

Preference Elicitation and Explanation in Iterative Planning

Lindsay Sanneman

Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, USA

1 Introduction

Planning for complex scenarios, particularly in which large teams of humans with distributed expertise and varying preferences share a set of resources, poses a number of challenges. While the team as a collective has full knowledge of the task requirements, constraints, and all existing preferences of individuals or subteams, no individual in the team knows the full model of the task and preferences. Such a scenario could be an ideal context to leverage an automated planning agent. However, in many complex domains, there exist context-dependent preferences and constraints that vary with each planning episode, so encoding a static model to represent the planning scenario is not possible.

[Smith, 2012] introduces the Mars Exploration Rover activity planning domain in which a group of science and engineering subteams with potentially competing preferences work together to develop the rover’s tactical activity plan. Throughout the planning process, information is aggregated through complex coordination structures between subteams and a resultant time-consuming iterative planning process. [Smith, 2012] highlights the need to integrate automated planning into an iterative process that begins before goals, objectives, and preferences are fully defined and outlines the technical implications for planning, including the need to naturally specify and utilize constraints in the planning process, generate multiple qualitatively different plans for analysis, and provide explanation of planning decisions.

Given the technical implications laid out by [Smith, 2012], we see three key pieces to solving the problem of providing autonomous assistance through a mixed-initiative planning system when working in a complex domain. First, preferences of individuals or subteams must be elicited for consideration in planning [Berry *et al.*, 2011]. Second, a plan must be generated that takes into account both hard constraints inherent to the problem and soft constraints elicited as preferences [Gerevini and Long, 2006]. Finally, when discrepancies between constraints or differing preferences occur, explanation as to the reason for infeasibility must be effectively communicated back to the humans in the loop such that they can more efficiently work towards replanning together with the autonomous system [Langley *et al.*, 2017]. This process of preference elicitation, optimization, and explanation can be integrated as an iterative process by which teams can converge on the ideal schedule.

Previous work by [Berry *et al.*, 2011] develops a personalized time management agent that iteratively elicits user scheduling preferences, integrates them into meeting scheduling, and improves user preference modeling through an online learning process. One limitation of this work is that preferences are learned over a fixed set of objective terms, which could limit expressiveness of true user preferences. Further, feedback to the user is given in the form of candidate schedules, and no information as to why certain preferences were not accounted for is provided.

We envision using Linear Temporal Logic (LTL) as a common language that provides a natural link between the three components of the iterative planning problem, facilitating both elicitation of expressive preferences and intelligible explanations of the system’s decision-making processes [Kim *et al.*, 2017]. Outputs of each of the preference elicitation, planning, and explanation pieces can be used as inputs to each next step in the process. For example, preferences elicited as LTL specifications are readily inputted to an automated planner as soft constraints for planning, and the plan generated by the planner in combination with the preferences elicited can be used to generate relevant explanations. The system can leverage learned preference information relevant to each team member in providing explanations, aiding the team in converging on a schedule much faster than they can in such processes today. Further, since LTL is readily understandable by both automated planners and humans interacting with the system, it can be used to both describe human preferences as constraints on the planning problem and explain planner decisions in a concise way. In this thesis, we plan to explore the three individual components proposed and the integration of these pieces into an iterative process.

2 Thesis Plan

2.1 Preference Elicitation

Preference elicitation is the first important step in the iterative planning process. While individuals might be able to explain broadly what their preferences entail, doing so in a way that exactly represents their full set of preferences as they relate to other tasks in the plan is difficult. Further, representing expressive preferences in a form that is legible to an autonomous planner poses an additional challenge.

Recent work explores learning of task specifications from

demonstrations [Shah *et al.*, 2018; Vazquez-Chanlatte *et al.*, 2018]. In particular, [Shah *et al.*, 2018] develops a system that uses Bayesian inference techniques to infer LTL task specifications from few demonstrations using specification templates. We will build on this work by taking as input each individual’s requested activities for the planning episode and a set of known domain constraints. We will then generate feasible sample plans in which an individual’s requested activities are scheduled at various times and their scheduled execution times are perturbed with respect to other activities in the plan. Finally, binary classifications of “acceptable plan” or “unacceptable plan” will be elicited from the user for each sample plan, and specifications representing preferences can be learned over the classified plans. After this learning process, users can verify whether the learned specification reflects a set of preferences that they have or not.

2.2 Preferences as Soft Constraints in Planning

In complex planning domains, while context-dependent dynamic constraints and preferences might exist, there often also exist a set of static hard constraints that are known a priori. For example, in the Mars rover domain, science preferences may change each day, but engineering requirements for rover maintenance will persist for every planning episode. In the second part of this work, we aim to develop a technique that incorporates both hard constraints and context-dependent preferences as soft constraints.

[Gerevini and Long, 2006] explore preferences and soft constraints in planning, and recent work by [Kim *et al.*, 2017] encodes high-level user strategies for planning as soft constraints and solves planning problems augmented with these strategies. Outputted plans were found to be of higher quality than plans generated without leveraging high-level user strategies, and plans were more similar to human-generated plans for the same domains than those generated by the planner alone. We see promise in using user preferences as soft constraints in planning problems, and in the second part of this work, we will take as input a planning problem specification and user preferences expressed as LTL specifications from all team members and output a plan that adheres to as many user preferences as possible. One interesting problem that we plan to explore further is how to assign relative weights to each of the team member preferences in the integrated planning process.

2.3 Preferences as Heuristics for Generating Explanations

Once a plan that takes into account individual preferences is generated, it is important both for user trust in the system and for potential plan modification to produce explanations about why the planner made a decision or left certain preferences out. [Langley *et al.*, 2017] motivates the importance of explainable agency and lays out four functional capabilities that explainable agents must have including “explaining decisions made during plan generation” and “communicating decisions and reasons”. Further, [Fox *et al.*, 2017] proposes six questions that users might ask of a system and that an explainable system should be able to answer and offers potential solutions to providing answers to these questions. While these

solutions provide a first step towards providing high-level explanations about system behavior to users, some, such as re-planning to explain why certain actions were not included in the plan, might be time consuming and lack detailed information about why an action was not taken.

We see the possibility to leverage learned user preferences as they relate to other actions in the plan to narrow the search space for an appropriate explanation and to provide relevant explanations to each user who submitted a preference to the planner. Keeping track of which soft constraints were dropped from the produced plan in conjunction with which related activities, as determined by the LTL specifications associated with the soft constraints, were included in the plan could serve as a heuristic to quickly determine what information is most relevant to each user.

3 Contributions and Timeline

This thesis pitch is based on my experiences observing the tactical planning process for the Mars Science Laboratory rover at the Jet Propulsion Lab. Once developed, the technique proposed in this pitch will be evaluated first through experiments with human participants in a synthetic domain, and then will be assessed in a real-world context. I am currently researching techniques for preference elicitation. In the second year of my PhD, I plan to explore integration of distributed preferences as soft constraints in planning. In the third year of my PhD, I will research explanation and integration of the iterative planning process.

References

- [Berry *et al.*, 2011] Pauline M Berry, Melinda Gervasio, Bart Peintner, and Neil Yorke-Smith. Ptime: Personalized assistance for calendaring. *ACM TIST*, 2(4), 2011.
- [Fox *et al.*, 2017] Maria Fox, Derek Long, and Daniele Magazzeni. Explainable planning. *arXiv preprint arXiv:1709.10256*, 2017.
- [Gerevini and Long, 2006] Alfonso Gerevini and Derek Long. Preferences and soft constraints in pddl3. In *ICAPS workshop on planning with preferences and soft constraints*, pages 46–53, 2006.
- [Kim *et al.*, 2017] Joseph Kim, Christopher J Banks, and Julie A Shah. Collaborative planning with encoding of users’ high-level strategies. In *AAAI*, 2017.
- [Langley *et al.*, 2017] Pat Langley, Ben Meadows, Mohan Sridharan, and Dongkyu Choi. Explainable agency for intelligent autonomous systems. In *IAAI*, 2017.
- [Shah *et al.*, 2018] Ankit Shah, Pritish Kamath, Julie A Shah, and Shen Li. Bayesian inference of temporal task specifications from demonstrations. In *NeurIPS*, pages 3808–3817, 2018.
- [Smith, 2012] David E Smith. Planning as an iterative process. In *AAAI*, 2012.
- [Vazquez-Chanlatte *et al.*, 2018] Marcell Vazquez-Chanlatte, Susmit Jha, Ashish Tiwari, Mark K Ho, and Sanjit Seshia. Learning task specifications from demonstrations. In *NeurIPS*, pages 5372–5382, 2018.