

CSC2542 Class Project

Sheila McIlraith
Department of Computer Science
University of Toronto
Fall 2010

Examples

- The next few pages contain examples of what you might need to do for various kinds of projects. These are just samples to get you started thinking ...

Example 1

- Suppose you want to compare algorithm X with algorithm Y
 - E.g., compare Satplan with HPLAN-P
- Has someone already done it? (check the literature)
- Theoretical comparison?
 - Expressivity: are there domains that one algorithm can represent/solve but the other can't (or can't represent as well)?
 - Worst-case complexity: does it mean much in practice?
 - Average-case analysis: probably pretty difficult
- Experimental study?
 - How can you be sure the experiments mean anything?
 - Randomly generated problems? What about built-in bias?
 - Problems from the planning competition?
 - Statistical significance? Confounding factors?
- Overall significance of the results?
 - Does it say something general, or just about one kind of problem?

Example 2

- Suppose you want to apply AI planning to some application X
 - Do you know enough about X to do a credible job?
 - Has someone already done it? (check the literature)
 - Analyze real-world requirements
 - Formalize the problem
 - Which algorithm? What kind of algorithm?
 - Domain-independent, domain- customized, domain-specific?
 - What makes your algorithm better than others for this problem?
 - Maybe compare two algorithms? (see example 1)
 - What will you do if the algorithm doesn't work well?

Example 3

- Suppose you want to extend algorithm X to include feature Y
- E.g., modify Satplan, FastForward or some other planner
 - To include preferences, state variables, temporal reasoning, resource management, control rules, or dependency-directed backtracking
 - To produce conditional or conformant plans
 - To work in multi-agent environments (cooperative? adversarial?)
- Has someone already done it? (check the literature)
- What is the motivation for the feature? (sometimes trivial to address)
- What characteristics does your algorithm have?
 - Soundness, completeness, efficiency?
 - General idea or just a hack?
- Experimental evaluation (see example 2)
- Overall significance?
 - How general is your approach? Could it be made to work for other algorithms as well?
 - Does your approach extend to other algorithms?
 - Good for just one kind of problem, or for many? Which ones?

Some things to think about...

- If you succeed in carrying out your idea, will the result be interesting?
 - Interesting to you? Interesting to others?
- What is needed to carry out your idea
 - Is it too hard to accomplish in the amount of time that you have?
 - Is it too easy to count as a "real" project?
- How do you ensure success regardless of how the result turns out?
 - If the only interesting/significant result is "yes it works," then you'll be in trouble if you can't actually get it to work. You should either
 - (1) think enough about it to be pretty sure you can get it to work, or
 - (2) reformulate it in such that any of the likely outcomes will be interesting as the main result of your project

Assessment Criteria

Your project doesn't have to necessarily make a new research contribution, though this would be nice! The primary objective is pedagogical. Nevertheless, as you train to be good researchers, it is good to evaluate whether your work is a research contribution. To that end, here are some of the questions that are often asked of reviewers assessing your work for publication (ICAPS-07 review form):

Rank (1=bad, 5=middle, 10=good)

- Reviewer familiarity *
- Relevance to the call of papers *
- Technical quality and soundness *
- Originality and novelty *
- Significance to theory and practice *
- Readability and organization *
- Overall recommendation *

Assessment Criteria (cont.)

- ORIGINALITY
 - Is the work described in the paper novel?
- SIGNIFICANCE
 - Is the work important?
 - [Research Track] Does the work present theoretical or experimental results that advance the current state of the art in planning and scheduling?
 - [Application Track] Does the work describe a high-impact application or use of planning and scheduling technologies in an operational setting?
 - [System Track] Does the paper describe a significant integration of diverse component technologies into complex planning, scheduling and execution systems?

Assessment Criteria (cont. 2)

- TECHNICAL QUALITY
 - Is the work technically sound?
 - Are the paper's arguments compelling?
 - Is there a compelling empirical evaluation?
 - [Application Track] Does the work clearly rationalize the use of planning/scheduling technologies in the target application and quantify the benefit over current practice?
 - [System Track] Does the work present a successful integration of planning, scheduling, learning, constraint satisfaction or other technologies into a system that solves a well-defined problem? Does the system represent a well-engineered solution to the problem?
- QUALITY OF PRESENTATION
 - Is the paper clearly written?
 - Does the paper motivate the research or application?
 - Is the paper well organized?

Specific Project Ideas

<we can discuss one on one>