}}[f]$, it holds

$$\mathbb{E}_{\mathbf{x}} \left[\hat{\mathfrak{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right] = \mathbb{E}_{\sigma, \mathbf{x}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \left((f(\mathbf{x}_i) - \mathbb{E}_{\mathcal{D}}[f]) + (\mathbb{E}_{\mathcal{D}}[f] - \hat{\mathbb{E}}_{\mathbf{x}}[f]) \right) \right| \right].$$

The subadditivity of the supremum and of the absolute value, and the linearity of the expectation allow us to split the r.h.s. into two summands and obtain

$$\begin{aligned} \mathbb{E}_{\mathbf{x}} \left[\hat{\mathfrak{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right] &\leq \mathbb{E}_{\sigma, \mathbf{x}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i (f(\mathbf{x}_i) - \mathbb{E}_{\mathcal{D}}[f]) \right| \right] \\ &\quad + \mathbb{E}_{\sigma, \mathbf{x}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbb{E}_{\mathcal{D}}[f] - \hat{\mathbb{E}}_{\mathbf{x}}[f]) \right| \right]. \end{aligned} \quad (51)$$

The first term on the r.h.s. is the RA of the *distributional* centralization of \mathcal{F} , i.e., it is $\mathfrak{R}_m(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathcal{D})$. The second term is the expectation w.r.t. \mathbf{x} of the ERA on \mathbf{x} of the family

$$\mathcal{Z}_{\mathbf{x}} \doteq \{x \mapsto \mathbb{E}_{\mathcal{D}}[f] - \hat{\mathbb{E}}_{\mathbf{x}}[f], f \in \mathcal{F}\}.$$

Each function in $\mathcal{Z}_{\mathbf{x}}$ is *constant*. Proceeding in exactly the same way as we did for the family $\mathcal{K}_{\mathbf{x}}$ in the proof of the upper bound, we can write

$$\hat{\mathfrak{R}}_m(\mathcal{Z}_{\mathbf{x}}, \mathbf{x}) = \text{SD}(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathbf{x})b(m). \quad (52)$$

Continuing from (51) and using (52) and the rightmost inequality of (6), we obtain

$$\mathbb{E}_{\mathbf{x}} \left[\hat{\mathfrak{R}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \right] \leq \mathfrak{R}_m(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) + 2\mathfrak{R}_m(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathcal{D})b(m) \leq (1 + 2b(m))\mathfrak{R}_m(\mathcal{C}_{\mathcal{D}}(\mathcal{F}), \mathcal{D}),$$

and our proof is complete. \square

Definition A.1.1. A function $Z \in \mathcal{X}^m \rightarrow \mathbb{R}$ is (α, β) -self-bounding with scale γ , for some $\alpha > 0$, $\beta \geq 0$, $\gamma \geq 0$ if for each $j = 1, \dots, m$, there exists a function $Z_j \in \mathcal{X}^m \rightarrow \mathbb{R}$ such that, for any $\mathbf{x} \in \mathcal{X}^m$ it holds that

1. $Z_j(\mathbf{x})$ does not depend on the j -th component \mathbf{x}_j of \mathbf{x} ; and
2. it holds $Z_j(\mathbf{x}) \leq Z(\mathbf{x}) \leq Z_j(\mathbf{x}) + \gamma$;

Additionally, the functions Z_j , $j = 1, \dots, m$, must be such that, for any $\mathbf{x} \in \mathcal{X}^m$, it holds $\sum_{j=1}^m \left(Z(\mathbf{x}) - Z_j(\mathbf{x}) \right) \leq \alpha Z(\mathbf{x}) + \beta$.

Theorem A.1.2. Let Z be a function from \mathcal{X}^m to \mathbb{R} that is (α, β) -self-bounding with scale γ , for $\alpha \geq 1/3$. Let $\delta \in (0, 1)$ and let \mathbf{x} be a collection of m i.i.d. samples from \mathcal{X} . With probability at least $1 - \delta$ over the choice of \mathbf{x} , it holds

$$\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] \leq Z(\mathbf{x}) + \alpha\gamma \ln \frac{1}{\delta} + \sqrt{\left(\alpha\gamma \ln \frac{1}{\delta} \right)^2 + 2\gamma(\alpha Z(\mathbf{x}) + \beta) \ln \frac{1}{\delta}}. \quad (53)$$

Additionally, when $\alpha = 1$, we may improve the constants to

$$\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] \leq Z(\mathbf{x}) + \frac{2}{3}\gamma \ln \frac{1}{\delta} + \sqrt{\left(\frac{1}{\sqrt{3}}\gamma \ln \frac{1}{\delta}\right)^2 + 2\gamma(Z(\mathbf{x}) + \beta) \ln \frac{1}{\delta}} . \quad (54)$$

Proof. In both cases, we will assume WLOG $\gamma = 1$. The results then hold by linearity, noting that if $Z(\cdot)$ is α - β self-bounding, with scale γ , then $\frac{1}{\gamma}Z(\cdot)$ is α - β/γ self-bounding, with scale 1; the general case thus follows by dividing out γ , obtaining a bound, and then multiplying through by γ .

We first show equation 53. Assume scale $\gamma = 1$. It is known that for $\gamma = 1$, we have for all $\alpha \geq \frac{1}{3}$, as described in [Boucheron et al., 2009, Thm. 1], which improves the earlier bounds of [Maurer, 2006]

$$\mathbb{P}\left(Z(\mathbf{x}) \leq \mathbb{E}_{\mathbf{x}}[Z(\mathbf{x})] - \varepsilon\right) \leq \exp\left(\frac{-\varepsilon^2}{2(\alpha \mathbb{E}_{\mathbf{x}}[Z(\mathbf{x})] + \beta)}\right) . \quad (55)$$

Now, taking δ equal to the RHS of (55), and solving for ε , this implies that with probability at least $1 - \delta$, we have

$$Z(\mathbf{x}) + \frac{\beta}{\alpha} \geq \mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \frac{\beta}{\alpha} - \sqrt{2(\alpha \mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \beta) \ln \frac{1}{\delta}} .$$

Note that this is a quadratic inequality in $\sqrt{\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \frac{\beta}{\alpha}}$, solving for which (via the quadratic formula) yields nondegenerate solution

$$\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] \leq Z(\mathbf{x}) + \alpha \ln \frac{1}{\delta} + \sqrt{\left(\alpha \ln \frac{1}{\delta}\right)^2 + 2\alpha(\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \beta) \ln \frac{1}{\delta}} .$$

Finally, in the general case, with γ -scaling, we have

$$\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] \leq Z(\mathbf{x}) + \gamma\alpha \ln \frac{1}{\delta} + \sqrt{\left(\gamma\alpha \ln \frac{1}{\delta}\right)^2 + 2\gamma\alpha(\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \beta) \ln \frac{1}{\delta}} .$$

We now show equation 54 (i.e., assume $\alpha = 1$). Again assume $\gamma = 1$. This result follows via identical logic to the above, this time using the *sub-gamma* form (see Boucheron et al. [2013, Ch. 2.1], section 2.1) of the stronger *sub-Poisson* 1 - β self-bounding function inequality [Boucheron et al., 2000, Thm. 1].

In particular, here we have that with probability at least $1 - \delta$,

$$Z(\mathbf{x}) \geq \mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \frac{1}{3} \ln \frac{1}{\delta} - \sqrt{2(\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \beta) \ln \frac{1}{\delta}} ,$$

which by the quadratic formula, yields

$$\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] \leq Z(\mathbf{x}) + \frac{2}{3} \ln \frac{1}{\delta} + \sqrt{\left(\frac{\gamma}{\sqrt{3}} \ln \frac{1}{\delta}\right)^2 + 2(\mathbb{E}_{\mathbf{x}} [Z(\mathbf{x})] + \beta) \ln \frac{1}{\delta}} .$$

The general result then follows via γ -scaling. □

Theorem 1.2.2. Suppose $m \geq 1$, and let $\chi \doteq 1 + 2b(m)$. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the choice of \mathbf{x} , it holds that

$$\mathbb{E}_{\mathbf{x}}[\hat{\mathbf{K}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})] \leq \hat{\mathbf{K}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) + \frac{2r\chi \ln \frac{1}{\delta}}{3m} + \sqrt{\left(\frac{r\chi \ln \frac{1}{\delta}}{\sqrt{3m}}\right)^2 + \frac{2r\chi(\hat{\mathbf{K}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) + rb(m)) \ln \frac{1}{\delta}}{m}} . \quad (9)$$

Proof. This proof proceeds by showing that $\hat{\mathbf{R}}_m(\hat{C}_x(\mathcal{F}), \mathbf{x})$ is a $(1, \text{rb}(m))$ -self-bounding function with scale $r \times /m$, then applying (54) from theorem A.1.2. First note that the result trivially holds for $m = 1$, as the empirically centralized ERA will always be 0, thus we assume $m \geq 2$ henceforth.

For any $\mathbf{x} \in \mathcal{X}^m$, let

$$Y(\mathbf{x}) \doteq \hat{\mathbf{R}}_m(\hat{C}_x(\mathcal{F}), \mathbf{x}),$$

and let $\mathbf{x}_{\setminus j}$ (resp. $\boldsymbol{\sigma}_{\setminus j}$) denote the $m - 1$ -dimensional vector of all but the j -th element of \mathbf{x} (resp. $\boldsymbol{\sigma}$). Define

$$Y_j(\mathbf{x}) \doteq \frac{m-1}{m} \hat{\mathbf{R}}_{m-1}(\hat{C}_{\mathbf{x}_{\setminus j}}(\mathcal{F}), \mathbf{x}_{\setminus j}) = \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1, i \neq j}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right].$$

We define these functions for convenience of notation. They will be handy when we later introduce the functions Z and Z_j , $j = 1, \dots, m$ that we want to show to be self-bounding.

We now show that $Y_j(\mathbf{x}) \leq Y(\mathbf{x}) + r/m\text{b}(m)$. Starting from the definition of $Y_j(\mathbf{x})$ and adding and subtracting $(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])/2m$ to the argument of the supremum, it holds

$$Y_j(\mathbf{x}) = \mathbb{E}_{\boldsymbol{\sigma}_{\setminus j}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \left(\sum_{\substack{i=1 \\ i \neq j}}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right) + \frac{1}{2} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]) - \frac{1}{2} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]) \right| \right].$$

Doubling and halving the sum in the argument of the expectation, and leveraging the subadditivity of the supremum and of the absolute value, we obtain

$$Y_j(\mathbf{x}) \leq \mathbb{E}_{\boldsymbol{\sigma}_{\setminus j}} \left[\left. \begin{aligned} & \frac{1}{2} \left(\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{\substack{i=1 \\ i \neq j}}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) + \left(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right) \\ & + \frac{1}{2} \left(\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{\substack{i=1 \\ i \neq j}}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) - \left(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right) \end{aligned} \right].$$

The two-term sum forming the argument of the outermost expectation is the expectation *w.r.t. only $\boldsymbol{\sigma}_j$* (i.e., *conditioned on $\boldsymbol{\sigma}_{\setminus j}$*) of the quantity

$$\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right|.$$

Thus, using the law of total expectation, we can write

$$Y_j(\mathbf{x}) \leq \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right].$$

By subtracting and adding $\hat{\mathbb{E}}_{\mathbf{x}}[f]$ to each term of the sum, and using the subadditivity of the supremum and of the absolute value, and the linearity of the expectation, we obtain

$$Y_j(\mathbf{x}) \leq \underbrace{\mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right) \right| \right]}_{=Y(\mathbf{x})} + \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(\hat{\mathbb{E}}_{\mathbf{x}}[f] - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right]. \quad (56)$$

The first term on the r.h.s. is $Y(\mathbf{x})$. The second term is the ERA of the sample-dependent family

$$\mathcal{W}_{\mathbf{x}} \doteq \left\{ y \mapsto \frac{1}{m} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]), f \in \mathcal{F} \right\} .$$

Each function in $\mathcal{W}_{\mathbf{x}}$ is *constant*. Using (7) and the linearity of expectation, like we did in the proof of lemma 1.2.1 for the family $\mathcal{K}_{\mathbf{x}}$ (see (50)), it holds

$$\hat{\mathbf{R}}_m(\mathcal{W}_{\mathbf{x}}, \mathbf{x}) = \frac{1}{m} \sup_{f \in \mathcal{F}} |f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]| b(m) \leq \frac{r}{m} b(m) .$$

Thus, continuing from (56) by incorporating the above fact, it holds

$$Y_j(\mathbf{x}) \leq Y(\mathbf{x}) + \frac{r}{m} b(m) . \quad (57)$$

We now show that $Y_j(\mathbf{x}) \geq Y(\mathbf{x}) - (1 + b(m))r/m$. Starting from the definition of Y_j and adding and removing

$$\frac{1}{m} \left(\sigma_j(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]) \right)$$

to the argument of the supremum, it holds

$$Y_j(\mathbf{x}) = \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \left(\sum_{\substack{i=1 \\ i \neq j}}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right) + \sigma_j(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]) - \sigma_j(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]) \right| \right] .$$

Then, from the triangle inequality and the fact that

$$\sup_{f \in \mathcal{F}} |\sigma_j(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])| \leq r,$$

we obtain

$$Y_j(\mathbf{x}) \geq \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right] - \frac{r}{m} .$$

From here, we add and subtract $\sigma_i \hat{\mathbb{E}}_{\mathbf{x}}[f]$ to each term of the sum, and then use the triangle inequality, the subadditivity of the supremum, and the linearity of expectation, to obtain

$$Y_j(\mathbf{x}) \geq \underbrace{\mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right) \right| \right]}_{=Y(\mathbf{x})} - \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \sigma_i \left(\hat{\mathbb{E}}_{\mathbf{x}}[f] - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right] - \frac{r}{m} .$$

The second term on the r.h.s. is again the ERA of a family of constant functions, each of them taking value at most r/m . Thus using (7), it follows that

$$Y_j(\mathbf{x}) \geq Y(\mathbf{x}) - (1 + b(m)) \frac{r}{m} .$$

Combining the above and (57), we obtain

$$Y(\mathbf{x}) - (1 + b(m)) \frac{r}{m} \leq Y_j(\mathbf{x}) \leq Y(\mathbf{x}) + \frac{r}{m} b(m) . \quad (58)$$

We now show that

$$\sum_{j=1}^m (Y(\mathbf{x}) - Y_j(\mathbf{x})) \leq Y(\mathbf{x}) . \quad (59)$$

Starting from the definition of the Y_j functions, and using the linearity of expectation and the subadditivity of the supremum

$$\begin{aligned} \sum_{j=1}^m Y_j(\mathbf{x}) &= \sum_{j=1}^m \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1, i \neq j}^m \boldsymbol{\sigma}_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right] \\ &\geq \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{j=1}^m \sum_{i=1, i \neq j}^m \boldsymbol{\sigma}_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right) \right| \right] . \end{aligned}$$

We rearrange the terms in the double sums, and use the linearity of expectation to obtain

$$\begin{aligned} \sum_{j=1}^m Y_j(\mathbf{x}) &\geq \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| (m-1) \sum_{i=1}^m \boldsymbol{\sigma}_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right) \right| \right] \\ &\geq (m-1) \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \left| \sum_{i=1}^m \boldsymbol{\sigma}_i \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right) \right| \right] , \end{aligned}$$

which completes our proof of (59), as the last expectation is $Y(\mathbf{x})$.

Define now the functions

$$Z(\mathbf{x}) \doteq Y(\mathbf{x}) \text{ and } Z_j(\mathbf{x}) \doteq Y_j(\mathbf{x}) - \frac{r}{m} \mathfrak{b}(m) \text{ for each } j = 1, \dots, m .$$

The value of $Z_j(\mathbf{x})$ clearly does not depend on the j -th component of \mathbf{x} . Also, from (58) it follows that

$$Z_j(\mathbf{x}) \leq Z(\mathbf{x}) \leq Z_j(\mathbf{x}) + (1 + 2\mathfrak{b}(m)) \frac{r}{m} \text{ for each } j = 1, \dots, m .$$

A consequence of (59) is finally that

$$\sum_{j=1}^m (Z(\mathbf{x}) - Z_j(\mathbf{x})) \leq Z(\mathbf{x}) + r\mathfrak{b}(m) .$$

Thus Z , i.e., $\hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x})$, is a $(1, r\mathfrak{b}(m))$ -self-bounding function with scale $(1 + 2\mathfrak{b}(m))r/m$. An application of (54) from theorem A.1.2 completes the proof. \square

Before proving theorem 1.3.1, we need the following lemma.

Lemma A.1.3. It holds

$$W(\mathcal{F}) \leq \frac{m}{m-1} \mathbb{E}_{\mathbf{x}}[\widehat{W}_{\mathbf{x}}(\mathcal{F})] .$$

Proof. Using Bessel's correction, we can rewrite the definition of wimpy variance to use the empirical expectation as

$$W(\mathcal{F}) = \sup_{f \in \mathcal{F}} \mathbb{E}_{\mathbf{x}} \left[\frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - \frac{\mathbb{E}[f]}{\mathcal{D}})^2 \right] = \sup_{f \in \mathcal{F}} \mathbb{E}_{\mathbf{x}} \left[\frac{1}{m-1} \sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f])^2 \right] .$$

An application of Jensen's inequality gives

$$W(\mathcal{F}) \leq \mathbb{E}_{\mathbf{x}} \left[\underbrace{\sup_{f \in \mathcal{F}} \frac{1}{m-1} \sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f])^2}_{= \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F})} \right]. \quad \square$$

Theorem 1.3.1. Suppose $m \geq 2$. Let $\delta \in (0, 1)$. With probability $\geq 1 - \delta$ over the choice of \mathbf{x} ,

$$W(\mathcal{F}) \leq \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F}) + \frac{r^2 \ln \frac{1}{\delta}}{m-1} + \sqrt{\left(\frac{r^2 \ln \frac{1}{\delta}}{m-1} \right)^2 + \frac{2r^2 \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F}) \ln \frac{1}{\delta}}{m-1}}. \quad (11)$$

Proof. This proof proceeds by showing that $\widehat{W}_{\mathbf{x}}(\mathcal{F})$ is a $(m/m-1, 0)$ -self-bounding with scale r^2/m , then applying lemma A.1.3, and finally (53) from theorem A.1.2.

Let $\mathbf{x}_{\setminus j}$ denote the vector \mathbf{x} with the j -th component removed, as we defined it also in the proof for theorem 1.2.2. Let $\hat{V}_{\mathbf{x}}[f]$ denote the (unbiased) sample variance of f over \mathbf{x} , i.e.,

$$\hat{V}_{\mathbf{x}}[f] \doteq \frac{1}{m-1} \sum_{i=1}^m \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right)^2.$$

Define

$$Z(\mathbf{x}) \doteq \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F}) = \sup_{f \in \mathcal{F}} \hat{V}_{\mathbf{x}}[f] = \sup_{f \in \mathcal{F}} \frac{1}{m-1} \sum_{i=1}^m \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right)^2$$

and

$$Z_j(\mathbf{x}) \doteq \sup_{f \in \mathcal{F}} \frac{1}{m-1} \sum_{i=1, i \neq j}^m \left(f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right)^2. \quad (60)$$

We first show that

$$Z_j(\mathbf{x}) = \sup_{f \in \mathcal{F}} \left[\hat{V}_{\mathbf{x}}[f] - \frac{1}{m} \left(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right)^2 \right], \quad (61)$$

as this form comes in handy many times. Starting from the definition of Z_j in (60), we add and subtract $\frac{1}{m-1} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2$ to the argument of the supremum, and then add and subtract $\hat{\mathbb{E}}_{\mathbf{x}}[f]$ to the argument of the sum, to obtain:

$$\begin{aligned} Z_j(\mathbf{x}) &= \sup_{f \in \mathcal{F}} \frac{1}{m-1} \left[\left(\sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right) - (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right] \\ &= \sup_{f \in \mathcal{F}} \frac{1}{m-1} \left[\left(\sum_{i=1}^m \left((f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f]) + (\hat{\mathbb{E}}_{\mathbf{x}}[f] - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]) \right)^2 \right) - (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right] \end{aligned}$$

By expressing the square in the argument of the sum, separating the three resulting terms in three distinct sums (associative property of the sum), and noticing that one of these sum is $\sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f]) = 0$, and another has argument $(\hat{\mathbb{E}}_{\mathbf{x}}[f] - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2$ independent from i , we obtain

$$Z_j(\mathbf{x}) = \sup_{f \in \mathcal{F}} \frac{1}{m-1} \left[\underbrace{\left(\sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f])^2 \right)}_{=(m-1)\hat{V}_{\mathbf{x}}[f]} + m(\hat{\mathbb{E}}_{\mathbf{x}}[f] - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 - (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right].$$

It holds $\hat{\mathbb{E}}_{\mathbf{x}}[f] = \frac{1}{m}f(\mathbf{x}_j) + \frac{m-1}{m}\hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]$, so we have

$$Z_j(\mathbf{x}) = \sup_{f \in \mathcal{F}} \frac{1}{m-1} \left[(m-1)\hat{V}_{\mathbf{x}}[f] + m \left(\frac{1}{m}f(\mathbf{x}_j) - \frac{1}{m}\hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] \right)^2 - (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right].$$

The identity in (68) then follows through simple algebraic steps.

We want to show that Z is a $(m/m-1, 0)$ -self-bounding function with scale r^2/m (see definition A.1.1). By definition of Z_j in (60), the value of $Z_j(\mathbf{x})$ does not depend on the j -th component of \mathbf{x} , as required by the first point in definition A.1.1.

We now show that, for any $j = 1, \dots, m$, it holds,

$$Z_j(\mathbf{x}) \leq Z(\mathbf{x}) \leq Z_j(\mathbf{x}) + \frac{r^2}{m} \text{ for any } \mathbf{x} \in \mathcal{X}^m, \quad (62)$$

as required by the second point in definition A.1.1. The leftmost inequality follows from the definitions of Z and Z_j . To show the rightmost inequality, we start from (68), and use the subadditivity of the supremum to obtain

$$Z_j(\mathbf{x}) \geq \left[\underbrace{\left(\sup_{f \in \mathcal{F}} \hat{V}_{\mathbf{x}}[f] \right)}_{=Z(\mathbf{x})} - \left(\sup_{f \in \mathcal{F}} \frac{1}{m} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right) \right].$$

The rightmost supremum is always smaller than r^2/m because $|f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f]| \leq r$, thus we have obtained the rightmost inequality in (62).

We now show that, for any $\mathbf{x} \in \mathcal{X}^m$, it holds

$$\sum_{i=1}^m (Z(\mathbf{x}) - Z_i(\mathbf{x})) \leq \frac{m}{m-1} Z(\mathbf{x}),$$

as in the last requirement of definition A.1.1. Starting again from (68) and using the subadditivity of the supremum, it holds

$$\sum_{j=1}^m Z_j(\mathbf{x}) = \sum_{j=1}^m \sup_{f \in \mathcal{F}} \left[\hat{V}_{\mathbf{x}}[f] - \frac{1}{m} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right] \geq \sup_{f \in \mathcal{F}} \sum_{j=1}^m \left[\hat{V}_{\mathbf{x}}[f] - \frac{1}{m} (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right].$$

By simple algebra we then get

$$\sum_{j=1}^m Z_j(\mathbf{x}) \geq \sup_{f \in \mathcal{F}} \left[m\hat{V}_{\mathbf{x}}[f] - \frac{1}{m} \sum_{j=1}^m (f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f])^2 \right].$$

From here, we use the fact that

$$\hat{\mathbb{E}}_{\mathbf{x}_{\setminus j}}[f] = \frac{1}{m-1} (m\hat{\mathbb{E}}_{\mathbf{x}}[f] - f(\mathbf{x}_j)),$$

to get

$$\sum_{j=1}^m Z_j(\mathbf{x}) \geq \sup_{f \in \mathcal{F}} \left[m\hat{V}_{\mathbf{x}}[f] - \frac{1}{m} \sum_{j=1}^m \left(\frac{m}{m-1} f(\mathbf{x}_j) - \frac{m}{m-1} \hat{\mathbb{E}}_{\mathbf{x}}[f] \right)^2 \right].$$

Now by simplifying some terms on the r.h.s., we obtain

$$\sum_{j=1}^m Z_j(\mathbf{x}) \geq \sup_{f \in \mathcal{F}} \left[m\hat{V}_{\mathbf{x}}[f] - \frac{m}{(m-1)} \underbrace{\frac{1}{m-1} \sum_{j=1}^m \left(f(\mathbf{x}_j) - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right)^2}_{=\hat{V}_{\mathbf{x}}[f]} \right].$$

Collecting terms and using the original definition of Z results in

$$\sum_{j=1}^m Z_j(\mathbf{x}) \geq \left(m - \frac{m}{m-1}\right) Z(\mathbf{x}) .$$

Thus,

$$\sum_{j=1}^m (Z(\mathbf{x}) - Z_j(\mathbf{x})) \leq mZ(\mathbf{x}) - \left(m - \frac{m}{m-1}\right) Z(\mathbf{x}) \leq \frac{m}{m-1} Z(\mathbf{x}),$$

which concludes our proof that Z is $(m/m-1, 0)$ -self-bounding with scale r^2/m .

We now use the above fact to prove the thesis. A consequence of lemma A.1.3 is

$$\mathbb{P}_{\mathbf{x}} \left(\widehat{W}_{\mathbf{x}}(\mathcal{F}) \leq W(\mathcal{F}) - \varepsilon \right) \leq \mathbb{P}_{\mathbf{x}} \left(\widehat{W}_{\mathbf{x}}(\mathcal{F}) \leq \frac{m}{m-1} \mathbb{E}_{\mathbf{x}}[\widehat{W}_{\mathbf{x}}(\mathcal{F})] - \varepsilon \right) .$$

From here, we use the definition

$$Z(\mathbf{x}) = \frac{m}{m-1} \widehat{W}_{\mathbf{x}}(\mathcal{F})$$

and apply (53) from theorem A.1.2 to obtain the thesis. \square

The constants in this bound are somewhat sub-optimal, as there is a significant gap between the best-known (sub-Poisson) tails for $(1, 0)$ -self-bounding and the best-known (sub-gamma) tails for $(1 + \varepsilon, 0)$ -self-bounding functions. We hope that future work leads to refined analysis of tail bounds for $(\alpha, 0)$ -self-bounding functions that decay gracefully as α exceeds 1.

Lemma 1.3.4. For any $\mathbf{x} \in \mathcal{X}^m$, it holds

$$\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) \geq \sqrt{\frac{\widehat{W}_{\mathbf{x}}^r(\mathcal{F})}{2m}} \text{ and } \hat{\mathbf{R}}_m(\hat{C}_{\mathbf{x}}(\mathcal{F}), \mathbf{x}) \geq \sqrt{\frac{\widehat{W}_{\mathbf{x}}(\mathcal{F})}{2m}} .$$

Furthermore, it holds

$$\lim_{m \rightarrow \infty} \sqrt{m} \hat{\mathbf{R}}_m(\mathcal{F}, \mathcal{D}) \geq \sqrt{\frac{2}{\pi} W^r(\mathcal{F})} \text{ and } \lim_{m \rightarrow \infty} \sqrt{m} \hat{\mathbf{R}}_m(C_{\mathcal{D}}(\mathcal{F}), \mathcal{D}) \geq \sqrt{\frac{2}{\pi} W(\mathcal{F})} .$$

Proof. From the subadditivity of the supremum, it holds that

$$\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) \geq \sup_{f \in \mathcal{F}} \mathbb{E}_{\boldsymbol{\sigma}} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i f(\mathbf{x}_i) \right\| \right] .$$

An application of Khintchine's inequality [Haagerup, 1982] gives

$$\hat{\mathbf{R}}_m(\mathcal{F}, \mathbf{x}) \geq \sup_{f \in \mathcal{F}} \frac{1}{\sqrt{2}} \sqrt{\frac{\|f(\mathbf{x})\|_2^2}{m^2}},$$

where $f(\mathbf{x})$ denotes the m -dimensional vector of values of f on \mathbf{x} . The proof of the leftmost inequality in the thesis ends by noting that

$$\widehat{W}_{\mathbf{x}}^r(\mathcal{F}) = \frac{\|f(\mathbf{x})\|_2^2}{m} .$$

The rightmost inequality is then a corollary, using the identity $\widehat{W}_{\mathbf{x}}^r(\hat{C}_{\mathbf{x}}(\mathcal{F})) = \widehat{W}_{\mathbf{x}}(\mathcal{F})$.

The asymptotic lower bounds follow by replacing the Khintchine's inequality step with an application of the central limit theorem. \square

Before proving theorem 1.3.7 we need to introduce an important technical result. For any $u \in \mathbb{R}$, let $h(u) \doteq (1+u)\ln(1+u) - u$, and let $(u)_+ \doteq \max(0, u)$.

Theorem A.1.4 (Samson's bound, [Boucheron et al., 2013, Thm. 12.11]). Let $\mathcal{Q}_1, \dots, \mathcal{Q}_m$ be possibly different probability distributions over a domain \mathcal{Y} . Let $\mathcal{G} \subseteq \mathcal{X} \rightarrow [-1, 1]$. Furthermore, assume that for each $g \in \mathcal{G}$ and $i \in \{1, \dots, m\}$, it holds $\mathbb{E}_{\mathcal{Q}_i}[g] = 0$. Now, for any $\mathbf{y} \in \mathcal{Y}^m$, let

$$Z(\mathbf{y}) \doteq \sup_{g \in \mathcal{G}} \sum_{i=1}^m g(y_i) \text{ and } S^2 \doteq \mathbb{E}_{\mathbf{y}} \left[\sup_{g \in \mathcal{F}} \sum_{i=1}^m \mathbb{E}_{y'_i \sim \mathcal{Q}_i} \left[((g(y_i) - g(y'_i))_+)^2 \right] \right].$$

Let $\mathbf{y} \in \mathcal{Y}^m$, with each $y_i \sim \mathcal{Q}_i$, independently (but not necessarily identically, since the distributions may be different). It holds¹

$$\mathbb{P}_{\mathbf{y}} \left(Z(\mathbf{y}) \leq \mathbb{E}_{\mathcal{Q}_{1:m}} [Z] - \varepsilon \right) \leq \exp \left(-\frac{S^2}{4} h \left(\frac{2\varepsilon}{S^2} \right) \right). \quad (63)$$

Theorem 1.3.7. Let $\boldsymbol{\sigma} \in (\pm 1)^{n \times m}$ be a matrix of i.i.d. Rademacher r.v.'s. Let $\delta \in (0, 1)$. With probability at least $1 - \delta$ over the choice of $\boldsymbol{\sigma}$, it holds

$$\hat{\mathbf{K}}_m(\mathcal{F}, \mathbf{x}) \leq \hat{\mathbf{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \frac{2\hat{q}_{\mathcal{F}}(\mathbf{x}) \ln \frac{1}{\delta}}{3nm} + \sqrt{\frac{4\widehat{W}_{\mathbf{x}}^r(\mathcal{F}) \ln \frac{1}{\delta}}{nm}}. \quad (14)$$

Proof. Without loss of generality, we assume that $\hat{q}_{\mathcal{F}}(\mathbf{x}) = 1$. The general case then follows via scaling.

Let

$$Z(\boldsymbol{\sigma}) \doteq nm \hat{\mathbf{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) = \sum_{j=1}^n \sup_{f \in \mathcal{F}} \left| \sum_{i=1}^m \boldsymbol{\sigma}_{j,i} f(\mathbf{x}_i) \right|.$$

It holds $\mathbb{E}_{\boldsymbol{\sigma}}[Z] = nm \hat{\mathbf{K}}_m(\mathcal{F}, \mathbf{x})$.

We first show that we can apply Samson's bound (theorem A.1.4) to Z , i.e., to the scaled MC-ERA. Consider the function family \mathcal{F}_{\pm} introduced in corollary 1.3.5, and consider the n -times Cartesian product of \mathcal{F}_{\pm} with itself

$$(\mathcal{F}_{\pm})^n = \underbrace{\mathcal{F}_{\pm} \times \dots \times \mathcal{F}_{\pm}}_{n \text{ times}}.$$

We use $\mathbf{f} = (f_1, \dots, f_n)$ to denote an element of $(\mathcal{F}_{\pm})^n$. Now, define the family

$$\mathcal{G} \doteq \{g(\boldsymbol{\sigma}_{j,i}) \doteq \boldsymbol{\sigma}_{j,i} f_j(\mathbf{x}_i), \mathbf{f} \in (\mathcal{F}_{\pm})^n\}.$$

The functions in \mathcal{G} have domain $\mathcal{Y} = \{-1, 1\}$ and values in $[-1, 1]$. It holds

$$Z(\boldsymbol{\sigma}) = \sup_{\mathbf{f} \in (\mathcal{F}_{\pm})^n} \sum_{j=1}^n \sum_{i=1}^m \boldsymbol{\sigma}_{j,i} f_j(\mathbf{x}_i) = \sup_{g \in \mathcal{G}} \sum_{(j,i) \in \{1, \dots, n\} \times \{1, \dots, m\}} g(\boldsymbol{\sigma}_{j,i}). \quad (64)$$

Thus Z has the form required by theorem A.1.4.

Let $\boldsymbol{\sigma}'$ denote a second $n \times m$ i.i.d. Rademacher matrix (like $\boldsymbol{\sigma}$), and define

$$\begin{aligned} S^2 &\doteq \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{\mathbf{f} \in (\mathcal{F}_{\pm})^n} \sum_{j=1}^n \sum_{i=1}^m \mathbb{E}_{\boldsymbol{\sigma}'_{j,i}} \left[((\boldsymbol{\sigma}_{j,i} f_j(\mathbf{x}_i) - \boldsymbol{\sigma}'_{j,i} f_j(\mathbf{x}_i))_+)^2 \right] \right] \\ &= n \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \mathcal{F}_{\pm}} \sum_{i=1}^m 2((\boldsymbol{\sigma}_{1,i} f(\mathbf{x}_i))_+)^2 \right]. \end{aligned}$$

¹To be precise, this is an immediate consequence of the statement of [Boucheron et al., 2013, Thm. 2.11], through an application of the Chernoff method to the moment generating function given therein.

It holds

$$S^2 \leq 2nm\widehat{W}_x^r(\mathcal{F}) . \quad (65)$$

For each $g \in \mathcal{G}$, $g(\sigma_{j,i})$ and $g(\sigma_{j',i'})$ are *independent*, though not necessarily *identically distributed*, for $(j,i) \neq (j',i')$, due to the dependence of $g(\sigma_{j,i})$ on indices (j,i) . It also holds, for each $g \in \mathcal{G}$, and indices (j,i) , that $\mathbb{E}_{\sigma_{i,j}}[g(\sigma_{i,j})] = 0$, simply due to multiplication by symmetric (Rademacher) r.v.'s.

Thus, we can use Samson's bound (theorem A.1.4) on \mathcal{G} , Z , and S^2 , although it is generally more convenient to work with \mathcal{F} and $(\mathcal{F}_\pm)^n$.

We now show the thesis. Fix $\varepsilon \in (0,1)$. It follows from Samson's bound that

$$\mathbb{P}_\sigma \left(\hat{\mathbf{K}}_m(\mathcal{F}, \mathbf{x}) \geq \hat{\mathbf{K}}_m^n(\mathcal{F}, \mathbf{x}, \sigma) + \varepsilon \right) = \mathbb{P}_\sigma \left(\mathbb{E}[Z] \geq Z(\sigma) + nm\varepsilon \right) \leq \exp \left(-\frac{S^2}{4} h \left(\frac{2nm\varepsilon}{S^2} \right) \right) .$$

The function

$$g(x) \doteq xh \left(\frac{2nm\varepsilon}{x} \right)$$

is monotonically decreasing in its argument. Thus, using (65) gives

$$\mathbb{P}_\sigma \left(\hat{\mathbf{K}}_m(\mathcal{F}, \mathbf{x}) \geq \hat{\mathbf{K}}_m^n(\mathcal{F}, \mathbf{x}, \sigma) + \varepsilon \right) \leq \exp \left(\frac{-nm\widehat{W}_x^r(\mathcal{F})}{2} h \left(\frac{\varepsilon}{\widehat{W}_x^r(\mathcal{F})} \right) \right) .$$

Now, for $u > -1/2$, define the function

$$h_1(u) \doteq 1 + u - \sqrt{1 + 2u} .$$

Using the fact (see Boucheron et al. [2013, Ch. 2.4]) that

$$h(u) \geq 9h_1 \left(\frac{u}{3} \right) \text{ for every } u \in (-1, +\infty),$$

we obtain

$$\mathbb{P}_\sigma \left(\hat{\mathbf{K}}_m(\mathcal{F}, \mathbf{x}) \geq \hat{\mathbf{K}}_m^n(\mathcal{F}, \mathbf{x}, \sigma) + \varepsilon \right) \leq \exp \left(-\frac{9}{2} nm\widehat{W}_x^r(\mathcal{F}) h_1 \left(\frac{\varepsilon}{3\widehat{W}_x^r(\mathcal{F})} \right) \right) .$$

The result for $\hat{q}_\mathcal{F}(\mathbf{x}) = 1$ is obtained by imposing that the r.h.s. be at most δ and solving for ε using standard sub-gamma inequalities. The general case then follows via linear scaling. \square

This bound is quite comparable to Bousquet's bound on the SD (see theorem 1.3.2). The variance factors $\widehat{W}_x^r(\mathcal{F})$ and $\widehat{W}_x(\mathcal{F})$ are convenient, as they depend only on sample variances, rather than true variances and expected supremum deviations.

Even if Samson's inequality introduces additional 2-factors on both the range and variance w.r.t. theorem 1.3.2, both are divided by MC-trial count n , so for $n \geq 2$ trials, the Monte-Carlo error terms become negligible.

A.2 Details on the Experimental Evaluation

As mentioned in the main text, lemma 1.4.1 is a consequence of [Shalev-Shwartz and Ben-David, 2014, Lemmas 26.11, 26.10], reported here for completeness.²

²The identities in the lemma are not reported in the original, but can be easily obtained through a slightly more refined proof than the one presented in the original. See the proof of lemma 1.4.1 for intuition.

Lemma A.2.1 (Shalev-Shwartz and Ben-David, 2014, Lemmas 26.11, 26.10). It holds

$$\hat{\mathbf{r}}_m(\mathcal{F}_1, \mathbf{x}) = \mathbb{E}_\sigma \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i \mathbf{x}_i \right\|_\infty \right] \leq \max_i \|\mathbf{x}_i\|_\infty \sqrt{\frac{2 \ln(2d)}{m}},$$

and

$$\hat{\mathbf{r}}_m(\mathcal{F}_2, \mathbf{x}) = \mathbb{E}_\sigma \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i \mathbf{x}_i \right\|_2 \right] \leq \max_i \|\mathbf{x}_i\|_2 \frac{1}{\sqrt{m}}.$$

We now show the centralized variants.

Lemma 1.4.1. Let $\bar{\mathbf{x}} \doteq \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i \in \mathbb{R}^d$. For the ℓ_1 norm, it holds

$$\hat{\mathbf{r}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}_1), \mathbf{x}) = \mathbb{E}_\sigma \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_\infty \right] \leq \max_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|_\infty \sqrt{\frac{2 \ln(2d)}{m}},$$

while for the ℓ_2 norm, it holds

$$\hat{\mathbf{r}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}_2), \mathbf{x}) = \mathbb{E}_\sigma \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_2 \right] \leq \max_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|_2 \frac{1}{\sqrt{m}}.$$

Proof. We show the ℓ_2 case in detail; the reasoning for the ℓ_1 case is essentially the same (see details at the end of the proof). The definition of $\hat{\mathbf{r}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}_2), \mathbf{x})$ is

$$\hat{\mathbf{r}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}_2), \mathbf{x}) = \mathbb{E}_\sigma \left[\sup_{w: \|w\|_2 \leq 1} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i (w \cdot \mathbf{x}_i - \hat{\mathbb{E}}_{\mathbf{x}}[w]) \right| \right],$$

where

$$\hat{\mathbb{E}}_{\mathbf{x}}[w] = \frac{1}{m} \sum_{i=1}^m (w \cdot \mathbf{x}_i) = w \cdot \bar{\mathbf{x}}.$$

Using linearity, we then get

$$\hat{\mathbf{r}}_m(\hat{\mathcal{C}}_{\mathbf{x}}(\mathcal{F}_2), \mathbf{x}) = \mathbb{E}_\sigma \left[\sup_{w: \|w\|_2 \leq 1} \left| w \cdot \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right| \right].$$

Now, for ease of notation, let $u \doteq \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}})$. The supremum is realized when

$$w = \frac{u}{\|u\|_2},$$

because in this case the vector w has the same direction as u , and the largest possible norm $\|w\|_2 = 1$. Since the two vectors w and u are collinear, the Cauchy-Schwarz inequality holds with equality, and we have

$$w \cdot u = \|w\|_2 \|u\|_2 = \|u\|_2 = \left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_2.$$

We thus obtain

$$\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}_2), \mathbf{x}) = \mathbb{E}_{\boldsymbol{\sigma}} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_2 \right] .$$

From here, we can proceed as in the second part of the proof of [Shalev-Shwartz and Ben-David, 2014, Lemma 26.10] to obtain the thesis.

By similar reasoning (now with Hölder’s inequality in place of the Cauchy-Schwarz inequality, and following the proof of Shalev-Shwartz and Ben-David [2014, Lemma 26.11]), we get that

$$\hat{\mathfrak{R}}_m(\hat{C}_x(\mathcal{F}_1), \mathbf{x}) = \mathbb{E}_{\boldsymbol{\sigma}} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i (\mathbf{x}_i - \bar{\mathbf{x}}) \right\|_{\infty} \right] \leq \max_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|_{\infty} \sqrt{\frac{2 \ln(2d)}{m}} . \quad \square$$

A.2.1 Data Generation

Our data distributions for both the ℓ_1 and ℓ_2 constrained linear family experiments are both randomized and parameterized by dimension d . Rademacher averages and wimpy variances depend on the randomization and d , and ranges may be bounded *a priori* in terms of d .

ℓ_1 Datasets In our ℓ_1 experiments, each \mathbf{x}_j is independently Beta-distributed, thus $\mathbf{x} \sim \text{B}(\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1) \times \cdots \times \text{B}(\boldsymbol{\alpha}_d, \boldsymbol{\beta}_d)$. The parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are themselves randomized, in particular, we sample $\boldsymbol{\alpha}_j$ and $\boldsymbol{\beta}_j$ from $\sqrt{\chi_j^2}$, where χ_k^2 is the χ^2 distribution with k degrees of freedom. In these datasets, $r = q = 1$.

ℓ_2 Datasets In our ℓ_2 experiments, we generate random *mean vector* $\boldsymbol{\mu} \in \mathbb{R}^d$ and *covariance matrix* $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$, then sample $\mathbf{x}' \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and finally obtain sample \mathbf{x} by projecting \mathbf{x}' to the nonnegative hyperquadrant of the radius \sqrt{d} ℓ_2 sphere; i.e.,

$$\mathbf{x} = \underset{\mathbf{x} \in \mathbb{R}^d: \|\mathbf{x}\|_2 \leq \sqrt{d} \wedge \mathbf{0} \leq \mathbf{x}}{\operatorname{argmin}} \|\mathbf{x} - \mathbf{x}'\|_2 .$$

Taking \mathbf{I}_d to be the identity matrix, we sample $\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{1}, \mathbf{I}_d)$, and taking $\mathbf{a} \sim \mathcal{U}(0, 1)^{d \times d}$, we let $\boldsymbol{\Sigma} \doteq \frac{\mathbf{a} \mathbf{a}^\top}{d} + \mathbf{I}_d$. In these datasets, $r = q = \sqrt{d}$.

A.2.2 Supplementary Plots

Figure A.1 shows the same results as figure 1.1 (in the main text), but without the scaling of the quantities by \sqrt{m} . Similarly, figure A.2 shows the same results as figure 1.2, sans scaling by \sqrt{m} . Additionally, both plots also include a *McDiarmid term* $3r\sqrt{\ln \frac{1}{\eta}}/2m$, representing the *additive error* incurred bounding the SD in terms of $\hat{\mathfrak{R}}_m^1(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma})$. We stress that this term *does not* include the MC-ERA itself, and thus is just one summand of the total McDiarmid SD bound. Nevertheless, the McDiarmid term alone asymptotically exceeds *all other bounds* in all experiments, except for the (loose) noncentralized analytical bound of \mathcal{F}_1 over \mathbb{R}^{256} . This sharply illustrates the benefit of *variance-sensitive* bounds over (range-only) McDiarmid bounds.

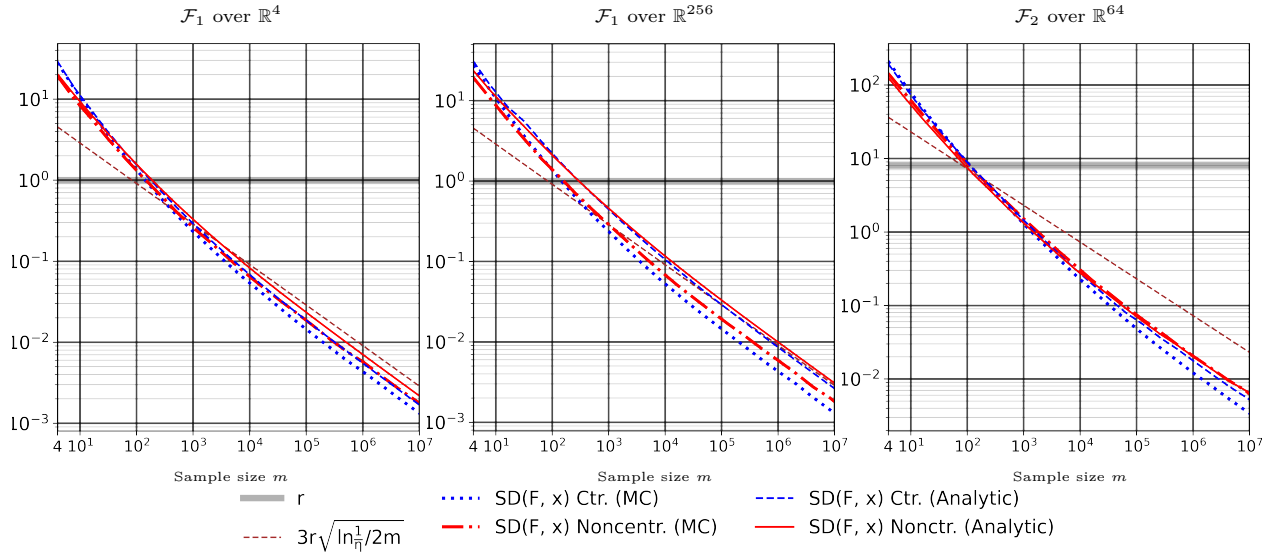


Figure A.1: Comparison of SD bounds as functions of the sample size m . See the main text for details.

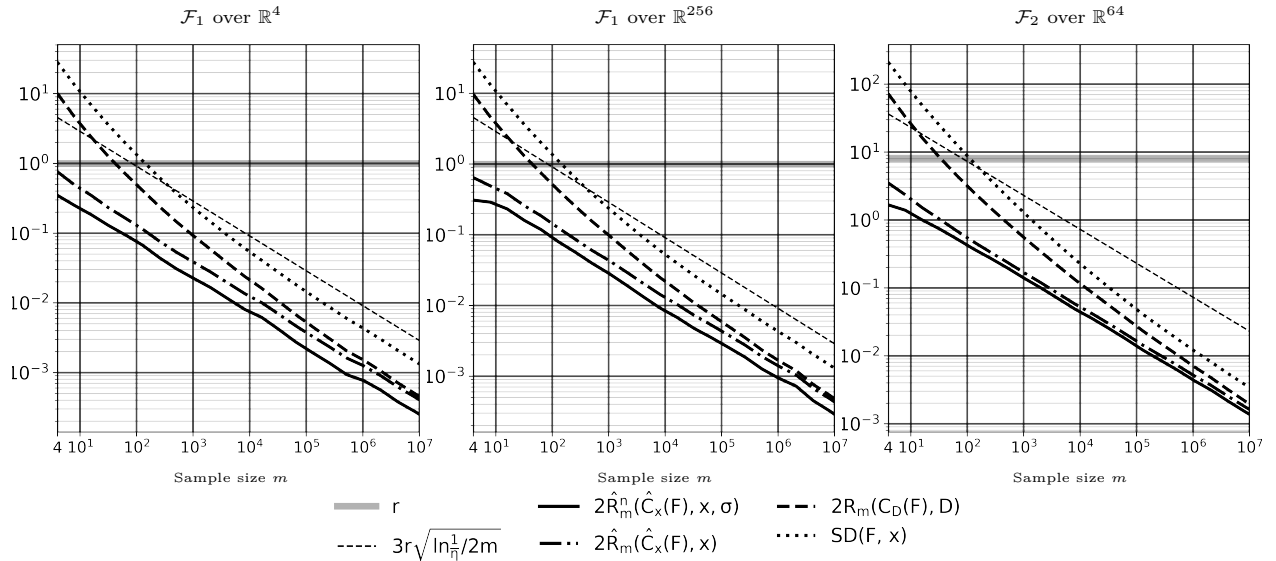


Figure A.2: Comparison of SD bounds as functions of the sample size m . See the main text for details.

Appendix B

Supplementary Material for Chapter 2

B.1 Deferred proofs

Proof of Lemma 2.4.1. For the sake of the proof, assume that we have two labeled set of samples of size m and m_L from $p_{\mathcal{X}\mathcal{Y}}$, call them respectively S and S_L . The set S represents our unlabeled sample, and the set S_L represents the labeled sample. For any $\delta \in (0, 1)$, we would like to find a $\gamma > 0$ such that with probability $1 - \delta$, for all $i \in 1, \dots, n$ (simultaneously), we have

$$\left| \frac{1}{m} \sum_{(x, \mathbf{y}) \in S} \ell(\varphi_i(x), \mathbf{y}) - \frac{1}{m_L} \sum_{(x, \mathbf{y}) \in S_L} \ell(\varphi_i(x), \mathbf{y}) \right| \leq \gamma . \quad (66)$$

The sample S represents the unlabeled data x_1, \dots, x_m we have access to. In fact, $\frac{1}{m} \sum_{(x, \mathbf{y}) \in S} \ell(\varphi_i(x), \mathbf{y}) = \hat{\mathbf{R}}(\varphi_i; X, \mathbf{Y}^*)$. The inequality (66) implies that for the true labeling of the unlabeled data x_1, \dots, x_m , for all $i \in 1, \dots, n$, it holds that:

$$\hat{\mathbf{R}}(\varphi_i; X, \mathbf{Y}^*) \in [\mu_i - \gamma, \mu_i + \gamma]$$

where $\mu_i = \frac{1}{m_L} \sum_{(x, \mathbf{y}) \in S_L} \ell(\varphi_i(x), \mathbf{y})$ is the empirical mean computed from the labeled sample S_L .

By using Hoeffding's inequality, we have that for a fixed i , it holds that:

$$\begin{aligned} & \mathbb{P}_{S, S_L} \left(\left| \frac{1}{m} \sum_{(x, \mathbf{y}) \in S} \ell(\varphi_i(x), \mathbf{y}) - \frac{1}{m_L} \sum_{(x, \mathbf{y}) \in S_L} \ell(\varphi_i(x), \mathbf{y}) \right| > \gamma \right) \\ & \leq 2 \exp \left(\frac{-2\gamma^2}{\sum_{j=1}^m \left(\frac{B}{m}\right)^2 + \sum_{j=1}^{m_L} \left(\frac{B}{m_L}\right)^2} \right) \\ & = 2 \exp \left(\frac{-2m_L m \gamma^2}{B^2(m + m_L)} \right) = \frac{\delta}{n} \end{aligned} \quad (67)$$

By taking a union bound and solving (67) with respect to γ , the statement follows.

Proof of Lemma 2.4.5. By invoking lemma 2.4.4, it is easy to see that the function $\mathbf{R}(\mathbf{h}_\theta; X, \mathbf{Y}')$ is a convex combination of convex functions with respect to θ , hence it is also convex in θ . Let $\mathbf{v} \in \partial \mathbf{R}(\mathbf{h}_\theta; X, \mathbf{Y}')$. By Fact 3, such a vector always exists. Then, we have that for any $\theta' \in \Theta$, it holds that:

$$\mathbf{R}(\mathbf{h}_{\theta'}, \mathbf{Y}') - \mathbf{R}(\mathbf{h}_\theta, \mathbf{Y}') \geq \mathbf{v}^T (\theta' - \theta)$$

As $f(\boldsymbol{\theta}') \geq R(\mathbf{h}_{\boldsymbol{\theta}'}, \mathbf{Y}')$, we have that:

$$f(\boldsymbol{\theta}') - f(\boldsymbol{\theta}) \geq \mathbf{v}^T(\boldsymbol{\theta}' - \boldsymbol{\theta})$$

which implies that \mathbf{v} is a subgradient of f at $\boldsymbol{\theta}$.

Proof of Theorem 2.4.6 We need to show that $f(\boldsymbol{\theta})$ is convex and L -Lipschitz continuous with respect to $\boldsymbol{\theta}$ to apply the standard convergence result for constant step size subgradient optimization [Bertsekas, 2015]

To show that $f(\boldsymbol{\theta})$ is convex it is straightforward to see that $\hat{R}(\mathbf{h}_{\boldsymbol{\theta}}, \mathbf{Y})$ is convex in $\boldsymbol{\theta}$ as it is the convex combination of convex functions in $\boldsymbol{\theta}$. For any $\lambda \in [0, 1]$, we have that:

$$\begin{aligned} f(\lambda\boldsymbol{\theta}_1 + (1-\lambda)\boldsymbol{\theta}_2) &= \max_{\mathbf{Y} \in \mathbb{Y}} L_c(h_{\lambda\boldsymbol{\theta}_1 + (1-\lambda)\boldsymbol{\theta}_2}, \mathbf{Y}) \\ &\leq \max_{\mathbf{Y} \in \mathbb{Y}} [\lambda L_c(\mathbf{h}_{\boldsymbol{\theta}_1}, \mathbf{Y}) + (1-\lambda)L_c(\mathbf{h}_{\boldsymbol{\theta}_2}, \mathbf{Y})] \\ &\leq \lambda \max_{\mathbf{Y} \in \mathbb{Y}} L_c(\mathbf{h}_{\boldsymbol{\theta}_1}, \mathbf{Y}) + (1-\lambda) \max_{\mathbf{Y} \in \mathbb{Y}} L_c(\mathbf{h}_{\boldsymbol{\theta}_2}, \mathbf{Y}) \\ &= \lambda f(\boldsymbol{\theta}_1) + (1-\lambda)f(\boldsymbol{\theta}_2) \end{aligned}$$

Also, $f(\boldsymbol{\theta})$ is L -Lipschitz in $\boldsymbol{\theta}$. In fact, it is straightforward to see that $\hat{R}(\mathbf{h}_{\boldsymbol{\theta}}; X, \mathbf{Y})$ is also L -Lipschitz continuous in $\boldsymbol{\theta}$. For any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$, we have that:

$$\begin{aligned} |f(\boldsymbol{\theta}_1) - f(\boldsymbol{\theta}_2)| &\leq \max_{\mathbf{Y} \in \mathbb{Y}} |\hat{R}(\mathbf{h}_{\boldsymbol{\theta}_1}; X, \mathbf{Y}) - \hat{R}(\mathbf{h}_{\boldsymbol{\theta}_2}; X, \mathbf{Y})| \\ &\leq L \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2 \end{aligned}$$

The subgradient of $f(\boldsymbol{\theta})$ in $\boldsymbol{\theta}$ is computed by using lemma 2.4.5.

Proof of Lemma 2.4.10. For each $i \in 1, \dots, n$, we have that:

$$\frac{\partial}{\partial \theta_i} \ell(\mathbf{h}_{\boldsymbol{\theta}}(x), \mathbf{e}) = 2 \left(\boldsymbol{\varphi}_i(x)^T \cdot \mathbf{h}_{\boldsymbol{\theta}}(x) - \boldsymbol{\varphi}_i(x)^T \cdot \mathbf{e} \right)$$

Therefore, we can bound the norm of the gradient of ℓ as:

$$\begin{aligned} \|\nabla_{\boldsymbol{\theta}} \ell(\mathbf{h}_{\boldsymbol{\theta}}(x), \mathbf{e})\|_2 &= 2 \sqrt{\sum_{i=1}^n (\boldsymbol{\varphi}_i(x)^T \cdot (\mathbf{h}_{\boldsymbol{\theta}}(x) - \mathbf{e}))^2} && \text{SEE ABOVE} \\ &\leq 2 \sqrt{\sum_{i=1}^n (1)^2} && \begin{aligned} &\|\boldsymbol{\varphi}_i(x)^T\|_1 = 1 \\ &\|\mathbf{h}_{\boldsymbol{\theta}}(x) - \mathbf{e}\|_{\infty} \leq 1 \end{aligned} \\ &= 2\sqrt{n} . && \text{HÖLDER'S INEQUALITY} \end{aligned}$$

This implies that the function $\ell(\mathbf{h}_{\boldsymbol{\theta}}(x), \mathbf{e})$ is $2\sqrt{n}$ -Lipschitz continuous with respect to $\boldsymbol{\theta}$.

Proof of Lemma 2.4.11. Without loss of generality, suppose that $e_i = 1$. We have that:

$$\ell(\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{x}), \mathbf{e}) = -\ln \left(\frac{\exp(\mathbf{w}_i^T \cdot \mathbf{x})}{\sum_{c=1}^k \exp(\mathbf{w}_c^T \cdot \mathbf{x})} \right)$$

It is easy to see that $\ell(\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{x}), \mathbf{e}) \geq 0$. By using Cauchy-Schwarz inequality, we have that:

$$\begin{aligned} \ell(\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{x}), \mathbf{e}) &= -\ln \left(\frac{\exp(\mathbf{w}_i^T \cdot \mathbf{x})}{\sum_{c=1}^k \exp(\mathbf{w}_c^T \cdot \mathbf{x})} \right) \\ &\leq -\ln \left(\frac{\exp(-B_w B_x)}{k \exp(B_w B_x)} \right) \\ &\leq 2B_w B_x + \ln k \end{aligned}$$

Proof of Lemma 2.4.12. For a fixed $(\mathbf{x}, \mathbf{e}) \in \mathcal{X} \times \mathcal{Y}$, consider the function $\omega(\mathbf{p}) : \mathbb{R}^k \rightarrow \mathcal{Y}^\diamond$ defined as

$$\omega(\mathbf{p}) \doteq - \sum_{c=1}^k e_c \cdot \ln(p_c) ,$$

and let

$$\mathbf{h}(\boldsymbol{\theta}) \doteq \left(\frac{\exp(\mathbf{w}_1^T \cdot \mathbf{x})}{\sum_{c=1}^k \exp(\mathbf{w}_c^T \cdot \mathbf{x})}, \dots, \frac{\exp(\mathbf{w}_k^T \cdot \mathbf{x})}{\sum_{c=1}^k \exp(\mathbf{w}_c^T \cdot \mathbf{x})} \right)^T ,$$

where $\boldsymbol{\theta} = (\mathbf{w}_1 \dots \mathbf{w}_k)^T$, and observe that $\ell(\mathbf{h}_\theta(x), \mathbf{e}) = \omega \circ \mathbf{h}(\boldsymbol{\theta})$.

It is well known that ℓ is $L_\omega L_h$ -Lipschitz continuous with respect to $\boldsymbol{\theta}$, where L_ω and L_h are the Lipschitz constants respectively of ω and \mathbf{h} . It is also a known result that $L_\omega \leq 1$ (see for example Proposition 4 of Gao and Pavel [2018]).

We now want to compute L_h . We will use the fact that $\max_{\boldsymbol{\theta} \in \Theta} \|J_h(\boldsymbol{\theta})\|_F \leq L_h$, where J_h denotes the Jacobian matrix of \mathbf{h} and $\|\cdot\|_F$ denotes the Frobenius norm.

For ease of notation, let $\mathbf{h}(\boldsymbol{\theta}) = \mathbf{p} = (p_1, \dots, p_k)^T$. We have that for each $i \in 1, \dots, k$, it holds that

$$\begin{aligned} \frac{\partial[\mathbf{h}(\boldsymbol{\theta})]_i}{\partial \mathbf{w}_j} &= p_i p_j \mathbf{x} \quad \text{for } j \neq i \\ \frac{\partial[\mathbf{h}(\boldsymbol{\theta})]_i}{\partial \mathbf{w}_i} &= (p_i - p_i^2) \mathbf{x} \end{aligned}$$

Therefore, we can bound the square of the Frobenius norm of the Jacobian matrix of \mathbf{h} with:

$$\begin{aligned} \|J_h(\boldsymbol{\theta})\|_F^2 &= \sum_{i,j} \left\| \frac{\partial[\mathbf{h}(\boldsymbol{\theta})]_i}{\partial \mathbf{w}_j} \right\|_2^2 \\ &\leq \|\mathbf{x}\|_2^2 \left(\sum_i [p_i(1-p_i)]^2 + \sum_{i \neq j} [p_i p_j]^2 \right) \\ &\leq \|\mathbf{x}\|_2^2 (k + k^2/2) \leq \|k\mathbf{x}\|_2^2 \end{aligned}$$

We can conclude that \mathbf{h} is kB_x -Lipschitz continuous, and the statement follows.

Proof of Theorem 2.4.8 From Chapter 14 of [Mitzenmacher and Upfal, 2017], we know that:

$$\mathbf{R}(\mathbf{h}_{\hat{\boldsymbol{\theta}}}) \leq \hat{\mathbf{R}}(\mathbf{h}_{\hat{\boldsymbol{\theta}}}; X, \mathbf{Y}^*) + 2\hat{\mathbf{R}}(\mathcal{L}; X, \mathbf{Y}^*) + \mathbf{O} \left(B \sqrt{\frac{\ln \frac{1}{\delta}}{m}} \right)$$

By definition of $f(\cdot)$, it holds that $\hat{\mathbf{R}}(\mathbf{h}_{\hat{\boldsymbol{\theta}}}; X, \mathbf{Y}^*) \leq f(\hat{\boldsymbol{\theta}})$. As $\hat{\boldsymbol{\theta}}$ is the optimal solution of (16), we have that $f(\hat{\boldsymbol{\theta}}) \leq f(\boldsymbol{\theta}^*)$. Let $\mathbf{Y}' \doteq \operatorname{argmax}_{\mathbf{Y} \in \mathcal{Y}^\diamond} \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y})$. It holds that:

$$\begin{aligned} f(\boldsymbol{\theta}^*) &= \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}') \\ &= \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}') + \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}^*) - \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}^*) \\ &= \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}^*) + \left| \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}') - \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}^*) \right| \end{aligned}$$

By using the fact that ℓ is bounded, and the definition of diameter $D_{\mathcal{Y}^\diamond}$, we have that:

$$\begin{aligned} \left| \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}') - \hat{\mathbf{R}}(\mathbf{h}_{\boldsymbol{\theta}^*}; X, \mathbf{Y}^*) \right| &= \left| \frac{1}{m} \sum_{j=1}^m \sum_{c=1}^k \ell(\mathbf{h}(x_j), \mathbf{e}_c) (\mathbf{y}'_{jc} - \mathbf{y}^*_{jc}) \right| \\ &\leq B \frac{1}{m} \sum_{j=1}^m \sum_{c=1}^k \left| \mathbf{y}'_{jc} - \mathbf{y}^*_{jc} \right| \leq B D_{\mathcal{Y}^\diamond} \end{aligned}$$

To wrap it up, it results that

$$\begin{aligned}
R(\mathbf{h}_{\hat{\theta}}) &\leq \hat{R}(\mathbf{h}_{\theta^*}; X, \mathbf{Y}^*) + BD_{\mathbb{Y}^\circ} + 2\hat{\mathfrak{R}}(\mathcal{L}; X, \mathbf{Y}^*) + \mathbf{O}\left(B\sqrt{\frac{\ln \frac{1}{\delta}}{m}}\right) \\
&\leq R(\mathbf{h}_{\theta^*}) + BD_{\mathbb{Y}^\circ} + 4\hat{\mathfrak{R}}(\mathcal{L}; X, \mathbf{Y}^*) + \mathbf{O}\left(B\sqrt{\frac{\ln \frac{1}{\delta}}{m}}\right) \\
&\leq R(\mathbf{h}_{\theta^*}) + BD_{\mathbb{Y}^\circ} + \sup_{Y \in \mathbb{Y}^\circ} 4\hat{\mathfrak{R}}(\mathcal{L}; X, Y) + \mathbf{O}\left(B\sqrt{\frac{\ln \frac{1}{\delta}}{m}}\right).
\end{aligned}$$

Proof of Theorem 2.4.9 This result is essentially a corollary of theorem 2.4.8. After bounding the minimax-risk, the next lines follow via elementary geometric inequalities and range assumptions.

Proof of Lemma 2.4.13 The proof is along the same lines of the proof of lemma 2.4.1, but we take a union bound with respect to all the nK intervals $\Delta_{i,c,\hat{c}}$ for $i \in 1, \dots, n$, $c \in 1, \dots, k$, and $\hat{c} = 1, \dots, k_i$. Moreover, as for any $j \in 1, \dots, m$, we have that $y_{j,c}[\varphi_i(x_j)]_{\hat{c}} \leq 1$, we take $B = 1$ during the proof (as in lemma 2.4.1).

B.2 Additional Experimental Details

We provide further information specifying the experimental setup used to generate our figures.

B.2.1 Weak Supervision Sources

We first build the weak supervision sources on our two datasets as follows.

Animals with Attributes Each class is annotated with a binary vector of attributes. For each attribute, we train a binary classifier by finetuning a ResNet-18 using labeled data from the seen classes. When we consider a classification task between two unseen classes, we use as weak supervision sources the classifiers for the attributes which are different between the animals of these two unseen classes. We report the results of the 4 binary classification tasks which have the lowest majority vote accuracy. We chose these particular tasks to demonstrate the abilities of our methods on the tasks that have the least accurate and most highly correlated weak supervision sources.

DomainNet We sample 5 classes among the 25 classes of DomainNet with the largest number of data. For each domain, we use 60% of the available data for those classes to fine tune a pretrained ResNet-18 network. We perform this procedure on two disjoint samples of test classes to illustrate our results on two distinct multi class classification tasks.

In our experiments, we use the pretrained ResNet-18 from PyTorch. We finetune this ResNet-18 network following the approach described in [He et al., 2016], using cross-entropy loss.

B.2.2 Algorithm Hyperparameters

The subgradient method (algorithm 1) used to train AMCL-CC and AMCL-LR uses the following hyperparameters:

AMCL-CC We set $\delta = 0.1$ and build the constraints as in lemma 2.4.1. We use $\varepsilon = 0.1$ and define the step size h and the number of iterations T as in corollary 2.4.7.

AMCL-LR In this case, the loss function is bounded as in lemma 2.4.11. Since this value could be potentially very large, which in turn it would result in large intervals and number of iterations, we use the value $B = 1$ in the experiments. We set δ to 0.1 and build the constraints as in lemma 2.4.1. We do not bound the set of weights Θ : in the experiments, the norm of the weights of the multinomial logistic regression model has never diverged. We run the subgradient algorithm for $1/\varepsilon^2$ iterations with step size ε , with $\varepsilon = 0.1$

B.3 Additional Figures

B.3.1 Animals with Attributes

We provide the remaining figures for our experiments on the Animals with Attributes dataset. The last two binary classification tasks are *bat* versus *rat* and *horse* versus *giraffe*.

From figure B.1, we note that our methods show similar results as the figures displayed in the main body of the paper. AMCL-LR matches or outperforms all other methods on both tasks, over all ranges labeled data. AMCL-CC is within a few accuracy points of the other baselines and AMCL-LR on these tasks.

B.3.2 DomainNet

We provide the remaining figures for our experiments on the DomainNet dataset. We provide histograms when using the other 4 domains as the target task and also provide histograms for results on another of the samples of 5 classes. The first sample of classes as mentioned in the main body of the paper is {sea turtle, vase, whale, bird, violin }. The second sample is {tornado, trombone, submarine, feather, zebra }.

From figures B.2–B.6, we note that in most domains our methods perform better than or match all other approaches, namely in both samples of *clipart*, *quickdraw*, *painting*, and the second sample of *sketch*. Our methods achieve slightly lower accuracy than the best performing baseline on the *real* domain and on the second sample of the *infograph* domain, although they are not beaten by a single baseline in all of these tasks. We believe that the combination of our theoretical guarantees and that our methods achieve similar or sometimes better empirical performance captures the benefits of AMCL.

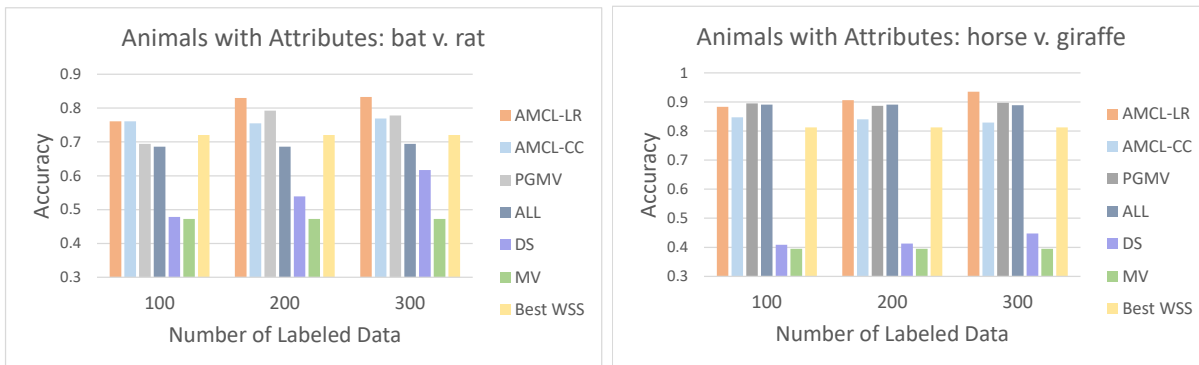


Figure B.1: Experimental results on the Animals with Attributes dataset for the binary classification tasks of *bat* versus *rat* (left) and *horse* versus *giraffe* (right) as we vary the amount of labeled data. Each method uses 347 unlabeled data points for bat versus rat and 1424 unlabeled data points for horse versus giraffe.

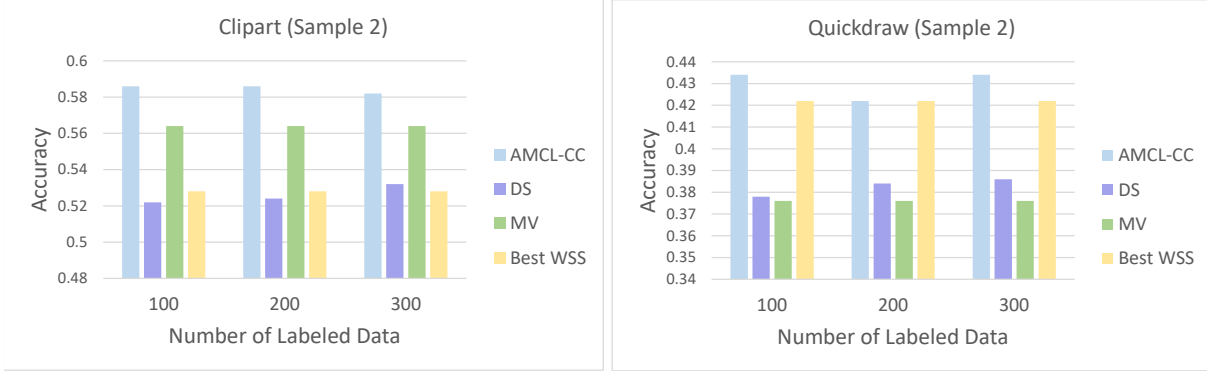


Figure B.2: Experimental results on the second sample of Domain Net for the *clipart* and *quickdraw* domains as we vary the amount of labeled data. Each method uses 500 unlabeled data points.

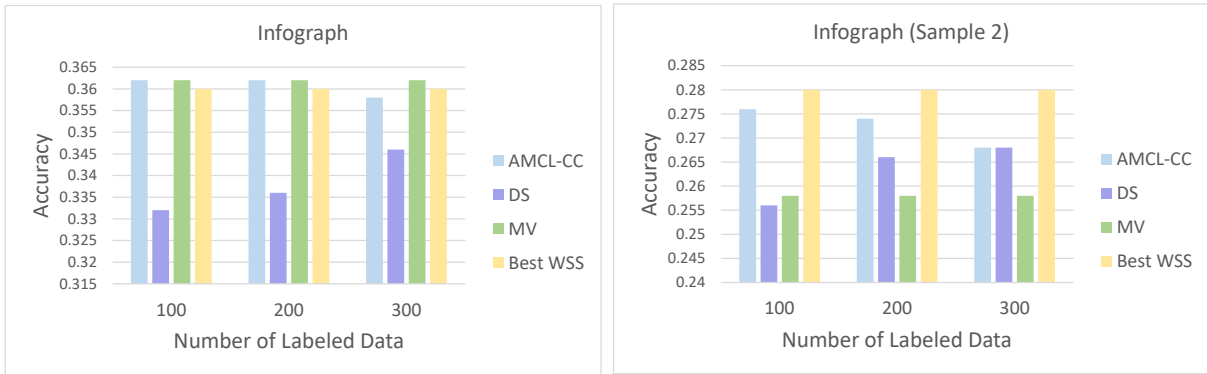


Figure B.3: Experimental results on both samples of Domain Net for the *infograph* domain as we vary the amount of labeled data. Each method uses 500 unlabeled data points.

B.4 Experiments on Synthetic Data

We run synthetic experiments to show that our method is robust with respect to the addition of correlated weak supervision sources. (Similar discussion has been done in Arachie and Huang [2019]).

We consider a multi class classification task over 5 classes, and 25 weak supervision sources $\varphi_1, \dots, \varphi_{25}$. In this classification task, each item of the domain \mathcal{X} has an unique true label. Given an item $x \in \mathcal{X}$, for $i \in 1, \dots, 10$, the weak supervision source φ_i returns the correct label with probability $\frac{1}{2}$, and a random label with probability $\frac{1}{2}$. The output of the weak supervision source φ_i is independent to the output of the weak supervision sources φ_j for $j \in \{1, \dots, 10\} \setminus \{i\}$. Therefore, the weak supervision source φ_i is correct with probability $\frac{1}{2}(1 + \frac{1}{k})$. For $i \in 11, \dots, 25$, the weak supervision sources φ_i outputs the same result than the weak supervision source φ_1 . Note that the weak supervision sources $\varphi_{11}, \dots, \varphi_{25}$ do not provide any additional information with respect to the target classification task, as they add redundant constraints to the set of feasible labelings \mathbb{Y}^\diamond . The majority vote of the weak supervision sources $\varphi_1, \dots, \varphi_{25}$ is highly affected by these dependencies, and it is very likely to provide the same answer as φ_1 , which is only $\frac{1}{2}(1 + \frac{1}{k})$ accurate on average. On the other hand, the majority vote of the weak supervision sources $\varphi_1, \dots, \varphi_{10}$ would improve upon the individual accuracy of the weak supervision sources, as their output is independent.

We use 500 unlabeled examples, run experiments varying the amount of labeled data, and show that our method AMCL-CC is robust against those dependencies. For the sake of these experiments, as we want to use very small amount of labeled data, we set $\gamma = 0$ when building the constraints for \mathbb{Y}^\diamond as in lemma 2.4.1.



Figure B.4: Experimental results on both samples of Domain Net for the *painting* domain as we vary the amount of labeled data. Each method uses 500 unlabeled data points.

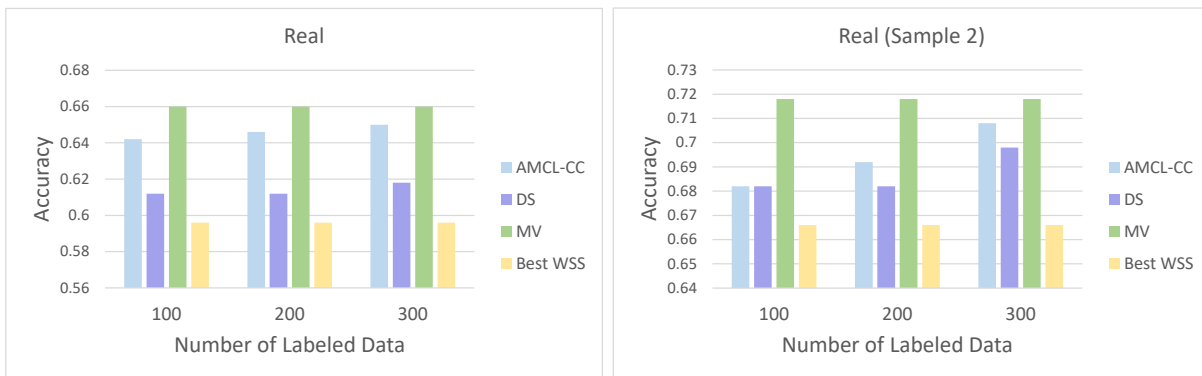


Figure B.5: Experimental results on both samples of Domain Net for the *real* domain as we vary the amount of labeled data. Each method uses 500 unlabeled data points.

The experimental results are reported in appendix B.4. The table shows that AMCL-CC is robust with respect to dependencies among weak supervision sources, whereas majority vote is greatly affected by them. In fact, in this case the majority vote does not improve upon the individual accuracy of the weak supervision sources, which is on average $\frac{1}{2}(1 + \frac{1}{k}) = \frac{3}{5}$.

Table B.1: We report the experimental results on the synthetic dataset. We report the accuracy obtained by our method AMCL-CC and the majority vote, when varying the amount of labeled examples (we report the average accuracy over 3 distinct runs).

Labeled Examples	AMCL-CC	Majority Vote
100	0.902	0.595
50	0.828	0.602
25	0.819	0.598

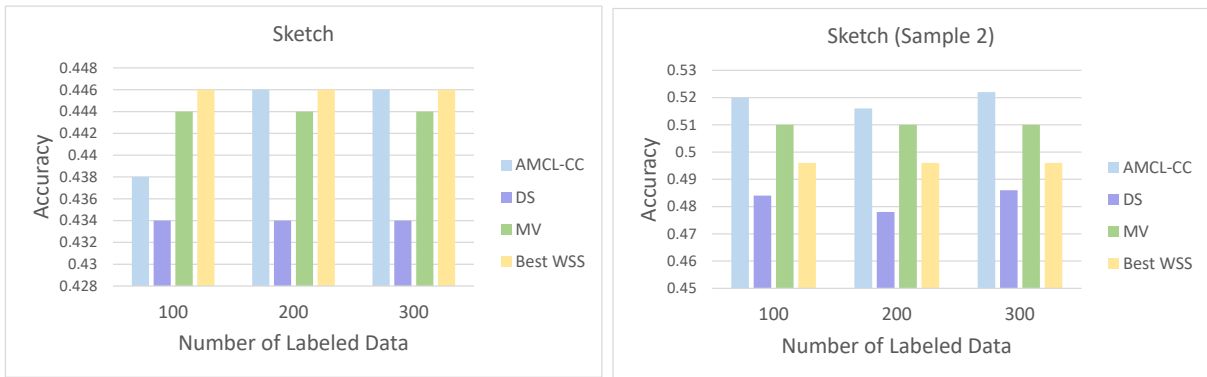


Figure B.6: Experimental results on both samples of Domain Net for the *sketch* domain as we vary the amount of labeled data. Each method uses 500 unlabeled data points.

Appendix C

Supplementary Material for Chapter 3

C.1 Welfare and Malfare

We now show theorem 3.3.7.

Theorem 3.3.7 (Properties of the Power-Mean). Suppose \mathcal{S}, ε are sentiment values $\Omega \rightarrow \mathbb{R}_{0+}$, and \mathbf{w} is a probability measure over Ω . Then

1. Monotonicity: $M_p(\mathcal{S}; \mathbf{w})$ is weakly-monotonically-increasing in p , and strictly if \mathcal{S} attains distinct $a, b \in \mathbb{R}$ with nonnegligible probability.
2. Subadditivity: $\forall p \geq 1 : M_p(\mathcal{S} + \varepsilon; \mathbf{w}) \leq M_p(\mathcal{S}; \mathbf{w}) + M_p(\varepsilon; \mathbf{w})$.
3. Contraction: $\forall p \geq 1 : |M_p(\mathcal{S}; \mathbf{w}) - M_p(\mathcal{S}'; \mathbf{w})| \leq M_p(|\mathcal{S} - \mathcal{S}'|; \mathbf{w}) \leq \|\mathcal{S} - \mathcal{S}'\|_\infty$.
4. Curvature: $M_p(\mathcal{S}; \mathbf{w})$ is concave in \mathcal{S} for $p \in [-\infty, 1]$ and convex for $p \in [1, \infty]$.

Proof. We omit proof of item 1, as this is a standard property of power-means (generally termed the *power mean inequality* [Bullen, 2013, chapter 3]).

We first show item 2. By the triangle inequality (for $p \geq 1$), we have

$$M_p(\mathcal{S} + \varepsilon; \mathbf{w}) \leq M_p(\mathcal{S}; \mathbf{w}) + M_p(\varepsilon; \mathbf{w}) .$$

We now show item 3 First take $\varepsilon \doteq \mathcal{S} - \mathcal{S}'$. Now consider

$$\begin{aligned} M_p(\mathcal{S}; \mathbf{w}) &= M_p(\mathcal{S}' + \varepsilon; \mathbf{w}) && \text{DEFINITION OF } \varepsilon \\ &\leq M_p(\mathcal{S}' + \varepsilon_+; \mathbf{w}) && \text{MONOTONICITY} \\ &\leq M_p(\mathcal{S}'; \mathbf{w}) + M_p(\varepsilon_+; \mathbf{w}) && \text{ITEM 2} \\ &\leq M_p(\mathcal{S}'; \mathbf{w}) + M_p(|\mathcal{S} - \mathcal{S}'|; \mathbf{w}) . && \text{MONOTONICITY} \end{aligned}$$

By symmetry, we have $M_p(\mathcal{S}', \mathbf{w}) \leq M_p(\mathcal{S}, \mathbf{w}) + M_p(|\mathcal{S} - \mathcal{S}'|; \mathbf{w})$, which implies the result.

We now show item 4. First note the special cases of $p \in \pm\infty$ follow by convexity of the maximum ($p = \infty$) and concavity of the minimum ($p = -\infty$).

Now, note that for $p \geq 1$, by concavity of $\sqrt[p]{\cdot}$, Jensen's inequality gives us

$$M_1(\mathcal{S}; \mathbf{w}) = \mathbb{E}_{\omega \sim \mathbf{w}}[\mathcal{S}(\omega)] = \underbrace{\mathbb{E}_{\omega \sim \mathbf{w}}[\sqrt[p]{\mathcal{S}^p(\omega)}]}_{\text{DEFINITION OF CONVEXITY}} \leq \sqrt[p]{\mathbb{E}_{\omega \sim \mathbf{w}}[\mathcal{S}^p(\omega)]} = M_p(\mathcal{S}; \mathbf{w}) ,$$

i.e., convexity, and similarly, for $p \leq 1$, $p \neq 0$, we have by convexity of $\sqrt[p]{\cdot}$, we have

$$M_1(\mathcal{S}; \mathbf{w}) = \mathbb{E}_{\omega \sim \mathbf{w}}[\mathcal{S}(\omega)] = \underbrace{\mathbb{E}_{\omega \sim \mathbf{w}}[\sqrt[p]{\mathcal{S}^p(\omega)}]}_{\text{DEFINITION OF CONCAVITY}} \geq \sqrt[p]{\mathbb{E}_{\omega \sim \mathbf{w}}[\mathcal{S}^p(\omega)]} = M_p(\mathcal{S}; \mathbf{w}) .$$

Similar reasoning, now by convexity of $\ln(\cdot)$, shows the case of $p = 0$. \square

We now show theorem 3.3.8.

Theorem 3.3.8 (Population Mean Properties). Suppose population-mean function $M(\mathcal{S}; \mathbf{w})$. If $M(\cdot; \cdot)$ satisfies (subsets of) the population-mean axioms (see definition 3.3.4), we have that $M(\cdot; \cdot)$ exhibits the following properties. For each, assume arbitrary sentiment-value function $\mathcal{S} : \Omega \rightarrow \mathbb{R}_{0+}$ and weights measure \mathbf{w} over Ω . The following then hold.

1. *Identity*: Axioms 6 & 7 imply $M(\omega \mapsto \alpha; \mathbf{w}) = \alpha$.
2. Axioms 1-5 imply $\exists p \in \mathbb{R}$, strictly-monotonically-increasing continuous $F : \mathbb{R} \rightarrow \mathbb{R}_{0+}$ s.t.

$$M(\mathcal{S}; \mathbf{w}) = F\left(\int_{\mathbf{w}} f_p(\mathcal{S}(\omega)) d(\omega)\right) = F\left(\mathbb{E}_{\omega \sim \mathbf{w}}[f_p(\mathcal{S}(\omega))]\right) , \quad \text{with } \begin{cases} p = 0 & f_0(x) \doteq \ln(x) \\ p \neq 0 & f_p(x) \doteq \text{sgn}(p)x^p \end{cases} .$$

3. Axioms 1-7 imply $F(x) = f_p^{-1}(x)$, thus $M(\mathcal{S}; \mathbf{w}) = M_p(\mathcal{S}; \mathbf{w})$.
4. Axioms 1-5 and 8 imply $p \in (-\infty, 1]$.
5. Axioms 1-5 and 9 imply $p \in [1, \infty)$.

Proof. Item 1 is an immediate consequence of axioms 6 & 7 (multiplicative linearity and unit scale).

We now note that item 2 is the celebrated Debreu-Gorman theorem [Debreu, 1959, Gorman, 1968], extended by continuity and measurability of \mathcal{S} to the weighted case.

We now show item 3. This result is essentially a corollary of item 2, hence the dependence on axioms 1-4. Suppose $\mathcal{S}(\cdot) = 1$. By item 1, for all $p \neq 0$, we have

$$\alpha = \alpha M(\mathcal{S}; \mathbf{w}) = M(\alpha \mathcal{S}; \mathbf{w}) = F\left(\mathbb{E}_{\omega \sim \mathbf{w}}[f_p(\alpha \mathcal{S}(\omega))]\right) = F\left(\mathbb{E}_{\omega \sim \mathbf{w}}[f_p(\alpha)]\right) = F(\text{sgn}(p)\alpha^p) .$$

From here, we have $\alpha = F(\text{sgn}(p)\alpha^p)$, thus $F^{-1}(u) = \text{sgn}(p)u^p$, and consequently, $F(v) = \sqrt[p]{\text{sgn}(p)v}$.

Taking $p = 0$ gets us

$$\alpha = \alpha M(\mathcal{S}; \mathbf{w}) = F\left(\mathbb{E}_{\omega \sim \mathbf{w}}[\ln(\alpha \mathcal{S}(\omega))]\right) = F(\ln a) ,$$

from which it is clear that $F^{-1}(u) = \ln(u) \Rightarrow F(v) = \exp(v)$.

For all values of $p \in \mathbb{R}$, substituting the values of f_p and $F(\cdot)$ into item 2 yields $M(\mathcal{S}; \mathbf{w}) = M_p(\mathcal{S}; \mathbf{w})$ by definition.

We now show 4 and 5. These properties follow directly from 2, wherein f_p are defined, and Jensen's inequality. \square

We now show corollary 3.4.2.

Corollary 3.4.2 (Statistical Estimation with Hoeffding and Bennett Bounds). Suppose fair power-mean malfare $M(\cdot; \cdot)$ (i.e., $p \geq 1$), loss function $\ell : \mathcal{X} \rightarrow [0, r]$, $\mathcal{S} \in [0, r]^g$ s.t. $\mathcal{S}_i = \mathbb{E}_{\mathcal{D}_i}[\ell]$, samples $\mathbf{x}_i \sim \mathcal{D}_i^m$, and take $\hat{\mathcal{S}}_i \doteq \frac{1}{m} \sum_{j=1}^m \ell(\mathbf{x}_{i,j})$. Then with probability at least $1 - \delta$ over choice of \mathbf{x} ,

$$\left| M_p(\mathcal{S}; \mathbf{w}) - M_p(\hat{\mathcal{S}}; \mathbf{w}) \right| \leq r \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}} .$$

Alternatively, again with probability at least $1 - \delta$ over choice of \mathbf{x} , we have

$$\left| M_p(\mathcal{S}; \mathbf{w}) - M_p(\hat{\mathcal{S}}; \mathbf{w}) \right| \leq \frac{r \ln \frac{2g}{\delta}}{3m} + \max_{i \in \{1, \dots, g\}} \sqrt{\frac{2 \mathbb{V}_{\mathcal{D}_i}[\ell] \ln \frac{2g}{\delta}}{m}} .$$

Proof. This result is a corollary of lemma 3.4.1, applied to ε , where we note that for $p \geq 1$, by theorem 3.3.7 item 3 (contraction) it holds that

$$M_p(\hat{\mathcal{S}} + \varepsilon; \mathbf{w}) \leq M_p(\hat{\mathcal{S}}; \mathbf{w}) + \|\varepsilon\|_\infty \quad \& \quad M_p(\mathbf{0} \vee (\hat{\mathcal{S}} - \varepsilon); \mathbf{w}) \leq M_p(\hat{\mathcal{S}}; \mathbf{w}) - \|\varepsilon\|_\infty .$$

Now, for the first bound, note that we take $\varepsilon_i \doteq r \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}}$, and by Hoeffding's inequality and the union bound, for $\Omega = \{1, \dots, n\}$, we have $\forall \omega : \mathcal{S}'(\omega) - \varepsilon(\omega) \leq \mathcal{S}(\omega) \leq \mathcal{S}'(\omega) + \varepsilon(\omega)$ with probability at least $1 - \delta$. The result then follows via the *power-mean contraction* (theorem 3.3.7 item 3) property.

Similarly, for the second bound, note that we take $\varepsilon_i \doteq \frac{r \ln \frac{2g}{\delta}}{3m} + \sqrt{\frac{2 \mathbb{V}_{\mathcal{D}_i}[\ell] \ln \frac{2g}{\delta}}{m}}$, which this time follows via Bennett's inequality and the union bound. Now, we again apply lemma 3.4.1, noting that $M(\varepsilon) \leq M_\infty(\varepsilon) = \|\varepsilon\|_\infty$ (by *power-mean monotonicity*, theorem 3.3.7 item 1), and the rest follows as in the Hoeffding case. \square

Appendix D

Supplementary Material for Chapter 4

D.1 Pseudocode and Ewok^s Algorithm Details

Algorithm 6 presents pseudocode for the Ewok^s. In particular, we describe the *optimization* algorithm CODECSELECTION(...), where the optimal $\hat{c}_W \in \binom{C}{k}$ is selected, as well as the routines STOREMEDIA(...), which adds a piece of media to the database, and QUERYMEDIA(...), which serves a user the best-available encoding for their loss function.

Algorithm 6 Pseudocode for EWOks Operations

```
1: procedure CODECSELECTION( $\mathcal{C}, k, \mathcal{L}, W, \mathbf{x}$ )  $\mapsto \hat{\mathbf{c}}_W$ 
2:   Input: Codec family  $\mathcal{C}$ , selection size  $k$ , loss distribution  $\mathcal{L}$ , welfare concept  $W$ , sample  $\mathbf{x} \in \mathcal{X}^m$ 
3:   Output: EWOks-optimal  $\mathbf{c} \in \binom{\mathcal{C}}{k}$ 
4:   Effect Sets global codec-selection  $\hat{\mathbf{c}}_W$  (for use with ENCODEMEDIA)
5:   for all  $i \in 1, \dots, |\mathcal{C}|$  do ▷ Precompute attribute values
6:     for all  $j \in 1, \dots, m$  do
7:        $\mathbf{A}_{i,j,:} \leftarrow \mathcal{C}_i(\mathbf{x}_j)$  ▷  $\mathbf{A} \in \mathbb{R}_+^{|\mathcal{C}| \times m \times a}$ 
8:     end for
9:   end for
10:   $\mathbf{I} \leftarrow \{i \in 1, \dots, |\mathcal{C}| \mid \nexists i' < i \text{ s.t. } \mathbf{A}_{i',:,} = \mathbf{A}_{i',:,}\}$  ▷ Prune Coincident Codecs
11:  if PRUNE MONOTONIC then ▷ Prune monotonically-dominated codecs (Optional)
12:     $\mathbf{I} \leftarrow \{i \in \mathbf{I} \mid \nexists i' \text{ s.t. } \mathbf{A}_{i',:,} \preceq \mathbf{A}_{i',:,}\}$ 
13:  end if
14:  if PRUNE GENERIC then ▷ Prune  $\mathcal{L}$ -dominated codecs (Optional)
15:     $\mathbf{I} \leftarrow \left\{ i \in \mathbf{I} \mid \nexists i' \neq i \text{ s.t. } \sup_{\ell \in \mathcal{L}} \max_{j \in 1, \dots, m} \ell(\mathbf{A}_{i',j,:}) - \ell(\mathbf{A}_{i,j,:}) < 0 \right\}$ 
16:  end if
17:   $\mathbf{I}' \leftarrow \underset{\mathbf{I}' \in \binom{\mathbf{I}}{k}}{\operatorname{argmin}} \hat{W}(\{i(j) \mapsto \mathbf{A}_{i,j,1:a} \mid i \in \mathbf{I}'\}; \mathcal{L}, 1, \dots, m)$  ▷ Combinatorial Optimization
18:   $\hat{\mathbf{c}}_W \leftarrow \{\mathcal{C}_i \mid i \in \mathbf{I}'\}$ 
19:  return  $\hat{\mathbf{c}}_W$ 
20: end procedure
21: procedure ENCODEMEDIA( $x, \text{ID}$ )  $\mapsto \emptyset$ 
22:   Input: Media  $x$ , media ID  $\text{ID}$ .
23:   Effect Stores each encoding  $\forall c \in \hat{\mathbf{c}}_W : c(x)$  of media  $x$  at index  $\text{ID}$ .
24:   for  $c \in \hat{\mathbf{c}}$  do
25:      $\text{STORE}_{\text{ID},c} \leftarrow c(x)$  ▷ Store encoded media and attributes
26:   end for
27: end procedure
28: procedure QUERYMEDIA( $\ell, \text{ID}$ )  $\mapsto (\text{CODEC} \times \text{ENCODING})$ 
29:   Input: Loss function  $\ell$ , some media ID  $\text{ID}$  that has previously been passed to ENCODEMEDIA( $\cdot, \cdot$ )
30:   Output: Optimal codec  $c \in \hat{\mathbf{c}}_W$  and encoding  $c(x)$ , for  $x$  the media referenced by  $\text{ID}$ 
31:    $c \leftarrow \underset{c' \in \hat{\mathbf{c}}}{\operatorname{argmin}} \ell(\text{STORE}_{\text{ID},c'})$ 
32:   return  $(c, \text{STORE}_{\text{ID},c})$  ▷ Return encoded media
33: end procedure
```

D.2 Supplementary Experiments and Methods

Here we present supplementary detail on our datasets, additional experiments in the video domain, and additional high-detail plots and extensions of experiments presented in the paper body.

D.2.1 Experimental Setup

As discussed in Section 4.3, we performed encoding experiments on audio data extracted from audio-books and music. In this section, we present additional detail on our audio and video datasets, as well as encoding parameters.

For each of our *music*, *speech*, and *video* datasets, we give the *sample count*, *sample length*, and detailed *format information* in table D.1. Each music dataset was ripped from CD, and the *au* dataset was downloaded from librivox.org [McGuire]. Each of the files in the original datasets are split into *sample count* contiguous

Table D.1: Sample count, length, rate, and depth of all music, speech, and video datasets. All numbers given are for the original (uncompressed) data formats.

Modality	Dataset	Count	Length	Sample Rate	Sample Depth	Information Rate
MUSIC	de	8741	2s	44100hz	16 bit Stereo (2ch)	1411200 b/s 172 KB/s
	ex	9443				
	lz	13163				
	TOTAL	31347				
SPEECH	au	16530	10s	44100hz	16 bit Monaural (1ch)	705600 b/s 86 KB/s
VIDEO	video	1033	1s	30fps	256 × 144 px 3 col (RGB) 8 b/col/px	26542080 b/s 3240 KB/s

nonoverlapping *sample length* segments, measured in *seconds* (s), with a possible shorter final “leftover segment” per file. In particular, we decompose the *information rate*, measured in *bit rate* bits per second (b/s), or, perhaps more conveniently, base-1000-kilobytes per second (KB/s; 1KB = 8000b), into *sample rate*, which gives the number of discrete samples per second, measured in *cycles / second* (Hertz or hz), or equivalently in *frames / second* (fps), and *sample depth*, which is the number of bits per sample. Sample depth is further decomposed by domain: audio is either *stereo*, which refers to 2-channel (left and right) sound (thus *doubling* the bitrate), or *monaural*, which contains a single channel, generally heard in both ears; similarly, *video samples* are *frames*, or *images*, each split into a resolution $x \times y$ array of *pixels* (px), each of which has 3 color values (RGB). Both audio waveform samples and color intensity values are represented as linearly-scaled unsigned fixed-width integers.

All audio was initially in the uncompressed `wav` format (see table D.1), and all `mp3` encoding was performed with the LAME [Hegemann et al., 2017] encoder. LAME is a recursive acronym for *Lame Ain’t an Mp3 Encoder*, and is known for its advanced psychoacoustic modelling, flexibility, and high perceptual-quality `mp3` encodings. Details on audio encoding are given in table 4.1. The video dataset was ripped from a youtube video (compressed), but we treated it as though the original rip were the ground-truth (uncompressed) video, and measured deviation from the original as distortion, though we compute *compression ratio* as a ratio of the original compressed video size. Our experiments utilize 22 `mp4` codec variants, each a *constant rate factor* (CRF) `mp4`, encoded by FFMPEG, with *quality* parameter ranging from $\{30, \dots, 51\}$.

D.2.2 Video Experiments

To demonstrate the flexibility of our approach, we also apply EWoks to the selection of video codecs. Perceptual quality metrics are an area of active research, particularly in the visual domain, and in our video experiments, we characterize distortion with the *peak signal to noise ratio* (PSNR) Erfurt et al. [2019] metric. Figure D.1 visualizes the Pareto-optimal frontier for 22 `mp4` video compression codecs, which showcases similar trade-offs as those demonstrated for audio compression in figure 4.1. We use `mp4` codecs with integer-valued *Constant Rate Factor* (quality parameter) between 30 and 51 with the FFMPEG compression software. The mean performance of 22 codecs is plotted, along with Pareto optimality curves for all $k \in \{1, 2, |\mathcal{C}|\}$.

Because this experiment is a proof of concept, our dataset is limited in size and complexity. Our set of media consists of a 17:13 minutes:seconds long, 144p 30fps video *youtube video* (<https://www.youtube.com/watch?v=310PF8ctRTM>) sliced into 1033 distinct 1-second segments. We removed the audio from the video, so the we need only compress 1-second intervals of frames of images.

We measure divergence in terms of the *peak signal-to-noise ratio* (PSNR), a common method for measuring how the *fidelity* of a compressed image (i.e., how much said image resembles its original form) [Gonzalez and Woods, 1992]. PSNR compares two images by taking the *logarithm* of the *maximum squared error* over the

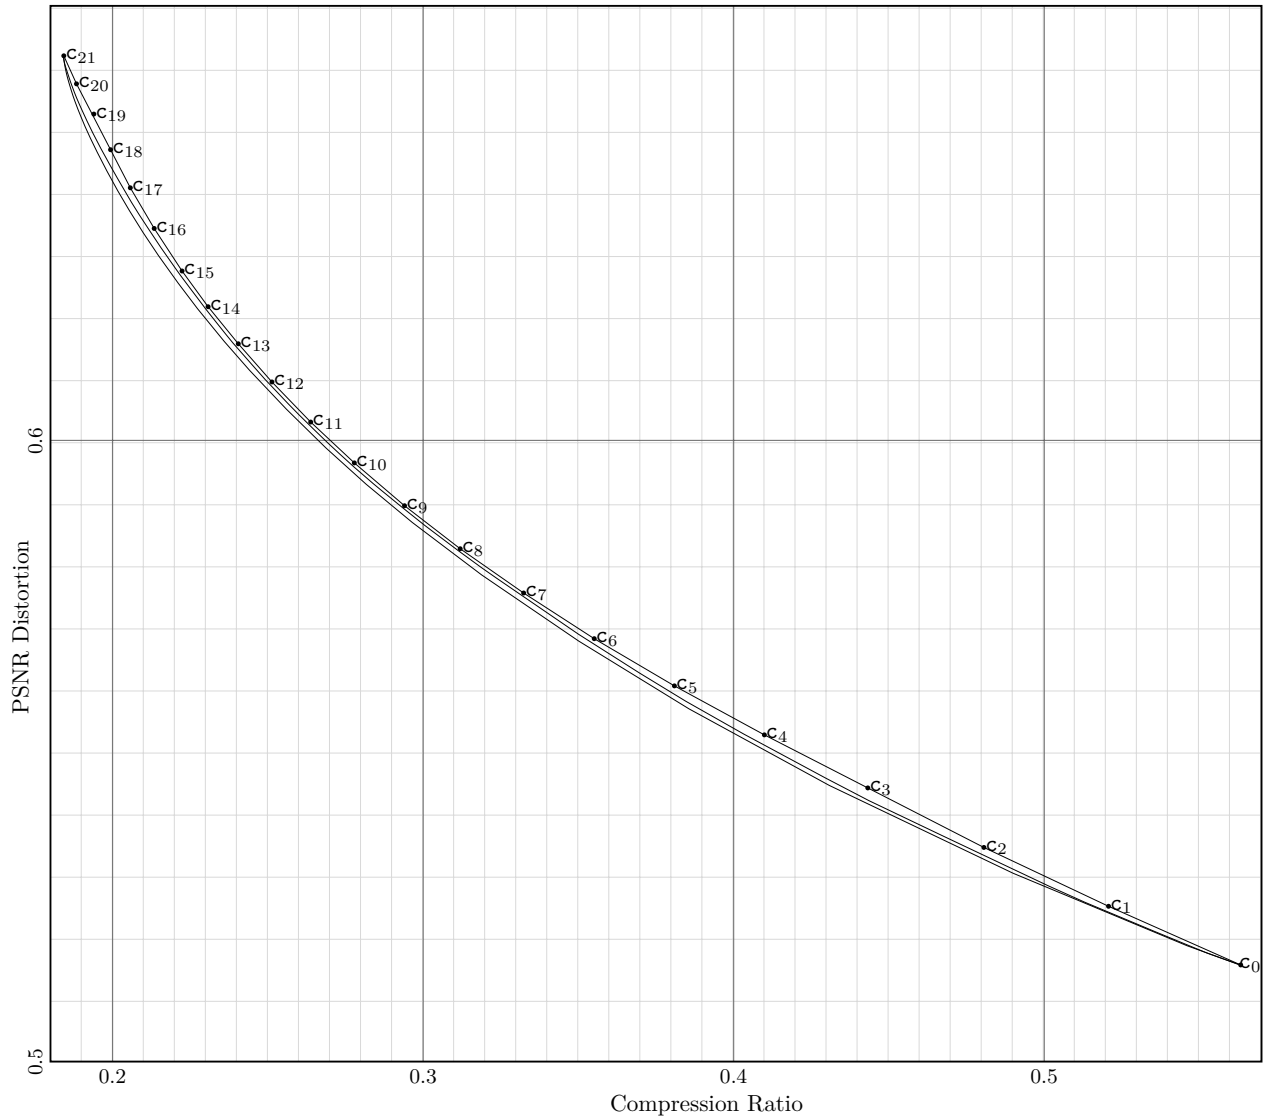


Figure D.1: Empirical Pareto-optimal frontiers for a selection of video codecs, for $k \in \{1, 2, |\mathcal{C}|\}$. Divergence measure $PSNR$ (y -axis) versus $compression\ ratio$ (x -axis) on the video dataset.

mean squared error. This scaling ensures that subtle and drastic changes to an image can be compared at the same order of magnitude with the same metric. This metric can also be applied to video compression by computing frame-by-frame differences between pixels. Because we need to restrict our attributes to the interval $[0, 1]$, and want a measure of *distortion* rather than *quality*, we take the $PSNR\ Divergence$ to be $1 - \frac{PSNR}{\max\ PSNR}$, where $\max\ PSNR$ refers to the largest PSNR divergence over all our codecs for the given fixed video segment.

D.2.3 Supplementary Audio Experiments

Here we show additional plots and experiments on the audio datasets discussed in section 4.3.

In particular, we first complement the Pareto-optimality visualization of figure 4.1 with figure D.2, which shows the same in much greater detail (matching the inset, with mean-attribute sets of all codec pairs).

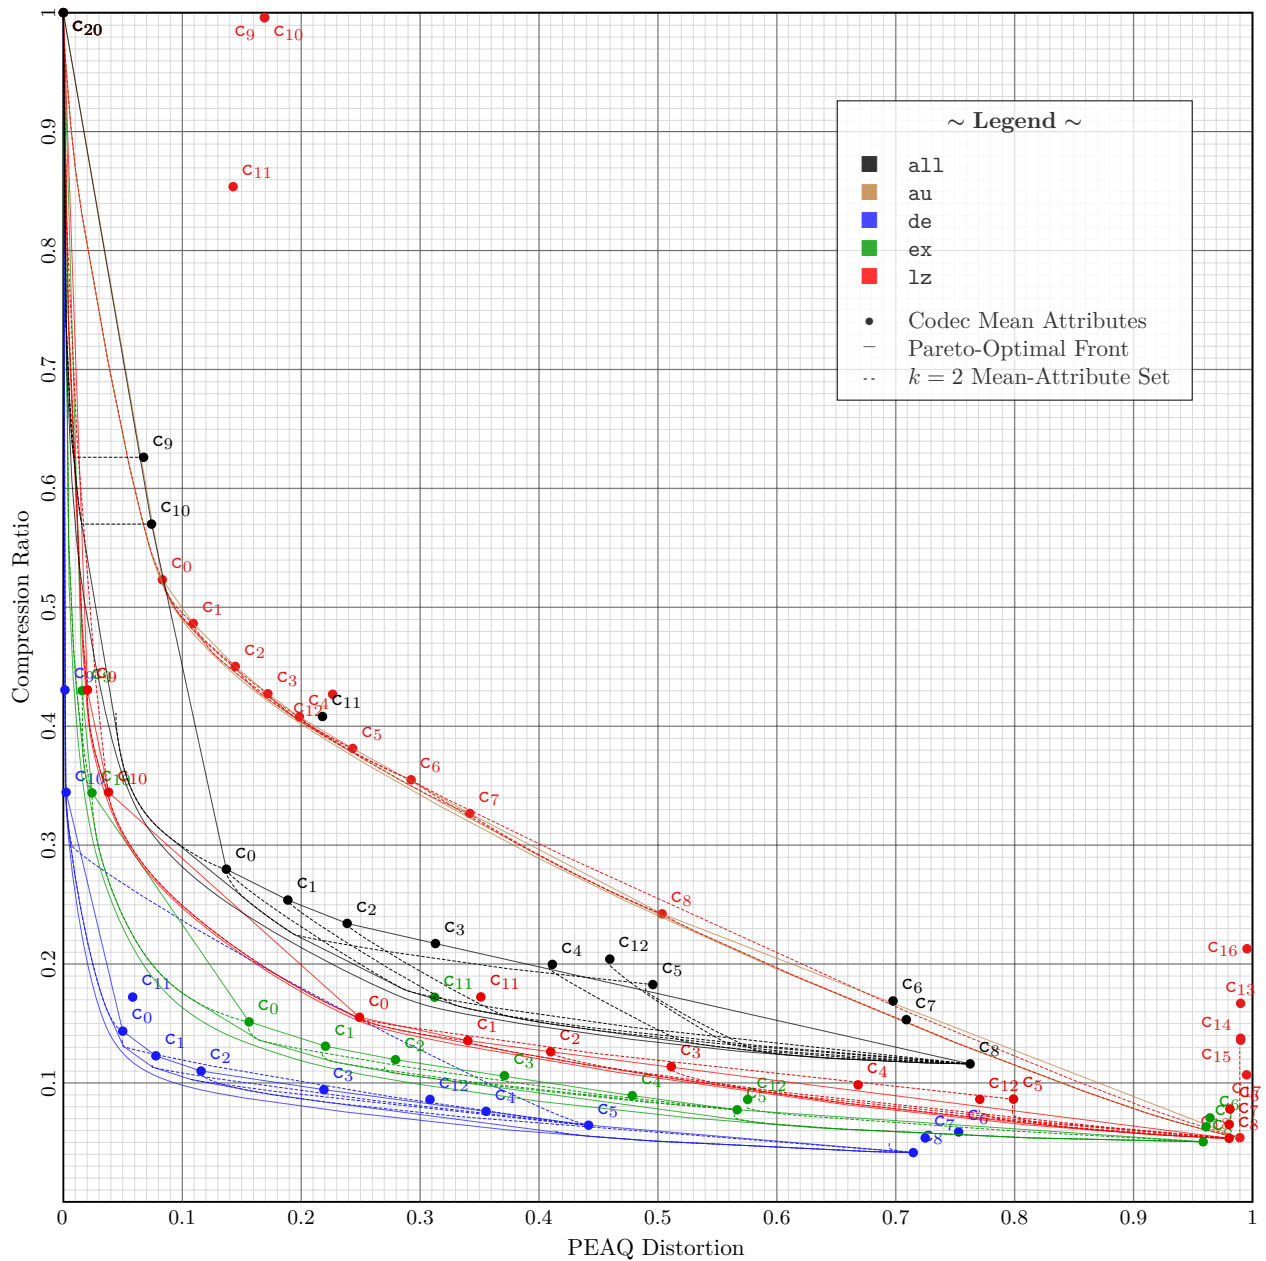
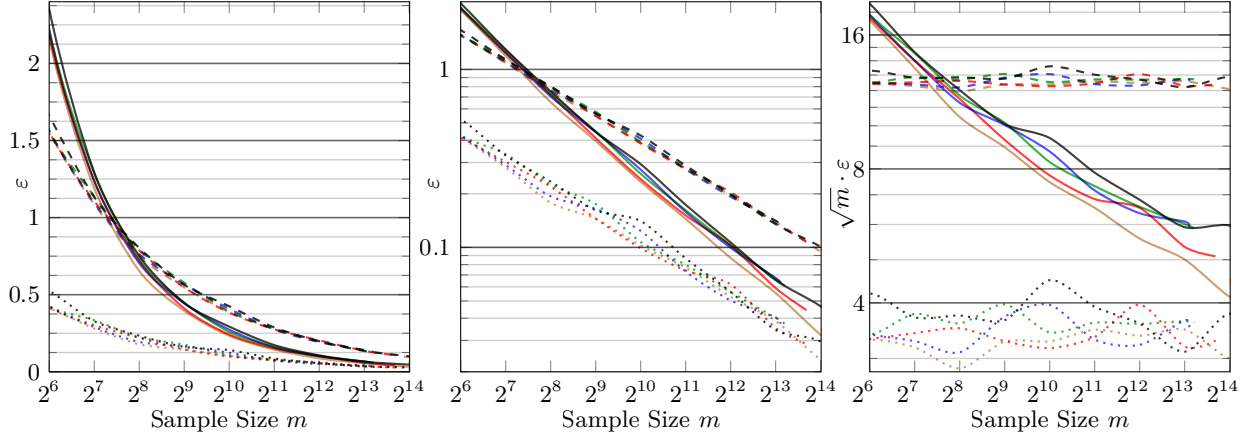
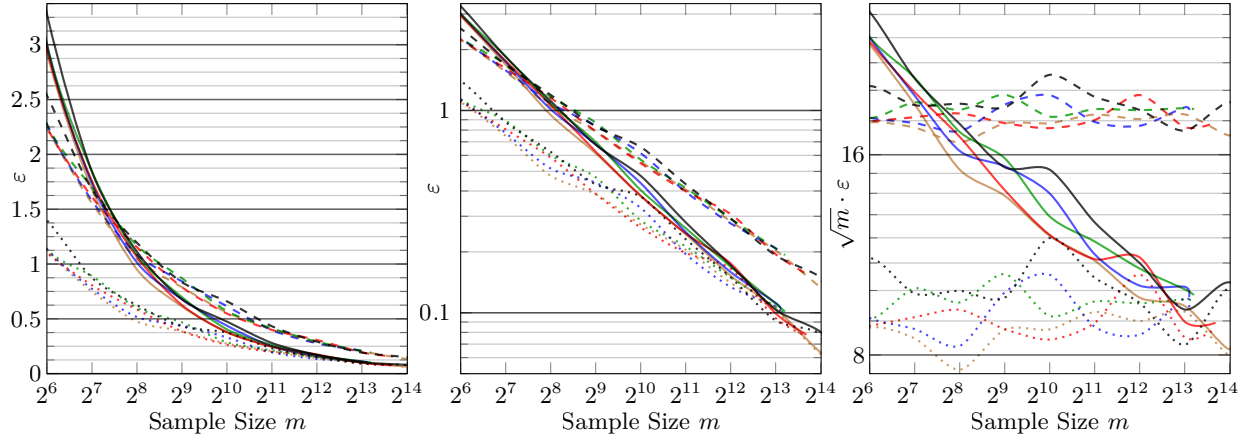


Figure D.2: Pareto optimality curves for $k \in \{1, 2, |\mathcal{C}|\}$, for all audio datasets. The figure matches figure 4.1, except here, the music dataset is split into its constituent datasets (de, ex, and lz). All codecs, Pareto-optimality curves, and Pareto-optimal pair mean-attribute-sets are color-coded by dataset (see legend in top right), and again k always increases from top-left to bottom-right.

W(\cdot)-Gap Bounds ε versus Sample Size m , $k = 3$, All Datasets



W(\cdot)-Gap Bounds ε versus Sample Size m , $k = 5$, All Datasets



~ Legend ~

Dataset	de	ex	lz	au	all
ε Bound	... $4\hat{\mathfrak{R}}_m^n(\mathcal{F}, \mathbf{x}, \sigma)$		--- $4\hat{\mathfrak{R}}_m^n(\mathcal{F}, \mathbf{x}, \sigma) + 6\sqrt{\frac{\ln \frac{1}{\delta}}{2m}}$		— $4\hat{\mathfrak{R}} + 2\varepsilon_{\text{SD}}$ (Theorem 4.2.2)

Figure D.3: Semilog, log-log, and normalized log-log plots of EWoks welfare-optimality gap bounds ε (x -axis) bounds versus sample size m (y -axis), in expectation with $4\hat{\mathfrak{R}}_m^n(\mathcal{F}, \mathbf{x}, \sigma)$, and w.h.p. with McDiarmid and theorem 4.2.2. The first row presents results for $k = 3$, and the second for $k = 5$. The left, center, and right plots in each row contain the same information, and vary only in presentation. The left column contains *semilog* (log- x , natural- y axes) plots, and the center column contains *log-log* plots. The right column also contains *log-log* plots, but in this plot, all bounds are scaled by \sqrt{m} , to visually characterize $\Theta\sqrt{1/m}$ rates as straight lines.

We also give in figure D.3 plots supplementing figure 4.3, now with semilog, log-log, and \sqrt{m} -scaled plots of the same data. We additionally present plots for $k = 3$ and $k = 5$ (only $k = 3$ is given in the paper body).

D.3 Proof Compendium

We now prove our main results stated in the paper body.

D.3.1 Proof of Theorem 4.2.2

The result is shown via a chain of union bounds on Rademacher averages and variances. We now show several lemmas in service of the main result.

Lemma D.3.1 (Empirical Wimpy Variance Expectation). Suppose $\mathbf{x} \sim \mathcal{D}^m$, and take

$$v \doteq \sup_{f \in \mathcal{F}} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\sum_{i=1}^m (f(\mathbf{x}_i) - \mathbb{E}_{\mathcal{D}}[f])^2 \right] \quad \& \quad \hat{v} \doteq \sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f])^2 .$$

Then $v \leq \frac{m}{m-1} \mathbb{E}_{\mathbf{x}}[\hat{v}]$.

Proof.

$$\begin{aligned} v &= \sup_{f \in \mathcal{F}} \mathbb{E}_{\mathbf{x}} \left[\frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - \mathbb{E}_{\mathcal{D}}[f])^2 \right] && \text{LINEARITY} \\ &= \sup_{f \in \mathcal{F}} \mathbb{E}_{\mathbf{x}} \left[\frac{1}{m-1} \sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f])^2 \right] && \text{BESSEL'S CORRECTION} \\ &\leq \mathbb{E}_{\mathbf{x}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m-1} \sum_{i=1}^m (f(\mathbf{x}_i) - \hat{\mathbb{E}}_{\mathbf{x}}[f])^2 \right] && \text{JENSEN'S INEQUALITY} \\ &= \frac{m}{m-1} \mathbb{E}_{\mathbf{x}}[\hat{v}] && \text{DEFINITION OF } \hat{v} \end{aligned}$$

□

Lemma D.3.2 (Empirical Wimpy Variance Concentration). Suppose as in lemma D.3.1, and furthermore $\mathcal{F} \subseteq \mathcal{X} \rightarrow [0, 1]$. Then

$$\mathbb{P} \left(v \geq \frac{m}{m-1} \hat{v} + \frac{2 \ln(\frac{1}{\delta})}{m} + \sqrt{\frac{2 \hat{v} \ln(\frac{1}{\delta})}{m-1}} \right) \leq \delta .$$

Proof. Suppose $\mathbf{X} \in \mathcal{X}^m$, let $\mathbf{X}_{\setminus j}$ denote the vector \mathbf{X} , missing the j th component, and take $Z(\mathbf{X}) \doteq \sup_{f \in \mathcal{F}} (m -$

$1) \hat{V}_{\mathbf{X}}[f] = \sup_{f \in \mathcal{F}} \sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2$, and $Z_j(\mathbf{X}) \doteq \sup_{f \in \mathcal{F}} (m-2) \hat{V}_{\mathbf{X}_{\setminus j}}[f] = \sup_{f \in \mathcal{F}} \sum_{\substack{i=1 \\ i \neq j}}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2$, where $\hat{V}_{\mathbf{X}}[f]$ denotes the (unbiased) sample variance of f over \mathbf{X} .

We first show that

$$Z_j(\mathbf{X}) = \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m} (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 ; \quad (68)$$

to see this, consider that

$$\begin{aligned}
Z_j(\mathbf{X}) &= \sup_{f \in \mathcal{F}} \sum_{\substack{i=1 \\ i \neq j}}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 = \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 \right) - (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{ADDITIVE IDENTITY} \\
&&& \text{DEFINITION OF } Z_j \\
&= \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m \left((f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f]) + (\hat{\mathbb{E}}_{\mathbf{X}}[f] - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f]) \right)^2 \right) - (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{ADDITIVE IDENTITY} \\
&= \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 + 2(f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])(\hat{\mathbb{E}}_{\mathbf{X}}[f] - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f]) + (\hat{\mathbb{E}}_{\mathbf{X}}[f] - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 \right) - (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{UNFACTORING} \\
&= \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) + m(\hat{\mathbb{E}}_{\mathbf{X}}[f] - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 - (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f]) = 0 \\
&= \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) + m\left(\frac{1}{m}f(\mathbf{X}_j) - \frac{1}{m}\hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f]\right)^2 - (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \mathbb{E}_{\mathbf{X}}[f] = \frac{1}{m}f(\mathbf{X}_j) + \frac{m-1}{m}\hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f] \\
&= \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m}(f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 . && \text{ALGEBRA}
\end{aligned}$$

The desideratum reduces to showing that $\frac{m}{m-1}Z(\cdot)$ is an $(\frac{m}{m-1}, 0)$ -self-bounding function (see Boucheron et al. [2009]). *Nonnegativity* is clear from the definition of $Z(\cdot)$. We now show *underestimation*; first consider that for any $j \in \{1, \dots, m\}$,

$$\begin{aligned}
Z_j(\mathbf{X}) &= \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m}(f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{EQ. 68} \\
&\geq \left(\sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) \right) - \left(\sup_{f \in \mathcal{F}} \frac{m-1}{m}(f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 \right) && \text{PROPERTIES OF SUPREMA} \\
&\geq Z(\mathbf{X}) - \frac{m-1}{m} . && \text{DEFINITION OF } Z(\cdot) \\
&&& \forall f \in \mathcal{F}: \text{image}(f) \subseteq [0, 1]
\end{aligned}$$

We may thus conclude that $\frac{m}{m-1}Z_j(\mathbf{X}) \geq \frac{m}{m-1}Z(\mathbf{X}) - 1$, which satisfies 1-*underestimation*. We now show that $Z(\cdot)$, and subsequently $\frac{m}{m-1}Z(\cdot)$, obey $(\frac{m}{m-1}, 0)$ -*self-boundedness*.

$$\begin{aligned}
\sum_{j=1}^m Z(\mathbf{X}) - Z_j(\mathbf{X}) &= mZ(\mathbf{X}) - \sum_{j=1}^m \sup_{f \in \mathcal{F}} \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m} (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{EQ. 68} \\
&\leq mZ(\mathbf{X}) - \sup_{f \in \mathcal{F}} \sum_{j=1}^m \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m} (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{SUBADDITIVITY} \\
&= mZ(\mathbf{X}) - \sup_{f \in \mathcal{F}} m \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m} \sum_{j=1}^m (f(\mathbf{X}_j) - \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f])^2 && \text{LINEARITY} \\
&= mZ(\mathbf{X}) - \sup_{f \in \mathcal{F}} m \left(\sum_{i=1}^m (f(\mathbf{X}_i) - \hat{\mathbb{E}}_{\mathbf{X}}[f])^2 \right) - \frac{m-1}{m} \sum_{j=1}^m \left(\frac{m}{m-1} f(\mathbf{X}_j) - \frac{m}{m-1} \hat{\mathbb{E}}_{\mathbf{X}}[f] \right)^2 \quad \hat{\mathbb{E}}_{\mathbf{X}_{\setminus j}}[f] = \frac{1}{m-1} (m \hat{\mathbb{E}}_{\mathbf{X}}[f] - f(\mathbf{X}_j)) \\
&= \frac{m(m-1)}{m-1} Z(\mathbf{X}) - \left(\frac{m(m-1)}{m-1} - \frac{m}{m-1} \right) Z(\mathbf{X}) = \frac{m}{m-1} Z(\mathbf{X}) && \text{DEFINITION OF } Z(\cdot)
\end{aligned}$$

It is now a straightforward matter to observe that $\frac{m}{m-1}Z(\cdot)$, also obeys $(\frac{m}{m-1}, 0)$ -self-boundedness, as

$$\sum_{j=1}^m Z(\mathbf{X}) - Z_j(\mathbf{X}) \leq \frac{m}{m-1} Z(\mathbf{X}) \Rightarrow \frac{m}{m-1} Z(\mathbf{X}) - \frac{m}{m-1} Z_j(\mathbf{X}) \leq \frac{m}{m-1} \cdot \left(\frac{m}{m-1} Z(\mathbf{X}) \right) .$$

With all three desiderata shown, we may now conclude that $\frac{m}{m-1}Z(\cdot)$ is a 0-overestimating, 1-underestimating $(\frac{m}{m-1}, 0)$ -self-bounding function. Now, observe that

$$\begin{aligned}
\mathbb{P}_{\mathbf{x}} \left(W \geq \frac{m}{m-1} \widehat{W} + \varepsilon \right) &= \mathbb{P}_{\mathbf{x}} \left(\frac{m}{m-1} \widehat{W} \leq W - \varepsilon \right) && \text{EQUIVALENT EVENTS} \\
&\leq \mathbb{P}_{\mathbf{x}} \left(\frac{m}{m-1} \widehat{W} \leq \frac{m}{m-1} \mathbb{E}[\widehat{W}] - \varepsilon \right) && \text{LEMMA D.3.1} \\
&= \mathbb{P}_{\mathbf{x}} \left(\frac{1}{m-1} Z(\mathbf{x}) \leq \frac{1}{m-1} \mathbb{E}[Z] - \varepsilon \right) && \widehat{W} = \frac{1}{m} Z(\mathbf{x}) \\
&= \mathbb{P}_{\mathbf{x}} \left(\frac{1}{m} \cdot \frac{m}{m-1} Z(\mathbf{x}) \leq \frac{1}{m} \cdot \frac{m}{m-1} \mathbb{E}[Z] - \varepsilon \right) . && \text{EQUIVALENT EVENTS} \\
&\leq \exp \left(\frac{-m\varepsilon^2}{2(\frac{m}{m-1} \hat{v} + \varepsilon)} \right) = \exp \left(\frac{-m\varepsilon^2}{2(\frac{m}{m-1} \hat{v} + \varepsilon)} \right) . && \text{SBF INEQUALITY}
\end{aligned}$$

The step marked SBF INEQUALITY follows via standard self-bounding function inequalities; see Boucheron et al. [2009]. An application of the quadratic formula (positive solution) then yields the result, as

$$\varepsilon = \frac{\ln(\frac{1}{\delta})}{m} + \sqrt{\left(\frac{\ln(\frac{1}{\delta})}{m} \right)^2 + \frac{2\hat{v} \ln(\frac{1}{\delta})}{m-1}} \leq \frac{2 \ln(\frac{1}{\delta})}{m} + \sqrt{\frac{2\hat{v} \ln(\frac{1}{\delta})}{m-1}} \Rightarrow \mathbb{P} \left(v \geq \frac{m}{m-1} \hat{v} + \varepsilon \right) \leq \delta .$$

□

With these results, we now show a variance-sensitive supremum-deviation bound in terms of the n -MCERA.

Lemma D.3.3 (Empirical-Variance Sensitive Uniform Convergence Tail Bounds). Suppose as in theorem 4.2.2. It then holds with probability at least $1 - \delta$ over choice of $\mathbf{x}, \boldsymbol{\sigma}$ that

$$\sup_{f \in \mathcal{F}} \left| \frac{\mathbb{E}[f]}{D} - \hat{\mathbb{E}}_{\mathbf{x}}[f] \right| \leq 2\tilde{\mathbf{R}} + \varepsilon_{\text{SD}} .$$

Proof. This result follows from a chain of 4 union bounds, each of probability $\frac{\delta}{4}$, where first v is bounded in terms of \hat{v} with lemma D.3.2, then $\hat{\mathfrak{K}}_m(\mathcal{F}, \mathbf{x})$ is bounded in terms of $\hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma})$ with McDiarmid's bounded difference inequality [McDiarmid, 1989], $\mathfrak{K}_m(\mathcal{F}, \mathcal{D})$ is bounded in terms of $\hat{\mathfrak{K}}_m(\mathcal{F}, \mathbf{x})$ with the self-bounding function inequality [Boucheron et al., 2000, Thm. 1] (see Oneto et al. [2013] for details), and finally the supremum deviation is bounded with the symmetrization inequality equation 4 [Boucheron et al., 2013, see Lemma 11.4] and the refinement of Talagrand's inequality for empirical processes due to Bousquet [2002, Thm. 2.3]. The aforementioned results implicitly assume $r = 1$, so we first consider that case, and then note that by linear scaling, we attain the general result. In particular, while \mathcal{F} , the SD, and each Rademacher average and tail bound thereof scale linearly, the *variance-like* quantities ε_v , v , and \hat{v} all scale *quadratically* with r^2 , which has a linear effect, since these always appear inside a root in the final bound. Note also that while these bounds do place some restrictions on the failure probability δ , our 4-term union bound holds vacuously when $\delta \geq \frac{1}{4}$, and said restrictions do not apply to $\delta < \frac{1}{4}$.

All but the Monte-Carlo bounds are discussed elsewhere, so we give further detail here. Note that this result generalizes a theorem of Bartlett and Mendelson [2002, thm. 11], which holds for the special case of $n = 1$. The probability space in play is the $n \times m$ Rademacher random variables of $\boldsymbol{\sigma}$, changing any one of which produces a change of no more than $\frac{2r}{nm}$ to the n -MCERA, which yields via McDiarmid's inequality

$$\mathbb{P}_{\boldsymbol{\sigma}} \left(\hat{\mathfrak{K}}_m(\mathcal{F}, \mathbf{x}) \geq \hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + 2r \sqrt{\frac{\ln(\frac{4}{\delta})}{2nm}} \right) \leq \frac{\delta}{4} .$$

This is sufficient to show the stated result. A more careful analysis can improve this bound by noting that, due to the presence of the absolute value, $\hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) = \hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, -\boldsymbol{\sigma})$, and there is essentially one fewer degree of freedom as a result.

We note also that the stated bound in each of the above references is in most cases sharper than what we present; this is because we are interested in the explicit form of ε as a function of δ , and thus use the convenient sub-gamma characterizations, even when sharper sub-Poisson bounds are available. With numeric inversion, these bounds can be used with fixed δ , and are appropriate when sharper bounds are required, but they come at the cost of computational overhead, algebraic inconvenience, and lack of interpretability. Similarly, the uniform union bound of probability $\frac{\delta}{4}$ is also suboptimal, but again is algebraically convenient, and the potential improvement from nonuniform division of δ is quite small. \square

With this result, we are now ready to show theorem 4.2.2.

Theorem 4.2.2 (Variance-Sensitive Bounds). Suppose *distribution* \mathcal{D} over \mathcal{X} , *training sample* $\mathbf{x} \sim \mathcal{D}^m$, *Monte-Carlo trial count* n , i.i.d. Rademacher matrix $\boldsymbol{\sigma} \in (\pm 1)^{n \times m}$, *codec family* \mathcal{C} , *loss family* $\mathcal{L} \subseteq \mathbb{R}_{0+}^a \rightarrow [-r, r]$, λ -Lipschitz *welfare* $W(\cdot)$, and *codec selection size* k . $\forall \delta \in (0, 1)$, let $\mathcal{F} \doteq \min \circ \mathcal{L} \circ (\overset{c}{k})$, take \hat{v} to be the *empirical wimpy variance* of \mathcal{F} over \mathbf{x} and \hat{v}^{raw} to be the *empirical raw wimpy variance*, and take

$$\begin{aligned} 1. \quad \varepsilon_{\text{ERA}} &\doteq \frac{4r \frac{4}{\delta}}{3nm} + \sqrt{\frac{4\hat{v}^{\text{raw}} \ln \frac{4}{\delta}}{nm}}; & 3. \quad \tilde{\mathfrak{K}} &\doteq \hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}} + \varepsilon_{\text{RA}}; \\ 2. \quad \varepsilon_{\text{RA}} &\doteq \frac{2r \ln \frac{4}{\delta}}{m} + \sqrt{\frac{2r(\hat{\mathfrak{K}}_m^n(\mathcal{F}, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}) \ln \frac{4}{\delta}}{m}}; & 4. \quad \tilde{v} &\doteq \frac{m}{m-1} \hat{v} + \frac{4r^2 \ln \frac{4}{\delta}}{m} + \sqrt{\frac{2r^2 \hat{v} \ln \frac{4}{\delta}}{m-1}}; \& \\ 5. \quad \varepsilon_{\text{SD}} &\doteq \frac{r \ln \frac{4}{\delta}}{3m} + \sqrt{\frac{2(\tilde{v} + 4r\tilde{\mathfrak{K}}) \ln \frac{4}{\delta}}{m}}. & & \end{aligned}$$

It then holds with with pr. $\geq 1 - \delta$ over $\mathbf{x}, \boldsymbol{\sigma}$ that

$$\begin{aligned} 1. \quad & \left| W(\hat{c}_W, \mathcal{L}, \mathcal{D}) - \hat{W}(\hat{c}_W, \mathcal{L}, \mathbf{x}) \right| && \leq 2\lambda\tilde{\mathfrak{K}} + \lambda\varepsilon_{\text{SD}} \quad \& \\ 2. \quad & W(\hat{c}_W, \mathcal{L}, \mathcal{D}) - W(c_W^*, \mathcal{L}, \mathcal{D}) && \leq 4\lambda\tilde{\mathfrak{K}} + 2\lambda\varepsilon_{\text{SD}} . \end{aligned}$$

Proof. This result follows via lemma D.3.3, followed by several straightforward deterministic manipulations. This result is also a probabilistic guarantee, and it suffices to show that the condition on welfare is satisfied under the assumptions and probabilistic guarantee of the aforementioned lemma. Henceforth we assume the conclusion of lemma D.3.3 holds, and thus the following hold only with probability $\geq 1 - \delta$.

By assumption, $W(\mathbf{c}, \mathcal{L}, \mathbf{x})$ is λ - ℓ_∞ -Lipschitz in the empirical risks (w.r.t. \mathcal{L}) of \mathbf{c} on \mathbf{x} . Consequently, a uniform bound of $2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}}$ on *empirical risks* (lemma D.3.3) implies a $\lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}})$ bound on *welfare*. It therefore holds that

$$|W(\hat{\mathbf{c}}_W, \mathcal{L}, \mathcal{D}) - \hat{W}(\hat{\mathbf{c}}_W, \mathcal{L}, \mathbf{x})| \leq \lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}}) ,$$

which yields item (1). Similarly, we have

$$|\hat{W}(\mathbf{c}_W^*, \mathcal{L}, \mathbf{x}) - W(\mathbf{c}_W^*, \mathcal{L}, \mathcal{D})| \leq \lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}}) ,$$

and consequently, by the triangle inequality (noting that $\hat{W}(\hat{\mathbf{c}}_W, \mathcal{L}, \mathbf{x}) \leq \hat{W}(\mathbf{c}_W^*, \mathcal{L}, \mathbf{x})$ by definition), we have

$$|W(\hat{\mathbf{c}}_W, \mathcal{L}, \mathcal{D}) - W(\mathbf{c}_W^*, \mathcal{L}, \mathcal{D})| \leq 2\lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}}) = 4\lambda\tilde{\mathfrak{R}} + 2\lambda\varepsilon_{\text{SD}} ,$$

which completes the result. Note that here, the largest value the LHS can attain occurs when the empirical welfare of \mathbf{c}_W^* is *overestimated* by $\lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}})$, while $\hat{\mathbf{c}}_W$ is *underestimated* by $\lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}})$, making $\hat{\mathbf{c}}_W$ *empirically optimal* over \mathbf{x} , and no worse than $2\lambda(2\tilde{\mathfrak{R}} + \varepsilon_{\text{SD}})$ -suboptimal over \mathcal{D} . \square

D.3.2 Proof of Theorem 4.2.3

Theorem 4.2.3 combines several well-understood properties of variance with a bound on the Rademacher averages of minima of function families. We first show lemmas regarding variances of linear families and Rademacher averages of minima. Each lemma in this subsection assumes

1. *dual norms*: p, q , s.t. $|p^{-1}| + |q^{-1}| = 1$;
2. *Codec family*: \mathcal{C} such that $\sup_{c \in \mathcal{C}} \|c(x)\|_q \leq 1$; and
3. *Linear loss family* $\mathcal{L}_p = \{ \ell(\mathbf{a}) \doteq \mathbf{w} \cdot \mathbf{a} \mid \mathbf{w} \in \mathbb{R}_{0+}^a \text{ s.t. } \|\mathbf{w}\|_p \leq 1 \}$ (as defined in eq. 34).

We first show bounds on the variances of singleton codec families.

Lemma D.3.4 (Variance of Linear Families).

$$\sup_{c \in \mathcal{C}} \sup_{\ell \in \mathcal{L}} \hat{\mathbb{V}}_{\mathbf{x}}[\ell \circ c] = \sup_{c \in \mathcal{C}} \left\| \hat{\mathbb{C}}_{\mathbf{x}}[c] \right\|_{q \rightarrow p}$$

Proof. First note that this result essentially holds componentwise over each $c \in \mathcal{C}$. For a single c , in the special case of $p = q = 2$, it is the standard Rayleigh decomposition argument used in *singular value decomposition* (SVD) and *principal component analysis* (PCA), and for arbitrary $p \neq q \neq 2$ dual norms, the spectral norm is generalized to an operator norm. We introduce the unit-sphere notation

$$\mathcal{S}_p^{a-1} \doteq \{ \mathbf{w} \in \mathbb{R}^a \mid \|\mathbf{w}\|_p = 1 \} ,$$

and now show the result proper:

$$\begin{aligned} \sup_{c \in \mathcal{C}} \sup_{\ell \in \mathcal{L}} \hat{\mathbb{V}}_{\mathbf{x}}[\ell \circ c] &= \sup_{c \in \mathcal{C}} \sup_{\mathbf{w} \in \mathcal{S}_p^{a-1}} \hat{\mathbb{V}}_{\mathbf{x}}[\mathbf{w} \cdot c(\cdot)] && \text{DEFINITION OF } \mathcal{L} \\ &= \sup_{c \in \mathcal{C}} \sup_{\mathbf{w} \in \mathcal{S}_p^{a-1}} \mathbf{w} \hat{\mathbb{C}}_{\mathbf{x}}[c] \mathbf{w}^\top && \text{VARIANCE PROPERTIES} \\ &= \sup_{c \in \mathcal{C}} \left\| \hat{\mathbb{C}}_{\mathbf{x}}[c] \right\|_{p \rightarrow q} . && \left\| \hat{\mathbb{C}}_{\mathbf{x}}[c] \mathbf{w}^\top \right\|_q \leq \|\mathbf{w}\|_p \left\| \hat{\mathbb{C}}_{\mathbf{x}}[c] \right\|_{p \rightarrow q} \\ &&& \text{HÖLDER'S INEQUALITY} \end{aligned}$$

Note that Hölder's inequality, which generalizes the Cauchy-Schwarz inequality to $\mathbf{a} \cdot \mathbf{b} \leq \|\mathbf{a}\|_p \|\mathbf{b}\|_q$, holds with equality for appropriate choice of \mathbf{w} , which the supremum selects, yielding (strict) equality. \square

We now show compositional bounds on the Rademacher averages of codec selection.

Lemma D.3.5 (Bounding Rademacher Averages of Minima). Suppose $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k \subseteq \mathcal{X} \rightarrow \mathbb{R}$, and let

$$\mathcal{F}_{\min} \doteq \left\{ \min_{i \in \{1, \dots, k\}} \mathbf{f}_i(x) \mid \mathbf{f} \in \mathcal{F}_1 \times \dots \times \mathcal{F}_k \right\} .$$

Then $\hat{\mathfrak{R}}_m(\mathcal{F}_{\min}, \mathbf{x}) \leq \sum_{i=1}^m \hat{\mathfrak{R}}_m(\mathcal{F}_i, \mathbf{x})$.

Proof. We first show the result for $k = 2$.

$$\begin{aligned} \hat{\mathfrak{R}}_m(\mathcal{F}_{\min}, \mathbf{x}) &= \mathbb{E}_{\sigma} \left[\sup_{f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \min(f_1(\mathbf{x}_i), f_2(\mathbf{x}_i)) \right| \right] && \text{DEFINITION} \\ &= \mathbb{E}_{\sigma} \left[\sup_{f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \left(\frac{f_1 + f_2 - |f_1 - f_2|}{2} \right) (\mathbf{x}_i) \right| \right] && \text{MINIMUM IDENTITY} \\ &\leq \mathbb{E}_{\sigma} \left[\sup_{\substack{f_1 \in \mathcal{F}_1 \\ f_2 \in \mathcal{F}_2}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \left(\frac{f_1 + f_2}{2} \right) (\mathbf{x}_i) \right| \right] + \mathbb{E}_{\sigma} \left[\sup_{\substack{f_1 \in \mathcal{F}_1 \\ f_2 \in \mathcal{F}_2}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \left(\frac{|f_1 - f_2|}{2} \right) (\mathbf{x}_i) \right| \right] && \text{TRIANGLE INEQUALITY} \\ &= \hat{\mathfrak{R}}_m\left(\frac{f_1 + f_2}{2}, \mathbf{x}\right) + \hat{\mathfrak{R}}_m\left(\frac{|\mathcal{F}_1 \ominus \mathcal{F}_2|}{2}, \mathbf{x}\right) && \text{DEFINITION} \\ &= \hat{\mathfrak{R}}_m\left(\frac{f_1 + f_2}{2}, \mathbf{x}\right) + \hat{\mathfrak{R}}_m(|\cdot| \circ \frac{\mathcal{F}_1 \ominus \mathcal{F}_2}{2}, \mathbf{x}) && \text{ALGEBRA} \\ &\leq \hat{\mathfrak{R}}_m\left(\frac{f_1 + f_2}{2}, \mathbf{x}\right) + \hat{\mathfrak{R}}_m\left(\frac{\mathcal{F}_1 \ominus \mathcal{F}_2}{2}, \mathbf{x}\right) && \text{TALAGRAN'D'S CONTRACTION INEQUALITY} \\ &\leq \hat{\mathfrak{R}}_m(\mathcal{F}_1, \mathbf{x}) + \hat{\mathfrak{R}}_m(\mathcal{F}_2, \mathbf{x}) && \text{DIRECT SUMMATION} \end{aligned}$$

The result now follows by induction. Note that the contraction inequality avoids the usual 2-factor gap, as the $|\cdot|$ function is *odd*. \square

With these ingredients, we now restate and show theorem 4.2.3.

Theorem 4.2.3 (Bounds for Linear Loss Families). Suppose p, q , s.t. $|p^{-1}| + |q^{-1}| = 1$, *codec family* \mathcal{C} , \mathcal{L}_p as in (34), and assume $\sup_{c \in \mathcal{C}, \mathbf{x} \in \mathcal{X}} \|c(\mathbf{x})\|_q \leq s$. For all $k \in \{1, \dots, |\mathcal{C}|\}$, take $\mathcal{F}_k \doteq \min \circ \mathcal{L}_p \circ \binom{\mathcal{C}}{k}$, and \hat{v}_k to be the *empirical wimpy variance* of \mathcal{F}_k on *sample* \mathbf{x} . Taking $\hat{\mathbb{C}}_{\mathbf{x}}[c] \in \mathbb{R}^{a \times a}$ to denote the *empirical covariance matrix* of the attributes of codec c over \mathbf{x} , and $\|\cdot\|_{p \rightarrow q}$ to be the ℓ_p - ℓ_q *operator norm*, we have

$$\hat{v}_1 = \sup_{c \in \mathcal{C}} \sup_{\ell \in \mathcal{L}} \hat{\mathbb{V}}[\ell \circ c] \leq \sup_{c \in \mathcal{C}} \|\hat{\mathbb{C}}[c]\|_{q \rightarrow p} , \quad \& \quad \hat{\mathfrak{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \sigma) \leq \frac{1}{n} \sum_{j=1}^n \sup_{c \in \mathcal{C}} \left\| \frac{1}{m} \sum_{i=1}^m \sigma_{j,i} c(\mathbf{x}_i) \right\|_q .$$

Furthermore, for any $a, b \in \mathbb{N}$ s.t. $k = a + b$, we have

$$\hat{v}_k \leq \hat{v}_a + \hat{v}_b \leq k \hat{v}_1 , \quad \& \quad \hat{\mathfrak{R}}_m(\mathcal{F}_k, \mathbf{x}) \leq \hat{\mathfrak{R}}_m(\mathcal{F}_a, \mathbf{x}) + \hat{\mathfrak{R}}_m(\mathcal{F}_b, \mathbf{x}) \leq k \hat{\mathfrak{R}}_m(\mathcal{F}_1, \mathbf{x}) .$$

Proof. We first consider the bounds on \hat{v}_1 and $\hat{\mathfrak{R}}_m^1(\mathcal{F}_1, \mathbf{x}, \sigma)$. Both are inequalities, as they are shown by defining a linear loss family without nonnegativity constraints $\mathcal{L}_{p \pm}$, and bounding the corresponding quantities for $\mathcal{L}_{p \pm}$. In particular, take

$$\mathcal{L}_{p \pm} \doteq \{ \ell(\mathbf{a}) \doteq \mathbf{w} \cdot \mathbf{a} \mid \mathbf{w} \in \mathbb{R}^a \text{ s.t. } \|\mathbf{w}\|_p \leq 1 \} ,$$

which differs from \mathcal{L}_p in that $\mathbf{w} \in \mathbb{R}_{0+}^a$ is relaxed to $\mathbf{w} \in \mathbb{R}^a$, thus $\mathcal{L}_p \subseteq \mathcal{L}_{p\pm}$. Now take $\mathcal{F}_{1\pm} \doteq \mathcal{L}_p \circ \mathcal{C}$, and note that $\mathcal{F}_1 \subseteq \mathcal{F}_{1\pm}$. It thus holds that the empirical wimpy variance and 1-ERA of $\mathcal{F}_{1\pm}$ are at least as great as the corresponding quantities for \mathcal{F}_1 .

Consequently, from the definition of operator norm, we have

$$\hat{v}_1 = \sup_{c \in \mathcal{C}} \sup_{\ell \in \mathcal{L}_p} \sup_{\mathbf{x}} \hat{\mathbb{V}}[\ell \circ c] \leq \sup_{c \in \mathcal{C}} \sup_{\ell \in \mathcal{L}_{p\pm}} \sup_{\mathbf{x}} \hat{\mathbb{V}}[\ell \circ c] = \sup_{c \in \mathcal{C}} \left\| \hat{\mathbb{C}}[c] \right\|_{q \rightarrow p} .$$

Similarly, via dual-norm properties, we have

$$\hat{\mathbf{R}}_m^1(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) \leq \hat{\mathbf{R}}_m^1(\mathcal{F}_{1\pm}, \mathbf{x}, \boldsymbol{\sigma}) = \sup_{c \in \mathcal{C}} \left\| \frac{1}{m} \sum_{i=1}^m \boldsymbol{\sigma}_i c(\mathbf{x}_i) \right\|_q .$$

Now, note that the property

$$\hat{v}_{a+b} \leq \hat{v}_a + \hat{v}_b \leq (a+b)\hat{v}_1$$

then follows by subadditivity of the variance of minima, and finally

$$\hat{\mathbf{R}}_m(\mathcal{F}_{a+b}, \mathbf{x}) \leq \hat{\mathbf{R}}_m(\mathcal{F}_a, \mathbf{x}) + \hat{\mathbf{R}}_m(\mathcal{F}_b, \mathbf{x}) \leq (a+b)\hat{\mathbf{R}}_m(\mathcal{F}_1, \mathbf{x})$$

follows from lemma D.3.5. □

Corollary D.3.6 (Bounds for Linear Loss Families). Suppose as in theorem 4.2.3, and also $p = q = 2$, and all attribute values are contained by the unit sphere. Then

$$\hat{v}_k \leq k \sup_{c \in \mathcal{C}} \left\| \hat{\mathbb{C}}[c] \right\|_{2 \rightarrow 2} ,$$

where $\|\cdot\|_{2 \rightarrow 2}$ is the *spectral norm* (largest eigenvalue). Additionally, with $\text{pr.} \geq 1 - \delta$ over $\mathbf{x}, \boldsymbol{\sigma}$, we have

$$\begin{aligned} \hat{\mathbf{R}}_m(\mathcal{F}_k, \mathbf{x}) &\leq k \left(\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \sqrt{\frac{\ln \frac{1}{\delta}}{2nm}} \right) \\ &\leq k \left(\frac{1}{n} \sum_{i=1}^n \sup_{c \in \mathcal{C}} \left\| \frac{1}{m} \sum_{j=1}^m \boldsymbol{\sigma}_{i,j} c(\mathbf{x}_j) \right\|_2 + \sqrt{\frac{\ln \frac{1}{\delta}}{2nm}} \right) . \end{aligned}$$

Consequently, theorem 4.2.2 holds with

$$\begin{aligned} \varepsilon_{\text{RA}} &\leq \frac{\ln \frac{4}{\delta}}{3m} + \sqrt{\frac{2k(\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}) \ln \frac{4}{\delta}}{m}}; \\ \tilde{\mathbf{R}} &\leq k(\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}) + \varepsilon_{\text{RA}}; \ \& \\ \tilde{v} &\leq \frac{m}{m-1} k \hat{v}_1 + \frac{2 \ln \frac{4}{\delta}}{m} + \sqrt{\frac{2k \hat{v}_1 \ln \frac{4}{\delta}}{m-1}} . \end{aligned}$$

Proof. This result follows directly from theorem 4.2.3, noting that $\|\cdot\|_{2 \rightarrow 2}$ is the *spectral norm*, and applying the McDiarmid argument of lemma D.3.3 to (probabilistically) upper-bound $\hat{\mathbf{R}}_m(\mathcal{F}_1, \mathbf{x})$ as $\hat{\mathbf{R}}_m^n(\mathcal{F}_1, \mathbf{x}, \boldsymbol{\sigma}) + \varepsilon_{\text{ERA}}$, and then noting that $\hat{\mathbf{R}}_m(\mathcal{F}_k, \mathbf{x}) \leq k \hat{\mathbf{R}}_m(\mathcal{F}_1, \mathbf{x})$. □

D.3.3 Justifications of Equations

We now justify or further elucidate several numbered equations found in the paper body.

Asymptotic Uniform Convergence Bounds We first consider equations 33 and 35. These follow from theorem 4.2.2 and corollary D.3.6, respectively. Equation 33 assumes $n \in \mathbf{\Omega}(\frac{r^2}{v})$, which is required to ensure that the contributions of the term ε_{ERA} are dominated by ε_{SD} , as

$$\varepsilon_{\text{ERA}} = \sqrt{\frac{r^2}{nm}} \in \mathbf{O}\sqrt{\frac{v}{m}} .$$

Note also that the ERA itself dominates the ε_{ERA} terms, as by the Khintchine inequality [see Haagerup, 1982], we have

$$\hat{\mathfrak{R}}_m(\mathcal{F}, \mathbf{x}) \geq \sqrt{\frac{\hat{v}}{m}} \xrightarrow{m \rightarrow \infty} \sqrt{\frac{v}{m}} ,$$

where the $\hat{v} \rightsquigarrow v$ is asymptotic in m for fixed r . Essentially the same arguments hold for equation 35, now assuming $n \in \mathbf{\Omega}(\frac{1}{v_1})$, as we assume (corollary D.3.6) $r^2 = r = 1$, and the analysis now operates on \mathcal{F}_1 , thus v_1 is the variance-concept of interest.

Pareto Optimality Equations We now consider equations 37 and 38. Note that we define the Pareto-optimal fronts as the $\text{LCH}(\cdot)$ of the *Pareto-optimal codecs / codec sets*, but the expressions given merely take the $\text{LCH}(\cdot)$ of *all codecs / codec sets*. This is because any Pareto-suboptimal codecs or codec sets included in the $\text{LCH}(\cdot)$ by definition have no impact on the $\text{LCH}(\cdot)$, thus the two formulations are identical.

Appendix E

Supplementary Material for Chapter 5

E.1 Efficient FPAC-Learning

We now show theorem 5.3.1.

Theorem 5.3.1 (Efficient FPAC Learning via Convex Optimization). Suppose each hypothesis space $\mathcal{H}_d \in \mathcal{H}$ is indexed by $\Theta_d \subseteq \mathbb{R}^{\text{Poly}(d)}$, i.e., $\mathcal{H}_d = \{h(\cdot; \theta) \mid \theta \in \Theta_d\}$, s.t. (Euclidean) $\text{Diam}(\Theta_d) \in \text{Poly}(d)$, and $\forall x \in \mathcal{X}, \theta \in \Theta_d$, $h(x; \theta)$ can be evaluated in $\text{Poly}(d)$ time, and $\tilde{\theta} \in \mathbb{R}^{\text{Poly}(d)}$ can be Euclidean-projected onto Θ_d in $\text{Poly}(d)$ time. Suppose also ℓ such that $\forall x \in \mathcal{X}, y \in \mathcal{Y} : \theta \mapsto \ell(y, h(x; \theta))$ is a *convex function*, and suppose Lipschitz constants $\lambda_\ell, \lambda_{\mathcal{H}} \in \text{Poly}(d)$ and some norm $\|\cdot\|_{\mathcal{Y}}$ over \mathcal{Y} s.t. ℓ is $\lambda_\ell \|\cdot\|_{\mathcal{Y}}$ -Lipschitz in \hat{y} , i.e.,

$$\forall y, \hat{y}, \hat{y}' \in \mathcal{Y} : |\ell(y, \hat{y}) - \ell(y, \hat{y}')| \leq \lambda_\ell \|\hat{y} - \hat{y}'\|_{\mathcal{Y}} \quad ,$$

and also that each \mathcal{H}_d is $\lambda_{\mathcal{H}} \|\cdot\|_2 \|\cdot\|_{\mathcal{Y}}$ -Lipschitz in θ , i.e.,

$$\forall x \in \mathcal{X}, \theta, \theta' \in \Theta_d : \|h(x; \theta) - h(x; \theta')\|_{\mathcal{Y}} \leq \lambda_{\mathcal{H}} \|\theta - \theta'\|_2 \quad .$$

Finally, assume $\ell \circ \mathcal{H}_d$ exhibits ε - δ *uniform convergence* with sample complexity $m_{\text{UC}}(\varepsilon, \delta, d) \in \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d)$.

It then holds that, for arbitrary initial guess $\theta_0 \in \Theta_d$, for any group distributions $\mathcal{D}_{1:g}$, group weights \mathbf{w} , and fair malfare function $\mathbb{M}(\cdot; \cdot)$, the algorithm (see algorithm 2)

$$\mathcal{A}(\mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}(\cdot; \cdot), \varepsilon, \delta, d) \doteq \mathcal{A}_{\text{PSG}}(\ell, \mathcal{H}_d, \theta_0, m_{\text{UC}}(\cdot, \cdot, d), \mathcal{D}_{1:g}, \mathbf{w}, \mathbb{M}(\cdot; \cdot), \varepsilon, \delta)$$

fair-PAC-learns (\mathcal{H}, ℓ) with sample complexity $m(\varepsilon, \delta, d, g) = g \cdot m_{\text{UC}}(\frac{\varepsilon}{3}, \frac{\delta}{g}, d)$, and (training) time-complexity $\in \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$, thus $(\mathcal{H}, \ell) \in \text{FPAC}_{\text{Poly}}^{\text{Agn}}$.

Proof. We now show that this subgradient-method construction of \mathcal{A} requires $\text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$ time to identify an ε - δ - $\mathbb{M}_p(\cdot; \cdot)$ -optimal $\tilde{\theta} \in \Theta_d$, and thus fair-PAC-learns (\mathcal{H}, ℓ) . This essentially boils down to showing that (1) the empirical malfare objective is *convex* and *Lipschitz continuous*, and (2) that algorithm 2 runs sufficiently many subgradient-update steps, with appropriate step size, on a sufficiently large training set, to yield the appropriate guarantees, and that each step of the subgradient method, of which there are polynomially many, itself requires polynomial time.

First, note that by theorem 3.3.8 items 3 and 5, we may assume that $\mathbb{M}(\cdot; \cdot)$ can be expressed as a p -power mean with $p \geq 1$; thus henceforth we refer to it as $\mathbb{M}_p(\cdot; \cdot)$. Now, recall that the empirical malfare objective (given $\theta \in \Theta_d$ and training sets $\mathbf{z}_{1:g}$) is defined as

$$\mathbb{M}_p(i \mapsto \hat{\mathbf{R}}(h(\cdot; \theta); \ell, \mathbf{z}_i); \mathbf{w}) \quad .$$

We first show that empirical malfare is convex in Θ_d . By assumption and positive linear closure, $\hat{R}(h(\cdot; \theta'); \ell, \mathbf{z}_i)$ is convex in $\theta \in \Theta_d$. The objective of interest is the composition of $\mathbb{M}_p(\cdot; \mathbf{w})$ with this quantity evaluated on each of g training sets. By theorem 3.3.7 item 4, $\mathbb{M}_p(\cdot; \mathbf{w})$ is convex $\forall p \in [1, \infty]$ in \mathbb{R}_{0+}^g , and by the monotonicity axiom, it is monotonically increasing. Composition of a monotonically increasing convex function on \mathbb{R}_{0+}^g with convex functions on Θ_d yields a convex function, thus we conclude the empirical malfare objective is convex in Θ_d .

We now show that empirical malfare is Lipschitz-continuous. Now, note that for any $p \geq 1$, \mathbf{w} ,

$$\forall \mathcal{S}, \mathcal{S}' : |\mathbb{M}_p(\mathcal{S}; \mathbf{w}) - \mathbb{M}_p(\mathcal{S}'; \mathbf{w})| \leq 1 \|\mathcal{S} - \mathcal{S}'\|_\infty ,$$

i.e., $\mathbb{M}_p(\cdot; \mathbf{w})$ is $1 \cdot \|\cdot\|_\infty$ -Lipschitz in *empirical risks* (see theorem 3.3.7 item 3), and thus by Lipschitz composition, we have Lipschitz property

$$\forall \theta, \theta' \in \Theta_d : \left| \mathbb{M}_p(i \mapsto \hat{R}(h(\cdot; \theta); \ell, \mathbf{z}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto \hat{R}(h(\cdot; \theta'); \ell, \mathbf{z}_i); \mathbf{w}) \right| \leq \lambda_\ell \lambda_{\mathcal{H}} \|\theta - \theta'\|_2 .$$

We now show that algorithm 2 FPAC-learns (\mathcal{H}, ℓ) . As above, take $m \doteq m_{UC}(\frac{\varepsilon}{3}, \frac{\delta}{g}, d)$. Our algorithm shall operate on a training sample $\mathbf{z}_{1:g, 1:m} \sim \mathcal{D}_1^m \times \dots \times \mathcal{D}_g^m$.

First note that evaluating a subgradient (via forward finite-difference estimation or automated subdifferentiation) requires $(\dim(\Theta_d) + 1)m$ evaluations of $h(\cdot; \cdot)$, which by assumption is possible in $\text{Poly}(d, m) = \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$ time.

The subgradient method produces $\tilde{\theta}$ approximating the empirically-optimal $\hat{\theta}$ such that [see Shor, 2012]

$$f(\tilde{\theta}) \leq f(\hat{\theta}) + \frac{\|\theta_0 - \hat{\theta}\|_2^2 + \Lambda^2 \alpha^2 n}{2\alpha n} \leq \frac{\text{Diam}^2(\Theta_d) + \Lambda^2 \alpha^2 n}{2\alpha n} ,$$

for $\Lambda \|\cdot\|_2$ -Lipschitz objective f , thus taking $\alpha \doteq \frac{\text{Diam}(\Theta_d)}{\Lambda \sqrt{n}}$ yields

$$f(\tilde{\theta}) - f(\hat{\theta}) \leq \frac{\text{Diam}(\Theta_d) \Lambda}{\sqrt{n}} .$$

As shown above, $\Lambda = \lambda_\ell \lambda_{\mathcal{H}}$, thus we may guarantee *optimization error*

$$\varepsilon_{\text{opt}} \doteq f(\tilde{\theta}) - f(\theta^*) \leq \frac{\varepsilon}{3}$$

if we take iteration count

$$n \geq \frac{9 \text{Diam}^2(\Theta_d) \lambda_\ell^2 \lambda_{\mathcal{H}}^2}{\varepsilon^2} = \left(\frac{3 \text{Diam}(\Theta_d) \lambda_\ell \lambda_{\mathcal{H}}}{\varepsilon} \right)^2 \in \text{Poly}(\frac{1}{\varepsilon}, d) .$$

As each iteration requires $m \cdot \text{Poly}(d) \subseteq \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$ time, the subgradient method identifies an $\frac{\varepsilon}{3}$ -empirical-malfare-optimal $\tilde{\theta} \in \Theta_d$ in $\text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$ time.

As m was selected to ensure $\frac{\varepsilon}{3} \cdot \frac{\delta}{g}$ uniform convergence, we thus have that by uniform convergence, and union bound (over g groups), with probability at least $1 - \delta$ over choice of $\mathbf{z}_{1:g}$, we have

$$\forall i \in \{1, \dots, g\}, \theta \in \Theta_d : \left| \mathbb{M}_p(i \mapsto \hat{R}(h(\cdot; \theta); \ell, \mathbf{z}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto R(h(\cdot; \theta); \ell, \mathcal{D}_i); \mathbf{w}) \right| \leq \frac{\varepsilon}{3} .$$

Combining *estimation* and *optimization* errors, we get that with probability at least $1 - \delta$, the approximate-EMM-optimal $h(\cdot; \tilde{\theta})$ obeys

$$\begin{aligned} \mathbb{M}_p(i \mapsto R(h(\cdot; \tilde{\theta}); \ell, \mathcal{D}_i); \mathbf{w}) &\leq \mathbb{M}_p(i \mapsto \hat{R}(h(\cdot; \tilde{\theta}); \ell, \mathbf{z}_i); \mathbf{w}) + \frac{\varepsilon}{3} \\ &\leq \mathbb{M}_p(i \mapsto \hat{R}(h(\cdot; \hat{\theta}); \ell, \mathbf{z}_i); \mathbf{w}) + \frac{2\varepsilon}{3} \\ &\leq \mathbb{M}_p(i \mapsto \hat{R}(h(\cdot; \theta^*); \ell, \mathbf{z}_i); \mathbf{w}) + \frac{2\varepsilon}{3} \\ &\leq \mathbb{M}_p(i \mapsto R(h(\cdot; \theta^*); \ell, \mathcal{D}_i); \mathbf{w}) + \varepsilon . \end{aligned}$$

We may thus conclude that \mathcal{A} fair-PAC learns \mathcal{H} with sample complexity $gm = g \cdot m_{\text{UC}}(\frac{\varepsilon}{3}, \frac{\delta}{g}, d)$. Furthermore, as the entire operation requires polynomial time, we have $(\mathcal{H}, \ell) \in \text{PAC}_{\text{Poly}}^{\text{Agn}}$. \square

We now work towards proof of theorem 5.3.2. We begin with a technical lemma deriving relevant properties of the cover employed in the main result.

Lemma E.1.1 (Group Cover Properties). Suppose loss function ℓ of bounded codomain (i.e., $\|\ell\|_\infty$ is bounded), hypothesis class $\mathcal{H} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$, and per-group samples $\mathbf{z}_{1:g,1:m} \in (\mathcal{X} \times \mathcal{Y})^{g \times m}$, letting $\bigcirc_{i=1}^g \mathbf{z}_i$ denote their *concatenation*. Now define

$$\hat{\mathcal{C}}_{\cup(1:g)} \doteq \bigcup_{i=1}^g \hat{\mathcal{C}}(\ell \circ \mathcal{H}, \mathbf{z}_i, \gamma) \quad \& \quad \hat{\mathcal{C}}_{\circ(1:g)} \doteq \hat{\mathcal{C}}\left(\ell \circ \mathcal{H}, \bigcirc_{i=1}^g \mathbf{z}_i, \frac{\gamma}{\sqrt{g}}\right) .$$

Then, letting $\hat{\mathcal{C}}$ refer generically to either $\hat{\mathcal{C}}_{\cup(1:g)}$ or $\hat{\mathcal{C}}_{\circ(1:g)}$, the following hold.

1. If $\hat{\mathcal{C}}$ is of *minimal cardinality*, then

$$\left| \hat{\mathcal{C}}_{\cup(1:g)} \right| \leq g \mathcal{N}(\ell \circ \mathcal{H}, m, \gamma) \quad \& \quad \left| \hat{\mathcal{C}}_{\circ(1:g)} \right| \leq \mathcal{N}(\ell \circ \mathcal{H}, gm, \frac{\gamma}{\sqrt{g}}) .$$

2. $\sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} \mathfrak{R}_m(\ell \circ \mathcal{H}, \mathcal{D}) \leq \inf_{\gamma \geq 0} \gamma + \|\ell\|_\infty \sqrt{\frac{\ln \mathcal{N}(\ell \circ \mathcal{H}, \gamma)}{2m}}$.

3. Suppose $\ln \mathcal{N}(\ell \circ \mathcal{H}, \gamma) \in \text{Poly}(\frac{1}{\gamma})$. Then the uniform-convergence sample-complexity of $\ell \circ \mathcal{H}$ over g groups obeys

$$\begin{aligned} m_{\text{UC}}(\ell \circ \mathcal{H}, \varepsilon, \delta, g) &\doteq \text{argmin} \left\{ m \left| \sup_{\mathcal{D}_{1:g} \text{ over } (\mathcal{X} \times \mathcal{Y})^g} \mathbb{P} \left(\max_{i \in \{1, \dots, g\}} \sup_{h \in \mathcal{H}} \left| \mathbb{E}_{\mathcal{D}_i}[\ell \circ h] - \hat{\mathbb{E}}_{\mathbf{z}_i}[\ell \circ h] \right| > \varepsilon \right) \leq \delta \right\} \\ &\leq \frac{8 \|\ell\|_\infty^2 \ln \left(\sqrt[4]{\frac{2g}{\delta}} \mathcal{N}(\ell \circ \mathcal{H}, \frac{\varepsilon}{4}) \right)}{\varepsilon^2} \\ &\in \mathbf{O} \left(\frac{\ln \frac{g \mathcal{N}(\ell \circ \mathcal{H}, \varepsilon)}{\delta}}{\varepsilon^2} \right) \subset \text{Poly} \left(\frac{1}{\varepsilon}, \exp \frac{1}{\delta}, \exp g \right) . \end{aligned}$$

4. For the sample \mathbf{z}_i associated with each group $i \in 1, \dots, g$, $\hat{\mathcal{C}}$ is a γ -uniform-approximation of *empirical risk* $\hat{\mathbb{R}}(h; \ell, \mathbf{z}_i)$, and a γ - ℓ_2 cover of the *loss family* $\ell \circ \mathcal{H}$, as

$$\max_{i \in \{1, \dots, g\}} \min_{h_\gamma \in \hat{\mathcal{C}}} \left| \hat{\mathbb{R}}(h; \ell, \mathbf{z}_i) - \hat{\mathbb{R}}(h_\gamma; \ell, \mathbf{z}_i) \right| \leq \max_{i \in \{1, \dots, g\}} \min_{h_\gamma \in \hat{\mathcal{C}}} \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2} \leq \gamma .$$

5. $\hat{\mathcal{C}}_{\circ(1:g)}$, but not necessarily $\hat{\mathcal{C}}_{\cup(1:g)}$, *simultaneously* (across all groups) γ -uniformly-approximates *empirical risk*, and is a γ - ℓ_2 cover of the *loss family* $\ell \circ \mathcal{H}$, as

$$\min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \max_{i \in \{1, \dots, g\}} \left| \hat{\mathbb{R}}(h; \ell, \mathbf{z}_i) - \hat{\mathbb{R}}(h_\gamma; \ell, \mathbf{z}_i) \right| \leq \min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \max_{i \in \{1, \dots, g\}} \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2} \leq \gamma .$$

Proof. We first show items 1 to 3, followed by a key intermediary relating risk values and ℓ_2 distances, and close by showing items 4 and 5.

We begin with item 1. Both bounds follow directly from the definition of uniform covering numbers.

We now show item 2. This result follows via a standard sequence of operations over the Rademacher average. In particular, observe

$$\begin{aligned}
\sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} \mathfrak{R}_m(\ell \circ \mathcal{H}, \mathcal{D}) &= \sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} \mathbb{E}_{z \sim \mathcal{D}^m} \left[\hat{\mathfrak{R}}_m(\ell \circ \mathcal{H}, z) \right] && \text{DEFINITION OF } \mathfrak{R} \\
&\leq \sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} \mathbb{E}_{z \sim \mathcal{D}^m} \left[\inf_{\gamma \geq 0} \gamma + \hat{\mathfrak{R}}_m(\mathcal{C}^*(\ell \circ \mathcal{H}, z, \gamma), z) \right] && \text{DISCRETIZATION} \\
&\leq \inf_{\gamma \geq 0} \gamma + \sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} \mathbb{E}_{z \sim \mathcal{D}^m} \left[\|\ell\|_\infty \sqrt{\frac{\ln |\mathcal{C}^*(\ell \circ \mathcal{H}, z, \gamma)|}{2m}} \right] && \text{MASSART'S INEQUALITY} \\
&\leq \inf_{\gamma \geq 0} \gamma + \|\ell\|_\infty \sqrt{\frac{\ln \mathcal{N}(\ell \circ \mathcal{H}, \gamma)}{2m}}, && \text{DEFINITION OF } \mathcal{N}
\end{aligned}$$

where the MASSART'S INEQUALITY step follows via *Massart's finite class inequality* [Massart, 2000, lemma 1], and the DISCRETIZATION step via *Dudley's discretization argument*.

We now show item 3. By the *symmetrization inequality*, and a 2-tailed application of McDiarmid's bounded difference inequality [McDiarmid, 1989], where changing any $z_{i,j}$ has bounded difference $\frac{\|\ell\|_\infty}{m}$, we have that

$$\forall i : \mathbb{P} \left(\sup_{h \in \mathcal{H}} \left| \mathbb{E}_{\mathcal{D}_i}[\ell \circ h] - \hat{\mathbb{E}}_{z_i \sim \mathcal{D}_i^m}[\ell \circ h] \right| > 2\mathfrak{R}_m(\ell \circ \mathcal{H}, \mathcal{D}_i) + \|\ell\|_\infty \sqrt{\frac{\ln \frac{2}{\delta}}{2m}} \right) \leq \delta$$

thus by union bound over g groups, we have

$$\mathbb{P} \left(\max_{i \in \{1, \dots, g\}} \sup_{h \in \mathcal{H}} \left| \mathbb{E}_{\mathcal{D}_i}[\ell \circ h] - \hat{\mathbb{E}}_{z_i \sim \mathcal{D}_i^m}[\ell \circ h] \right| > \sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} 2\mathfrak{R}_m(\ell \circ \mathcal{H}, \mathcal{D}) + \|\ell\|_\infty \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}} \right) \leq \delta.$$

Now, let *estimation error* bound $\varepsilon_{\text{est}} \doteq \sup_{\mathcal{D} \text{ over } \mathcal{X} \times \mathcal{Y}} 2\mathfrak{R}_m(\ell \circ \mathcal{H}, \mathcal{D}) + \|\ell\|_\infty \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}}$, and observe that via item 2,

$$\varepsilon_{\text{est}} \leq \inf_{\gamma \geq 0} 2\gamma + 2\|\ell\|_\infty \sqrt{\frac{\ln \mathcal{N}(\ell \circ \mathcal{H}, \gamma)}{2m}} + \|\ell\|_\infty \sqrt{\frac{\ln \frac{2g}{\delta}}{2m}}.$$

From here, we solve for an upper-bound on sample-size m to get

$$\begin{aligned}
\text{m}_{\text{UC}}(\ell \circ \mathcal{H}, \varepsilon, \delta, g) &\leq \inf_{\gamma \geq 0} \frac{4\|\ell\|_\infty^2 \ln \mathcal{N}(\ell \circ \mathcal{H}, \gamma) + \|\ell\|_\infty^2 \ln \frac{2g}{\delta}}{2(\varepsilon - 2\gamma)^2} = \frac{2\|\ell\|_\infty^2 \ln \left(\sqrt[4]{\frac{2g}{\delta}} \mathcal{N}(\ell \circ \mathcal{H}, \gamma) \right)}{(\varepsilon - 2\gamma)^2} \\
&\leq \frac{8\|\ell\|_\infty^2 \ln \left(\sqrt[4]{\frac{2g}{\delta}} \mathcal{N}(\ell \circ \mathcal{H}, \frac{\varepsilon}{4}) \right)}{\varepsilon^2} && \text{SET } \gamma = \frac{\varepsilon}{4} \\
&\in \mathbf{O} \left(\frac{\ln \frac{g \mathcal{N}(\ell \circ \mathcal{H}, \varepsilon)}{\delta}}{\varepsilon^2} \right) \subset \text{Poly} \left(\frac{1}{\varepsilon}, \exp \frac{1}{\delta}, \exp g \right). && \mathcal{N}(\ell \circ \mathcal{H}, \varepsilon) \in \text{Poly} \frac{1}{\varepsilon}
\end{aligned}$$

We now show an intermediary which immediately implies the left inequalities of both items 4 and 5. In particular, we may relate these empirical risk gaps to (size-normalized) ℓ_2 distance, as (for each i) we have

$\forall h \in \mathcal{H}, h_\gamma \in \hat{\mathcal{C}}$ that

$$\left| \hat{R}(h; \ell, \mathbf{z}_i) - \hat{R}(h_\gamma; \ell, \mathbf{z}_i) \right| \leq \frac{1}{m} \sum_{j=1}^m |(\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j})| \leq \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2}.$$

Here the last inequality holds since we divide m *inside the* $\sqrt{\cdot}$. The opposite inequality holds for standard ℓ_1 and Euclidean distance, where m is not divided, essentially because the ℓ_1 and ℓ_2 distances differ by up to a factor \sqrt{m} , but this form may be familiar as the relationship between the *mean* and *root mean square* errors. The unconvinced reader may note that this size-normalized ℓ_2 distance is in fact the (unweighted) $p = 2$ power-mean, and thus this step follows via theorem 3.3.7 item 1.

We now show the right inequality of item 4. Note that for the case of $\hat{\mathcal{C}}_{\cup(1:g)}$, the result is almost tautological, as it holds per group by the union-based construction of $\hat{\mathcal{C}}_{\cup(1:g)}$. The case of $\hat{\mathcal{C}}_{\circ(1:g)}$ is more subtle, but we defer its proof to the final item, as it then follows as an immediate consequence of the *max-min inequality*, i.e., $\forall h \in \mathcal{H}$,

$$\max_{i \in 1, \dots, g} \min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2} \leq \min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \max_{i \in 1, \dots, g} \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2}.$$

We now show item 5. Suppose (by way of contradiction) that there exists some $h \in \mathcal{H}$ such that

$$\min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \max_{i \in 1, \dots, g} \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2} > \gamma.$$

One then need only consider the summands associated with a maximal i to observe that this implies

$$\min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \sqrt{\frac{1}{mg} \sum_{i=1}^g \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2} > \frac{\gamma}{\sqrt{g}},$$

thus $\hat{\mathcal{C}}_{\circ(1:g)}$ is not a $\frac{\gamma}{\sqrt{g}}$ - ℓ_2 cover of $\bigcirc_{i=1}^g \mathbf{z}_i$, which contradicts its very definition. We thus conclude

$$\forall h \in \mathcal{H} : \min_{h_\gamma \in \hat{\mathcal{C}}_{\circ(1:g)}} \max_{i \in 1, \dots, g} \sqrt{\frac{1}{m} \sum_{j=1}^m ((\ell \circ h)(\mathbf{z}_{i,j}) - (\ell \circ h_\gamma)(\mathbf{z}_{i,j}))^2} \leq \gamma.$$

□

With lemma E.1.1 in hand, we are now ready to show theorem 5.3.2.

Theorem 5.3.2 (Efficient FPAC-Learning by Covering). Suppose loss function ℓ of bounded codomain (i.e., $\|\ell\|_\infty$ is bounded), and hypothesis class sequence \mathcal{H} , s.t. $\forall m, d \in \mathbb{N}$, $\mathbf{z} \in (\mathcal{X} \times \mathcal{Y})^m$, there exist

1. a γ - ℓ_2 cover $\mathcal{C}^*(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$, where $|\mathcal{C}^*(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)| \leq \mathcal{N}(\ell \circ \mathcal{H}_d, \gamma) \in \text{Poly}(\frac{1}{\gamma}, d)$; and
2. an algorithm to enumerate a γ - ℓ_2 cover $\hat{\mathcal{C}}(\ell \circ \mathcal{H}_d, \mathbf{z}, \gamma)$ of size $\text{Poly} \mathcal{N}(\ell \circ \mathcal{H}_d, m, \gamma)$ in $\text{Poly}(m, \frac{1}{\gamma}, d)$ time.

Then $(\ell, \mathcal{H}) \in \text{FPAC}_{\text{Poly}}$, where (1) is sufficient to show that (ℓ, \mathcal{H}) is fair-PAC learnable with polynomial *sample complexity*, and (2) is required only to show polynomial *training time complexity*.

Proof. We now constructively show the existence of a fair-PAC-learner \mathcal{A} for (ℓ, \mathcal{H}) over domain \mathcal{X} and codomain \mathcal{Y} . As in theorem 5.3.1, we first note that by theorem 3.3.8 items 3 and 5, under the conditions of

FPAC learning, this reduces to showing that we can learn any malfare concept $\Lambda_p(\cdot; \cdot)$ that is a p -power mean with $p \geq 1$.

We first assume a training sample $\mathbf{z}_{1:g,1:m} \sim \mathcal{D}_1^m \times \cdots \times \mathcal{D}_g^m$, i.e., a collection of m draws from each of the g groups. In particular, we shall select m to guarantee that the *estimation error* for the malfare does not exceed $\frac{\varepsilon}{3}$ with probability at least $1 - \delta$, i.e., we require that with said probability,

$$\varepsilon_{\text{est}} \doteq \left| \Lambda_p(i \mapsto \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i); \mathbf{w}) - \Lambda_p(i \mapsto \mathbf{R}(h; \ell, \mathcal{D}_i); \mathbf{w}) \right| \leq \frac{\varepsilon}{3} .$$

Now, note that by theorem 3.3.7 item 3 (contraction), we have

$$\left| \Lambda_p(i \mapsto \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i); \mathbf{w}) - \Lambda_p(i \mapsto \mathbf{R}(h; \ell, \mathcal{D}_i); \mathbf{w}) \right| \leq \max_{i \in \{1, \dots, g\}} \sup_{h \in \mathcal{H}_d} \left| \mathbf{R}(h; \ell, \mathcal{D}_i); \mathbf{w} \right| - \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i); \mathbf{w} \right| ,$$

and by lemma E.1.1 item 3, a sample of size

$$m = \left\lceil \frac{81 \|\ell\|_\infty^2 \ln \left(\sqrt[4]{\frac{2g}{\delta}} \mathcal{N}(\ell \circ \mathcal{H}, \frac{\varepsilon}{12}) \right)}{\varepsilon^2} \right\rceil \in \mathbf{O} \left(\frac{\ln \frac{g \mathcal{N}(\ell \circ \mathcal{H}, \varepsilon)}{\delta}}{\varepsilon^2} \right) \subset \text{Poly} \left(\frac{1}{\varepsilon}, \exp \frac{1}{\delta}, \exp g \right)$$

suffices to ensure that

$$\mathbb{P} \left(\max_{i \in \{1, \dots, g\}} \sup_{h \in \mathcal{H}_d} \left| \mathbf{R}(h; \ell, \mathcal{D}_i); \mathbf{w} \right| - \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i); \mathbf{w} \right| > \frac{\varepsilon}{3} \right) \leq \delta ,$$

thus guaranteeing the stated estimation error bound.

With our sample size and *estimation error* guarantee, we now define the *learning algorithm* and bound its *optimization error*. Take *cover precision* $\gamma \doteq \frac{\varepsilon}{3\sqrt{g}}$. By assumption, for each $d \in \mathbb{N}$, we may enumerate a γ -cover $\hat{\mathcal{C}}(\ell \circ \mathcal{H}_d, \bigcirc_{i=1}^g \mathbf{z}_i, \gamma)$, where $\bigcirc_{i=1}^g \mathbf{z}_i$ denotes the *concatenation* of each \mathbf{z}_i , in $\text{Poly}(gm, \frac{1}{\gamma}, d) = \text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, d, g)$ time. For the remainder of this proof, we refer to this cover as $\hat{\mathcal{C}}$.

Now, we take the learning algorithm to be *empirical malfare minimization* over $\hat{\mathcal{C}}$. Let

$$\hat{h} \doteq \underset{h_\gamma \in \hat{\mathcal{C}}}{\text{argmin}} \Lambda_p(i \mapsto \hat{\mathbf{R}}(h_\gamma; \ell, \mathbf{z}_i); \mathbf{w}) \quad \& \quad \tilde{h} \doteq \underset{h \in \mathcal{H}_d}{\text{argmin}} \Lambda_p(i \mapsto \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i); \mathbf{w}) ,$$

where ties may be broken arbitrarily. Note that via standard covering properties, that the *optimization error* is bounded as

$$\begin{aligned} \varepsilon_{\text{opt}} &\doteq \Lambda_p(i \mapsto \hat{\mathbf{R}}(\hat{h}; \ell, \mathbf{z}_i); \mathbf{w}) - \Lambda_p(i \mapsto \hat{\mathbf{R}}(\tilde{h}; \ell, \mathbf{z}_i); \mathbf{w}) && \text{DEFINITION} \\ &= \sup_{h \in \mathcal{H}_d} \min_{h_\gamma \in \hat{\mathcal{C}}} \Lambda_p(i \mapsto \hat{\mathbf{R}}(h_\gamma; \ell, \mathbf{z}_i); \mathbf{w}) - \Lambda_p(i \mapsto \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i); \mathbf{w}) && \text{PROPERTIES OF SUPREMA} \\ &\leq \sup_{h \in \mathcal{H}_d} \min_{h_\gamma \in \hat{\mathcal{C}}} \max_{i \in \{1, \dots, g\}} \left| \hat{\mathbf{R}}(h_\gamma; \ell, \mathbf{z}_i) - \hat{\mathbf{R}}(h; \ell, \mathbf{z}_i) \right| && \text{ITEM 3 (CONTRACTION)} \\ &\leq \sqrt{g} \gamma = \frac{\varepsilon}{3} . && \text{LEMMA E.1.1 ITEM 5} \end{aligned}$$

This controls for *optimization error* between the true and approximate EMM solutions \tilde{h} and \hat{h} .

We now combine the optimization and estimation error inequalities, letting

$$h^* \doteq \underset{h \in \mathcal{H}_d}{\text{argmin}} \Lambda_p(i \mapsto \mathbf{R}(h; \ell, \mathcal{D}_i); \mathbf{w}) ,$$

denote the true *malware optimal* solution, over *distributions* rather than *samples*, breaking ties arbitrarily. We then derive

$$\begin{aligned}
& \mathbb{M}_p(i \mapsto R(h^*; \ell, \mathcal{D}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto R(\hat{h}; \ell, \mathcal{D}_i); \mathbf{w}) \\
&= \left(\mathbb{M}_p(i \mapsto R(h^*; \ell, \mathcal{D}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto \hat{R}(h^*; \ell, \mathbf{z}_i); \mathbf{w}) \right) && \leq \varepsilon_{\text{est}} \\
&+ \left(\mathbb{M}_p(i \mapsto \hat{R}(h^*; \ell, \mathbf{z}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto \hat{R}(\tilde{h}; \ell, \mathbf{z}_i); \mathbf{w}) \right) && \leq 0 \\
&+ \left(\mathbb{M}_p(i \mapsto \hat{R}(\tilde{h}; \ell, \mathbf{z}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto \hat{R}(\hat{h}; \ell, \mathbf{z}_i); \mathbf{w}) \right) && \leq \varepsilon_{\text{opt}} \\
&+ \left(\mathbb{M}_p(i \mapsto \hat{R}(\hat{h}; \ell, \mathbf{z}_i); \mathbf{w}) - \mathbb{M}_p(i \mapsto R(\hat{h}; \ell, \mathcal{D}_i); \mathbf{w}) \right) && \leq \varepsilon_{\text{est}} \\
&\leq \varepsilon_{\text{opt}} + 2\varepsilon_{\text{est}} = \frac{\varepsilon}{3} + \frac{2\varepsilon}{3} = \varepsilon . && \text{SEE ABOVE}
\end{aligned}$$

We thus conclude that this algorithm produces an ε - $\mathbb{M}_p(\cdot; \cdot)$ optimal solution with probability at least $1 - \delta$, and furthermore both the sample complexity and time complexity of this algorithm are $\text{Poly}(\frac{1}{\varepsilon}, \frac{1}{\delta}, g, d)$. Hence, as we have constructed a polynomial-time fair-PAC learner for (\mathcal{H}, ℓ) , we may conclude $(\mathcal{H}, \ell) \in \text{FPAC}_{\text{Poly}}$. \square

Appendix F

Supplementary Material for Chapter 7

F.1 A Compendium of Missing Proofs

Definition F.1.1. Consider $\mathbf{X}_{1:T} : X_1, X_2, \dots, X_T$. We define C_i to be the autocovariance of two steps of \mathcal{M} (at stationarity) being i apart. i.e., $C_i \doteq \text{Cov}(X_1, X_{1+i})$.

Note that for i.i.d samples the autocovariance of any pair of samples X_i and X_j is zero. We do not have independence here nevertheless for reversible Markov chains (See for example Equation 12.9 from Levin and Peres [2017]) we have $C_i \leq \lambda^i \sqrt{\mathbb{V}_\pi[f] \mathbb{V}_\pi[f]} = \lambda^i v$. The following lemmas will be used throughout:

Lemma F.1.2. Suppose $\mathbf{X}_{1:T} : X_1, X_2, \dots, X_T$ is a trace of length T of \mathcal{M} . using the above definition for C_i , the trace variance are related as

$$v_T = \frac{1}{T} v_\pi + \frac{2}{T^2} \sum_{i=1}^{T-1} (T-i) C_i . \quad (69)$$

Proof. The trace variance is

$$v_T = \mathbb{E} \left[\left(\frac{1}{T} \sum_{i=1}^T (f(X_i) - \mu) \right)^2 \right] = \mathbb{E} \left[\frac{1}{T^2} \sum_{i=1}^T \sum_{j=1}^T (f(X_i) - \mu)(f(X_j) - \mu) \right] = \frac{1}{T} v_\pi + \frac{2}{T^2} \sum_{i=1}^{T-1} (T-i) C_i.$$

□

Lemma F.1.3. For any lazy reversible chain \mathcal{M} , we have $v_\pi \leq 2v_2 \leq 3v_3 \leq \dots$, i.e, for any T and T' satisfying $T < T'$ we have, $Tv_T \leq T'v_{T'}$.

Lemma F.1.4. For $1 \leq i \leq \frac{n}{2}$, let f_i be defined as in lemma 7.2.3, we have $v_T(f_i) \leq \Theta(i^2/2)$.

Proof. We first prove that $v_T(f_i) \leq \Theta(i^2)/T$. Note that the function $x \mapsto x \bmod 2i$ partitions the set $[n]$ to $2i$ partitions, and the relaxation time of the projection chain \mathcal{M}_{f_i} is $\Theta(i^2)$.

We first assume WLOG that i divides 3: this is so we may analyze an convenient integer-length trace, but the remaining cases hold with constant-factor differences. Note that $\tau_{\text{rel}} \in \Theta(n^2)$, and our proof operates by analyzing I -traces, for $I \doteq i^2/9$, and then applying trace-variance inequalities to draw the desired conclusions.

Assume X_1, X_2, \dots, X_{i^2} is a trace of the unbiased walk on the cycle, and define $Y_k = X_{k+1} - X_k$. Thus, $X_k = \sum_{j=1}^{k-1} Y_j + s_0$ where s_0 is the starting point. Note that Y_j is a symmetric random variable (thus has equal mean and median), which we shall use to apply Lévy's inequality. The proof strategy here is to

lower-bound $v_{i^2/9}$ by showing that a constant fraction of stationary traces in $\mathcal{M}^{(i^2/9)}$ see only one color. We call such traces *homogeneous*, and note that for such traces \mathbf{X} , we have $(f_{\text{avg}}(\mathbf{X}) - \mu)^2 = \frac{1}{4}$ (as $f_{\text{avg}}(\mathbf{X}) \in \{0, 1\}$, and $\mu = \frac{1}{2}$), and from there we bound trace variances as appropriate.

We begin with a key step in deriving a lower-bound on the proportion of such homogeneous traces. In particular, take $S_0 \doteq \{i/3 + 1, i/3 + 2, \dots, 2i/3\} \subseteq [n]$, and similarly take $S_k \doteq \{ki + i/3 + 1, ki + i/3 + 2, \dots, ki + 2i/3\} \subseteq [n]$, in other words each S_k is the *middle third* of the k th contiguous color region. Finally, take $S \doteq \bigcup_{k=1}^{n/i} S_k$. These regions are depicted graphically in figure 7.2.

Let $\text{SD}(\mathbf{Y}) \doteq \max_{k \in 1, \dots, i^2/9} \left| \sum_{j=1}^I Y_k \right|$, i.e., $\text{SD}(\mathbf{Y}) = \max_{k \in 1, \dots, i^2/9} \Delta_{\circ}(X_1, X_k)$, for Δ_{\circ} the shortest-path distance on the cycle, and observe

$$\begin{aligned} \mathbb{P}\left(\text{SD}(\mathbf{Y}) \geq \frac{i}{3}\right) &\leq 2 \mathbb{P}\left(\left|\sum_{j=1}^{i^2/9} Y_j\right| \geq \frac{i}{3}\right) && \text{LÉVY'S INEQUALITY} \\ &\leq 4 \exp\left(\frac{-2\left(\frac{i}{3}\right)^2}{i^2/9}\right) && \text{HOEFFDING'S INEQUALITY} \\ &= 4 \exp(-2) . \end{aligned}$$

From here, we decompose the trace variance and bound it as

$$\begin{aligned} v_I &= \mathbb{E}_{\mathbf{Y}} \left(\frac{1}{i^2/9} \sum_{j=1}^{i^2/9} Y_j - \mu \right)^2 && \text{BY DEFINITION} \\ &\geq \mathbb{P}(X_1 \in S) \mathbb{P}\left(\text{SD}(\mathbf{Y}) \leq \frac{i}{3} \mid X_1 \in S\right) \mathbb{E}_{\mathbf{Y} \mid \text{SD}(\mathbf{Y}) \leq \frac{i}{3}} \left(\frac{1}{i^2/9} \sum_{j=1}^{i^2/9} Y_j - \mu \right)^2 && \text{LAW OF TOTAL EXPECTATION} \\ &\geq \frac{1}{3} (1 - 4 \exp(-2)) \frac{1}{4} > 0.03822 . && \begin{aligned} \mathbb{P}(X_1 \in S) &= \frac{1}{3} \\ \text{SEE ABOVE} \\ (0 - \frac{1}{2})^2 &= (1 - \frac{1}{2})^2 = \frac{1}{4} \end{aligned} \end{aligned}$$

We thus conclude $v_I \in \Theta(1)$, therefore via lemma F.1.3, $v_T \in \Theta(i^2/T)$. □

F.1.1 Correctness and Efficiency of DynaMITE

A Note on Nonstationarity

Note that theorems 7.2.5 and 7.2.6 assume *stationarity* i.e., they consider traces $\mathbf{X}_{1:m} : X_1, X_2, \dots, X_m$ assuming $X_1 \sim \pi$. This assumption is often prohibitive, as drawing even a single such sample can be computationally infeasible. We overcome this problem using the following equation for $\mathbf{X}_{1:m} : X_1, X_2, \dots, X_m$, $X_1 \sim \nu$.

$$\mathbb{P}_{\mathbf{X}_{1:m}} \left(|\hat{\mu} - \mu| \geq \varepsilon \right) \leq \left\| \frac{\partial \nu}{\partial \pi} \right\|_{\pi, \infty} \mathbb{P}_{\mathbf{X}_{1:m}} \left(|\hat{\mu} - \mu| \geq \varepsilon \right) ; \quad (70)$$

see, e.g., [Fan et al., 2018], proof of theorem 2.3. Note that by eq. 70, it is sufficient to have $\left\| \frac{\partial \nu}{\partial \pi} \right\|_{\pi, \infty} = \text{ess sup}_{\omega \in \Omega} \left| \frac{\nu(x)}{\pi(x)} \right| \in \mathcal{O}(1)$. This can generally be accomplished straightforwardly with a *warm-start* by selecting an arbitrary fixed $\omega \in \Omega$, taking ν to be the distribution reached after running \mathcal{M} for $\tau_{\text{rel}}(\mathcal{M}) \ln(1/\pi_{\min})$ steps (see, e.g., Levin and Peres [2017]). With this in mind, for simplicity we assume stationarity, knowing that our proofs and algorithm generalize with trivial modifications.

Stationarity and mixing time of trace chain $\mathcal{M}^{(T)}$

Lemma F.1.5 (Trace Chain Mixing). Having a Markov chain \mathcal{M} with stationary distribution π and mixing time τ_{mix} , the trace chain $\mathcal{M}^{(T)}$'s stationary distribution is $\pi^{(T)}$ where $\pi^{(T)}(\mathbf{a}) \doteq \pi(a_1) \prod_{i=1}^{T-1} \mathcal{M}(a_i, a_{i+1})$ and its mixing time is bounded as $\tau_{\text{mix}}(\mathcal{M}^{(T)}) \leq 1 + \frac{1}{T} \tau_{\text{mix}}(\mathcal{M})$. In particular for $T \geq \tau_{\text{mix}}(\mathcal{M})$, we have that the second largest eigenvalue of $\mathcal{M}^{(T)}$ is at most $4/5$.

Proof of Lemma F.1.5. We first remark that it is not difficult to prove by induction that

$$\left(\mathcal{M}^{(T)} \right)^j(\mathbf{a}, \mathbf{b}) = \mathcal{M}^{(j-1)T+1}(a_T, b_1) \prod_{i=1}^{T-1} \mathcal{M}(b_i, b_{i+1}).$$

It is not difficult to see that with the above definitions, for any \mathbf{a} a trace of length T we have: $\pi^{(T)}(\mathbf{a}) = \sum_{\mathbf{b} \in \Omega^T} \pi^{(T)}(\mathbf{a}, \mathbf{b}) \mathcal{M}^{(T)}(\mathbf{a}, \mathbf{b})$.

We now show that $\mathcal{M}^{(T)}$ is close to $\pi^{(T)}$ after $(\tau(\varepsilon)/T) + 1$ steps. Let X_0, X_1, X_2, \dots , be the trace of the original random walk, thus each X_i is a random variable having distribution $X_i \sim \mathcal{M}^i \mu$. We partition this trace into blocks of length T as follows: $B^j = (X_j, X_{j+1}, \dots, X_{j+T-1})$. Thus the trace of $\mathcal{M}^{(T)}$ will be B^0, B^1, B^2, \dots .

Since the mixing time of \mathcal{M} is τ_{mix} , we have: for any starting distribution ν over Ω , and $\tau \geq \tau_{\text{mix}} \log(\varepsilon)$, $\text{TVD}(\mathcal{M}^\tau(\nu), \pi) \leq \varepsilon$. Assume $\mathcal{M}^{(T)}$ has started at initial distribution ν' , looking at the distribution of $B_{\frac{\tau(\varepsilon)}{T}+1}$ we will have:

$$\begin{aligned} \text{TVD}(B_{\frac{\tau}{T}+1}, \pi^{(T)}) &= \frac{1}{2} \sum_{\mathbf{a} \in \Omega^T} \left| (\mathcal{M}^{(T)})^{\frac{\tau}{T}+1} \nu'(a) - \pi^{(T)}(a) \right| \\ &= \frac{1}{2} \sum_{\mathbf{a} \in \Omega^T} \left| (\mathcal{M}^\tau)(\nu'_T, a_1) \prod_{i=1}^{T-1} \mathcal{M}(a_i, a_{i+1}) - \pi^{(T)}(a) \right| \\ &= \frac{1}{2} \sum_{\mathbf{a} \in \Omega^T} \left| (\mathcal{M}^\tau)(\nu'_T, a_1) \prod_{i=1}^{T-1} \mathcal{M}(a_i, a_{i+1}) - \pi(a_1) \prod_{i=1}^{T-1} \mathcal{M}(a_i, a_{i+1}) \right| \\ &\leq \varepsilon \prod_{i=1}^{T-1} \mathcal{M}(a_i, a_{i+1}) \leq \varepsilon. \end{aligned}$$

The bound on λ then follows since $(\tau_{\text{rel}}(\mathcal{M}) - 1) \ln(2) \leq \tau_{\text{mix}}(\mathcal{M}) \leq \tau_{\text{rel}}(\mathcal{M}) \ln\left(\frac{2}{\sqrt{\pi_{\min}}}\right)$.

□

Trace Variance Estimation

Lemma 7.3.5 (Trace Variance Estimation). Let \hat{v} be defined as in definition 7.3.4, it then holds that, for any $\delta \in (0, 1)$, we have

$$\mathbb{P} \left(v_\pi \geq \hat{v} + \frac{(11 + \sqrt{21})(1 + \frac{\lambda}{\sqrt{21}})r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m} + \sqrt{\frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} \right) \leq \delta .$$

Proof. The strategy here is to apply a Bernstein-type inequality to the *empirical variance estimate* \hat{v} . As Bernstein's inequality bounds an *expectation*, we construct a *tensor product chain* $\mathcal{M} \otimes \mathcal{M}$, and define a function $g : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$, such that $\mathbb{E}_{\mathcal{M} \times \mathcal{M}}[g] = \mathbb{V}_{\mathcal{M}} f$.

In particular, we first note that by the tensor product chain rule (see Ex. 12.6 of [Levin and Peres, 2017]), the *spectral gap* $(1 - \lambda)$ of $\mathcal{M} \otimes \mathcal{M}$ equals that of \mathcal{M} . We now define the function $g : \mathcal{X} \times \mathcal{X} \mapsto [0, \frac{1}{2}r^2]$ as $g(x_1, x_2) \doteq \frac{1}{2}(x_1 - x_2)^2$.

We first show that g is an *unbiased estimator* of the variance of f , i.e., $\mathbb{E}_{\mathcal{M} \times \mathcal{M}}[g] = \mathbb{V}_{\mathcal{M}}[f]$:

$$\begin{aligned} \mathbb{E}[g] &= \mathbb{E} \left[\frac{1}{2}(X_1 - X_2)^2 \right] && \text{DEFINITION OF } g \\ &= \frac{1}{2} \mathbb{E} \left[(X_1 - X_2)^2 \right] && \text{STATIONARITY} \\ &= \frac{1}{2} \left(\mathbb{E}[X_1^2 + X_2^2] - 2\mathbb{E}[X_1 X_2] \right) && \text{LINEARITY} \\ &= \mathbb{E}[X_1^2] - (\mathbb{E}[X_1])^2 && \text{INDEPENDENCE} \\ &= \mathbb{V}[X_1] = v_\pi && \text{VARIANCE PROPERTIES} \end{aligned}$$

We now seek to apply the Bernstein bound to g on $\mathcal{M} \otimes \mathcal{M}$. Note that the *range* of g is $[0, \frac{1}{2}r^2]$. We require also a bound on the *variance of the variance* $\mathbb{V}[g]$. We break the infinite regress here by noting that

$$\begin{aligned} \mathbb{V}[g] &= \mathbb{E}[g^2] - (\mathbb{E}[g])^2 \\ &\leq \mathbb{E}[g \cdot g] \\ &\leq \mathbb{E}[g \frac{1}{2}r^2] \\ &= \frac{1}{2}r^2 \mathbb{E}[g] . \\ &= \frac{1}{2}r^2 v_\pi . \end{aligned}$$

Note that this is effectively the argument of Bernstein's inequality that removes higher moments (beyond the second) from the exponential-sum decomposition in the Chernoff-MGF bound.

Applying the Bernstein inequality (theorem 7.2.6) then yields

$$\mathbb{P}(v_\pi \geq \hat{v} + \varepsilon) \leq \delta ,$$

for

$$\begin{aligned} \varepsilon &\leq \frac{10r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\frac{2(1 + \lambda) \mathbb{V}[g] \ln \frac{1}{\delta}}{(1 - \lambda)m}} && \text{Range}(g) = \frac{1}{2}r^2 \\ &\leq \frac{5r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m} + \sqrt{\frac{(1 + \lambda)r^2 v_\pi \ln \frac{1}{\delta}}{(1 - \lambda)m}} && \mathbb{V}[g] \leq \frac{1}{2}r^2 v_\pi \end{aligned}$$

We now seek a form that depends only on the *empirical variance*, which we derive via the quadratic formula.

$$\begin{aligned}
\varepsilon &\leq \frac{(11 + \lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\left(\frac{(1 + \lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m}\right)^2 + \frac{5r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m} \cdot \frac{(1 + \lambda)r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m} + \frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} && \text{QUADRATIC FORMULA} \\
&= \frac{(11 + \lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\left(\frac{(1 + \lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m}\right)^2 + \left(\frac{\sqrt{5}(1 + \lambda)r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m}\right)^2 + \frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} && \text{ALGEBRA} \\
&= \frac{(11 + \lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\left((2\sqrt{5})^2 + (1 + \lambda)\right) \left(\frac{\sqrt{1 + \lambda}r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m}\right)^2 + \frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m} \left((2\sqrt{5})^2 + (1 + \lambda)\right) = 21 + \lambda} \\
&\leq \frac{(11 + \lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\left(\frac{(\sqrt{21} + \frac{11}{\sqrt{21}})\lambda r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m}\right)^2 + \frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} && \sqrt{21 + \lambda}\sqrt{1 + \lambda} \leq \sqrt{21} + \frac{11}{\sqrt{21}}\lambda \\
&\leq \frac{(11 + \sqrt{21} + (1 + \frac{11}{\sqrt{21}})\lambda)r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} && \sqrt{a^2 + b} \leq a + \sqrt{b} \\
&= \frac{((11 + \sqrt{21})(1 + \frac{\lambda}{\sqrt{21}})r^2 \ln \frac{1}{\delta}}{2(1 - \lambda)m} + \sqrt{\frac{(1 + \lambda)r^2 \hat{v} \ln \frac{1}{\delta}}{(1 - \lambda)m}} . && \text{ALGEBRA}
\end{aligned}$$

This concludes the result. \square

As an intermediary, we obtain the following, which shall prove useful in deriving efficiency guarantees.

Corollary F.1.6 (Trace Variance Estimation).

$$\mathbb{P} \left(v_\pi \geq \hat{v} + \frac{5r^2 \ln \frac{1}{\delta}}{(1 - \lambda)m} + \sqrt{\frac{(1 + \lambda)r^2 v_\pi \ln \frac{1}{\delta}}{(1 - \lambda)m}} \right) \leq \delta .$$

DynaMITE's correctness and efficiency

Theorem 7.3.1 (Correctness of DYNAMITE). Consider a Markov chain \mathcal{M} over Ω , and assume its second absolute eigenvalue is less than Λ . Assume we have function $f : \Omega \rightarrow [a, b]$ with $r \doteq b - a$, and we want to estimate *true mean*, $\mu \doteq \mathbb{E}_\pi[f]$ with precision $\varepsilon > 0$, and confidence $(1 - \delta) \in (0, 1)$.

If we start at stationarity, i.e., $(x_0, x_1) \sim \pi(\mathcal{M} \otimes \mathcal{M})$, we take *mean estimate* $\hat{\mu}$ as either

1. $\hat{\mu} \leftarrow \text{MCMCPRo}((x_0, x_1), \mathcal{M}, \Lambda, f, \varepsilon, \delta)$; or
2. $\hat{\mu} \leftarrow \text{DYNAMITE}((x_0, x_1), \mathcal{M}, \Lambda, f, \varepsilon, \delta)$ for lazy \mathcal{M} .

More generally, for a nonstationary start from any $\omega \in \text{Support}(\pi(\mathcal{M}))$, given minimum supported stationary probability at least π_{\min} , we may take

3. $\hat{\mu} \leftarrow \text{WARMSTARTDYNAMITE}(\omega, \mathcal{M}, \Lambda, \pi_{\min}, f, \varepsilon, \delta)$ for lazy, reversible \mathcal{M} .

In all cases, with probability at least $1 - \delta$, we have $|\hat{\mu} - \mu| \leq \varepsilon$.

Proof. We first show (1), i.e., correctness of MCMCPRo, and then show that correctness of WARMSTARTDYNAMITE easily follows from that of DYNAMITE, which follows from that of MCMCPRo.

We now show claim (1). First note that in initialization (independent of any sampling), MCMCPRo computes *iteration count* N and *initial sample size* α (line 4), which determine the *schedule of sample sizes* and

probabilistic bounds. Over the course of the algorithm, at each of N timesteps, a 1-tail (upper) bound on *variance* \mathbf{u}_i is computed (line 12), and a 2-tail bound on *mean* $\boldsymbol{\mu}_i$ is computed (line 13), for a total of $3N$ tail bounds. Each tail is bounded with probability $1 - \frac{\delta}{3N}$, thus by union bound, all hold simultaneously with probability $1 - \delta$. We assume henceforth that app tail bounds hold, thus all conclusions are thus qualified as holding *with probability* $\geq 1 - \delta$. Now, note that termination (line 14) occurs for one of two reasons: either $i = N$, or $\hat{\boldsymbol{\epsilon}}_i \leq \varepsilon$. We analyze these cases separately.

In case 1, we have termination at $i < N$, i.e., $\hat{\boldsymbol{\epsilon}}_i \leq \varepsilon$. As assumed above, all tail bounds at each iteration hold by union bound. Thus for the sample drawn at iteration i , in particular the *variance bounds* (line 12) hold, as $\hat{\mathbf{v}}_i$ is an *unbiased estimate* of v_T , and by lemma 7.3.5, w.h.p., $v_T \leq \mathbf{u}_i$. Similarly, the *mean bounds* (13) hold via theorem 7.2.6 (noting that, by *averaging over* a pair of *independent chains*, the variance proxy of interest is $\frac{1}{2}v_T$). Note that while both tail bounds are taken over the tensor-product chain $\mathcal{M} \otimes \mathcal{M}$, it holds that $\lambda(\mathcal{M} \otimes \mathcal{M}) = \lambda(\mathcal{M})$, so the bound remains valid.

Consequently, when it holds that $\hat{\boldsymbol{\epsilon}}_i \leq \varepsilon$, the algorithm returns the estimate $\hat{\boldsymbol{\mu}}_i$ (line 15), which by the above is sufficiently accurate to satisfy the stated guarantees.

We now consider case 2, wherein we have termination at step $i = N$. When this occurs, it holds that

$$m_i = m_N = \left\lceil \alpha 2^N \right\rceil \geq m_H(\lambda, r, \varepsilon, \frac{\delta}{3T}) ,$$

and thus by Hoeffding's inequality for mixing processes (theorem 7.2.5), we have $|\mu - \hat{\boldsymbol{\mu}}_N| \leq \varepsilon$, (i.e., the schedule was selected exactly to ensure m_N samples would be sufficient, regardless of early termination and variance.

We now proceed to show claim (2). To see this result, note that

$$\lambda(\mathcal{M}^{(T)}) \leq \lambda^T(\mathcal{M}) ,$$

and thus claim (2) follows directly from claim (1).

Finally, note that claim (3) follows from claim (2), paired with Eq. 70. In particular, note that the uniform mixing time is selected such that we have

$$\forall \omega : \left\| \frac{\partial \mathcal{M}^{\tau_{\text{unif}}}(\omega)}{\partial \pi} \right\|_{\pi, \infty} \leq 2 ,$$

i.e., \mathcal{M} is uniformly mixed. This further implies that the trace chain is uniformly mixed, as both π_{\min} and Λ are exponential in the trace-length T , which cancels out in the uniform-mixing bound (see Eq. 12.13 of [Levin and Peres, 2017]). However, we still apply nonstationarity correction $\frac{\delta}{4}$ instead of $\frac{\delta}{2}$, to account for the fact that the *tensor product chain* $(\mathcal{M} \otimes \mathcal{M})$ *uniformly mixes* slightly slower than \mathcal{M} , even though it *relaxes* and *mixes* as quickly as \mathcal{M} . \square

Before showing efficiency, we state a lemma that allows us to relate Tv_T with $\tau_{\text{rel}}v_{\tau_{\text{rel}}}$.

Lemma F.1.7 (Trace Variance of Large Traces). Suppose a lazy reversible Markov chain \mathcal{M} . The product of trace-size and trace-variance, i.e., Tv_T , is *asymptotically equivalent* for $T \geq \tau_{\text{rel}}$. In other words, we have

$$\Theta(\tau_{\text{rel}}v_{\tau_{\text{rel}}}) = \Theta(Tv_T) = \Theta\left(\lim_{t \rightarrow \infty} tv_t\right) .$$

Proof. We show this result by deriving a bidirectional inequality chain. We begin with the upper-bounds. By monotonicity, we have

$$\tau_{\text{rel}}v_{\tau_{\text{rel}}} \leq Tv_T \leq \lim_{t \rightarrow \infty} tv_t .$$

We now show the lower-bounds. This direction is a bit subtler. Assume for now that $T = n\tau_{\text{rel}}$ for some integer n . We then have

$$\begin{aligned}
Tv_T(\mathcal{M}) &= n\tau_{\text{rel}}v_n(\mathcal{M}^{(\tau_{\text{rel}})}) && \text{IDENTITIES} \\
&\leq n\tau_{\text{rel}}\frac{2\tau_{\text{rel}}(\mathcal{M}^{(\tau_{\text{rel}})})v_{\tau_{\text{rel}}}}{n} && \text{EQ. 3.14 OF [PAULIN, 2015]} \\
&= n\tau_{\text{rel}}\frac{2\tau_{\text{rel}}(\mathcal{M}^{(\tau_{\text{rel}})})v_{\tau_{\text{rel}}}(\mathcal{M})}{n} && v_\pi(\mathcal{M}^{(\tau_{\text{rel}})}) = v_{\tau_{\text{rel}}}(\mathcal{M}) \\
&= 2\tau_{\text{rel}}\tau_{\text{rel}}(\mathcal{M}^{(\tau_{\text{rel}})})v_{\tau_{\text{rel}}}(\mathcal{M}) && \text{ALGEBRA} \\
&\in \Theta(\tau_{\text{rel}}v_{\tau_{\text{rel}}}(\mathcal{M})) . && \tau_{\text{rel}}(\mathcal{M}^{(\tau_{\text{rel}})}) \in \Theta(1)
\end{aligned}$$

Correcting for non-integer n is simple, as we may apply the result to the nearest integer n , and the ratio $\frac{v_T}{v_{T'}}$ is bounded when $\frac{T}{T'}$ is near 1.

Finally, note that

$$\lim_{t \rightarrow \infty} tv_t \in \Theta(\tau_{\text{rel}}v_{\tau_{\text{rel}}})$$

follows similarly via Eq. 15 of [Paulin, 2015]. □

We now present and prove an extended statement of theorem 7.3.2, which provides finite-sample and asymptotic sample complexity bounds to MCMCPRO, DYNAMITE, and WARMSTARTDYNAMITE.

Theorem F.1.8 (Efficiency of DYNAMITE). Suppose as in theorem 7.3.1, and take $N \doteq \log_2\left(\frac{r}{2\varepsilon}\right)$ and $T \doteq \lceil \frac{1+\Lambda}{1-\Lambda} \ln \sqrt{2} \rceil$. Then with probability at least $1 - \frac{\delta}{3N}$, each mean-estimation algorithm runs for no more than \hat{m} steps (individually), where \hat{m} is

1. for MCMCPRO:

$$\begin{aligned}
\hat{m} &\leq 4T \ln \frac{3N}{\delta} \left(\frac{5(6+\Lambda)r}{2(1-\Lambda)\varepsilon} + \frac{(1+\Lambda)v_T}{(1-\Lambda)\varepsilon^2} \right) \\
&\in \mathcal{O} \left(\frac{1}{1-\Lambda} \log \left(\frac{\log(r/\varepsilon)}{\delta} \right) \left(\frac{r}{\varepsilon} + \frac{v_\pi}{\varepsilon^2} \right) \right) ;
\end{aligned}$$

2. for DYNAMITE:

$$\begin{aligned}
\hat{m} &\leq 2T \ln \frac{3N}{\delta} \left(\frac{65r}{\varepsilon} + \frac{12v_T}{\varepsilon^2} \right) \\
&\in \mathcal{O} \left(T \log \left(\frac{\log(r/\varepsilon)}{\delta} \right) \left(\frac{r}{\varepsilon} + \frac{v_T}{\varepsilon^2} \right) \right) \\
&= \mathcal{O} \left(\log \left(\frac{\log(r/\varepsilon)}{\delta} \right) \left(\frac{r}{(1-\Lambda)\varepsilon} + \frac{\tau_{\text{rel}}v_{\tau_{\text{rel}}}}{\varepsilon^2} \right) \right) ; \&
\end{aligned}$$

3. for WARMSTARTDYNAMITE:

$$\begin{aligned}
\hat{m} &\leq 2 \lceil \frac{\ln \frac{1}{\pi_{\min}}}{\ln \frac{1}{\lambda}} \rceil + 2T \ln \frac{12N}{\delta} \left(\frac{65r}{\varepsilon} + \frac{12v_T}{\varepsilon^2} \right) \\
&\in \mathcal{O} \left(\frac{\ln \frac{1}{\pi_{\min}}}{\ln \frac{1}{\lambda}} + T \log \left(\frac{\log(r/\varepsilon)}{\delta} \right) \left(\frac{r}{\varepsilon} + \frac{v_T}{\varepsilon^2} \right) \right) \\
&= \mathcal{O} \left(\frac{\ln \frac{1}{\pi_{\min}}}{\ln \frac{1}{\lambda}} + \log \left(\frac{\log(r/\varepsilon)}{\delta} \right) \left(\frac{r}{(1-\Lambda)\varepsilon} + \frac{\tau_{\text{rel}} v_{\text{rel}}}{\varepsilon^2} \right) \right) .
\end{aligned}$$

Proof. The strategy here is to derive a sample size m' , dependent on $r, v_T, \varepsilon, \delta$, s.t. w.h.p., each algorithm will terminate after drawing a sample of at least m' traces. We then bound the total number of samples drawn over the course of this process, and make some substitutions to derive the result. For brevity, throughout this result we take $\eta \doteq \ln \frac{3N}{\delta}$.

We show the result for MCMC_{PRO}, using second absolute eigenvalue bound λ , as it immediately implies the corresponding results for DYNAMITE and WARMSTARTDYNAMITE.

We first show that, with high probability, the *empirical variance* is not much larger than the true variance, and thus with high probability, the variance-bounds used by MCMC_{PRO} are not loose. Let

$$\varepsilon_{v,1} \doteq \frac{5r^2\eta}{(1-\Lambda)m}, \quad \varepsilon_{v,2} \doteq \sqrt{\frac{(1+\Lambda)r^2v_\pi\eta}{(1-\Lambda)m}} .$$

Now, note that by corollary F.1.6,¹ we have for any sample size m that

$$\mathbb{P}(\hat{\mathbf{v}}_i \geq v_\pi + \varepsilon_{v,1} + \varepsilon_{v,2}) \leq \frac{\delta}{3N} .$$

We now consider the first iteration i such that $m_i \geq m'$, letting $m \doteq m_i$ and $\hat{v} \doteq \hat{\mathbf{v}}$. On line 12 of MCMC_{PRO}, we have (w.h.p.)

¹Note that here we take a lower-tail bound, rather than an upper-tail bound; the constants are identical and the result similarly follows from the Bernstein inequality.

$$\begin{aligned}
\mathbf{u}_i &\leq \hat{v} + \frac{(11 + \sqrt{21} + (1 + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m} + \sqrt{\frac{(1 + \Lambda)r^2\hat{v}\eta}{(1 - \Lambda)m}} \\
&\leq v_\pi + \varepsilon_{v,1} + \varepsilon_{v,2} + \frac{(11 + \sqrt{21} + (1 + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m} + \sqrt{\frac{(1 + \Lambda)r^2(v_\pi + \varepsilon_{v,1} + \varepsilon_{v,2})\eta}{(1 - \Lambda)m}} \quad \hat{v} \leq v_\pi + \varepsilon_{v,1} + \varepsilon_{v,2} \text{ (w.h.p.)} \\
&= v_\pi + \varepsilon_{v,2} + \frac{(16 + \sqrt{21} + (1 + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m} + \sqrt{\frac{(1 + \Lambda)r^2(v_\pi + \varepsilon_{v,1} + \varepsilon_{v,2})\eta}{(1 - \Lambda)m}} \\
&= v_\pi + \varepsilon_{v,2} + \frac{(16 + \sqrt{21} + (1 + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m} + \sqrt{\frac{(1 + \Lambda)r^2\eta}{(1 - \Lambda)m}} \sqrt{v_\pi + \varepsilon_{v,1} + \varepsilon_{v,2}} \quad \text{ALGEBRA} \\
&< v_\pi + \varepsilon_{v,2} + \frac{(16 + \sqrt{21} + (1 + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m} + \sqrt{\frac{(1 + \Lambda)r^2\eta}{(1 - \Lambda)m}} (\sqrt{v_\pi} + \sqrt{\varepsilon_{v,1}}) \quad \sqrt{a + b + 2\sqrt{ab}} = \sqrt{a} + \sqrt{b} \\
&\leq v_\pi + \varepsilon_{v,2} + \frac{(16 + \sqrt{5} + \sqrt{21} + (1 + \frac{\sqrt{5}}{2} + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m} + \sqrt{\frac{(1 + \Lambda)r^2v_\pi\eta}{(1 - \Lambda)m}} \\
&= v_\pi + \underbrace{\frac{(16 + \sqrt{5} + \sqrt{21} + (1 + \frac{\sqrt{5}}{2} + \frac{11}{\sqrt{21}})\Lambda)r^2\eta}{(1 - \Lambda)m}}_{\doteq \varepsilon_{v,3}} + 2 \underbrace{\sqrt{\frac{(1 + \Lambda)r^2v_\pi\eta}{(1 - \Lambda)m}}}_{\doteq \varepsilon_{v,4}}.
\end{aligned}$$

Substitution into the Bernstein bound (line 13), and similar algebra, gives us

$$\begin{aligned}
\hat{\varepsilon}_i &= \frac{10r\eta}{(1 - \Lambda)m_i} + \sqrt{\frac{(1 + \Lambda)\mathbf{u}_i\eta}{(1 - \Lambda)m_i}} \\
&\leq \frac{10r\eta}{(1 - \Lambda)m_i} + \sqrt{\frac{(1 + \Lambda)(v_\pi + \varepsilon_{v,3} + \varepsilon_{v,4})\eta}{(1 - \Lambda)m_i}} \quad \text{SEE ABOVE} \\
&< \frac{5(6 + \Lambda)r\eta}{2(1 - \Lambda)m_i} + \sqrt{\frac{(1 + \Lambda)v_\pi\eta}{(1 - \Lambda)m_i}} \quad \sqrt{a + b + 2\sqrt{ab}} = \sqrt{a} + \sqrt{b} \\
&\quad \text{ALGEBRAIC BOUNDS}
\end{aligned}$$

We terminate when $\hat{\varepsilon}_i \leq \varepsilon$, thus this implies sufficient sample size

$$m' \leq +\eta \left(\frac{5(6 + \Lambda)r}{2(1 - \Lambda)\varepsilon} + \frac{(1 + \Lambda)v_\pi}{(1 - \Lambda)\varepsilon^2} \right).$$

Now, due to the doubling geometric grid, we must have $m_i \in [m', 2m']$, thus we have

$$m_i \leq 2\eta \left(\frac{5(6 + \Lambda)r}{2(1 - \Lambda)\varepsilon} + \frac{(1 + \Lambda)v_\pi}{(1 - \Lambda)\varepsilon^2} \right).$$

Now, each step of the tensor-product chain $(\mathcal{M} \otimes \mathcal{M})$ requires two steps of \mathcal{M} , so we conclude (1) by noting that $4m'$ samples suffice.

Now, to get (2), note that in DYNAMITE, we take $T \doteq \lceil \frac{1+\Lambda}{1-\Lambda} \ln \sqrt{2} \rceil$, and thus $\lambda(\mathcal{M}^{(T)}) \leq \lambda^T \leq \frac{1}{2}$. The finite-sample bound then follows from (1) as DYNAMITE simply calls MCMCPRO (see line 24, multiplying total sample complexity by T , as each step in $(\mathcal{M}^{(T)} \otimes \mathcal{M}^{(T)})$ (i.e., the tensor-product trace-chain) takes T steps in \mathcal{M} for every step in the $\mathcal{M} \otimes \mathcal{M}$ chain of MCMCPRO. Finally, applying lemma F.1.7 yields the result.

We now show (3). In WARMSTARTDYNAMITE, the nonstationarity correction appears as a factor-4 increase on $\frac{1}{\delta}$ (see line 31), and an additive sampling cost for the warm start (see line 29). \square