

# Hybrid Hierarchical Learning for Adaptive Persuasion in Human-Robot Interaction

Shane Saunderson, *Student Member, IEEE* and Goldie Nejat, *Member, IEEE*

## SUPPLEMENTARY MATERIAL

In Section V-D, we reference other ablation experiments that were run to find the most competitive benchmark to compare the HHL to. Here, we further describe that experiment and the results.

We also considered using different ablations of the HHL as a comparative benchmark: 1) exclusively considering static user states (i.e. NfC), and 2) exclusively considering dynamic user states (i.e. affect). The trials were run with the default persuade (10%) and transition (40%) rates and all other parameters were the same as experiment #1. The static-only ablation uses a CB approach to attempt to learn which abstract action to choose (i.e. central or peripheral) that aligns with the user's static NfC preference, however, must randomly select from one of six possible primitive actions, as it has no capacity to adapt to dynamic affective states. The dynamic-only ablation uses QL to select one of 12 primitive actions in response to the user's dynamic affective state, however, does not consider the user's static NfC preference. The results are presented in Figure 7 below.

The static-only method performs poorly and never accumulates a positive cumulative reward. This is likely because, while this approach may learn a user's static NfC preference, its inability to consider dynamic affective states means that each interaction has, at best, a one in six chance of a successful persuasive attempt. On the other hand, the dynamic-only method performs similarly to (though slightly worse than) the POMDP approach. The two methods are both able to learn successful dynamic affective state and primitive action pairings. However, the POMDP method (which uses a belief MDP approach) performs slightly better likely due to its attempts to form a belief about the user's NfC preference, whereas the dynamic-only method only considers user affective state. Based on this brief experiment, the POMDP approach was therefore used in this study as the most competitive benchmark for the HHL.

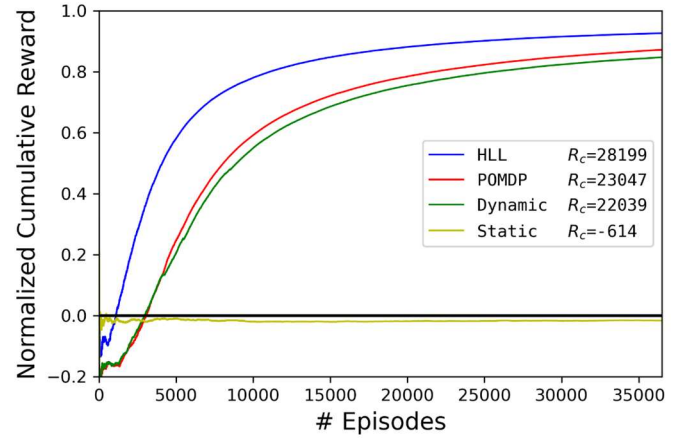


Fig. 7. Convergence graphs for HHL (blue), POMDP (red), dynamic-only (green), and static-only (yellow) methods