

BROWN UNIVERSITY
Department of Computer Science
Master's Project
CS-94-M12

“Optimal Policies for Partially Observable Markov Decision Processes”

by

Anthony Rocco Cassandra

**Optimal Policies
for
Partially Observable Markov Decision
Processes**

by

Anthony Rocco Cassandra

Department of Computer Science
Brown University

Submitted in partial fulfillment of the requirements for the
Degree of Master of Science in the Department of
Computer Science at Brown University

May 1994



Professor Leslie Pack Kaelbling
Advisor

Optimal Policies for Partially Observable Markov Decision Processes

Anthony R. Cassandra*

Department of Computer Science

Brown University

Providence, RI 02912

May 17, 1994

1 Introduction

Decision making is an important task in everyday life. Making the correct decision in some circumstances can help utilize resources better or even be the difference between life and death. There are many applications in which automated decision making is desirable (automated stock trading) and others in which it is necessary (e.g., real time flight controls of a supersonic jet). Unfortunately, at this time, the process of human decision making procedures for many types of problems is difficult or impossible to formalize. Given this limitation, we would still like to have automated decision making perform as well as possible.

An often used and sometimes convenient way to make a machine aware of the dynamics of the domain in which it will be making decisions is to provide the machine with a model of the world. For some domains this is quite acceptable, since the dynamics of the world may be well defined. For other domains, it may be impossible to specify the dynamics fully. Even for

*This work was done jointly with Leslie P. Kaelbling and Michael L. Littman. Anthony Cassandra's work was supported in part by National Science Foundation Award IRI-9257592.

algorithms that do not rely upon being given a model [17], there is a structure imposed upon the environment which acts as a somewhat more abstract model. In this paper, we will be exploring a specific model-based scheme in which decisions need to be made. Even when we cannot assume we have the model, we can use techniques, [2], that allow us to approximate the model and then apply these model-based schemes. The many problems associated with such models will be outlined in a subsequent section. Throughout this discussion, the term *agent* will refer to the automated process that has to make decisions. A convenient example of an agent is that of an autonomous robot trying to survive in a real world environment. However, the agent can simply be a computer program such as one that does medical diagnosis. In this case, the model of the world might be based upon statistics and medical research.

The model we will explore assumes that the agent exists in a unique state at any time point. Furthermore, there are only a finite number of different possible states it can ever be in. For the robot example, these states can be different locations in the world or even the agent's state of knowledge at a given time. Time is assumed to pass in discrete increments (ticks) of some clock and the agent must choose some action to perform at each tick of the clock (it could be to do nothing). Each time the agent performs an action, it will move to a new state, though it is possible for it to remain in the same state. An alternative, but equivalent, view of the agent and world is one where the world exists in some state and the agent interacts with it.

If a machine (computer) is going to make decisions by itself, it needs to have some metric that it can use to differentiate between the good and bad choices. This is another potential trouble spot, because it is often difficult for people to articulate exactly what the differences are between good and bad choices. Nevertheless, we must do so if we desire to have the machine choose between a number of options. The choice in our model is to assume that each state has an associated reward for performing each possible action in that state. This reward reflects the immediate desirability of performing a particular action given that we are currently in some state.

These rewards do not necessarily have to be something the agent actually gets as it operates. We define these rewards as a way of assigning relative values to different states of the environment. Given these abstract values, the agent can then attempt to get to the locations that it knows have higher value. Although we can view the rewards as abstract entities, they might




Figure 1: Sample environment showing why one step greedy strategy is undesirable

well be grounded in some real phenomenon that the agent can perceive.

The model described thus far has dynamics as follows: The agent starts in a state i at time t , performs action a , receives an immediate reward for that action in that state and then finds itself in state j at time $t + 1$. The dynamics of an environment specified with such a model can be described with just a couple of tables. One table specifies the next state based upon the current state and the action chosen. The other specifies immediate rewards for all combinations of states and actions.

Although this model is much simpler than the one that is actually treated in this paper, it is useful to think about how an agent would use this information to make decisions at each time step. One fairly intuitive notion is that the agent wants to get as much reward as possible after all, these rewards are our way of informing the agent what state-action pairs we think are desirable. But exactly what is meant by “as much reward as possible?” Do we simply want to perform the action that gets us the most reward for this particular step? The answer is probably not. Figure 1 is an environment in which we can get a reward of 2 now, but no more rewards afterwards no matter what actions we decide to take. Also in this environment there is another action we can perform which will only give us a reward of 1 now, but after we take this action we will be able to get a reward of 5 for the next action (and even more with subsequent actions).

The above situation shows an example in which we would probably want to take into account the rewards we might receive in the future, not just the immediate rewards. If we knew in advance how many time steps (call

this number k) we had to make decisions for, then we could simply look at the model and decide which sequence of k actions (from a given starting state) will give us the maximum reward. Of course, there could be many possible sequences to consider, but there would always be a finite number of them. However, what if you don't know how many decisions the agent will have to make? In this case we need to consider the future rewards in some other manner. Dealing with this is a somewhat controversial topic. For this paper we will consider an infinite horizon with a discounting factor. The justifications for using this are given later, as are the drawbacks of discounting. Here we will just present the basic concept behind this approach.

The term *infinite horizon* is used to indicate that we will consider rewards for an infinite number of time steps. This is why we need a discount factor. If we merely try to sum rewards forever, we will get an ever increasing sum and all actions will look as if we can get an infinite amount of reward. By adding a discount factor we can cause these infinite sums to converge to a number which we can then use in a comparison to decide which actions are best.

Although we wouldn't expect the agent to make decisions forever, the infinite horizon with discounting is a good model for the instances in which we do not know for how long the agent will have to make decisions. If we assume that the agent will stop making decisions at each step with some probability, then it makes sense not to count rewards that could be received far into the future as much as rewards that could be received closer to the present. The discount factor does this by discounting rewards based upon how far in the future they could occur.

We now come to the point where we would like to make our model more robust. In the previous model description we assumed that each action will always have the same output given a specific starting state. Many domains of interest, however, do not possess this property. More specifically, an agent that is a mechanical device is subject to physical limitations and tolerances. Trying to move a robot 10 inches by applying the exact amount of voltage for exactly the proper time is a difficult thing to accomplish and even if this is overcome, there may be slippage in the wheels depending upon the floor surface, humidity, dirt, etc. So if you've modeled an environment in which one of the actions is for a robot to move 10 inches and another is to move 9 inches, the action to move 10 inches might not succeed. Although the results of the robot's actions are in a sense deterministic, the amount of knowledge




Figure 2: Sample environment showing stochastic actions

that we would have to have to predict it is overwhelming (e.g., where each dirt particle was, coefficient of friction of carpet for the current atmospheric conditions, etc.) Instead of getting bogged down in all these details, the standard approach is to model the actions probabilistically. Instead of saying that performing action a in state i will result in our next state being j , we say that 90% of the time this will happen, but that 10% of the time we could remain in state i . This is shown pictorially in Figure 2.

You can add these action probabilities into the model and get a more robust model. The nice part about this is that you can still figure out what the best action to take is. It is a little more complex than models without probabilistic actions, but solution techniques exist. The technical name for this type of model is a Markov Decision Process (MDP) and it is a very thoroughly explored model in mathematics, operations research, computer science, and related fields.

There is one extremely important element that we have assumed up to this point, but which is not always a good assumption. When we are choosing the best action to take, we first need to know what state we are in. How do we know what state we are in? Before we had probabilities associated with the actions, we always knew the next state after we took an action, so as long as we knew where we started we could tell exactly where we were. Now that we have introduced uncertainty in the outcome of actions, how do we know what state we actually did end up in? The easy answer is the one that is assumed in the previous model: Just look. The problem is that it is not always possible or realistic to assume you can tell exactly what state you are in at a given time. For instance, the robot example demonstrated the physical limitations of machines. These limitations apply not only to the actions they can perform, but also to the things they can observe in the world. If a robot is 18 inches from a wall and decides to move 10 inches, how will it know whether it is in the state of being 8 inches or 9 inches from the wall? Even if the robot has some sensors that can very accurately tell the distance

to the wall, there is still some probability that they could malfunction. The typical case is that mobile robots have crude sensors which are nowhere near 100% reliable. Even if there were perfect or near perfect sensors, the agent will rarely be able to see the whole world at once.

Dealing with uncertainty in observation is the main issue which this paper explores. The previously described MDP model can be made more general by introducing uncertainty in observations. This type of model is referred to as a partially observable markov decision process (POMDP) and is the focus of this work. Many existing techniques utilize the MDP model [7, 3] and the assumption of complete (reliable) observability is one of the major drawbacks of these techniques and those that build upon these. This paper explores existing techniques for POMDPs, presents a new algorithm and shows the usefulness and limitations of the POMDP model in general.

2 The Model

2.1 The MDP Model

A Markov decision process, MDP, is a fairly simple model with a fancy name. It consists of a finite set of states, a finite set of actions and a reward structure defined for each state-action pair. For a robot navigation problem, the states can be viewed as the location of the robot in the environment. Though most environments are really a continuum of states, it is often convenient to discretize these continuous spaces. The actions are the things the agent can do, such as move forward, turn left, pick up an object, etc. Associated with performing each action is an immediate reward. Because the effectiveness of performing an action can depend upon what state the agent is in, the MDP model actually assigns an immediate reward for each combination of states and actions. For example, if the robot is trying to get to a specific location, then performing the action *move forward* when it is next to and facing that location should result in a high reward. If the robot was directly in front of a descending stairway (and it did not know how to walk down stairs), then we would expect the action *move forward* to give much less reward.

Formally, an MDP is a quadruple, $\langle S, \mathcal{A}, \mathcal{P}, \mathcal{W} \rangle$. S is a finite set of states with state i denoted by s_i . \mathcal{A} is a finite set of actions with action i denoted by a_i or sometimes just a when we refer to a specific action. \mathcal{W} is

the reward structure, which is slightly more general than the one described previously. Here, w_{ij}^a represents the immediate reward the agent will get if it performs action a while in state s_i and moves to state s_j . Although this model is more general, we can simplify the reward structure to the one described previously by taking a weighted sum over all possible next states. The weight used in this sum is the probability that we actually move to each of the next states.

$$q_i^a = \sum_j p_{ij}^a w_{ij}^a \quad (1)$$

Alternatively, we could just describe the q_i^a rewards directly. The w_{ij}^a are the actual rewards that would be received, but the q_i^a are the *expected* rewards representing what we would expect to receive on average in the long term.

2.2 Policies

A *policy* is a function or mapping that tells the agent what action to execute. It typically depends upon the state that the agent is in. For instance, this current state could be its physical location. A policy completely specifies the appropriate action for each possible situation (state) that occurs in the model. A policy should not be confused with a plan. A plan is a sequence of actions to perform, and does not necessarily specify the appropriate action for each possible situation. In some sense, a policy is the most general possible plan, since it completely specifies the action to take for each possible state of the environment.

A policy can be either *deterministic* or *stochastic*. A deterministic policy is one that specifies a single action to take in each state. A stochastic policy specifies a number of possible actions to execute in each state.. In addition to a set of possible actions, the agent is given a probability distribution over the set of actions. The agent will then stochastically choose one action according to those probabilities. For this paper we will only concern ourselves with deterministic policies. Work in stochastic policies include [14]. Note that deterministic policies are really a subset of stochastic policies where the probability distribution assigns one action probability 1.

Policies (both deterministic and stochastic) can also be categorized as *stationary* or *non-stationary*. A stationary policy is independent of time; the same policy is applied regardless of when the policy is used. Thus, the policy only depends on the state of the agent and/or environment. A non-stationary

policy is dependent on time; the time the policy is to be used will affect what action will be taken. In this case, the time and state and/or environment determine what action the agent should execute.

Policies can also be grouped by the amount of memory required. As an agent moves about its environment performing actions, it is building up a history of its movements. In some situations it might be desirable to remember this history and perform future actions based upon this. A *memory-less* policy is one that uses no history at all. The choice of action depends only on an immediate situation of the agent. A *k*-memory policy chooses actions based upon the last *k* pieces of the agent's history.

There are also policies that require a finite amount of state, but the amount of state doesn't directly correspond to the last *k* pieces of the agent's history. For example, a policy that depends upon whether a button was ever pressed, requires only a single bit of information, but this single bit does not correspond to *k*-memory for any fixed *k*.

In this paper we will focus on deterministic stationary policies with memory. We will denote the entire space of these deterministic stationary policies as Δ where a specific deterministic stationary policy is denoted by δ . We will denote a deterministic non-stationary policy as a set of stationary policies, $\delta = \{\delta_0, \delta_1, \dots, \delta_t\}$ where $\delta_i \in \Delta$ is the stationary policy to use when there are *i* steps remaining.

2.3 Infinite vs. Finite horizon

In this paper we will discuss two ways to construct optimal policies for POMDP models. Both of these also apply to MDP models and, again for simplicity, it is in the MDP framework that we first present them.

We first consider problems in which the agent only has to make a known finite number of decisions. We refer to this class of problems as *finite horizon* problems; a *k-horizon* problem is one in which the agent will make decisions for *k* time steps. This period of *k* time steps is usually referred to as the agent's *lifetime* or the size of the horizon. The actual structure of a finite horizon policy consists of a sequence of deterministic policies, one for each time step (i.e., a deterministic non-stationary policy). The policy, δ_i , is a complete mapping from all states to actions, $\delta_i : \mathcal{S} \rightarrow \mathcal{A}$. The policy takes the current state of the agent as an argument, so the action specified by the policy for state s_i is denoted by $\delta_i(s_i)$.




Figure 3: Example showing need for non-stationary policies

$t = 2$		$t = 1$	
state	action	state	action
0	A	0	B
1	A	1	A
2	A or B	2	A or B
3	A or B	3	A or B

Table 1: Policies for horizons $t = 2$ and $t = 1$

To see why a non-stationary policy is necessary for the finite horizon case consider Figure 3. Assume that our horizon is $k = 2$ time steps where $k = 2$ means that there are exactly two time steps for which the agent needs to make decisions. The optimal policies for $t = 1$ and $t = 2$ are shown in table 1. Notice that each is a policy (complete mapping from states to actions), but that they are different. Since the model is small it should be easy to convince yourself that any other policy would not do as well for each of the two time steps. Therefore the true optimal policy for a $k = 2$ horizon is to apply the $t = 2$ policy first and the $t = 1$ policy next. Note that optimality here is being measured as total accumulated reward for the two time steps.

The other class of problems we consider are termed *infinite horizon* problems. Although we wouldn't expect any agent to have to make decisions forever, we use this type of model when we do not know in advance how long the agent's lifetime will be, which in many circumstances is more realistic than the finite horizon models. At first, attempting to consider an

infinite number of time steps might seem impossible. Although it does add some complexity, it can be and is often done, especially in the MDP model. The trick that needs to be incorporated is the addition of a discount factor to the rewards, so that rewards that are received further in the future are discounted more than rewards received closer to the present time. To show why the addition of the discount factor can make things work out nicely, imagine that we have some divine knowledge of exactly which state we will be in at each step and which action will be executed at each step. We will represent this knowledge as the functions $S(k)$ and $A(k)$ respectively. With this knowledge we can write the value, in terms of rewards, as

$$V_n(s_i) = \sum_{k=0}^{\infty} \beta^k q_{S(k)}^{A(k)}.$$

This formula is associating a value for starting in a given state, s_i , with n time steps remaining assuming we will know exactly what states we are in, $S(k)$, and which actions we perform, $A(k)$, at each time step. Here k represents the time when there are $n - k$ steps remaining.

This formulation, usually referred to as the value function, is a power series of the immediate rewards. Since we ultimately want to discuss and compare the value of different policies we would like this series to converge. A well known property of power series of this form is that they will converge for variables with an absolute value of less than one. In this case the variable of interest is the discount factor β . Since it is difficult to interpret a negative discount factor in our problems, we will only consider discount factors in the range: $0 \leq \beta < 1$. A main motivation for using a discount factor is that it allows us to talk about and compare these infinite sums. The niceness of the mathematical formulation and convergence are hardly appropriate justifications for using a discount factor; however, more motivated interpretations and justifications of the discount factor will be presented later (as will the problems with using a discount factor.) There are also techniques for solving undiscounted infinite horizon problems [7, 13], but this paper does not attempt to treat these problems.

You will recall that, for the finite horizon problem, a policy had the form $\delta = \{\delta_0, \delta_1, \dots, \delta_t\}$ since we (potentially) needed a different policy for each time step. We have thus far formulated our infinite horizon problem the same way, but trying to represent an infinite horizon policy this way leads to

some real problems. First, how could we ever calculate an infinite number of policies and, even if we figure out how to do that, how could we put it into a form that we could use that wouldn't take up an infinite amount of space? This is one place where the infiniteness can actually help us out. Imagine for a moment that you knew δ_i , that is, the optimal policy to use at time $t = i$ when there is an unbounded number of decisions to be made in the future. Well then, at time $t = i + 1$ there is also an unbounded amount of time left as there is for any particular time. Therefore, the policy that is best must be δ_i since we defined δ_i to be the optimal policy to use when there is an infinite amount of time remaining. Since our infinite horizon solution looks like this: $\delta = \{\delta_i, \delta_i, \delta_i, \dots\}$, we can simply refer to the policy as δ because at any time step we will always be using the same policy. This is what we previously termed a *stationary* policy, because at every time step the policy is the same. Solving the infinite horizon problem now looks a lot more promising, since we only need to search for and represent a single policy. Also, we have developed a way to ensure that all potential policies have a finite value which we can compare to one another.

2.4 Policies in MDPs

The goal of the techniques discussed in this paper is to derive a computational procedure for finding an optimal deterministic *policy* for a given POMDP model. Working toward that end, we first discuss policies in the MDP model and a traditional method (policy iteration) for finding optimal policies in these models. We will later generalize both of these to the POMDP case.

Since our aim is to find the optimal policy, we need a metric that gives us a measure of each policy's usefulness. This metric allows us to compare policies with one another. As a start in this direction, we define $V_n^{\delta_n}(s_i)$ to be the expected reward that the agent can accumulate in its lifetime if it executes policy $\delta = \{\delta_0, \delta_1, \dots, \delta_n\}$ for n steps. When it is currently in state s_i , there are n steps for which the agent needs to make decisions given that it executes actions according to the policy δ_m at time m .

$$V_n^{\delta_n}(s_i) = \sum_j p_{ij}^{\delta_n(s_i)} [w_{ij}^{\delta_n(s_i)} + \beta V_{n-1}^{\delta_{n-1}}(s_j)]$$

Since our MDP model has uncertainty in the outcome of its actions, we use the probabilities of moving to all possible next states in a weighted sum

of the rewards we would get if we actually moved into that state. The reward for moving to a state is the immediate reward received for the current action, current state and next state, w_{ij}^a , plus the value of state s_j with one less step remaining, $V_{n-1}^{\delta_{n-1}}(s_j)$. We discount the value of the next step by β , which will become necessary when we discuss the infinite horizon problems. For the undiscounted finite horizon problems we can just set $\beta = 1$. Factoring and substitution of formula 1 allows us to simplify this expression to

$$V_n^{\delta_n}(s_i) = q_i^{\delta_n(s_i)} + \beta \sum_j p_{ij}^{\delta_n(s_i)} V_{n-1}^{\delta_{n-1}}(s_j).$$

Note that for the non-stationary policies of finite horizon problems, this formula requires the value for the last $n - 1$ steps before we can compute the value for the n^{th} step. Thus, the value function is computed with dynamic programming. The values for $V_0^{\delta_0}(\cdot)$ depend only on the q_i^a values. For the infinite horizon problem, the policy, $\delta(\cdot)$, that gets executed at any time is the same. This results in a value function:

$$V^\delta(s_i) = q_i^{\delta(s_i)} + \beta \sum_j p_{ij}^{\delta(s_i)} V^\delta(s_j).$$

This is a system of $|\mathcal{S}|$ equations with $|\mathcal{S}|$ unknowns. As long as $0 \leq \beta < 1$, this system of equations will have a solution. Thus, whether we are solving the finite or infinite horizon problem, if we are given a policy (stationary or non-stationary), the value function formulas allow us to easily compute a metric for the policy. These value functions will allow us to compare different policies and eventually to prove that a given policy is optimal. These formulas show how we judge a policy's usefulness by the actions and states that the agent is led through.

In an MDP, any mapping from states to actions is a policy, but what we are concerned with is finding the optimal policy. Optimality is considered with respect to the value of the state-action pairs which are derived from the reward structure \mathcal{W} . We would like a policy that performs better than any other possible policy in terms of the value of all states the agent passes through. We will denote the optimal policy by δ^* and its associated value function by $V_n^*(\cdot)$. More formally, an optimal policy is one where, for all states, s_i , and all other policies, δ , $V^{\delta^*}(s_i) \geq V^\delta(s_i)$. It is non-trivial, [7, 3], that such a δ^* exists.

The distinction between the policy and the value function of a policy is a fairly important one. Some of the algorithms discussed in this paper construct both the value function and the policy, whereas others construct only the policy. Typically, it is easier to construct the policy than it is to construct the policy and its value function. However, given a policy the value function can be derived by the formulas above. Conversely, a policy can be constructed from the value function.

2.5 The POMDP Model

A partially observable Markov decision process, POMDP, is defined by a hextuple, $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \Theta, \mathcal{R}, \mathcal{W} \rangle$. \mathcal{S} is a finite set of states with state i denoted by s_i . \mathcal{A} is a finite set of actions with action i denoted by a_i or sometimes just a when we refer to a specific action. \mathcal{P} is the transition probabilities for each action in each state and defines the Markov process that the agent actually operates within, though it typically does not have access to this *core process*. p_{ij}^a denotes the probability that the agent moves to state s_j given that it was in state s_i and it just performed action a . Θ is a finite set of observations where θ_i denotes observation i . Although there is an underlying Markov process, the agent does not directly observe it. \mathcal{R} is the observation model, in which $r_{j\theta}^a$ denotes the probability that we observe θ when we are in state s_j at time t and when our last action (at $t - 1$) was a . \mathcal{W} specifies the immediate rewards, $w_{ij\theta}^a$ denotes the immediate reward received at time $t + 1$ for performing the action a in state s_i at time t , moving to state s_j at time $t + 1$, and making observation θ . These immediate rewards are essential in our quest for optimal solutions, since it is these values that inform us how useful different strategies are.

There is another notation that will prove convenient in our later formulations. We let q_i^a be the immediate reward of performing action a in state i . It can easily be derived from the immediate rewards $w_{ij\theta}^a$ weighted by the transition and observation probabilities for the different actions and observations:

$$q_i^a = \sum_{j,\theta} p_{ij}^a r_{j\theta}^a w_{ij\theta}^a. \quad (2)$$

This will simplify many of the formulas that appear later in the paper. We have also used the q_i^a notation in describing the MDP model. The use is the

same in that we would like it to stand for the immediate reward received for performing action a in state s_i , but the calculation differs since, in the POMDP model, we must factor in the uncertainty (probabilities) of the observations. Further use of the value q_i^a will not be confusing if you remember which model is being discussed.

There are actually two ways a POMDP problem can be formulated. In the first, which is the one presented above and used throughout this paper, observations are made after an action is taken. Alternatively, we could define a POMDP in which we first make an observation and then perform the action. These two formulations are essentially equivalent, since a problem in either one of these forms can be converted to a problem in the other. However, to avoid confusion we will always discuss the former since this is the model most often used in the existing literature.

2.6 Belief States

When we moved from a deterministic action model to a stochastic action model, we could still use the same basic machinery to solve the two types of problems. This was possible because we assumed the agent always knew where it was (or could be) in the finite state space. Thus, to find a policy the agent merely needed to decide what the appropriate action was for each state in this finite set. However, in the POMDP model, adding partial observability creates a real problem. If the agent can never know for sure what state it is in, how can it possibly know what action to take? One observation we can exploit is that the agent will be more likely to be in some states than others based upon the actions it has taken and the observations it has made. For instance, in Figure 4, even if we do not know whether we started in state $s = 0$ or $s = 1$, if we take action A and make observation $\theta = 1$, we are more likely to believe that we are in state $s = 3$ than $s = 2$, since the probability that we observe 1 in state $s = 2$ is so low.

It turns out that we can keep track of how likely we are to be in each of the states. We call this probability distribution over the state space the *belief state*. We denote the belief state $\pi = \{\pi_0, \pi_1, \dots, \pi_{|\mathcal{S}|}\}$ where $|\mathcal{S}|$ is the number of states in the model and π_i represents the probability that we are currently in state s_i . An important point (and drawback) about maintaining a belief state is that we must know the model of the world if we wish to compute it.




Figure 4: Belief state example

After each action and observation, we can update our belief state with a simple application of Bayes' rule:

$$\pi'_j = \frac{\sum_i \pi_i p_{ij}^a r_{j\theta}^a}{\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a}. \quad (3)$$

The new belief state will be a vector of probabilities computed according to the above formula. We define a belief transformation function:

$$\pi' = T(\pi|a, \theta).$$

The first formula is more enlightening since it more explicitly states that we just consider all possible ways the agent could have ended up in a state weighted by the probabilities of those ways actually occurring for that action-observation combination. The second formula merely simplifies the notation for subsequent formulas. An important result, as shown in [15], is that this belief state captures all the necessary information for any sequence of actions and observations. Therefore, by constantly updating this belief state, we are implicitly saving the relevant part of our past history of actions and observations. Since we cannot know our location with certainty, this belief state seems like the next best thing.

A diagram of the basic dynamics of a POMDP is shown in Figure 5. The action is generated as a function of the belief state, and the observation is generated by the environment. Recall that we are using the POMDP model in which observations follow the actions.




Figure 5: Control dynamics for a POMDP

2.7 Policies for POMDPs

In the regular MDP model, a policy is a mapping from states to actions and since the state space is finite, the policy and/or the value function are both easy to calculate and to represent (e.g., in a table). In the POMDP model, since we never know the true underlying state, our policies must now map belief states into actions. The number of belief states is infinite, and therefore, storing the policy or value function in tables is no longer feasible. This means we must find some other representation for both the policies and the value functions of policies.

For now, we only consider the finite horizon case of a POMDP problem since this is the focus of much of this paper. A later section will discuss infinite horizon solutions.

Imagine that we need to solve the finite horizon for a POMDP problem with the horizon $k = 1$. In this case, the agent gets to choose to execute a single action and no more. Here, the observation it makes after that action doesn't really matter, because once it chooses an action the game is over. In order for this problem to be well defined, we need to define a terminating reward for the process. This terminating reward represents the value of the agent ending up in each of the environment states when $t = 0$ (i.e., there is no time left to take actions). We will denote the terminating reward for state s_i as q_i^0 . Recall that there are also immediate rewards for each state-action pair. For $k = 1$, the agent will accrue one of these immediate rewards (depending on the state it was in and the action it chose to execute) plus the terminating reward for the state it ended up in. However, neither the starting nor ending state will be directly observable to the agent. All the agent will have access to is the probability distribution over the states (i.e., belief state) it starts out in. Our task is to derive a policy that obtains, based upon the belief state, the maximum expected reward for a single action.

Without knowledge of the actual underlying states of the agent, it might seem peculiar that we could ever find the best action to take. However, since the belief state represents the likelihood of being in each state, we can discuss the expected accumulated rewards by taking an expected value. For instance, assume that the agent's last belief state is π' , then the expected terminating reward, or the value of ending up (at $t = 0$) in belief state π' is

$$V_0(\pi') = \sum_i \pi'_i q_i^0. \quad (4)$$

Now that we know the value of ending up in a particular belief state, we move on to discuss finding the best action for each belief state, π , given that we will only be executing one action. We know that the ending belief state will be $\pi' = T(\pi|a, \theta)$. The agent will have control over the action a , but not over the observation θ . Because it cannot control the resulting observation, the agent must use the model of observation probabilities to weigh all the possibilities:

$$V_1^*(\pi) = \max_{a \in \mathcal{A}} \left[\sum_i \pi_i [q_i^a + \sum_{j, \theta} p_{ij}^a r_{j\theta}^a V_0(T(\pi|a, \theta))] \right].$$

By substituting for $V_0(\cdot)$ and using formula 3 for $T(\pi|a, \theta)$, after some cancellation we get

$$V_1^*(\pi) = \max_{a \in \mathcal{A}} \left[\sum_i \pi_i q_i^a + \sum_{i, j, \theta} \pi_i p_{ij}^a r_{j\theta}^a q_i^0 \right].$$

With this formula we can determine the best action to take for any given belief state by performing the maximization shown. Thus, the optimal policy for $t = 1$ can be defined as

$$\delta_1^*(\pi) = \arg \max_{a \in \mathcal{A}} \left[\sum_i \pi_i q_i^a + \sum_{i, j, \theta} \pi_i p_{ij}^a r_{j\theta}^a q_i^0 \right].$$

Although the belief space is continuous, we are able to represent a policy over its entirety with this small finite formula, though there is some extra work required to perform the maximization. Notice here that the policy was derived from a value function. We show, shortly, that if you can represent the value function in this nice way then you also have a way to represent the policy since it is just the argument used in the maximization procedure. Indeed, this is exactly what we show for the general case next.

2.7.1 Value functions for POMDPs

To move to the general k -horizon case of the value function, we use the same ideas demonstrated for the 1-horizon case. First, let us assume we know the optimal value function when there are $n - 1$ time steps remaining, $V_{n-1}^*(\cdot)$. The basic form is the same as the MDP, but here we have a few extra complications to consider. First, we must take the weighted average of all the states, with the weighting of a state's value coming directly from the belief state. Second, we need to factor in all the possible observations we could possibly make and take a weighted sum of the values using the observation probabilities. Lastly, when we consider the value for the next step we must transform the belief vector based upon the current belief state, the action and the observation. Folded together and taking the maximum for all possible actions we get

$$V_n^*(\pi) = \max_a \left[\sum_i \pi_i \sum_j p_{ij}^a \sum_{\theta} r_{j\theta}^a \left(w_{ij\theta}^a + V_{n-1}^*[T(\pi|a, \theta)] \right) \right].$$

Just as we did with the MDP model we can use factoring and substitution of formula 2 for the q_i^a values to simplify this to

$$V_n^*(\pi) = \max_a \sum_i \pi_i q_i^a + \sum_{i,j,\theta} \pi_i p_{ij}^a r_{j\theta}^a V_{n-1}^*[T(\pi|a, \theta)]. \quad (5)$$

This formula can appear rather complex at first. With its recursive definition and abundant summations it would seem that trying to solve for such a value function would get messy. However, if $V_n^*(\pi)$ is piecewise linear and convex (we show shortly that it is) it can be written much simpler as

$$V_n^*(\pi) = \max_k \sum_i \pi_i \alpha_i^k(n)$$

for some set of vectors $\alpha(n) = \{\alpha^0(n), \alpha^1(n), \dots\}$. This is the crucial point for understanding the remainder of the paper and the algorithms. Unfortunately, this is also the easiest point to get confused. The simultaneous introduction of terms *piecewise linear and convex* and the $\alpha_i^k(n)$ notation can easily make your head spin at first. Going through each piece slowly, a piecewise linear function is one that is composed solely of line segments or hyperplanes. There may be many linear segments that together combine to make up this function,

but at any point there is one line segment (hyperplane) that covers that point. A linear function of variables x_i can always be written with coefficients α_i as

$$\sum_i \alpha_i x_i = \alpha_0 x_0 + \alpha_1 x_1 + \dots + \alpha_N x_N.$$

A piecewise linear function can consist of one or more linear segments of this form and so we superscript the coefficients of each to indicate which linear segment they come from so that the k^{th} linear segment will have the equation:

$$\sum_{i=0}^N \alpha_i^k x_i.$$

We will refer to the vectors $\alpha^k = [\alpha_0^k, \alpha_1^k, \dots, \alpha_N^k]$ as α -vectors in the remainder of the paper and each vector will represent the coefficients of one of the linear pieces of a piecewise linear function. These piecewise linear functions will be the value functions for each step in the finite horizon POMDP problem. As a result it will be convenient to index the α -vectors by the number of time steps remaining where $\alpha^k(n)$ represents one of the linear pieces of the value function $V_n^*(\cdot)$.

We now develop the value functions for the finite horizon POMDP problem and show that they are indeed piecewise linear. With this result we can see how to represent a value function over a continuum of points (i.e., the belief space) with a finite number of items (i.e., α -vectors). This result is also shown in [16, 15]. This simpler representation is the key element to the POMDP solution techniques. Remember that the agent will only know the belief state, which is a probability distribution over the states of the model. With the MDP model there was a finite number of states and because the agent knew which state it was in, it could store the value function as a table. In the POMDP framework, the belief states are continuous. There are an uncountably infinite number of belief states which makes representing the value function by tables impossible. With this $\alpha(n)$ vector representation we have found a way to represent the value function for this continuous belief state space.

As an example Figure 6 shows a value function over an $|\mathcal{S}| = 2$ belief simplex. It consists of 4 vectors, which in this case define 4 lines in a plane. Since our value function takes a maximum dot product of these for all belief




Figure 6: Sample piecewise linear and convex value function for $|\mathcal{S}| = 2$

points, the value function is the upper-most line segment at each point. This maximization always results in a convex function of this type for $|\mathcal{S}| = 2$. This convexity generalizes for higher dimension state spaces as will be shown shortly.

Our value function above is defined recursively, but our vector representation of the value function is not, it merely shows the representation as if we knew the various α -vectors for that value function. To be useful, we would like a formulation that will define the vectors for time step n based upon the vectors from time step $n - 1$. This is achieved by simple substitutions and the introduction of an indexing function to keep track of how we use the $n - 1^{st}$ vectors, $\alpha(n - 1)$. This will also demonstrate the inductive step in the proof that the value function for a finite horizon POMDP problem is piecewise linear and convex.

Our inductive hypothesis is that $V_{n-1}^*(\cdot)$ is piecewise linear and convex and thus can be written with some set of α -vectors, $\{\alpha^0(n - 1), \alpha^1(n - 1), \dots, \alpha^m(n - 1)\}$ as

$$V_{n-1}^*(\pi) = \max_k \left[\sum_i \pi_i \alpha_i^k(n - 1) \right].$$

However, when we are interested in the value for $t = n - 1$, we will actually be dealing with a transformed belief vector so that we are really interested

in

$$V_{n-1}^*(T(\pi|a, \theta)) = \max_k \left[\sum_i \pi'_i \alpha_i^k (n-1) \right].$$

Substituting for the transformed belief vector we get

$$V_{n-1}^*(T(\pi|a, \theta)) = \max_k \left[\frac{\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k (n-1)}{\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a} \right].$$

Every evaluation of this formula, for a specific π , a , and θ , can result in a different $\alpha^k(n-1)$ being selected (from the maximization process. Since we will soon need to put this definition of $V_{n-1}^*(\cdot)$ into the formula for $V_n^*(\cdot)$ we need a way to keep track of which of the $\alpha^k(n-1)$ was the maximum. To this end we define the function $\iota(\pi, a, \theta)$ to be the index of the α -vector that maximizes $V_{n-1}^*(T(\pi|a, \theta))$.

$$\iota(\pi, a, \theta) = \arg \max_k \left[\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k (n-1) \right] \quad (6)$$

Notice that the π argument is actually being transformed by the function $T(\pi|a, \theta)$ and that the denominator is not relevant for the maximization procedure since it will be the same for all k . This is how the $V_{n-1}^*(\cdot)$ formula looks with our new index notation:

$$V_{n-1}^*(T(\pi|a, \theta)) = \frac{\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi, a, \theta)} (n-1)}{\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a}.$$

The next step is to substitute this into our recursive definition of the value function given in formula 5:

$$V_n^*(\pi) = \max_a \sum_i \pi_i q_i^a + \sum_{\theta} \frac{[\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a] [\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi, a, \theta)} (n-1)]}{\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a}$$

After some summation manipulation, cancellations and factoring we are left with

$$V_n^*(\pi) = \max_a \sum_i \pi_i \left[q_i^a + \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi, a, \theta)} (n-1) \right]. \quad (7)$$

The terms are arranged this way to show that we still have the value function in terms of linear segments. In this case the vector portion is defined by the inner bracketed quantity:

$$\alpha_i(n) = q_i^a + \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{(\pi,a,\theta)}(n-1).$$

This completes the inductive step of the proof that the finite horizon POMDP problem value function is piecewise linear and convex. To complete the proof, we use the base case referred to earlier. The terminating rewards we defined in formula 4 show that the value function for $t = 0$ is also piecewise linear and convex. This proof was first shown in [16] as was the observation that the value function could be represented by a finite set of vectors.

If we know all the $\alpha^k(n-1)$ vectors for the $V_{n-1}^*(\cdot)$ value function, then by simple calculations using formulas 6 and 7 we can compute the above term to arrive at a linear segment, $\alpha(n)$, of the $V_n^*(\cdot)$ value function.

Another very important thing to notice about this new $V_n^*(\cdot)$ formulation is that for a given belief state π , in order to find its value we must do a maximization over all possible actions. According to the form above each action will result in a new vector. Each of these α -vectors for time step n has an associated action that gave rise to it. When we perform these maximizations in forming the $V_n^*(\cdot)$ value function we should also keep track of the specific action, a , that gave rise to it for each $\alpha(n)$ vector. With this combination of vector-action pairs we will have a compact way to represent the optimal policy for the finite horizon POMDP problem, namely just perform the maximization over all $\alpha(n)$ and take the associated action of that vector.

2.7.2 Geometric Interpretation of Value Function

The formulas and language of the previous section can serve to obscure the nice properties of the value function. In this section we attempt to improve the intuition behind the previously mention concepts of piecewise linear and convex and to provide a geometric interpretation of the value function. This representation is most valuable when trying to describe and understand the various algorithms discussed in this paper and in any other POMDP paper.

Note that with all of our previous definitions of value functions we have assumed that higher values are better than lower values. This stems from the fact that we chose the $w_{ij\theta}^a$ to represent rewards. However, the $w_{ij\theta}^a$




Figure 7: Simplex of belief space for $|\mathcal{S}| = 2$

could be defined to represent costs. This formulation is found in some of the existing POMDP literature and results in the maximization procedures being replaced by minimization procedures. Other than this change the rest of the formulation remains unchanged except that typically the value functions $V^*(\cdot)$ are referred to as cost functions, $C^*(\cdot)$. We will continue to adopt the reward/value interpretation.

For all the following examples and diagrams we start with the space of beliefs. Since beliefs are probability distributions over the underlying model states we see the belief simplex for the 2 and 3 dimensional cases in Figures 7 and 8 respectively.

The value function is a function of the number of components in the belief state and is thus a function in $|\mathcal{S}| + 1$ space where $|\mathcal{S}|$ is the number of states in the model. Because we would like to demonstrate these geometric ideas for the cases $|\mathcal{S}| = 2$ and $|\mathcal{S}| = 3$, we need a slightly different representation of the belief simplexes. This will prove to be both easier to understand as well as easier to produce on the printed page. The key to this other representation lies in the observation that one component of the belief state can always be omitted since the components (being a probability distribution) must sum to one. With this we can represent the 2 dimensional belief space as a single line or axis as in Figure 9. Here the distance from the left axis is the first component π_0 and the distance from the right can be used to compute the second component π_1 . Our value function can now be represented by drawing it above this line. In Figure 9 we see a value function over the belief




Figure 8: Simplex of belief space for $|\mathcal{S}| = 3$

space which consists of two linear segments. Since we assume that we are concerned with getting the maximum reward the actual value function is always the larger of the two linear segments.

Notice also that the value function for this case imposes a partition of the π_0 axis and thus is a partition of the $|\mathcal{S}| = 2$ belief space. This partition divides the belief space according to which of the linear segments (α -vectors) is maximum. The partition is shown with small vertical lines on the axis.

We turn the $|\mathcal{S}| = 3$ case into a two dimensional triangle as shown in Figure 10. In reality, the belief space is a two-dimensional triangle lying in three-dimensional space as shown in Figure 8, but in Figure 10 we view it simply as a two dimensional object where the value of a point in the belief space can be obtained by the perpendicular distance to sides of the triangle each distance representing a different component of the belief vector.

The value function for and $|\mathcal{S}| = 3$ problem can be thought of as a three dimensional surface lying above this triangle as shown in Figure 11. The surface is comprised of hyperplanes (i.e., it is piecewise linear and convex) and we can also view it as imposing a partition on the belief space. The borders of the partition are the projections of the intersection lines of the




Figure 9: Sample value function for $|\mathcal{S}| = 2$




Figure 10: Example belief for $|\mathcal{S}| = 3$, $\pi = [0.1, 0.3, 0.6]$




Figure 11: Sample value function for $|\mathcal{S}| = 3$

hyperplanes.

As shown previously, the value function is piecewise linear and convex. Geometrically this means that the value function is composed solely of straight lines (for $|\mathcal{S}| = 2$), planes (for $|\mathcal{S}| = 3$) or hyperplanes (for $|\mathcal{S}| > 3$). Each α -vector represents the coefficients for one of these lines or planes. The highest plane (or line) at a point is the α -vector that represents the value function for this point. A key point with these pictures is that each α -vector defines a region over the simplex. These regions represent a set of belief states. Since there only a finite number of α -vectors, there are only a finite number of regions defined over the simplex.

2.8 Alternate Value Function Interpretation

There is another way in which we could go about determining the best policy for a k -horizon POMDP problem. This method comes from the observation that at each step there are only a finite number of things we can choose to do




Figure 12: Example policy tree

and a finite number of things we can possibly observe. Since we have a finite horizon, we could actually enumerate each possible policy we could follow as a decision tree. As an example consider the tree shown in Figure 12. This is one of the potential policy trees for the case where we have 4 possible actions, 3 possible observations and a finite horizon of $k = 3$. The nodes represent an action decision and the branch taken from a node depends on the observation made.

Knowing the initial belief state allows us to calculate the expected value for each of these trees. To do this we merely take an expectation for the immediate rewards (weighed by the belief state) for the action in the current belief state and then add the expectation for the subsequent actions. This requires us to compute the transformed belief states and weight each possibility by the probability of each observation. There are a finite number of trees and so we could just enumerate all of them. This would actually specify a policy since for a given belief state, we could evaluate all the possible trees to decide which is the best action to take. However, there are 4^{13} possible policy trees for this example since each of the nodes represents a decision point which could have any one of 4 possible values. The number of nodes in the tree will be determined by the formula

$$\sum_{i=0}^k |\Theta|^i = \frac{|\Theta|^{k+1} - 1}{|\Theta| - 1}.$$

and thus the number of possible trees is

$$|\mathcal{A}|^{\frac{|\Theta|^{k+1} - 1}{|\Theta| - 1}}.$$

This is too many to expect to be able to generate. However, many of these policy trees will be useless in the sense that there will be no belief state for which they will be the best one. If we could somehow only generate the useful trees we would have reduced the complexity greatly. An additional savings could be made if we could collapse similar parts of the useful policy trees into the same sub-tree.

The the previous sections we have seen a way to construct a piece of the value function $V_n^*(\pi)$ given that we know the value function $V_{n-1}^*(\cdot)$. However, this only gives us the $\alpha(n)$ vector in the $V_n^*(\cdot)$ for a single specific belief state π . In order to create the entire $t = n$ value function we would need to apply formulas 6 and 7 at every single belief point in Π . There are too many of these points (an uncountably infinite number) for this naive algorithm (of applying these formulas to all points) to work. However, there will only be a finite number of possible $\alpha(n)$ vectors for any given horizon k . We will see that these $\alpha(n)$ vectors have a direct correspondence with the useful policy trees.

3 POMDP Solutions

We will first explore the structure of the solutions to finite horizon POMDP problems and their relationship to the solutions to infinite horizon problems. In the next section we will present a number of algorithms for computing these solutions. The general method for the finite horizon problem iterates over the time steps. We first derive the optimal policy for the case when there is only one time step remaining. Given this solution we now find the optimal policy for two remaining time steps and this process is repeated until we have arrived at the solution to the the finite horizon size we are interested in. It is essentially a dynamic programming problem where the optimal solution for time step t is phrased in terms of the optimal solution for time step $t - 1$. For the finite horizon case we will actually be constructing the value function and the policy. For this, we will utilize the recursive formula for the value function described previously, and repeated below. Recall that the value function at each step can be represented by a set of vectors:

$$V_n^*(\pi) = \max_a \sum_i \pi_i \left[q_i^a + \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{(\pi,a,\theta)}(n-1) \right].$$

With this form we see that if we know how to compute the value function and optimal policy for one time step (given the value function for the previous time step) then we know how to solve the finite horizon problem for any horizon since we can continually use the current output as the input to another iteration. Therefore when we discuss these finite horizon methods we will only need to discuss the task of computing the value function, $V_n^*(\cdot)$, for one step given that we have the previous value function, $V_{n-1}^*(\cdot)$ as input.

One of the unpleasant parts of the finite horizon solution is that our resulting policy is potentially very large. If we decide to solve a k -horizon policy and perform the above mentioned iterative procedure, then the final policy and value function we get out (the one for k remaining time steps) is only a part of our answer. Each policy and value function at each step is actually part of the correct k -horizon non-stationary policy. This becomes clearer when we consider how the agent would use this policy knowing it only has k time steps remaining. At first there are k time steps remaining and so it can look at the policy it generated for this case (the last one computed from above). This will tell it the optimal action to take, but immediately after it executes the proper action there are $k - 1$ time steps remaining. The policy it just used is not relevant for this case, however, in computing the k horizon policy we had to compute the $k - 1$ horizon policy. It is this policy that the agent needs to use now. This same argument repeats and so all of the intermediate work in solving the k horizon policy must be retained.

As a quick aside to try and head off a possibly confusing part of the notation we need to emphasize exactly what the t time step represents. Here the t represents how many more decisions (actions) will the agent have left to take. It is easy to confuse this t with how many steps the agent has taken. This becomes especially confusing when one starts to think of the agent executing these computed policies. When we think of the agent carrying out these policies we would expect time to proceed forward (i.e., $t, t + 1, \dots$), but the policies it will execute will be proceeding in decreasing order of t since after each step the agent has one less decision (action) to make.

There is a certain class of POMDP problems for which we can use this iterated finite horizon approach to solve the discounted infinite horizon problem. This doesn't say that all infinite horizon problem can be computed this way, but for a certain class (called *finitely transient* and discussed later) we can perform the finite horizon solution technique until we find two consecutive time steps that have exactly the same answer. When this occurs, we know

that further iterations will continue to give this exact same answer no matter how large the horizon is made. This follows from the fact that at each step you derive the optimal solution based solely on the previous optimal policy. If the previous optimal policies are the same for any two time steps, then the computed optimal policy for these two times steps must be the same. The pleasant part of the infinite horizon case is that, unlike the finite horizon, we can discard all the intermediate results. It is only the final repeating optimal policy that needs to be retained since the optimal policy is a stationary one. Of course in order for the value function to converge, we must add a discount factor to the value, otherwise if we merely kept accumulating the rewards from all the previous time steps, the value function could grow without bound and thus never converge. The addition of a discount factor β is a simple change to the value function:

$$V_n^*(\pi) = \max_a \sum_i \pi_i \left[q_i^a + \beta \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi,a,\theta)}(n-1) \right].$$

We will return to discuss the infinite horizon case in more detail later, but we thought it instructive to draw the connection between the finite and infinite horizon problems here.

An important point which has not yet been mentioned, though might already be obvious, is the use of the value function. For one step of the iterated solution procedure we assume that we have the previous value function $V_{n-1}^*(\cdot)$ with which to construct the new value function $V_n^*(\cdot)$. Our value function, $V_{n-1}^*(\cdot)$, will take the form of a set of vectors $\alpha(n-1) = \{\alpha^0(n-1), \alpha^1(n-1), \dots, \alpha^m(n-1)\}$ where each one represents one of the linear pieces that comprise this value function. As we saw previously we can use the recursive value function equation to give us a new $\alpha(n)$ vector, but the key point here is that in order to construct this α -vector we must first know a particular belief state π . Using this the value function formula for any valid belief state will give you an actual linear piece of the $V_n^*(\cdot)$ value function, however, it does not let you make any claims about the region (set of belief states) for which this is optimal (other than it includes the belief state we used to generate it).

The bottom line here is that given a point in the belief space we can find a piece of the value function, but the problem is that there are just too many (uncountably infinite) belief states for us to crank through the formula this

way. One solution is how to just choose points at random in the state space to use in the formula. One would hope that with enough points we wouldn't miss any of the true α -vectors that comprise the value function. This turns out to be a poor solution for many reasons. The first problem is to decide how many points to choose and this is directly related to the dimensionality of the belief space. The larger the belief space the more points you would expect to have to choose. The worst problem (from the theoretical perspective) is that we are never guaranteed to find the true value function (and optimal policy) with this method since we never know if some new point would generate another α -vector. We have run this approach on a number of problems and found that empirically it performs poorly even for small problems.

The key to finding optimal policies (i.e., true value functions) is to develop a systematic and terminating algorithm to explore the entire continuous space of beliefs. A finite method is possible since we know that there will only be a finite number of α -vectors to discover. The method used in exploring the belief space is where most of the algorithms discussed in this paper differ.

3.1 Structure of POMDP Solutions

The previous sections have discussed the mathematical foundations for solutions to POMDP problems. In this section we look at the solutions from a practical view point. Although we do not discuss solution procedure until the next section, here we will begin to show how these solutions can be utilized. We discuss the finite horizon problem which will then lead us directly to the infinite horizon problem.

3.1.1 Finite Horizon

The raw output of the algorithms we later discuss is a set of vectors, $\alpha(t) = \{\alpha^0(t), \alpha^1(t), \dots, \alpha^m(t)\}$, one set for each decision time, t . This refers to the finite horizon case where there is a non-stationary policy. In addition, each of these vectors will have a single associated action with it. To decide what action to take (i.e., use the policy) we take the set of vectors for the current time period and find the vector in this set which gives the largest dot product with the current belief state. The action associated with this vector is the action that should be executed. In this way, the policy can be stored as the




Figure 13: Example POMDP environment

vectors that comprise the value function along with an associated action for each vector.

In order to better motivate the discussion that follows we introduce a simple POMDP example shown in Figure 13. In this example, imagine that we are standing in front of two closed doors. Behind one of the doors is a tiger and behind the other is a large reward. If the door with the tiger is opened, then a large penalty is received (presumably as some amount of bodily injury). Most likely, we would prefer to open the door with the large reward instead. Aside from opening one of the two doors we have another action we can take, namely to *listen*. We choose to listen hoping that we will be able to hear which door the tiger is behind, but listening is not free, there is a cost associated with it. Unfortunately, listening is not entirely accurate and there is a possibility that we will get the wrong information when we listen.

The complete set of parameters for this problem is given in Tables 2 through 4. State s_0 represents the state of the world when the tiger is on the left and s_1 represents the state when the tiger is on the right. Action $a = 0$ is the action of listening and $a = 1$ and $a = 2$ are, respectively, the actions of opening the left door and opening the right door. The reward for opening the correct door is $+10$ whereas there is a penalty of -100 for choosing the door with the tiger behind it. The cost of listening is -1 . The two observations possible are to hear the tiger on the left or right. We have constructed this problem so that immediately after opening a door, the problem resets with the tiger randomly placed behind one of the two doors. Aside from resetting

p_{ij}^0	$j = 0$	$j = 1$
$i = 0$	1.0	0.0
$i = 1$	0.0	1.0

p_{ij}^1	$j = 0$	$j = 1$
$i = 0$	0.5	0.5
$i = 1$	0.5	0.5

p_{ij}^2	$j = 0$	$j = 1$
$i = 0$	0.5	0.5
$i = 1$	0.5	0.5

Table 2: Example transition probabilities

the tiger, the parameters have been constructed so that the belief state will be reset to the uniform distribution, $\pi = (0.5, 0.5)$, when either door is chosen, regardless of the belief state at the time the door is opened.

If we are given the policy (i.e., the vectors and their associated actions) for each time step, then we can determine the proper action to take, but only if we have the current belief state. Let us begin with the policy for the time step $t = 1$, which is the policy when the agent will only get to make a single decision. There are three different actions the agent might choose to execute, open the right or left door or listen. If our belief state was such that we had a high probability of being in the state “tiger-left”. then we would imagine that the best thing to do would be to open the right door. In the symmetric case of having a high probability of “tiger-right”, we would expect opening the left door to be the best action to do. But what if we are highly uncertain about where the tiger is at this time? The best thing to do would probably be to listen. To see this notice that guessing wrong will cost us a penalty of -100 , whereas guessing correctly only rewards us with 10 . When we have no strong beliefs either way we would expect to guess wrong as often as we guess right. Thus, the expected reward would seem to be $\frac{-100+10}{2} = -55$. Listening always costs us -1 and so, in the long run we would expect listening to cost us less than opening one of the doors in the

$r_{j\theta}^0$	$\theta = 0$	$\theta = 1$
$j = 0$	0.85	0.15
$j = 1$	0.15	0.85

$r_{j\theta}^1$	$\theta = 0$	$\theta = 1$
$j = 0$	0.5	0.5
$j = 1$	0.5	0.5

$r_{j\theta}^2$	$\theta = 0$	$\theta = 1$
$j = 0$	0.5	0.5
$j = 1$	0.5	0.5

Table 3: Example observation probabilities

$w_{ij\theta}^0$	$\theta = \{0, 1\}$
$i = 0, j = \{0, 1\}$	-1
$i = 1, j = \{0, 1\}$	-1

$w_{ij\theta}^1$	$\theta = \{0, 1\}$
$i = 0, j = \{0, 1\}$	-100
$i = 1, j = \{0, 1\}$	+10

$w_{ij\theta}^2$	$\theta = \{0, 1\}$
$i = 0, j = \{0, 1\}$	+10
$i = 1, j = \{0, 1\}$	-100

Table 4: Example rewards




Figure 14: Tiger example policy for $t = 1$

case where we are uncertain about the tiger's location.

The previous paragraph gives the intuitive argument for what the policy for $t = 1$ should be and it turns out that this is exactly the optimal policy for this case. Figure 14 shows this policy pictorially. In this figure we show each of the vectors as a node in a graph (which currently has no edges). The actual vectors are shown above each node and below each node is the belief interval ¹ over which it is the best vector. These three vectors cover the entire belief space (i.e., partition the belief space) and each has an associated action (shown inside the node circle) and so this specifies a policy for $t = 1$.

We now move on to the case where the agent has two decisions to make, $t = 2$. This is shown in Figure 15 and has an interesting property; it always chooses to listen. There is a logical reason for this. If the agent were to open one of the doors at $t = 2$, then, due to the way the problem has been formulated, the tiger is randomly placed behind one of the doors and the agent's belief state will get reset to $\pi = (0.5, 0.5)$ (i.e., it has no information about where the tiger is). So after opening a door the agent is left with no information about the tiger and one action left to take. We just saw that for the case where $t = 1$ and $\pi = (0.5, 0.5)$ the best thing to do was to listen. Therefore no matter what happens, whether it opens a door or listens, one of the two decisions will result in a listen action. In a way, the agent knows this and so chooses to listen first since this will give it information about the tiger. This way the agent will be better informed (since it has listened once) when it makes its next decision.

The other interesting aspect of Figure 15 is the number of vectors that exist. Although it will always choose to take the action listen, there are several vectors that have that action associated with it. These vectors are actually partitioning the belief space into pieces that have some structural similarity. The similarity the belief states in a partition element share is

¹The belief interval is in terms of the component π_0 only since π_1 can be found by subtracting π_0 from 1.




Figure 15: Tiger example policy for $t = 2$




Figure 16: Belief state mapping from $t = 2$ to $t = 1$

that when they are transformed, via $T(\pi|a, \theta)$ in Formula 3, the resulting belief states will all lie in the same partition defined by the policy for $t = 1$. In other words, every single belief state in a particular partition for $t = 2$ will, for the same action, a , and observation, θ , be transformed to exactly the same partition imposed by the policy for $t = 1$. We have shown this relationship pictorially in Figure 16. Notice that the edges only need to be labeled with the observations since the action used in the transformation of the belief state is dictated by the node corresponding to that belief state.

Interestingly, the optimal policy for $t = 3$ also consists solely of nodes with the listen action. The nature of the problem as it is set up is such that if we start from the uniform belief state, $\pi = (0.5, 0.5)$, listening once cannot give us enough information (i.e., change our belief state enough) to make choosing a door more rewarding than listening. As mentioned before, choosing a door will always reset our belief state to this uniform distribution. Therefore, just like the argument for $t = 2$, if we chose a door with $t = 3$ steps remaining, we would be guaranteed to listen for the next two steps. Since the agent knows that this is the case, it would prefer to do the two listening steps up front since that will give it the most information when it finally decides what to do at $t = 1$. You can see in Figure 16 that starting with belief state $\pi = (0.5, 0.5)$ we will choose to listen when $t = 2$ ($\alpha^2(2)$), and then, no matter what we observe, move to a belief state that lies in the

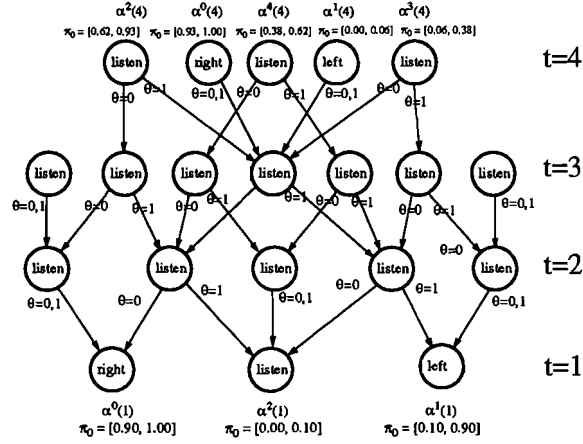


Figure 17: Policies and relationship for $t = 1$ through $t = 4$

listen node for $t = 1$ ($\alpha^2(1)$).

This argument for doing the listening up front no longer applies after $t = 3$ and all the optimal policies for $t > 3$ will choose to open a door for some initial belief state. In Figure 17 we have shown the structure that emerges for the optimal policies from $t = 1$ to $t = 4$. Note that the belief state partitions imposed by these policies is only shown for the first and last policy.

Figure 17 shows many of the interesting structures of a finite horizon POMDP solution. Notice that at the $t = 3$ level there are a couple of nodes ($\alpha^0(3)$ and $\alpha^1(3)$) that do not have any incoming arcs from the $t = 4$ level. This is interesting because it is showing that no matter what belief state you start out in at time $t = 4$, there is no action you can take and observation you can make that will leave you in a belief state that lies in the partitions defined by those unused nodes at $t = 3$. (i.e., Certain belief states are not possible at this time step no matter what the agent's belief state was on the previous time step.)

Perhaps the most interesting aspect of Figure 17 is the way in which the belief states behave. If we choose one of the belief partitions for $t = 4$, and then execute this policy, the actions and observations the agent makes will precisely define the nodes in the graph that are traversed (i.e., which actions to take at each time step). Regardless of the actual starting belief state, as long as it lies in that particular partition, the action sequence will be exactly




Figure 18: Two possible policy trees for tiger example

the same for a given sequence of observations.

The policy trees discussed in a previous section are very much related to this solution structure. Recall that a policy tree was a specific sequence of actions to take based solely on the observations made. In that section we discussed the possibility of generating every possible tree and then evaluating each to see which was the best for a particular belief state. The graph shown in Figure 17 actually has some of these policy trees embedded in it. One major difference between this solution structure and the policy trees is that the solution structure does not have all possible policy trees in it, only the useful ones. The other major difference is that, within the solution structure, policy trees with similar sub-components will have these sub-components collapsed. For instance consider the two policy trees for $t = 4$ shown in Figure 18. These are two of the 2^{15} possible policy trees. Both of these policy trees are embedded in the graph of Figure 17 and are therefore two of the useful policy trees. Notice how nodes of the tree are collapsed in the resulting figure. Figure 19 shows one of the policy trees that is not useful. There is no starting belief state where this tree will be the best thing to do.

3.1.2 Infinite Horizon

With the finite horizon problems, the non-stationary policies generated can be quite cumbersome since there is a complete policy for each and every time step of interest. Without discounting future rewards, these policies will generate a different set of vectors (i.e., value function) for each and every time step.

When we add a discount factor to decrease the value of future rewards,




Figure 19: A policy tree that is not useful

the structure of the finite horizon POMDP problem changes slightly. As the time, t , gets larger the effect of the rewards received for early times (e.g., $t = 1$, $t = 2$, etc.) will begin to have negligible influence on the policy for these later time steps. As $t \rightarrow \infty$, the vectors output by the discounted finite horizon algorithms begin to converge on a fixed set of values. As a result, for large t , the policy looks much the same as the policy for $t-1$. Figure 20 shows the solution to the discounted finite horizon of the tiger POMDP example for large values of t . Notice that the structure of the graph is exactly the same from one time to the next. The actual corresponding vectors for each of the nodes (which together comprise the value function) in the graph differ only in the fifteenth decimal places. This shows how the value function is converging. The structure of the graph remains unchanged, even though the values of the underlying vectors are slightly different from one time step to the next. This structure first appears at time step $t = 56$ and remains constant all the way up to $t = 105$. When $t = 105$ the precision of the algorithm used to calculate the policy can no longer discern the difference between the vectors' values for succeeding intervals.

If we solve the finite horizon problem for larger and larger t , the value function (and its policy) for the largest time step gets closer and closer to the value function for the infinite horizon problem. The two converge when $t \rightarrow \infty$. In this way we can use algorithms for the finite horizon problem to get arbitrarily close answers to the infinite horizon problem. Since there is finite precision in most of the algorithm implementations, there will be a point at which the solution appears to converge. Although the actual value function only converges in the limit, for practical purposes, the solution the




Figure 20: Structure of solution for large t

algorithms converge on for the discounted finite horizon are close enough to the actual infinite horizon solution. The advantage of infinite horizon solutions is that we only need the last policy generated and can throw away all the other time steps' policies. (i.e., The optimal solution for the infinite horizon problem is a stationary policy.)

3.2 Policy graphs

One major drawback of the POMDP approach is that the agent must maintain a belief state. While in theory (since we assume we know the model), this is not a significant problem, in practice it can become a major problem when the agent might not have the resources (time or hardware) to perform the updating of belief states. Fortunately there is a way to encode the policy in a graph such that no explicit belief states need to be maintained. We refer to such graph as *policy graphs* and they appear first in [16]. The policy graphs we discuss apply to the infinite horizon solutions.

Recall Figure 20 where we have converged upon an infinite horizon policy. If we continued to run the algorithm longer and longer, we would continue to get the exact same structure. Because we have the same structure at every level, we can re-draw the edges from one level to itself. Although this is erroneous in the finite horizon problem (because eventually we will get to lower values of t where the graph is not the same), for the infinite horizon, we can generate as many of these identical levels as desired. This rearrangement of edges is shown in Figure 21.




Figure 21: Rearranging edges for infinite horizon




Figure 22: Policy graph for tiger example

In Figure 22 we have redrawn this graph into what we refer to as a *policy graph*. Here “TL” and “TR” represent the observations “tiger-left”, $\theta = 0$, and “tiger-right”, $\theta = 1$, respectively.

The policy graph in Figure 22 has the interesting property that there are some nodes of the graph that will never be visited once either the open left or right door actions are taken. This results from the resetting of the belief state to $\pi = (0.5, 0.5)$ induced by our particular problem. If we imagine that the agent always starts in a state of complete uncertainty, then it will never be in a belief state that lies in a partition of these non-reachable nodes. This results in a simpler representation of the policy graph, shown in Figure 23.

Recall that the vectors that each node represents defines a partition of the belief space and, furthermore, that these all beliefs within a particu-




Figure 23: Trimmed policy graph for tiger example

lar partition will map to a node on the next level exactly contained in the partition imposed by the vectors of that next level. In the case where we collapse the levels into one, we are guaranteeing that the partitions have a one-to-one mapping for the transformed belief states (for a particular action and observation). This policy graph representation allows us to forget about maintaining belief states because as long as we are guaranteed that the belief states behave in this manner, the exact belief state is not important. The only thing that is important is which partition the belief state lies in and, by the properties of the policy graph, all the subsequent belief states will lie in partitions solely determined by the following the actions on the graph and transitioning to another node based upon the observation seen.

3.3 Finitely Transient Policies

In our discussion about using finite horizon algorithms for finding solutions to discounted infinite horizon problems, we ignored a few important aspects about infinite horizon solutions. Solving the finite horizon for larger and larger horizon sizes will give solutions (value functions) that get closer and closer to the optimal infinite horizon problem. However, only in the limit do the two actually converge.

We developed the idea of policy graphs assuming that the solutions for each time step would converge. This will happen for only a certain class of problems. When this does converge, then we can construct the policy graph, but we cannot do this when the solutions do not converge.

Finite Transience is formally defined in [16] and is the property of a policy

not of the POMDP problem. Here, we opt for a more intuitive understanding. Recall that a particular policy will define a partition of the belief space and that each node of the policy graph represents one of these regions. Moving around in the policy graph (based upon the resulting observations), is actually the same as moving around in the belief space since the belief state will be transformed by $T(\pi|a, \theta)$ at each step (whether or not we choose to actually compute it or not). With the policy graph we are implicitly assuming that all the belief states within a particular partition element will be transformed to another element under $T(\pi|a, \theta)$ for a particular a and θ . This held from the properties of the POMDP formulation and the structure of the resulting optimal policy. However, not all optimal policies for POMDPs will define partitions with this property, but the ones that do are termed finitely transient policies. For non-finitely transient policies we will not be able to construct policy graphs of the type described.

4 Finite Horizon Algorithms

4.1 POMDP History

POMDP research grew directly out of the MDP research, both of which began to flourish in the 1960's. [3] and [7] provided much of the basic framework and solution procedures for MDP models. Among the earliest work dealing with partial observability are [5] and [1]. Although none of these presents algorithmic solutions to the general POMDP model, each provided some of the groundwork for treating the general problem. In this section we will give just an overview of some of the POMDP researchers and their work. Their algorithms are treated in more detail in the subsequent sections.

The first researcher to give a detailed algorithm for finding optimal policies for the general POMDP model is E. J. Sondik [16]. In this work he gives an algorithm for finding the exact solution to finite horizon problems and solutions to infinite horizon problems that are arbitrarily close to optimal. Sondik treats both the discounted and undiscounted infinite horizon, though this paper focuses only on discounted/undiscounted finite horizon and discounted infinite horizon. Unfortunately, Sondik's finite horizon algorithm is difficult to follow as well as to implement and is slow for all but the smallest of problems. An attempt is made in a later section to present Sondik's

algorithm in a language that is more comprehensible than his thesis or his subsequently published article [15].

Sondik's main contribution was to observe that since there were a finite number of linear segments (regions over the belief simplex) in the value function we could determine all of these segments by iteratively finding a particular segment and the belief space region for which it was optimal. Once this region was known you would know exactly where points in neighboring regions would lie, namely, on the borders of the current region. The algorithm was guaranteed to terminate since there are only a finite number of these regions in the finite horizon case.

Monahan [11] presented a much simpler algorithm for computing optimal policies for the finite horizon. Although it is simpler to understand, it appears more inefficient than Sondik's method, because it is an exhaustive algorithm, but actually for most problems it turns out to actually be more efficient than Sondik's method. The reasons for this are discussed in more detail in a later section. Monahan's approach is to enumerate all possible linear segments that could exist for $t = i$ and then go through them one by one to see which were relevant. Monahan's technique differs from most of the others presented since it does not explore the belief space in any way. His insight was that with a finite set of previous $\alpha(n-1)$ vectors and finite action and observation spaces, there were only a finite number of possible ways the value function could give different answers. Thus Monahan decides to enumerate all possible $\alpha(n)$ vectors and check each for validity afterwards.

Eagle [6] presented the details of an optimization of this method suggested by Monahan, that reduces the work needed to solve the problem. Eagle's uses Monahan's observation that for many of the linear segments (α -vectors) enumerated by Monahan's method, many could immediately be dismissed if they were component-wise dominated by a previously enumerated one. However, even with this optimization optimal policies to small problems can still require a substantial amount of computational resources.

Cheng introduced two algorithms [4] the first of which (Relaxed Region) is very similar to Sondik's except that he defines regions that are typically larger than Sondik's and thus, the algorithm tends to be more efficient. Cheng's other algorithm (Linear Support) uses successive approximations of the value function to converge on the actual value function by comparing the current approximation with the true value function at a set of points. Cheng guarantees that if there is any difference between the two at all, then one of those

points will show the maximum difference. An additional advantage of the Linear Support algorithm is that it can also be used to generate approximate solutions. Both of Cheng’s algorithms require the use of an algorithm for finding all the corner points of the regions (i.e., interior points of convex polyhedra) which can result in poor running times.

Our Witness algorithm is similar to Cheng’s Linear Support algorithm except that instead of having to discover all of the corner points of a region, we define a linear program to find points where the approximation is within ϵ of the true value function. Use of a linear programming formulation has resulted in better theoretical and empirical running times. The effect the ϵ term has upon the resulting answers is discussed after the algorithm is presented.

At this time it appears that seeking exact solutions to POMDPs is not practical. There are many techniques for determining approximate methods and many of these are discussed in [8].

4.2 Overview of Algorithms

In the following sections we will give detailed descriptions of some of the better known algorithms for solving finite horizon POMDPs. We have ordered the algorithms in a slightly out-of-chronological order. Our purpose here is to attempt to present the simplest two algorithms (Monahan’s and Eagle’s) first (both of which actually succeeded Sondik’s method). After these we present, in chronological order Sondik’s one-pass algorithm, both of Cheng’s algorithms (Relaxed Region and Linear Support) and finally our Witness Algorithm. It is hoped that the presentations of these algorithms are clearer and more complete than some of the original descriptions of these algorithms. We conclude this section with a comparison of these existing techniques.

All of these algorithms have a few aspects in common. They all are performed iteratively (i.e., as a dynamic program). Additionally, all of the algorithms try to find the set of vectors that define both the value function and the optimal policy at each time step. The description of the algorithms is limited to the process of completing one iteration since each iteration requires the exact same technique. They all start with the finite set of vectors, $(\alpha(n-1) = \{\alpha^0(n-1), \alpha^1(n-1), \dots, \alpha^{M-1}(n-1)\})$, for the previous time step (i.e., the piecewise linear segments of the convex value function, $V_{n-1}^*(\cdot)$). They all produce another finite set of vectors,

($\alpha(n) = \{\alpha^0(n), \alpha^1(n), \dots, \alpha^{N-1}(n)\}$), which represent the piecewise linear value function, $V_n^*(\cdot)$, for the next successive time step. In addition, each of the $\alpha(n)$ vectors will have a single control action a associated with it. The $\alpha(n)$ comprise the value function, $V_n^*(\cdot)$, whereas the combination of the vectors and associated actions are used in defining the policy, $\delta_n^*(\cdot)$.

All of the algorithms are presented as discounted finite horizon algorithms. The discounting factor β is not necessary for finite horizon problems, but is included for the case where we want to use the iterated finite horizon to solve for the discounted infinite horizon (in hopes that it converges). The discount factor is a trivial addition and does not make the algorithms any more or less complex. In fact, when $\beta = 1$, this is exactly the same as an undiscounted finite horizon problem.

It is easy to set up the iteration for all of these algorithms by feeding the resulting $\alpha(n)$ vectors back through the algorithm to get the $\alpha(n+1)$ vectors. Note that the associated actions for the input $\alpha(n-1)$ vectors do not come into play for the algorithms when they are computing the $\alpha(n)$ vectors. They are only required to define the policy $\delta_{n-1}^*(\cdot)$. Aside from the POMDP model and the initial set of $\alpha(n-1)$ vectors, no other input is needed for these algorithms (except the discount factor if one is being used).

4.3 Finding a Vector for a Single Point

This is one step that is common to all of the algorithms and so we have chosen to deal with it in a separate section. Given a single point in the belief state space we can use

$$\alpha_{a,i}(n) = q_i^a + \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{(\pi,a,\theta)}(n-1) \quad (8)$$

to immediately generate a vector for each possible action. For these vectors (referred to as $\alpha_a(n)$) we can use the formula

$$V_n^*(\pi) = \max_a \left[\sum_i \pi_i \alpha_{a,i}(n) \right]$$

to find out which of these will be the true vector at the belief point π . This maximal $\alpha_a(n)$ then is one of the true linear segments of the value function. Note that this is just the inner quantity from the recursive value equation from formula 7.

Generating the $\alpha(n)$ vector immediately might be a bit of an over statement, it really involves a few steps, but each step is straight-forward and requires nothing more than simple calculations. However, we will step through this exact procedure since it is a crucial step in all of the algorithms.

The basic scheme is that we will need to try all possible values of action a in the formula and then choose the one that gives the maximum value. The key point here is that it is not enough to merely save the action a and its value (from formula 7), you must also save the inner bracketed quantity (which is actually a vector) since once you find the action that maximizes the value function it is the corresponding inner quantity that is the α -vector for the particular belief π .

In the value function above the quantities q_i^a , p_{ij}^a , $r_{j\theta}^a$ and β are directly available from the POMDP model and the belief state π is assumed to be given, for the moment. The only unspecified term appears to be $\alpha^{\iota(\pi, a, \theta)}(n-1)$. If you recall, the function $\iota(\pi, a, \theta)$ was merely a convenience that was used to more succinctly represent the previous value function, $V_{n-1}^*(\cdot)$. It was defined as

$$\iota(\pi, a, \theta) = \arg \max_k \left[\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k(n-1) \right].$$

We see in this formula that everything we need to compute $\iota(\pi, a, \theta)$ is readily available: p_{ij}^a and $r_{j\theta}^a$ from the model and $\alpha_j^k(n-1)$ as the input to the current iteration. So we see that given a particular belief π , we need only to crank it through the above two formulas and arrive at a piece (one of the linear segments) of the true value function, $V_n^*(\cdot)$. Since we know there are only a finite number of these pieces (linear segments or α -vectors), if we could somehow know the exact set of points that would generate each of these pieces, we could construct the entire true value function as a set of α -vectors from a finite set of points. All the algorithms discussed except one, try to be clever about finding this set of points. Monahan's algorithm, which is conceptually much simpler, is the one exception.

4.4 Monahan's Algorithm

We start with the easiest algorithm of the bunch, though it was presented after the landmark work of Sondik ([16] [15]). Monahan actually credits this algorithm to Sondik and, though there are similarities, Monahan's description is different enough to warrant a separate treatment. This algorithm is both

easy to understand as well as easy to implement and it will provide a lot of the groundwork that will be essential for understanding the other algorithms. We begin again with our value function:

$$V_n^*(\pi) = \max_a \sum_i \pi_i \left[q_i^a + \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi,a,\theta)}(n-1) \right].$$

The key insight behind Monahan's method is noticing that there are only a finite number of possible α -vectors (i.e., inner bracketed quantities) that can be constructed from the above formula, despite the fact that there is a continuum of belief states. The actual values (components) of the α -vectors generated are not at all dependent upon the value of the belief state chosen. The belief state only serves to decide which of the $\alpha^k(n-1)$ vectors to include in the summation over the observations. Since there is only a finite number of the $\alpha^k(n-1)$ vectors and a finite number of both actions and observations we can simply list all the possible ways that the value formula could be evaluated. Sure, there can be a good number of them², but the number is a finite one and we could enumerate and calculate all of them, given enough time. Notice that since we have abandoned worrying about the belief state π we cannot really attempt a maximization over the different actions. Instead we just factor this into the enumeration scheme so that now we consider for each possible action all the ways the $\alpha^k(n-1)$ vectors could be combined in a summation over the observations, θ .

Here we will try to be more explicit about Monahan's enumeration scheme so that no confusion can arise. For a given action, as we do a summation over the observations, θ , we get a choice of one of the previous $\alpha(n-1)$ vectors for each θ . This amounts to filling in a table (e.g., $\theta_0 = \alpha^4(n-1)$, $\theta_1 = \alpha^2(n-1)$, $\theta_2 = \alpha^9(n-1)$, etc.) with one of the $\alpha(n-1)$ indices (of which there are M possible choices). This means that there are $M^{|\Theta|}$ ways we could fill out this table and that there are at most $M^{|\Theta|}$ ways to do the summation over observations. However, each summation can apply for each action and so there are at most $|\mathcal{A}|M^{|\Theta|}$ possible $\alpha(n)$ vectors we could generate. The first step of Monahan's method is to create all of them.

The biggest problem now is that our value function probably doesn't require all of these $\alpha(n)$ vectors. We have enumerated them without concern for whether or not they really provide a maximal value for some belief state.

² $|\mathcal{A}|M^{|\Theta|}$, where M is the number previous $\alpha(n-1)$ vectors

We could just remember them all, and each time we need to find what action to perform, maximize over all of these and select its appropriate action. However, this is a lot of excess baggage to carry around since we could waste a significant amount of time checking vectors that will never be maximum regardless of the belief state chosen. For time critical applications, it would not be very desirable to have the policy stored this way. An even worse problem arises in the finite horizon problem, since we iterate many times sending the results from one iteration to the next. If we never trim away useless vectors, the number of vectors in the next iteration might be an order of magnitude larger than if we had kept only the relevant ones. Even when we do the trimming, the number of vectors at each iteration grows quickly, without trimming this problem is greatly magnified. A later section provides some analysis about how large these sets can grow, here we are content to say “there are lots of them”³.

To trim away the non-relevant vectors generated during the enumeration phase Monahan’s method utilizes linear programming (LP). Linear programming is also a common theme throughout the rest of the algorithms discussed. What it does is more important than how it does it, but for the reader interested in how LPs work see [18]. There are many levels that linear programming can be interpreted. We opt for the geometric view, since that is how we have been presenting the belief space and value functions.

The observation that will let us use LPs to trim the set of vectors is that for a vector to be a true part of the value function, there must be at least one belief state π for which it gives a larger value than all the other newly created $\alpha(n)$ vectors. Imagine that we want to check if $\alpha^j(n)$ was a true vector. To do this we set up a linear program with a constraints for each other $\alpha^k(n)$ we have generated:

$$\sum_i \pi_i \alpha_i^k(n) \leq \sum_i \pi_i \alpha_i^j(n), \quad \forall k.$$

This says that we are only interested in the region of the belief space where $\alpha^j(n)$ will provide a better value than all other $\alpha^k(n)$. The variables in this LP are the components of the state vector, π_i , and a rewriting of this

³ $|\mathcal{A}|^{\frac{|\Theta|^{k+1}-1}{|\Theta|-1}}$

in a more useful way is:

$$\sum_i \pi_i (\alpha_i^k(n) - \alpha_i^j(n)) \leq 0, \quad \forall k.$$

We have to add one additional constraint to these, namely

$$\sum_i \pi_i = 1.$$

This makes sure that the point lies on the belief space simplex (i.e., it is a probability distribution). Most LP packages assume that all of the variables are non-negative. If this is not the case then we must also add the constraints for the ranges of the variables

$$\pi_i \geq 0, \quad \forall i.$$

We can use virtually any objective function we want with this LP, since here we are only concerned with whether or not there is any feasible solution to the LP. If there is any solution at all, then there must be some point on the belief state simplex that satisfies all of these constraints and thus $\alpha^j(n)$ is indeed part of the true value function. If the LP is infeasible, then not only have we found that it is not part of the true value function, we no longer need to keep it around for subsequent LPs.

The linear programming formulation and solving described above only checks a single vector. This entire process must be repeated for each of the $\alpha(n)$ vectors enumerated. The only pleasing part of this procedure is that the size of these LPs will diminish as we are able to trim more and more extraneous vectors away since once we have determined a vector to be extraneous, we don't have to generate a constraint for it when checking the other vectors. The whole algorithm is then:

- (1) Generated all possible vectors from formula 8.
- (2) Add each vector to list and mark as undiscovered.
- (3) Choose an undiscovered vector from list, if none then we are done, list contains all useful vectors.
- (4) Construct LP for that vector (see above) and mark the vector as discovered.
- (5) If the LP is not feasible, then remove the vector

from the list.
(6) Go to step (3).

4.5 Eagle’s Variant of Monahan’s Algorithm

Monahan [11] mentions that dominated vectors could be removed to help reduce the number of linear programs that need to be solved. In [6] this optimization is made more explicit. The optimization occurs in the phase where we need to enumerate all of the possible new $\alpha(n)$ vectors. Eagle’s shows that if, in our enumeration process, we ever come across a vector whose components are completely dominated by another vector’s components (one already generated by the enumeration process), then we can immediately discard it since it is impossible for it to be a true vector in the value function. In other words, if we just generated $\alpha^j(n)$ and the following condition holds:

$$\exists k, \forall i, \quad \alpha_i^k(n) \geq \alpha_i^j(n)$$

then we can immediately discard $\alpha^j(n)$. This is true because all the belief components π_i are non-negative. Finding a single existing vector that is component-wise larger than another implies that for any belief π at all, the former will result in a larger value, thereby rendering the latter extraneous.

Other than this check during the enumeration, Eagle’s algorithm works exactly the same as Monahan’s: enumerate all possible vectors and verify each one with a linear program. Eagle also mentions that we can skip the linear program verification step if we are willing to pay the computation price that will be incurred by keeping the extraneous vectors around. Since each extra vector in one time step can lead to many extra vectors in the next step, this seems to be an impractical solution for any problem. It is unlikely that the any LP package would be slow enough to prefer eliminating this step.

4.6 Sondik’s One-Pass Algorithm

4.6.1 Background

Edward J. Sondik [16] presented the first solution techniques for finding optimal policies for general POMDP problems. This was the seminal work from

which all the other algorithms described in this paper were derived. His algorithm for the finite horizon case is also described in [15]. This latter reference can be both a help and a hindrance. The formulas and notation in this article are much easier to follow than that of his thesis, and this article version is devoid of the many theorems and proofs. However, a few typographical errors, a couple of vague sentences and one serious oversight [12, 8] make this exactly the wrong primary source to use for an attempted implementation.

4.6.2 The One-Pass Algorithm

This algorithm begins where all the other ones begin; with the recursive value function:

$$V_n^*(\pi) = \max_a \sum_i \pi_i \left[q_i^a + \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{\ell(\pi,a,\theta)} (n-1) \right]$$

Here, unlike the exhaustive enumeration of Monahan, Sondik actually develops a method for finding the proper set of belief states to plug into this formula to get all the necessary vectors. Furthermore, his algorithm guarantees that this set of belief states is finite and that they generate all the true vectors comprising the value function (i.e., none are missed). This algorithm amounts to a search through the belief space by moving from one region to another until every region is found. We will soon discuss what exactly these regions represent, but here the important point is that his algorithm is guaranteed to only visit a finite number of regions and that the union of all these regions is equal to the entire belief space (i.e., these regions are a partition of the belief space.)

The basic idea behind the algorithm is not too complicated to state, but trying to implement this algorithm exposes many subtle problems. Dealing with these problems is what adds the real complications to the implementation of this algorithm. Here is an over-simplified version of Sondik's algorithm.

- (1) Initialize a search list of belief states to contain any single point.
- (2) Remove a point from the search list, if the list is empty then we are finished.
- (3) Find the true vector (and its associated action) for this point. (Use formulas 6 and 7)

- (4) Define a region around this point where this vector is guaranteed to be the true vector.
- (5) Select points that lie on the edges of this region and add them to the search list.
- (6) Go to step (2).

Although this is an over-simplification, it is nonetheless an important abstraction to understand since it is the very heart of this and the subsequently described algorithms. Before we discuss the complications that can arise, we will first expand upon the steps as they are currently shown.

The first two steps are straightforward and require nothing more than maintaining a list-type data structure to store belief vectors (which are just tuples of size $|\mathcal{S}|$). A previous section describes, in detail, how to perform step (3), just remember that in addition to generating the vector you need to store the associated action that was used to create it, since this is what is used to construct the policy.

Steps number (4) and (5) are the most interesting. These are where Sondik provided real insight into the POMDP problem. As mentioned earlier, when we generate a vector from a single belief point, we are guaranteed nothing about what the true vectors are for all the other points in the belief state space. However, this isn't really quite true. Sondik, by examining the steps with which the vector is created, is able to define a region (volume) around this belief point where that vector is guaranteed to be the true linear portion of the value function. His method is to state a series of constraints on the belief states that, when satisfied, assure that the belief states satisfying those constraints have the generated vector as the piece of the true value function. This has the effect of defining a volume of belief states instead of the single point we started with. Describing the region (step (4)) is necessary for generating the new belief states to add to the search list as is done in step (5). We first describe the former step.

A little extra notation is required in order to develop the constraints that define the region. We use π to be the point selected from the search list. Recall that to find the true maximum $\alpha(n)$ vector we need to do a maximization over all actions. What this really gives us is $|\mathcal{A}|$ vectors for the point π . ($|\mathcal{A}|$ is the cardinality of the set of actions, i.e. number of possible actions.) We will refer to each of the vectors generated in this step as $\alpha_a(n)$

where a is the action that generated the vector. We use $\alpha_{a,i}(n)$ to represent the i^{th} component of those vectors. Of these $\alpha_a(n)$ vectors, one of them will give the maximal value for the particular belief π . We will refer to this vector as $\alpha^*(n)$. Additionally, the action for this best vector will be referred to as a^* . Note that $\alpha^*(n) = \alpha_{a^*}(n)$.

We now must think about how the belief state π could vary, while still resulting in $\pi\alpha^*(n)$ being larger than all $\pi\alpha_a(n)$. This exact condition is actually one of the constraints on the region we are looking for:

$$\sum_i \pi_i \alpha_i^*(n) \geq \sum_i \pi_i \alpha_{a,i}(n), \quad \forall a.$$

Notice that this formula expresses exactly our desire to constrain the belief states to satisfy the condition stated above, one constraint for each $\alpha_a(n)$. However, this is just the tip of the iceberg. All of the $\alpha_a(n)$ were generated dependent upon the actual value of π . Although the above constraints hold true when all the $\alpha_a(n)$ stay unchanged, changes in π can cause changes in the $\alpha_a(n)$ component values. Therefore, the above set of constraints is too liberal since it might include belief states where the $\alpha_a(n)$ have completely different values. We need to constraint the belief state further to ensure that the actual values of the $\alpha_a(n)$ stay the same. To do this we must examine the formulas that gave rise to these $\alpha_a(n)$ values.

There are two fundamental ways in which our $\alpha^*(n)$ vector could result in a lower value than one of the other $\alpha_a(n)$ for a particular π . The first is if the components of the $\alpha^*(n)$ vector change. If these components change, we can no longer be sure that it will dominate all of the other $\alpha_a(n)$ vectors, it will depend upon exactly how the components changed and by how much. The other way $\alpha^*(n)$ can lose its domination is if one or more of the other $\alpha_a(n)$ vectors have their components change. In this case, although the quantity $\sum_i \pi_i \alpha_i^*(n)$ stays the same, the quantities we compared it to have changed. Again, whether or not $\alpha^*(n)$ still provides the best value all depends upon how the $\alpha_{a,i}(n)$ components changed. There is also the case where both $\alpha^*(n)$ and the $\alpha_a(n)$ vectors change simultaneously, but handling the individual cases will also handle this case.

What affects the $\alpha_a(n)$ and $\alpha^*(n)$ values? The $\iota(\pi, a, \theta)$ function does. Recall that this function was a convenience for specifying the index of the previous time step's $\alpha^k(n-1)$ vector to use in the value formula, $V_n^*(\cdot)$. The $\iota(\pi, a, \theta)$ function represented a maximization of the transformed belief state,

$T(\pi|a, \theta)$, over all the $(n-1)^{st}$ vectors. All we need to do is to constrain the belief states so that this maximization works out the same, and we will have ensured that the $\iota(\pi, a, \theta)$ function stays the same. The $\iota(\cdot)$ function stays the same when the $\alpha_a(n)$ vectors that are generated in the maximization remain the same, which is the condition we are interested in. We restate the formula for the maximization over previous vectors so that the origins of the constraints to follow are easier to see:

$$V_{n-1}^*(T(\pi, a, \theta)) = \max_k \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k(n-1).$$

Let ι be the index k that makes this maximum for a particular point, π , action, a , and observation, θ . Then the constraint

$$\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota}(n-1) \geq \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k(n-1), \quad \forall k \neq \iota.$$

satisfies the condition of all belief states where ι is the index that maximizes the formula above. Notice that for a single pair of a and θ we require a constraint for each of the previous $\alpha^k(n-1)$ vectors. The above constraint only restricts one pair and does not ensure that the entire function $\iota(\pi, a, \theta)$ remains unchanged. To make sure that the whole function stays the same, we need a set of these constraints for each combination of a and θ . By adding all of these constraints we have ensured that neither $\alpha^*(n)$ or any of the $\alpha_a(n)$ will change. Due to the nature of the problem, the constraints generated can only specify a region that ensures things do not change. It cannot have any information about how things change and as a result, things can change without really affecting the optimal vector generated. This leads to defining regions that can be smaller than the actual regions formed by the various $\alpha(n)$ vectors in the optimal value function, $V_n^*(\cdot)$.

Aside from the constraints thus far outlined, there are also the constraints that restrict the belief states to lie on the belief state space simplex, namely

$$\sum_i \pi_i = 1$$

and

$$\pi_i \geq 0, \quad \forall i.$$

No doubt, this looks like a formidable number of constraints and, it is for this reason that this algorithm has not found wide spread use in solving

POMDP problems. The complexity of this algorithm is discussed further in a subsequent section, but the worst case number of regions for one time step is $|\mathcal{A}|M^{|\mathcal{A}||\Theta|}$ where M is the number of $\alpha(n-1)$ vectors and each region requires a linear program of size $|\mathcal{A}||\Theta|M + |\mathcal{A}| + 1$ constraints and $|\mathcal{S}|$ variables. Here M is the number of previous vectors in the $\alpha(n-1)$ set.

After much work, step (4) is now complete; we have successfully defined a region of the belief space where $\alpha^*(n)$ is guaranteed to be the true value function's vector. This is especially pleasing because we have, in effect, eliminated every belief state within this region from consideration in plugging it into our value function. We now know that for every belief state in this region, if we crank it through the value function we will get exactly $\alpha^*(n)$, and so there is no reason to consider these points further.

Although most of the conceptual work was done in step (4), the majority of the computational work lies in step (5). In this step we want to find belief states that are guaranteed not to be in the region we defined in step (4). If we can find a point not in the step (4) region, then we have found a point that must lie in some other region. With this point, we can do exactly as we did in step (4) and find its region, and so on until a complete partition of the belief space is found. When we have been to every region in the partition, we are sure to have generated all of the true vectors in the value function.

Given the region in step (4) how can we find points in other regions? The answer here is linear programming. We have already built up the constraints of a linear program (notice that they were all linear constraints in the π_i variables.) We rearrange the constraints above into more standard LP constraints and cluster them in one place here for convenience:

$$\begin{aligned} \sum_i \pi_i (\alpha_{a,i}(n) - \alpha_i^*(n)) &\leq 0, \quad \forall a \neq a^* \\ \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a (\alpha^k(n-1) - \alpha^{\iota(\pi,a,\theta)}(n-1)) &\leq 0, \quad \forall a, \theta, k \neq \iota(\pi, a, \theta) \\ \sum_i \pi_i &= 1 \\ \pi_i &\geq 0, \quad \forall i. \end{aligned}$$

What we are interested in are points lying on the edge of the region defined by these constraints. These linear constraints will define a piecewise linear convex region (volume) of the belief space. The neighboring regions




Figure 24: A region and its neighbors

lie directly adjacent to this region and actually share a common border (hyperplane). Points on the border of the currently defined region will also be points on the border of neighboring regions. Figure 24 shows this graphically.

Since our current region will typically border many other regions, we need to find a point on each region border. Although we have a large number of constraining equations, there is a small subset of them that actually define (are binding on) the region and the broken lines are constraints that are superfluous (i.e., there are other constraints that restrict the size of the region more severely than it). The problem is that we have no direct way to ascertain which constraints are binding. The method proposed by Sondik to accomplish this is to solve a series of linear programs. There are actually two ways this can be done. We present both, first the one proposed by Sondik and then an alternative.

In Sondik's approach we solve a single LP for each constraint in the LP (aside from the simplex constraints). The constraints for all of these LPs are identical (the ones shown previously), the only thing that changes is the objective function. For each constraint we use the actual constraint itself (without the right hand side, 0) as the objective function. Note that this constraint remains in the set of constraints for the LP. Solving this LP as a maximization problem will result in an answer with values for all the π_i variables. However, just because this LP returned an answer does not necessarily mean that we have found a new belief vector to add to our list. We

only want to add this point if the constraint currently under consideration is binding in this LP. Fortunately, the solution to the LP gives us that information as well. In an LP, each constraint has an associated slack variable which indicates how much tolerance we have in changing this constraint without affecting the solution. If the constraint that was inserted into the objective function has no slack (i.e., its slack variable has value zero), then the constraint is binding and thus we can insert the resulting value of π_i s into our search table. In summary, for each constraint:

- (1) Make constraint the objective function.
- (2) Solve the LP trying to maximize the objective function.
- (3) Check the slack variable of row that the constraint in the objective function appears.
- (4) If it is zero then the solution (values for all π_i) should be added to the search list, otherwise the constraint is non-binding.

A slightly simpler method, that still utilizes one LP for each constraint, is to slightly change the constraints as each LP is computed. In this method the actual objective function doesn't matter. What we do for each constraint (again, except for the belief simplex constraints) is to change its inequality into an equality. The result of solving this LP gives both pieces of information desired at the same time. If the LP has any solution at all, then the constraint must be one of the binding constraints because a non-binding constraint can not pass through the region defined by all the other constraints. Additionally, any solution it returns must lie directly on the border that this constraint makes for the region. With this method it doesn't even matter if we perform a maximization or minimization of the objective function and we do not need access to the dual variables (which aren't always easily available from an LP package anyway.) The only trick here is to remember to change the equality back into an inequality when you move on to try the next constraint.

One aspect of the solutions to the LPs that needs to be pointed out, is the nature of the solutions it returns. With an LP formulation, the solution can consist of no points, one point or an infinite number of points. The latter is different from an unbounded LP. In an unbounded LP the objective




Figure 25: Points found by linear programs

value function can be made arbitrarily large. An LP solution with an infinite number of solutions happens when all these points result in the same finite value of the objective function. The nature of the LPs we solve, except in a rare instance, all have an infinite number of solutions (if it has any at all), however the answer the LP returns (assuming that it uses some form of the simplex method) will be a single point and furthermore it will be one of the corner points or vertices of the region.

To illustrate this, in Figure 25 we see a region that has two adjacent regions that it shares a border with. In this case there will be two binding constraints we will call AB and BC . When we check constraint AB we find that there are really an infinite number of points that lie on the border (the entire line segment AB). However, the LP will return either A or B . Exactly which one is returned depends upon the exact nature of the LP and the way the LP package was implemented. For the constraint BC we could get either point B or point C . At first this might seem troublesome since it will be possible for these two constraints to result in the same point B , but this in itself is not a problem since all we need to guarantee is that we get a point that is on the border of a neighboring region. Point B satisfies this criterion for both the neighboring regions. This raises some serious questions about handling these borderline cases. This is all discussed further in the next section.

This completes the description of Sondik's basic algorithm. It may or may not have raised some questions at each step. Hopefully, many of these questions (which have, thus far, been deliberately ignored) are answered in the next section. Trying to deal with the complications while describing the algorithm would have been difficult to do without adding much confusion. However, trying to implement the algorithm based solely on the previous description is not recommended. There are too many unanswered questions that will necessarily arise when an implementation is attempted.

4.6.3 The Complications

For all the lengthy description, the basic ideas behind Sondik's algorithm are really not that complex. However, there are some fairly nasty subtleties lurking beneath the surface. They all emanate from overlooking the case where we get identical values as we perform the various maximizations in the value function while calculating the $\alpha(n)$ vectors. Recall that there are actually two places we need to perform a maximization: once over all the actions in the actual value function itself (over all the various $\alpha_a(n)$ vectors); and the other as we determine the function $\iota(\pi, a, \theta)$ (over all the previous $\alpha^k(n-1)$ vectors). When there are ties, we cannot simply choose an arbitrary one from among the candidates, this will not work.

We will first try to present what is happening geometrically for the two cases discussed above. Let us first suppose that we have generated all of the $\alpha_a(n)$ for a particular belief state π . Our next task is to perform the maximization to find which a and $\alpha_a(n)$ give the largest value for $\sum_i \pi_i \alpha_{a,i}(n)$. Let a' and a'' be the two actions that both give the same maximum value. This indicates that we are on the border of two separate regions as shown in Figure 26. We might be tempted to arbitrarily choose one (say $\alpha_{a'}(n)$) with the rationale that the $\alpha_{a''}(n)$ region is a neighbor and thus will be found when we set up the region for $\alpha_{a'}(n)$.

But what happens if our LP returns π as the point on the border between the two regions? Trouble. We would either see that we have used this point already and throw it out, or we might perform the same calculation and again choose the $\alpha_{a'}(n)$ region. In the former case the algorithm ends without finding the $\alpha_{a''}(n)$ region and in the latter case it could cycle forever. Although you might think that a little extra bookkeeping can keep you out of this trouble, consider the point shown in Figure 27. This is the case of a point




Figure 26: A point on the border of two regions

on the boundary of four regions. This requires more than a little bookkeeping, and as the dimensionality of the problem increases, this problem can get arbitrarily complicated. Constructing a truly general solution to handle these cases is possible, but it is cumbersome and prone to errors.

The above examples only dealt with the maximization over the actions in the value function. We still must deal with the maximization over the previous vectors that give rise to the $\iota(\pi, a, \theta)$ function. The regions depicted above were somewhat misleading since they seemed to show the regions of the actual value function, $V_n^*(\cdot)$ being constructed. In reality, the regions defined by Sondik's constraints are usually only a subset of the actual value function regions. To discuss the next complication, we need to examine these regions more closely.

The constraints that define the region specify exactly the conditions necessary for the $\alpha_a(n)$ and the $\alpha^*(n)$ vectors to change. The unpleasant part is that these conditions are not necessarily sufficient for a change. What this says is that we can tell when some of the vector components might change, but we cannot determine whether they actually will change or if this change will have any noticeable affect. A further unpleasant consequence is that even if the components of some of the $\alpha_a(n)$ vectors change they may change in such a way that $\alpha^*(n)$ is still the proper piece of the value function at these points. The bottom line here is that the constraints built up by Sondik's




Figure 27: A point on the border of four regions

methods are typically too conservative. The regions (or volumes) defined are typically small subsets of the actual regions (volumes) for a given vector. Figure 28 shows the typical relationship between the true $\alpha^*(n)$ value function region and the region defined by the constraints.

So, in reality, when we draw the regions that the algorithm generates we cannot really label it as the $\alpha^*(n)$ region since in all likelihood it is merely a subset of the true region of the value function. How then, should we label these regions? To answer this we consider how many possible ways could we construct a set of constraints since this represents the maximum number of regions we could define. The constraints that arise from a region are fully determined by the optimal action and the $\iota(\pi, a, \theta)$ function. Notice, the actual value of the belief state π never appears in the LP for the region (though it does help determine what the $\iota(\pi, a, \theta)$ function should be.) This shows that Sondik's algorithm can potentially explore a very large number⁴ of regions even when there are a small number of actual regions in the value function. This is discussed further in a subsequent section, but here we needed just to demonstrate the connection between the $\iota(\pi, a, \theta)$ function and the LP regions. Now we have established that each instance of an $\iota(\pi, a, \theta)$ function can potentially specify a different LP region. It might be the case that many instances of this function specify the same region, but the only

⁴ $|\mathcal{A}|M|\mathcal{A}||\Theta|$




Figure 28: Sondik's region vs. actual region

way to determine that is by actually constructing and comparing the regions.

We mentioned before the difficulties associated when ties are obtained while performing the maximization over the actions in the value function. We now turn to the similar case of ties in the maximization procedure while determining the $\iota(\pi, a, \theta)$ function. If we find two different previous $\alpha(n-1)$ vectors (call them $\alpha^j(n-1)$ and $\alpha^k(n-1)$) that give the same value in the formula

$$\iota(\pi, a, \theta) = \arg \max_k \left[\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k(n-1) \right],$$

then we have two equally valid and distinct $\iota(\pi, a, \theta)$ functions. As just previously shown, these could represent two completely different regions. So if we chose a belief state π right on the borders of these two regions (assuming for now that they are indeed different), we have two choices for which region to go with. The problem here is that you actually need to consider both and for exactly the same reason that we had to consider both in the previous case of ties. The same problems with arbitrarily picking one of these apply here; we might miss a region or we might loop around forever repeatedly generating the same π and choosing the same region.

The bookkeeping required to keep all of this straight is an order of magnitude more complex than in the previous case. In the previous case, we could only have as many ties as there were actions, so we could imagine keeping

a list of the ties and handling them one after the other. However, in this case there are potential for ties for each possible combination of a and θ . Every time there is a tie, the number of possible $\iota(\pi, a, \theta)$ functions, that give the same value, doubles. As if this weren't enough, for each of those many places where a tie can occur, there could be as many ties as there are previous $\alpha^k(n-1)$ vectors. Implementing the routines to handle this, though possible, is a major headache.

Individually, the two types of ties that can occur in the maximization procedures are bad enough, but they actually interact in an unpleasant way as well. The previously discussed ties for the $\alpha_a(n)$ vectors assumed that they were all tied while using the same $\iota(\pi, a, \theta)$ function. Now we see that we may end up with bunch of these functions and each one of that bunch has the potential to have ties in the $\alpha_a(n)$ vectors it produces.

All of this might appear to be needless theoretical worry over cases that aren't likely to happen in practice, however just the opposite is true. These problems of ties occur precisely on the borders between regions. What is the likelihood that we get points on these borders? The nature of the algorithm actually guarantees that we get points on the borders. Recall that all but the first point are generated from the LPs, which can only give us points directly on the borders of the regions. Therefore we spend all of our time directly on the borders and in the corners (vertices) where the ties are liable to happen. Thus handling ties is a very real problem that must be dealt with when implementing this algorithm.

The last complication we discuss is not nearly as intimidating as the previous problems, but nonetheless, it is a difficulty that must be dealt with. With all these regions being explored and with all the belief states being churned out by the LPs, we are bound to get duplicate points. If we have gone through all of the trouble to make sure we have covered every base (handling the ties) for a particular belief state π , then there is no reason to work on this point again. This requires us to keep track of all the π we have dealt with. The major trouble here is that now we are dealing with floating point comparisons. If we strictly compare the numbers, we might be duplicating a lot of effort since the same points could be slightly off from machine/algorithm inherent rounding errors. However, if we allow too much freedom in deciding which points are the same we might miss small regions of the value function. As the problem sizes grow larger, the likelihood that there will be these small regions increases.

Sondik's algorithm as described is quite inefficient and in the following sections we present three ways that it could be made less wasteful. The first two of these were proposed by Sondik himself.

4.6.4 Neighbor Optimization

The geometric interpretation of the value function and the regions they impose upon the belief states allow us to exploit the geometric properties of these regions. When we define a region and then look for the neighboring regions, we are finding each regions' adjacent neighbors. If we store this neighbor information, we can exploit it to restrict the number of constraints in the linear programs on the next time step. Recall the set of constraints:

$$\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a (\alpha^k(n-1) - \alpha^{\iota(\pi,a,\theta)}(n-1)) \leq 0, \quad \forall \theta, k \neq \iota(\pi, a, \theta).$$

This says that we must construct a constraint for every previous vector $\alpha^k(n-1)$, but actually we only need to consider the $\alpha^k(n-1)$ that are neighbors to $\alpha^{\iota(\pi,a,\theta)}(n-1)$. If we kept this region adjacency information from the previous iteration, we know exactly which $\alpha^k(n-1)$ are truly constraining the current $\alpha^{\iota(\pi,a,\theta)}(n-1)$.

4.6.5 LP Optimization

Even with the above optimization, the number of constraints in the linear programs will be large. The following insight is directly from [15]:

The procedure can be made more efficient if, for each iteration of the linear programming problem with the k^{th} inequality as the objective function, all other constraints are tested as objective functions to see if they are optimized at the current feasible solution. If a constraint is optimized at any point, then either this constraint forms a boundary of the region (a zero slack variable) or is a superfluous constraint (a nonzero slack variable). Once a constraint has been optimized, it need not be used as the objective function. We have found that this procedure typically decreases the number of linear programming iterations by approximately 50 percent.

We have no experience with this optimization and therefore have offered it in their own words. The geometric intuition behind this method is that as an LP solution progresses, it moves from corner point to corner point of the constrained region. Each corner point lies on one of the boundaries of the region and all this technique does is to see which border point it is on at each iteration by substituting each possible border (i.e., constraint). However, this requires an intimate relationship with the LP package. Unless you have written the code specifically for this application, it is unlikely to be easy and may be impossible for you to get this functionality out of your LP package.

4.6.6 Dominated Constraints

This is the simplest of all the optimizations and just requires that we compare each constraint, as it is generated, with all the others to determine whether it is component-wise dominated by some other constraint. Since our variables of the LP ($\pi_i s$) are constrained to be non-negative, this simple scheme works. Whether or not this checking is worth the trouble depends on your particular implementation and the speed of the LP package being used.

4.7 Cheng's Algorithms

In [4] two new algorithms are presented. They are both heavily based upon Sondik's One Pass algorithm, but typically require less computation time. One change present in Cheng's algorithms is the elimination of the use of linear programming. Instead, Cheng opts for the interior point method for convex polytopes (see [9] and [10]) to find the corner points of the regions define on the belief space. The reasons for this are discussed shortly.

4.7.1 Relaxed Region Algorithm

This algorithm is exactly the same as Sondik's One Pass algorithm except with fewer constraints. In each step of Sondik's algorithm, a region (for the newly discovered $\alpha(n)$ vectors of the value function, $V_n^*(\cdot)$) is defined that is sure to be no bigger than the actual region the $\alpha(n)$ vector occupies in the true value function. This is the conservative approach and thus the actual regions for each of the true vectors of the value function will have to be built up out of a bunch of these smaller regions that Sondik defines. Cheng

reverses this approach and defines regions that will typically be larger than the regions actually are (i.e., relaxed regions).

The algorithm can proceed exactly the same as Sondik's with two minor changes. The first is the set of constraints and the other is the use of the interior point method instead of normal linear programming. Recall that we will keep a search list of belief points and that for each belief point we will find the true vector $\alpha^*(n)$, its associated action a^* and all the other $\alpha_a(n)$ vectors that were the losers in the maximization process over actions. We also have the set of all the previous $\alpha(n-1)$ vectors. The only modification necessary is in the construction of the regions or set of constraints for the LP. The following is the set of constraints for the relaxed regions of Cheng:

$$\begin{aligned} \sum_i \pi_i (\alpha_{a,i}(n) - \alpha_i^*(n)) &\leq 0, \quad \forall a \neq a^* \\ \sum_{i,j} \pi_i p_{ij}^{a^*} r_{j\theta}^{a^*} (\alpha^k(n-1) - \alpha^{\iota(\pi, a^*, \theta)}(n-1)) &\leq 0, \quad \forall \theta, k \neq \iota(\pi, a^*, \theta) \\ \sum_i \pi_i &= 1 \\ \pi_i &\geq 0, \quad \forall i. \end{aligned}$$

If you compare this set of constraints with Sondik's it might take a while to even notice that they were at all different. The only change is that the second constraint is no longer defined over all actions. This constraint now only applies for each θ and previous $\alpha^k(n-1)$ vector with the action being fixed as the action associated with $\alpha^*(n)$.

By using an interior point algorithm to ensure that every corner point of one of these relaxed regions is discovered, Cheng can guarantee that all of the true value function vectors will eventually be discovered in a finite number of steps. Before we discuss how this can be guaranteed, we should point out the crucial difference between utilizing an interior point method and a linear programming method. If you recall from the discussion of Sondik, the linear programs will typically have an infinite number of potential points they can return and, by the nature of the LP method (usually simplex), we will get only one of the corner points. This works fine for Sondik's method as was previously discussed, but can lead to incorrect results if combined with the relaxed region algorithm.




Figure 29: Points and relaxed regions

For example in Figure 29 we see two regions defined by $\alpha^*(n)$ which was generated for belief state π ; Sondik's region is bounded by $ABCD$ and the relaxed region by ABE . Let us assume that the true region of the value function for $\alpha^*(n)$ is $ABCD$ (which just so happens to be the one found by Sondik's method). Sondik's algorithm will show two neighboring regions. For neighboring region CDE the LP method will return either point C or D and for neighboring region $ABGF$ either A or B . Even though we do not know exactly which point the LP might return, either one is fine since they will both lead to the other regions. In contrast, look at the relaxed region ABE , this one looks to only have one neighboring region and so the LP method will return either point A or B . This is not sufficient, because we will never choose a point in the region DCE which is a different region of the true value function that must be found. However, if we always make sure we find each corner point of the region, then we are fine, since the point E will lead to the region CDE .

Cheng demonstrates that by employing the interior point method, his algorithm will always uncover all the regions. Furthermore, since his regions are larger than those that Sondik provides, it should execute faster. This algorithm only needs to find a single relaxed region for each actual value function vector which can be substantially fewer than Sondik's One Pass method. Cheng has experimental results that verify that this is actually the

case. Cheng’s work is also explicit about handling the case of ties, whereas Sondik’s work lacked discussion about these cases.

4.7.2 Linear Support Algorithm

Cheng proposes another algorithm which employs the same idea of defining larger regions. However, this Linear Support algorithm has the advantage that it can be used to find near optimal answers as well. With the previous Relaxed Region algorithm, we start at a point and move from one region to another, slowly exploring the space of beliefs. The Linear Support algorithm, on the other hand, defines an approximate value function over the entire belief space and slowly works to refine this approximation until it reaches the true value function. The name of the algorithm comes from Cheng’s terminology for the linear segments of the value function which we have been calling α -vectors. In Cheng’s thesis he refers to these as the linear supports of the value function.

This algorithm starts by initializing a search list with the extreme points on the belief simplex (e.g., $[1, 0, 0, \dots]$, $[0, 1, 0, 0, \dots]$, $[0, 0, 1, 0, \dots]$, etc.) and an empty set of vectors, \hat{V} . For each of these points the true $\alpha(n)$ vector is calculated and added to \hat{V} (calculated from the usual recursive value function formula.) These vectors now form an approximation to the value function which we will call $\hat{V}_n(\cdot)$ and define some partition on the belief space (since for each point of the belief space one of these vectors must give the maximum value when compared to the rest.) For example Figure 30 shows the true value function as a dotted line and the approximation generated from the corner points as solid lines. Notice that close to the extreme points of the belief simplex the approximation is exact. This is for a case where $|\mathcal{S}| = 2$ and shows the belief space partition generated along the horizontal axis. This figure is an example where there are four linear regions that comprise the actual optimal value function. These four linear segments impose a partition on the underlying belief states as shown with the dotted lines across the π_0 axis. When we only calculate the linear segments for the extreme points of the simplex, we get parts of the true value function (near the extreme points), but we also under estimate the value function at many of the interior points since we have not yet found the true vectors for these points.

The key insight behind this algorithm, is that no matter what the true value function is, the largest difference between the approximation, $\hat{V}_n(\cdot)$,




Figure 30: First step of Linear Support algorithm

and the true value function, $V_n^*(\cdot)$ will always occur at one of the corner points of the approximation. This follows from the properties of piecewise linear functions. Notice that the true value function must always lie equal to or above the approximation. Because both the true function and the approximation are piecewise linear and convex the largest difference must occur at a corner point. This is proved in [4].

With this handy piece of information, Cheng then finds all the corner points of the regions in the partition induced by the approximation. Here again, Cheng utilizes an interior point algorithm instead of linear programming since he must ensure all corner points are generated. For Figure 30 this will consist of three points, the two edges of the simplex and one interior point. Since we have already handled the simplex corner points we only need to consider the single interior point. This generalizes so that of all the corner points found from all the regions, we disregard those we have seen before and add those we haven't to the search list. Note that this requires you to keep track of all points you have used previously.

The next step is to pick a point out of this search list and find the true $\alpha(n)$ vector (support) for that point. If this vector is different from all the other ones in our current approximation, then we add it to \hat{V} to arrive at a new approximation, $\hat{V}_n(\cdot)$. Figure 31 shows the succeeding approximation for the example shown previously.




Figure 31: Second step of Linear Support algorithm

We can now repeat this whole procedure with the new approximation: find corners of region; add new point to search list; generate $\alpha(n)$ vectors; add new vectors to \hat{V} to get a new approximation, $\hat{V}_n(\cdot)$. You can see in Figure 31 that calculating the new vector at the new corner point will lead to the only remaining undiscovered linear segment (α -vector).

Although this algorithm will work, as described, it is a bit wasteful. Each successive approximation imposes a new partition on the belief space. At each iteration we do not have to find all the corners of all the regions of the current partition. Cheng shows that if we just find the corner points for the region of the newly added vector we can get the same result as before with much less computation. This also guarantees that the algorithm only examines one region for each true vector (support) in the value function. This is its the big advantage over Sondik's method which typically uncovers more regions than are actually in the value function.

This is similar to the Relaxed Region algorithm because, as Figures 30 and 31 show, the region over the belief space in an approximation tends to be larger than it ultimately will be in the true value function. There is another point Cheng makes about uncovering the interior points of the regions which say that you only need to find corner points for all but one of the regions. This follows because the corner points of any single region imposed by a piecewise linear convex function also must be corner points of some other




Figure 32: Example of why one region doesn't need to have its corners checked

region. Figure 32 shows that if we generate all the corner points for the non-shaded regions, we will also have generated all the corner points of the shaded region.

Cheng employs the Linear Support algorithm as an approximation scheme as well. To do this he merely calculates the actual difference between the true value function and his approximation at the points generated at the vertices of the approximation, $\hat{V}_n(\cdot)$. By taking the largest difference in the two values he gets a bound on the current approximation. (Remember the largest difference must occur at one of these vertices.) He can stop his algorithm any time this maximum error difference is within some tolerable range.

4.8 The Witness Algorithm

There is a shortcoming of Cheng's Linear Support algorithm which this next algorithm addresses. Cheng's algorithm is guaranteed to only find the corner points (i.e., vertices) of one region for each true vector there is in the value function, $V_n^*(\cdot)$. Each of these regions is a convex polytope and as such can have an exponential number of vertices. The number of vertices is exponential

in the number of faces in the polytope and the number of faces is dictated by either the number of vectors in the current approximation, up to N , or the number of variables in the state space, $|\mathcal{S}|$. Since the approximation tends toward the optimal, Cheng's algorithm might require time exponential in the number of actual vectors in the value function, up to N or states, $|\mathcal{S}|$, in the worst case. This is a result of his use of an interior point algorithm that insists on finding all of the vertices.

The Witness algorithm was derived directly from Cheng's Linear Support algorithm and employs much of the same machinery. The main difference is in how it goes about finding more points to add to the search list. Like Cheng, we start with the extreme points of the belief simplex and find their actual $\alpha(n)$ vectors. Since we know that this approximation, $\hat{V}_n(\cdot)$, will differ from the true value function at one of the corner points of the current partition's regions (induced by $\hat{\mathcal{V}}$), we construct a linear program that explicitly states this criterion.

This algorithm starts with an empty set $\hat{\mathcal{V}}$ and begins building up the set with $\alpha(n)$ vectors in $V_n^*(\cdot)$ until we have the entire set of vectors comprising $V_n^*(\cdot)$. As mentioned, the first step is to determine the $\alpha(n)$ vectors for the extreme points of the simplex and add them to $\hat{\mathcal{V}}$.

Like Cheng's Linear Support algorithm, the Witness algorithm uses the property that the maximum difference between our current approximation and the true value function will occur at a corner point of our approximation.

We then iterate over the $\hat{\alpha}(n)$ in $\hat{\mathcal{V}}$. As we add vectors to $\hat{\mathcal{V}}$, we must ensure that we choose each vector to work on exactly once. Each time we add a vector to $\hat{\mathcal{V}}$ we will also set a flag that indicates whether we have defined its region (initially set to false). At each iteration of the algorithm, we choose one of the vectors, $\hat{\alpha}(n)$, from $\hat{\mathcal{V}}$ that has its flag set to false. We then set this flag to true and construct a linear program over the current region that $\hat{\alpha}$ defines. (This LP is discussed below.) The result of solving this LP will either give us zero or one new point in belief space. If it results in no points, then we just repeat this for another vector in $\hat{\mathcal{V}}$. If we get a new point, then this means that our approximation, $\hat{V}_n(\cdot)$ is not the same as $V_n^*(\cdot)$ at that point. We calculate the true $\alpha(n)$ at that point, add it to $\hat{\mathcal{V}}$ and then repeat the process on $\hat{\alpha}(n)$ again, but with this new $\hat{\mathcal{V}}$. As long as $\hat{\alpha}(n)$ generates a new point we continue to use it. When $\hat{\alpha}(n)$ generates no new points, we know that the region it currently defines is exactly correct for the optimal value function and so we then move on to another vector in

$\hat{\mathcal{V}}$ that we haven't worked on.

The heart of the algorithm is to construct a linear program that will have a solution at a point where our approximation, $\hat{V}_n(\cdot)$, is different from the true value function, $V_n^*(\cdot)$. In doing this we require three types of constraints. Note that the algorithm as described here, is the way it is currently implemented; however we have since found an easier way to set up these linear programs which combines the LPs for all regions into a single LP. This combining step will not improve the complexity, but makes the implementation conceptually easier. We will first define the three types of constraints needed and then tie them together to produce our algorithm.

4.8.1 Improvement Constraint

We will iterate over all the $\hat{\alpha}(n)$ in $\hat{\mathcal{V}}$. Pick a previously unselected $\hat{\alpha}(n)$ from $\hat{\mathcal{V}}$. We now want to know if there exists an $\tilde{\alpha}(n)$ such that

$$\sum_i \pi_i \tilde{\alpha}_i(n) \geq \sum_i \pi_i \hat{\alpha}_i(n).$$

Even though we do not know $\tilde{\alpha}(n)$, we do know that

$$\tilde{\alpha}_i(n) = q_i^a + \beta \sum_{j,\theta} p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi,\theta,a)}(n-1).$$

This will be true for some (currently) unknown a and π . The value at this point is simply

$$V_n^*(\pi) = \sum_i \pi_i \tilde{\alpha}_i(n).$$

A simple substitution gives us the standard value function equation

$$V_n^*(\pi) = \sum_i \pi_i q_i^a + \beta \sum_{i,j,\theta} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi,\theta,a)}(n-1).$$

So we choose to define a new variable M_θ for each θ such that there are as many variables as observations. These we define to be

$$M_\theta = \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi,\theta,a)}(n-1).$$

Substituting this into our value function equation gives us

$$V_n^*(\pi) = \sum_i \pi_i q_i^a + \beta \sum_\theta M_\theta.$$

Remember, we do not know the π (or a) where this proposed $\tilde{\alpha}(n)$ produces the true value (if there are any at all). If there is a π , then the following equation will hold true:

$$V_n^*(\pi) \geq \sum_i \pi_i \hat{\alpha}_i(n).$$

Substituting our variant of the value formula we get

$$\sum_i \pi_i q_i^a + \beta \sum_{\theta} M_{\theta} \geq \sum_i \pi_i \hat{\alpha}_i(n).$$

Rewriting gives us

$$\sum_i \pi_i (q_i^a + \hat{\alpha}_i(n)) + \beta \sum_{\theta} M_{\theta} \geq 0.$$

This constraint is the first step in creating an LP formulation (in $|\mathcal{S}| + |\Theta|$ variables) to determine a point where the vectors in our current approximation differ from a vector in the true value function.

4.8.2 M_{θ} Constraints

The previous formula gives us one constraint that the π vectors must satisfy if such a $\tilde{\alpha}(n)$ exists. By itself it doesn't help much, but it is a first step to our full formulation. One problem with this formula is that there are no constraints on the M_{θ} variables. We need to relate these variables to the quantities they represent and for values which can actually occur. Since we know all of the previous control intervals' vectors, $\alpha(n-1)$, we add the constraint

$$M_{\theta} \geq \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k(n-1).$$

For each M_{θ} we need a constraint for each a and k pair. For the answers we are interested in, this constraint must hold because if it does not, then we would be able to find a larger M_{θ} (namely, the a and k where it did not hold). M_{θ} must be equal to this sum for at least one particular a and k , since we are assuming that the M_{θ} can actually be generated. Note that this inequality only restricts the M_{θ} variables from being too small. They can be too large, but we will try to minimize these variable values to try to force the equalities to hold. This is discussed shortly.

4.8.3 Region Constraints

Another constraint that must be imposed is to ensure we do not generate points that will lead to one of the $\alpha(n)$ vectors already in $\hat{\mathcal{V}}$. To accomplish this we just need to restrict the values of the π we are interested in to those in the region defined by the current vector, $\hat{\alpha}(n)$, under consideration. These constraints are

$$\sum_i \pi_i \hat{\alpha}_i(n) \geq \sum_i \pi_i \alpha_i(n)$$

for all $\alpha(n) \neq \hat{\alpha}(n)$ in $\hat{\mathcal{V}}$.

4.8.4 Objective Function

As formulated, there will always be sufficiently large values of the M_θ variables that will satisfy the constraints. To alleviate this problem, we put all of these variable into the objective function and try to minimize their values. Since all of these variables are independent of each other, the LP will, in effect, minimize all of the variables individually. This has the effect of attempting to force one of the M_θ inequality constraints to be an equality. However, it does not ensure that an equality will hold for each M_θ . This means that we will have to examine the answers we get out of the LP, as discussed later.

4.8.5 Constraint Adjustment

Currently, our LP formulation, for a given $\hat{\alpha}(n)$ with variables π_i , can be summarized as follows:

$$\begin{aligned} \min : & \sum_{\theta} M_{\theta} \\ & \sum_i \pi_i (q_i^a + \hat{\alpha}_i(n)) + \beta \sum_{\theta} M_{\theta} \geq 0 \\ M_{\theta} \geq & \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k (n-1), \quad \forall \theta \in \Theta, a \in \mathcal{A}, k \in V_{n-1}^* \\ & \sum_i \pi_i \hat{\alpha}_i(n) \geq \sum_i \pi_i \alpha_i(n), \quad \forall \alpha(n) \in \hat{\mathcal{V}}. \end{aligned}$$

An adjustment is necessary because the improvement constraint requires that the actual value function be greater-than or equal to our approximation.

Since we only add a vector to our $\hat{\mathcal{V}}$ set when there is at least one point where it is equal to the true value function, we are guaranteed to be able to find a π that satisfies this inequality. More precisely, the corners of the current $\hat{\alpha}(n)$ region satisfy this criterion. We need to ensure that we only consider points where the true value function is strictly larger than $\pi\hat{\alpha}(n)$ since we want to know where our approximation is wrong, not where it agrees with the value function, $V_n^*(\cdot)$.

$$V_n^*(\pi) > \sum_i \pi_i \hat{\alpha}_i(n)$$

This transforms the improvement constraint into

$$\sum_i \pi_i (q_i^a + \hat{\alpha}_i(n)) + \beta \sum_{\theta} M_{\theta} > 0.$$

Because we have chosen an LP formulation, we cannot specify such a constraint and instead introduce a small value ϵ and so the actual constraint becomes

$$\sum_i \pi_i (q_i^a + \hat{\alpha}_i(n)) + \beta \sum_{\theta} M_{\theta} \geq \epsilon.$$

The effect of this ϵ is discussed later. In summary, we are still able to find exact solutions to finite horizon problems.

4.8.6 Solution Checking

Unfortunately, the LP defined will frequently return a solution which does not define a point of interest. This happens when the objective function did not minimize all of the M_{θ} variables enough to ensure at least one of the inequalities was actually an equality. If no equalities hold for a particular M_{θ} , then this indicates that there is no point within the currently defined region where our approximation is different from the optimal.

This requires a check after the LP has been solved to verify that it is a point we are interested in. This results when the LP does not drive down the values of the variables M_{θ} in the objective function a sufficient amount to make them actual values that could occur. Recall that these variables were defined as

$$M_{\theta} = \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi, \theta, a)} (n-1).$$

Since we have only constrained these variables with an inequality, it may be the case that none of the inequalities tightly constrain the solution. When none of the inequalities are tightly constraining, then the values for M_θ are not possible (i.e., M_θ are strictly greater than $\sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k (n-1)$ for all a and k) and so an LP solution with these values must be discarded.

There are a number of ways we can detect this condition, but here we describe the two ways that have actually been implemented. The first and simplest way is to make use of the LP's slack variables. We have a series of inequalities for each of the M_θ . As long as one of these, for each M_θ , is binding, then we know we have a solution that we are interested in. If a constraint is binding in an LP, then its corresponding slack variable will be zero. Therefore we simply ensure that at least one slack variable is zero for the set of constraints defined on each M_θ .

An alternative to the above approach needs to be used when slack variables are not available. Some of our LP packages did not make these slack values readily available and so we had to adopt an alternative strategy. In effect, it does the same thing as the slack variable technique, but requires some extra computation. With the slack variable technique when a slack variable is zero it means that the equality holds. Therefore, without slack variables we can simply check to see if the equality

$$M_\theta = \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^{\iota(\pi,\theta,a)} (n-1)$$

holds by computing the right hand side quantity for each θ and a . Again, we only require one of the equalities to hold for each θ .

4.8.7 Witness Discussion

The Witness algorithm can be summarized as follows:

- (1) Using formulas 6 and 7 add the true vectors for the corner points of the belief simplex into $\hat{\mathcal{V}}$.
- (2) Mark all vectors in $\hat{\mathcal{V}}$ as undiscovered.
- (3) Choose an undiscovered vector, $\hat{\alpha}(n)$ from $\hat{\mathcal{V}}$. If there are none then return $\hat{\mathcal{V}}$.
- (4) Mark $\hat{\alpha}(n)$ as discovered.
- (5) For $\hat{\alpha}(n)$ construct the LP (see below).

- (6) Verify that solution to LP is valid. If not, then got to step (3).
- (7) Using the values from LP for all π_i , construct the true value function vector for this point (with formulas 6 and 7).
- (8) Add this vector to $\hat{\mathcal{V}}$ as undiscovered.
- (9) Go to step (5)

The full LP formulation for a vector $\hat{\alpha}(n) \in \hat{\mathcal{V}}$ is

$$\begin{aligned}
& \min : \sum_{\theta} M_{\theta} \\
& \sum_i \pi_i (q_i^a + \hat{\alpha}_i(n)) + \beta \sum_{\theta} M_{\theta} \geq \epsilon \\
& M_{\theta} \geq \sum_{i,j} \pi_i p_{ij}^a r_{j\theta}^a \alpha_j^k (n-1), \quad \forall \theta \in \Theta, a \in \mathcal{A}, k \in V_{n-1}^*(\cdot) \\
& \sum_i \pi_i \hat{\alpha}_i(n) \geq \sum_i \pi_i \alpha_i(n), \quad \forall \alpha(n) \in \hat{\mathcal{V}}.
\end{aligned}$$

We solve one of these linear programs for each combination of vector, in $\hat{\mathcal{V}}$, and action. This is similar to how Cheng finds the corner points for all the regions in the current approximation. As a result we only need to solve a number of linear programs equal to twice the actual number of vectors in the value function, N . The size of these LPs are polynomial in the number of observations and current vectors in the approximation; there are $|\mathcal{S}| + |\Theta|$ variables and $N + |\Theta||\mathcal{A}|M + 2$ constraints. Here M is the number of vectors in the previous time step's value function and N the number of vectors in the current step's value function.

As formulated above, most LP packages will not be able to handle this formulation because they require all the LP variables to be non-negative. Although the π_i variables are guaranteed to be non-negative, the M_{θ} variables are not. However, we can employ the standard substitution trick used in LPs and substitute $M'_{\theta} - M''_{\theta}$, for each variable M_{θ} . We then just need to recombine the two substituted variables to get the real value of M_{θ} . We chose not to make this substitution explicit since it adds nothing but more variables and confusion to the discussion.

Although the simplex method and its variants are not a polynomial time algorithm, the class of LP problems can be solved in polynomial time. Though we use simplex methods in our implementations, the use of one of the guaranteed polynomial time algorithms (i.e., Karmarkar’s method) would ensure that the Witness algorithm is polynomial in the number of states, actions, observations and actual $\alpha(n)$ vectors.

4.8.8 Epsilon

In the Witness algorithm formulation, we are required to find differences between the current approximation and the true value function of at least ϵ . This has some effects on the solutions that the Witness algorithm produces. As discussed, we would actually prefer a strict inequality, but in order to fit it into the linear programming framework we have to settle for a somewhat less pleasing formulation.

This means that for each of the k time steps in a k -horizon POMDP solution we actually construct an approximation to the value function. This complicates matters, because this approximation is then used to find an approximation for the subsequent time steps. Therefore, the error can get compounded. Although at each step we are guaranteed to obtain a value function that is no worse than ϵ from the true value function, this assumes that we are working with the $\alpha(n - 1)$ vectors of the true value function, $V_{n-1}^*(\cdot)$. But if our value function for time step $n - 1$ is just an approximation, our approximation for the $V_n^*(\cdot)$ value function could be worse than ϵ at some point.

However, the error at each point is bounded and so we can actually bound the error for each approximation in the k steps. The fact that we use discounting helps because our errors at each point will also be discounted.

If we are interested in obtaining a policy that is within a specific error bound, we can work backwards to decide what value of ϵ to use. This is only a theoretical argument because, in practice, the actual values of ϵ that can be used will be a function of the software and hardware used in such a system.

We have previously compared the Witness algorithm to other algorithms which provide exact answers to POMDP problem. This might seem a bit unfair since we have thus far been dealing with the Witness algorithm as an approximation algorithm. However, there is a theoretical argument that

actually shows that the Witness algorithm can be used to find exact solutions.

Because we will need to state the POMDP formulation with a finite number of bits (i.e., rationals), we can actually determine the the largest amount of bits required to represent the solution to a specific problem. Given this information we can then determine the value of ϵ to use so that our value function is guaranteed to be less than the actual precision of the true value function which implies that the approximation will indeed be the true value function. Again, this is a theoretical argument that has little to do with actually providing exact solutions.

4.9 Algorithm Analysis

Littered throughout the discussion of the algorithms is discussion of the worst case running times for the algorithms. In this section we will focus our attention on this topic as well as explain how these arise. We will progress through the algorithms in the same order in which we discussed them previously. All of the analysis deals with producing one set of vectors from the previous set. The actual complexity of each algorithm on a complete k -horizon problem is actually much more complicated to analyze since the complexity of each iteration depends upon the complexity of the previous one. We can say that the worst case scenario for the general k -horizon POMDP problem can be as large as:

$$|\mathcal{A}|^{\frac{|\Theta|^{k+1}-1}{|\Theta|-1}}$$

This is the total number of possible policy trees as discussed earlier. If we could construct a diabolical problem where all of these trees were useful for some belief state, then any algorithm we could possibly derive would have to, at a minimum, generate this many vectors (since each vector represents a useful policy tree beginning at the k^{th} time step).

Before we begin the analysis we will just review some notation; M is the minimum number of vectors that are required exactly specify the value function, $V_{n-1}^*(\cdot)$ for the previous time step and N is the same except for the current time step's value function, $V_n^*(\cdot)$.

4.9.1 Monahan's Algorithm

The first step of Monahan's algorithm is to generate all the possible vectors that could be formed from the value function (formula 7). As mentioned there are $|\mathcal{A}|M^{|\Theta|}$ vectors that will be generated. The $M^{|\Theta|}$ term is the number of ways we could fill in the M previous $\alpha(n-1)$ vectors in a summation over all θ . For each of these summations there are $|\mathcal{A}|$ possible actions that can be used to generate a vector. The algorithm will always require generating this many, this is not a worst case analysis.

The next step in Monahan's algorithm is to produce and solve a linear program for each of these to determine whether or not it is actually useful (i.e., part of the optimal value function). The size of these LPs will vary over time since it will eliminate useless vectors as it proceeds. Nevertheless, the first LP will have $|\mathcal{S}|$ variables and $|\mathcal{A}|M^{|\Theta|} + 1$ constraints. For ease of discussion we will let $X = |\mathcal{A}|M^{|\Theta|} + 1$ and note that $N \leq X$. There will be a constraint for each of the generated vectors plus one for the simplex constraint, $\sum_i \pi_i = 1$. Since each vector must be evaluated in an LP, this algorithm must solve X linear programs. For each LP if we discover that the vector is part of the true value function then we will need to include this as a constraint in all the subsequent LPs. However, if there is no belief state for which this vector can be optimal we can discard it and it will no longer generate a constraint in subsequent LPs. Although we know we will have to specify X LPs, we cannot say exactly how big (i.e., number of constraints) each one will be because it will depend on the order in which we choose the vectors.

The best case is when we try all the useless vectors first. This trims down the number of constraints as fast as possible (one fewer constraint per LP). The trimming will eventually stop when we are left with all of the real vectors, but we still need to try them all since we have no *a priori* knowledge of how many there are. These last N LPs will then all have $N + 1$ constraints. One way we could look at the best and worst case complexity of the LP solving process is to count the total number of constraints generated. Then, in the best case, the total number of constraints is

$$\sum_{i=N+2}^{X+1} i + \sum_{i=1}^N (N + 1),$$

which simplifies to

$$\frac{X^2 + 3X - 3N^2 - 5N}{2}$$

In the worst case we will check all the N true vectors first and the number of constraints will not change until we start trying the useless vectors. The total number of constraints in the X LPs in the worst case is:

$$\sum_{i=1}^N (X + 1) + \sum_{i=N+1}^{X+1} i$$

which is:

$$\frac{X^2 + 3X + 2NX - N^2 + N + 2}{2}$$

In both the best and worst case, the overwhelming factor is the X^2 term which shows that for large problems the dominating factor will be $M^{|\Theta|}$.

Eagle's variant of Monahan's algorithm tries to reduce both the number of LPs that have to be solved and the number of constraints in each. Unfortunately, it becomes very hard to analyze the savings since it depends on the specific problem. This algorithm relies on being able to eliminate vectors that are component-wise dominated by previously or subsequently generated vectors. The number of such vectors that will exist cannot be determined in advance since it will depend on the specific parameters of the POMDP problem. Notice that even with this version all the X vectors must still be generated.

Let us assume that when we generate all of the vectors there will be Y of them that are component-wise dominated by others. Right away we get savings in the linear programming step because we now only have $X - Y$ vectors. We can substitute $X - Y$ for X in the previously shown formulas to get new measures for the complexity in this case. However, we also pay a computation price in the comparisons that needed to be done to check for the domination. The actual complexity that is added will depend on the specific way the domination is checked. There are many clever standard search and data structure tricks from computer science that could be used to make this fairly efficient, but the cost must still be factored in.

The last thing to mention about Eagle's optimization is that there do not have to exist any dominated vectors at all. Dominated vectors are useless vectors, but the converse is not true.

4.9.2 Sondik's One Pass Algorithm

By far, the portion of this algorithm that consumes the most time is the linear programming step. Maintaining the search list of belief states and generating vectors from a specific belief state are trivial computations compared to the vast amount of time that must be spent on the large number of LPs this algorithm generates. Therefore we restrict the complexity analysis to the number and size of the LPs that need to be solved.

As discussed, the belief space regions Sondik's algorithm generates are smaller than the true regions in the resulting value function. Therefore, since this algorithm sets up an LP for each of these regions, the number of regions will be the deciding factor. Recall that a region is defined so that all of the $\alpha_a(n)$ and $\alpha^*(n)$ remain unchanged. For this to occur, the index function $\iota(\cdot)$ must remain unchanged. The worst case for this algorithm is when it generates a region for every possible instantiation of an $\iota(\cdot)$ function for each action. This function can be specified as an $|\mathcal{A}|$ by $|\Theta|$ table where each entry in the table can be one of M possible choices (the previous vector indices). Thus, there are

$$M^{|\mathcal{A}||\Theta|}$$

possible $\iota(\cdot)$ functions. However, for each of these functions there may be areas in the belief space where different actions are optimal (i.e., the index function is unchanged, but a different $\alpha_a(n)$ is larger than the rest).

So the number of possible regions is

$$|\mathcal{A}|M^{|\mathcal{A}||\Theta|}.$$

For each of these regions, a single LP is set up, but many LPs are actually solved for each one that is set up. Remember that for each set of constraints generated for an area, we need to check each constraint to see whether or not it is binding on the region. In this way the number of constraints and the number of LPs solved will be exactly the same. This ignores any of the optimizations discussed previously. The number of constraints (and number of LPs solved) is:

$$|\mathcal{A}| + |\mathcal{A}||\Theta|M - |\mathcal{A}||\Theta|.$$

The first term is for the $\alpha_a(n)$ constraints, though there is one less (for the $\alpha_a(n)$ that is equal to $\alpha^*(n)$), we need to add one for the simplex constraint anyway so we can leave it as $|\mathcal{A}|$. The last two terms are for the $\alpha(n-1)$

constraints. The second term is the total number of ways these constraints can be formed and the last term corrects for the cases when $\alpha^k(n-1) = \alpha^{\iota(\pi, a, \theta)}(n-1)$.

So for a single region, the total number of constraints is the the number of constraints (above) squared. This number is large and can be significantly trimmed by implementing some of the clever optimizations discussed earlier, but the sheer number of regions that could be generated will still dominate this and so for large problems we should expect worst case times that are predominantly dictated by $M^{|\mathcal{A}||\Theta|}$ which is actually worse than Monahan's "exhaustive" algorithm which is dominated by only $M^{|\Theta|}$.

4.9.3 Cheng's Relaxed Region Algorithm

This algorithm is the one that is very close to Sondik's and so we choose to only discuss the differences in complexity from that algorithm. Since Cheng defines regions which are typically larger than they need be, he is able to stay away from the explosive growth of regions that Sondik's algorithm suffers from. The Relaxed Region algorithm will only require N regions to be defined which is a significant savings. The number of constraints on these regions is also less, but not significantly. For each region the number of constraints is

$$|\mathcal{A}| + |\Theta|M - |\Theta|.$$

The next step in Sondik's algorithm would be to find the binding constraints, but Cheng opts to find all of the corner points of the region. The region defined by these constraints form a convex polytope in $|\mathcal{S}|$ dimensions. Cheng uses an interior point method to find these vertices, and this is where Cheng's solution suffers a little. The worst case time for finding all corner points in a convex polytope is exponential in the number of faces. The number of faces in the polytope will be dictated by the smaller of either the number of constraints or the number of variables. Although this could be bad, the huge savings in the number of regions makes this algorithm much more practical to use than either Sondik's or Monahan's.

4.9.4 Cheng's Linear Support Algorithm

Like the Relaxed Region algorithm, this one will only need to specify N different regions. The constraints on the regions will always be the number

of vectors in the current approximation plus the simplex constraint. It is difficult to specify how many constraints there will be at each point since it will depend on the nature of the problem and the order in which the regions are processed in. Even the number of constraints on the first region will not be known before hand since the $|\mathcal{S}|$ belief simplex corner points can generate anywhere from 1 to $|\mathcal{S}|$ vectors in the initial approximation. However, at all times we do know an upper bound since we will never have more than N vectors in our approximation. By definition, if we have N vectors in our approximation, \hat{V} , then we have all of the vectors necessary to exactly specify $V_n^*(\cdot)$.

Also, like the Relaxed region algorithm, Linear Support uses an interior point method for finding all of the vertices of the region (i.e., convex polytope) and as a result, inherits the same drawbacks. As stated before, the number of vertices can be exponential in either the number of constraints, at most N , or the number of variables, $|\mathcal{S}|$.

4.9.5 Witness Algorithm

This algorithm's main advantage over Cheng's Linear Support algorithm is that it is not exponential in any of the variables of interest. The difference between this algorithm and Linear Support is that for each of the N regions two LPs are solved. Although for any one $\hat{\alpha}(n)$ we may set up many LPs, we can guarantee that we will only do at most $2N$ LPs. This is guaranteed because when we solve an LP we either generate a new vector in $V_n^*(\cdot)$ or else we discover that $\hat{\alpha}(n)$ currently defines its actual region. We can view this as solving two LPs for each vector: one to discover it and one to verify that it when it defines the true region. It may not be the case that $\hat{\alpha}(n)$ leads to exactly two LPs, but if it leads to more than two, then there must be other vectors that do not lead to any.

The number of constraints in each LP is, in the worst case,

$$|\mathcal{A}||\Theta|M + N + 2.$$

The first term is derived from the M_θ constraints; the second from the region constraints (the number of vectors in the current approximation); and the last term is the improvement constraint and the simplex constraint.

Since this is formulated as an LP, it can be solved in a time that is polynomial in the input constraint size. Recall that the interior point methods

could take time that is exponential in the input size. The one problem with the Witness algorithm is that there is the extra solution verification step. In the worst case it will involve calculating $|\mathcal{A}||\Theta|M$ formulas for each LP (when slack variables aren't available). Also worth noting is that the number of variables for the Witness algorithm is $|\mathcal{S}| + |\Theta|$ whereas for all the other algorithms there were only $|\mathcal{S}|$.

For small problems, the comparative execution speed between Linear Support and Witness will only depend on which is implemented better, but as the problem sizes increase, the Witness algorithm's savings will make it the faster of the two.

4.10 Complications

This section presents some of the implementation complications that we have encountered in our implementations which are not specific to any one algorithm.

As the problem size increases, the value function will typically be comprised of a larger number of vectors. As this number increases the difference between vectors becomes harder to distinguish. The vectors begin to get closer and closer to each other as the solution requires minute differences to differentiate between different regions of the policy. The floating point precision of the existing computers will impose an upper limit on the size of problems that can be solved. After so many decimal places, two slightly different vectors will appear equal to the precision of the machine. Even for modest sized problems this becomes an issue, because there are many instances in the algorithms where these vectors must be compared to each other. Unless extreme care is taken during implementation, this could result in poor algorithms. Round off errors can lead to duplicate slightly different vectors when there should only be one. But if you try to make your implementation distinguish these errors and ignore them, you might be classifying two different vectors as being the same since their difference is not that great.

The heavy dependence on linear programming techniques presents many difficulties. None of these algorithms should be attempted without a stable, robust LP package. Among the more serious problems encountered are those of precision and scaling.

The precision that the LP software uses internally needs to be consistent with those used and assumed by the rest of the implementation. This is

closely related to the machine precision problems.

The LPs generated by the algorithms have constraint coefficients that are based upon the differences between two vectors. As the problem size increases, these vectors will get closer and closer in value making the differences extremely small. This results in LPs with a wide dynamic range of coefficients which causes the LPs to be extremely unstable. Good LP packages will scale the constraints to avoid this instability. If the package used does not provide this functionality, erratic and even incorrect answers can result.

5 Conclusions

5.1 Advantages of POMDPs

In this section we try to show why the POMDP model is worth exploring and why deriving optimal solutions is desirable. Obviously, for the latter, optimality is always more desirable (all else being equal) than sub-optimal policies. Even when the computational costs of finding optimal policies is prohibitive, knowing the form and methods for finding the optimal policy can help guide the process of exploring the space of approximate solution techniques.

Another way in which exact solution techniques might be useful is in sub-components of a larger system. For large world models or environments, finding optimal solutions for every decision process will not be feasible. However, there might be small sub-problems that can be cast in a compact POMDP model, for which we will be able to obtain the optimal policy.

5.2 Drawbacks/Assumptions of POMDPs

In this section we attempt to list the major problems with POMDP models and exact solutions procedures for these models. One of the most glaring deficiencies is expecting an agent to know the model. The POMDP model requires knowing a vast amount of information (e.g., transition probabilities, observation probabilities, rewards, etc.) and all of this supposes that the agent knows the number of states, actions and observations. This is not realistic for many problems and it is highly desirable to have algorithms that perform well without having access to such complete information about the

environment.

Another limitation is that of finiteness of the model. The POMDP model assumes finiteness of the states, actions and observations, but many problems are better modelled with continuous quantities for these. A robot that has a `turn-left` action, can usually be given any angle not a finite set of angles. The same would be true for a robot's location and its sensory readings.

The solution procedures for POMDPs assume that the model is static. It does not permit the model to change over time. Actually, we can model change over time, but it requires exploding the state space to deal with all the possible changes that could occur. While theoretically feasible, this state space expansion is impractical and leads to a model of the world which is very unintuitive.

While the notion of a belief state proves to be convenient for some aspects of the model, it requires time to compute this at run time. For certain applications, the resources to maintain and update belief states might not be available.

The most depressing news of solving the POMDP model is that it is impossible to compute the optimal policy for anything but small problems.

5.3 Contributions

The existing literature on POMDPs suffers from a few problems which this work has addressed. The papers that present specific algorithms [16, 15, 6, 4] are difficult to follow unless the reader has an intimate knowledge of the area to begin with. The survey articles [11, 8] are at too high level to give much insight into the details of POMDPs and the solution procedures. To compound this problem, each author has adopted different notation and terminology than the others. This makes exploring this field very difficult for persons not already familiar with this area.

This paper has tried to present some of the major work in the POMDP area in consistent notation and terminology. The paper's length is a direct result of our attempt to explain the algorithms and intuitions behind them in language that can be understood by persons not already familiar with the area. We have presented the structure of POMDP solutions in many forms and have tried to relate each of these to the others. We believe our approach is helpful for understanding POMDPs and have not seen it done this way in any other work to date. This paper has also attempted to sort out and clarify

many of the existing bugs in the literature.

We have also presented a new algorithm, Witness, for computing optimal solutions to finite horizon POMDPs which does not suffer from the problems of previous solutions. Empirically we have been able to solve problems much larger than any other previously presented work.

The Operations Research community, where most of the POMDP work has come from, does not address the complexity issues when comparing algorithms. In this work we have attempted to classify a number of finite horizon algorithms by their complexity. This analysis identifies the limiting elements of the running times of each and allows comparison between them. This theoretical analysis coincides with the existing empirical performance results, based upon our results and those of [4].

References

- [1] K. J. Astrom. Optimal control of markov decision processes with incomplete state estimation. *J. Math. Anal. Appl.*, 10:174–205, 1965.
- [2] L. E. Baum. An inequality and associated maximization technique in statistical estimation for probabilistic functions of a markov process. *Inequalities 3*, pages 1–8, 1972.
- [3] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, New Jersey, 1957.
- [4] Hsien-Te Cheng. *Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, University of British Columbia, British Columbia, Canada, 1988.
- [5] A. W. Drake. *Observation of a Markov Process Through a Noisy Channel*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1962.
- [6] James N. Eagle. The optimal search for a moving target when the search path is constrained. *Operations research*, 32(5):1107–1115, 1984.
- [7] Ronald A. Howard. *Dynamic Programming and Markov Processes*. The MIT Press, Cambridge, Massachusetts, 1960.
- [8] William S. Lovejoy. A survey of algorithmic methods for partially observed markov decision processes. *Annals of Operations Research*, 28(1):47–65, 1991.
- [9] T. H. Mattheis. An algorithm for determining irrelevant constraints and all vertices in systems of linear inequalities. *Operations Research*, 21:247–260, 1973.
- [10] T. H. Mattheis and David S. Rubin. A survey and comparison of methods for finding all vertices of convex polyhedral sets. *Mathematics of Operations Research*, 5(2):167–185, 1980.
- [11] George E. Monahan. A survey of partially observable markov decision processes: Theory, models, and algorithms. *Management Science*, 28(1):1–16, 1982.

- [12] H. Mukherjee, N. Shahabuddin, and K. Setk. Optimal control policies for partially observable markov processes - a corrected and improved algorithm. Technical report, Indian Institute of Technologu - Delhi, Delhi, India, 1985.
- [13] Anton Schwartz. A reinforcement learning method for maximizing undiscounted rewards. In *Proceedings of the Tenth International Conference on Machine Learning*, Amherst, Massachusetts, 1993. Morgan Kaufmann.
- [14] Satinder Pal Singh, Tommi Jaakkola, and Michael I. Jordan. Model-free reinforcement learning for non-markovian decision problems. In *Proceedings of the Machine Learning Conference*, 1994. To appear.
- [15] Richard D. Smallwood and Edward J. Sondik. The optimal control of partially observable markov processes over a finite horizon. *Operations Research*, 21:1071–1088, 1973.
- [16] Edward J. Sondik. *The Optimal Control of Partially Observable Markov Processes*. PhD thesis, Stanford University, Stanford, California, 1971.
- [17] C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3):279–292, 1992.
- [18] Wayne L. Winston. *Introduction to Mathematical Programming: Applications and Algorithms*. PWS-KENT, Boston, Massachusetts, 1991.