Responsive Planning and Recognition for Closed-Loop Interaction

Richard G. Freedman **Yi Ren Fung** and **Roman Ganchin** and **Shlomo Zilberstein**
Smart Information Flow Technologies, LLC
rfreedman@sift.net
University of Massachusetts Amherst
College of Information and Computer Sciences
{freedman, yfung, rganchin, shlomo}@cs.umass.edu

Abstract

Many intelligent systems currently interact with others using at least one of fixed communication inputs or preset responses, resulting in rigid interaction experiences and extensive efforts developing a variety of scenarios for the system. Fixed inputs limit the natural behavior of the user in order to effectively communicate, and preset responses prevent the system from adapting to the current situation unless it was specifically implemented. Closed-loop interaction instead focuses on dynamic responses that account for what the user is currently doing based on interpretations of their perceived activity. Agents employing closed-loop interaction can also monitor their interactions to ensure that the user responds as expected. This demonstration implements an assistive interactive agent that integrates planning, plan recognition, and intent recognition to predict what the user is trying to accomplish and autonomously decide on actions to take in response to these predictions. The interaction will take place in a turn-based simulated game.

1 Introduction

From entertainment to personal assistance, intelligent systems are interacting with people in a variety of applications. However, even when these systems appear to act autonomously and allow the user free will, there is usually extensive back-end development to engineer the interactive experience. Though not as restrictive as expert systems with hand-coded tables of what to exactly do in every considerable situation, there is usually a fixed set of inputs or outputs that is mapped from or to artificial intelligence algorithms. For example, natural language interfaces might perform speech-to-text and then map that text to a set of expected inputs through parsing or machine learning. Likewise, embodied agents might have a preprogrammed finite state machine that specifies what output behavior to perform, and task and motion planning algorithms determine how to execute those behaviors given the current environment's configuration.

Even though these intelligent systems exhibit artificial intelligence and account for the environment and stimuli, they are not actually interacting with an understanding of the user. People act with purpose, explore their environment, make mistakes, and will sometimes change their mind in the middle of doing something. Closed-loop interaction addresses this by *modeling users and making decisions with respect to those models*.

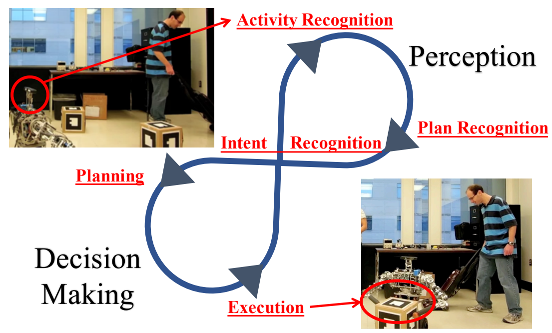

Figure 1: The PRETCIL framework's general flow for how perception and decision making affect each other.

We thus introduce the Planning and Recognition Together Close the Interaction Loop (PRETCIL) framework as a cognitive architecture, which is illustrated in Figure 1. Similar frameworks that integrate planning and recognition for closed-loop interaction either rely on a library of precomputed plans for less robust recognition and execution monitoring (Levine and Williams 2014; 2018) or require negotiation with the user to confirm the agent's understandings and actions (Geib et al. 2016). Instead, PRETCIL iterates indefinitely to *update the recognized intents and plans using its perception and expected responses of the user* while also *revising its decisions of how to act based on these updates*. As a general framework, any appropriate algorithms can be applied to PRETCIL; this demonstration implements PRETCIL using responsive planning (Freedman and Zilberstein 2017) and recognition as planning (Ramírez and Geffner 2010) as its primary components.

2 Perception via Planning as Recognition

When the interactive experience begins, the assistive agent has no model of the user. This also means that the agent is not aware of what the user wants to do and must first observe them in order to make any informed decisions. The user will play a turn-based game in a simulated environment with the freedom to select from a set of completion criteria—successfully satisfying any one of them will result in winning the game. The user will act on the first turn, which provides some information to the agent running our implementation of the PRETCIL framework about which criteria they intend to complete. Then the agent will act on the second turn if it received sufficient information to decide how to respond (in this case, do something that will help the user win the game).

Due to the simulated game setting of this demonstration, user inputs will be limited to discrete button presses and mouse clicks that are easily identifiable without any raw sensor data. Thus this implementation of PRETCIL simply performs activity recognition as a mapping from the input to the game's corresponding action.

The plan and intent recognition components will receive these actions as observations for probabilistic recognition as planning (Ramírez and Geffner 2010). This class of algorithms runs a generative planner to simulate the user solving a variety of problems and then compares their solutions in order to identify which of the completion criteria are most likely. The key assumption applied in recognition as planning is that the user *is acting as optimally as possible to achieve their goal*. This means that the observations either lead to completing the criteria (optimal to perform) or are out of the criteria's way (not optimal to perform). The more likely criteria will have a greater difference between solutions for these two cases.

The plans generated for these comparisons serve as the output for plan recognition, providing information about what the user is expected to do by themself when satisfying each completion criteria. The distribution over the different criteria, computed using the costs of these recognized plans, is the output for intent recognition because it identifies how likely each criteria is motivating the user's actions.

3 Decision Making via Responsive Planning

When deciding how to respond to the user's possible intents, it is important to consider the long-term interaction as much as the current action being taken. This is especially important at the beginning of the interactive experience because the user's initial actions are often relevant to completing multiple criteria. Furthermore, assisting the user towards the completion criteria that they did not select can hinder the experience and reduce the user's trust and willingness to work with the assistive agent.

Our implementation of the PRETCIL framework accounts for this by identifying the necessities (Freedman and Zilberstein 2017), which are shared features between the user's more likely goals. With respect to the distribution over the possible criteria that the intent recognition component provides, this is the weighted sum over the parts of each completion criteria. It is rarely the case that different intents are mutually exclusive of each other; so the agent can assist the user by completing the common tasks that progress towards all the likely completion criteria until the user performs some action that disambiguates their intent.

The necessities generate an intermediate goal for the agent, and the planner used for probabilistic recognition as planning in Section 2 can find a sequence of actions that will accomplish it. The planner generates the plan from current state and assigns actions to both the agent and the user each turn until the intermediate goal is accomplished. This joint solution is not revealed to the user, but is necessary for the agent to realize how the state might change after the user's turns. While this plan is active, the agent executes the next action in response to the user. Like with activity recognition, this demonstration's simulated game setting allows us to simply map the action to animation and update the state.

4 Execution Monitoring

Although the joint plan derived in Section 3 is not conveyed to the user, the actions assigned to the user are assumed to take place in order for the agent's future actions to execute successfully without uncertainty. Our implementation of the PRETCIL framework thus uses this plan as the second use of intent recognition to predict how the user will respond to the agent's actions each turn. If the user's action returned from the activity recognition component matches, then we can assume that the interaction is going smoothly and execute the agent's next action in the joint plan. If the user's action does not match, then there is a chance that the agent recognized incorrectly and needs to reassess the possible completion criteria with this newest user action. This execution monitoring system completes the interaction loop.

5 Conclusion

For less structured interactions between users and intelligent systems, closed-loop interaction that perceives what people do and decides how to appropriately respond is necessary. We introduced the PRETCIL framework as a cognitive architecture for such interaction and explained an implementation of an assistive agent for a game. The demonstration will give users flexibility to play as they choose, and the agent will adapt to their choices using this integration of recognition and responsive planning.

References

Freedman, R. G., and Zilberstein, S. 2017. Integration of planning with recognition for responsive interaction using classical planners. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 4581–4588.

Geib, C.; Weerasinghe, J.; Matskevich, S.; Kantharaju, P.; Craensen, B.; and Petrick, R. P. A. 2016. Building helpful virtual agents using plan recognition and planning. In *Twelfth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 162–168.

Levine, S., and Williams, B. 2014. Concurrent plan recognition and execution for human-robot teams. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling*, 490–498.

Levine, S. J., and Williams, B. C. 2018. Watching and acting together: Concurrent plan recognition and adaptation for human-robot teams. *Journal of Artificial Intelligence Research* 63:281–359.

Ramírez, M., and Geffner, H. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 1121–1126.