

Most research papers in human-computer interaction contain an experiment and/or evaluation as part of the work. The readings we do in class will explain correct experimental procedure, but these are often not followed by the authors. Published papers are peer-reviewed but not necessary for flaws in the experiments or evaluations. This assignment will expose you to the experiments conducted in papers published at the top-tier conferences in human-computer interaction. If you find an error in a particularly influential HCI paper, you may become famous! Here is an example article that points out numerous flaws in published journal papers: <http://phys.org/pdf235130717.pdf>

You will review a total of 6 full papers published since 2007 from either the proceedings of CHI or UIST that contain experiments and/or quantitative evaluations¹. Select only full papers (typically 10 pages); short papers (typically 4 pages) are okay but two short papers only count as one full paper. Do not review workshop papers, work-in-progress / poster papers, panel papers, extended abstracts, etc. You may need to use the Brown VPN or a computer at Brown to access articles that are not online (try a Google Scholar search first). The list of papers for each conference year is available from the conference program or in the ACM Digital Library, for example: <http://www.chi2007.org/attend/program.php> or <http://dl.acm.org/citation.cfm?id=1240624> under the Table of Contents tab. When you select a paper, quickly skim the paper to see if they match our criteria so you do not waste time reading theoretical papers, heavily qualitative papers, or papers describing methods or systems without an evaluation. At most two students can review a single paper, which you will claim on a first-come-first-served basis by writing your name and the title of the paper on <http://goo.gl/msOg1G>.

Keep a journal (lab notes) of your work as you go. For each paper you review, describe the procedure the authors followed in your own words (about 1 paragraph). Then use the in-class readings to look for errors that invalidate some results in the paper. If you find an error, describe the error in detail and how you would fix it as the experimenter. If you suspect an error but there is insufficient information to fully prove this, such as a bias that could potentially affect the objectivity of the work, please describe that in detail as well. If you are unable to find any errors or problems in the procedure, discuss whether you think the procedure was flawless or if it is simply described without enough detail to validate.

Write a final report by organizing your journal notes and make it public but copying to `/pro/web/web/courses/cs2951-1/[yourcsid]/bughunt.html`. Our goal as a class is to find out whether it is possible to find experimental errors occurring in peer-reviewed published papers, and if so, what are the pitfalls to avoid them ourselves.

This is a big assignment! You have over three weeks for this assignment, but start early. On February 28 (midpoint), you should have reviewed the experimental procedure and results for at least three papers and be prepared to share errors found or your thoughts in class. The assignment grade will be 95% based on your report, and 5% based on the errors you identify.

¹Problematic qualitative work is more difficult to prove and are beyond the scope of this assignment.