Planning With Others in Mind

Nastaran Arfaei (na2481@nyu.edu)
Department of Psychology, New York University, 2-6 Washington Place
New York, NY 10003 USA

Wei Ji Ma (weijima@nyu.edu)
Center for Neural Science and Department of Psychology, New York
University, 2-6 Washington Place New York, NY 10003 USA

Abstract

Planning is rarely done in isolation, but in the presence of other agents who affect the environment and the future action space. This shared nature of environment, cost, and reward is even more crucial to consider when planning collaboratively, where reward maximization is dependent on the actions of all collaborating agents. How do people incorporate the future actions of others in their planning process? We developed a collaborative dyadic turn-taking task with decision sequences up to length 16 to answer this question. We found that people can effectively plan in this context and evidence that they incorporate future potential moves of their partner in their planning process. We constructed and tested computational models of the behavior, among which, a collaborative heuristic search algorithm that simulated and evaluated the future actions of the partner fit the data the best. We also showed specific shortcomings of the competing models.

Keywords: Planning, Collaboration; Collaborative planning; Social Planning; Heuristic Tree Search; Computational; Modeling

Background and Introduction

People live and plan in social environments with intricate interpersonal dynamics, and planning, as a critical cognitive process, is rarely done in total isolation. Since planning requires a prospective consideration of future actions and states (Ho, Saxe, & Cushman, 2022), planning effectively in a social environment requires prediction about the future actions of others (Ho et al., 2022) and the integration of these predictions into the planning process. Previous work has shown that people can infer the beliefs and desires of others (Baker & Saxe, 2011), and predict their actions (Joiner, Piva, Turrin, & Chang, 2017; Khalvati et al., 2019), the process of collaborative planning has been understudied (but see: (Strachan & Török, 2020)). It is not clear whether and how predictions of actions by other agents are incorporated into the planning process. Here we have explored this question by using a novel collaborative decision making task with large state spaces and deep decision trees. We have constructed computational models of collaborative planning that integrate action prediction into the planning process tested these models in predicting human behavior.

Task and research design. We designed a task called "Collaborative Road Construction" (Fig. 1A), in which participants planned a route to collectively connect as many cities as possible within a given shared budget (length of path available). Players started from different map locations (cities indicated as grey dots) and took turns either to connect a city or skip their turn, allowing their collaborator to make the next move. Participants played the single-player version of this game alone, making decisions for both players. We recruited 30 dyads. Each subject played the multiplayer and the single player version of the task in two counterbalanced sessions.

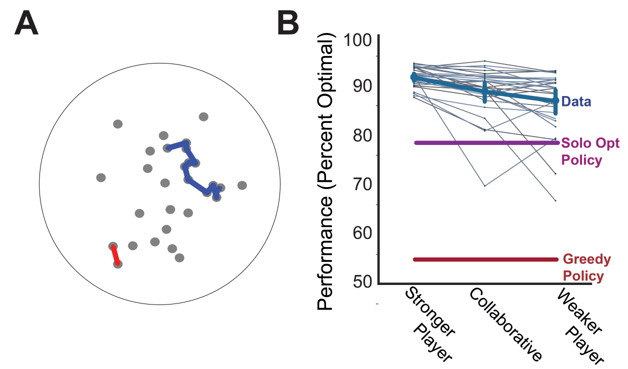

Figure 1: A. Collaborative road construction task. An example of a path taken by a dyad during a trial. B. Performance of participants and policies as a percentage of the optimal solution.

Results

Performance was calculated as the percent ratio of score from optimal. Dyads played well but not optimally, with a performance average between the performance of the weaker and stronger player and better than the best-performing policy that did not consider the collaborator at all (Solo Optimal)($t = 11.27$, $p < 0.001$). Skipping can be interpreted as the strategic preservation of action and state-space, and evidence towards taking the partner's future moves into account during planning. Dyads skipped more than the minimum necessary for an optimal solution ($t = 6.59$, $p < 0.001$). We defined "optimal contribution gap" as the average difference between the number of cities collected by self and other during optimal solutions of each state. Data showed a sharp decline in skip rate and near symmetrical trends in response time at zero contribution gap (Fig. 2.A), suggesting that people account for their partner's available gameplays. We also observed a decline in accuracy in states where more moves were left in the optimal path (Fig. 2.B). The U trend in reaction times implied more deliberation at the start and end of each trial.

Models of Collaborative Planning.

We explored the performance of some decision-making models with and without planning and with and without consideration of the future actions of the collaborator.

Greedy Model assigns value to each reachable city based on the budget it has left after traveling to that city (*b*), plus a normally distributed noise factor (η). It then travels to the highest value city or skips with a given uniform probability (P_{skip}) or makes a random move based on a lapse rate (λ). This model does not take any assumptions about its collaborator into account.

$$V(a) = b + \sigma \cdot \eta$$

Jointly Greedy Model assumes the value of its collaborator's greedy action in addition to its own. It then takes an action if the value of its own greedy action is higher than the collaborator's greedy action, and otherwise skips.

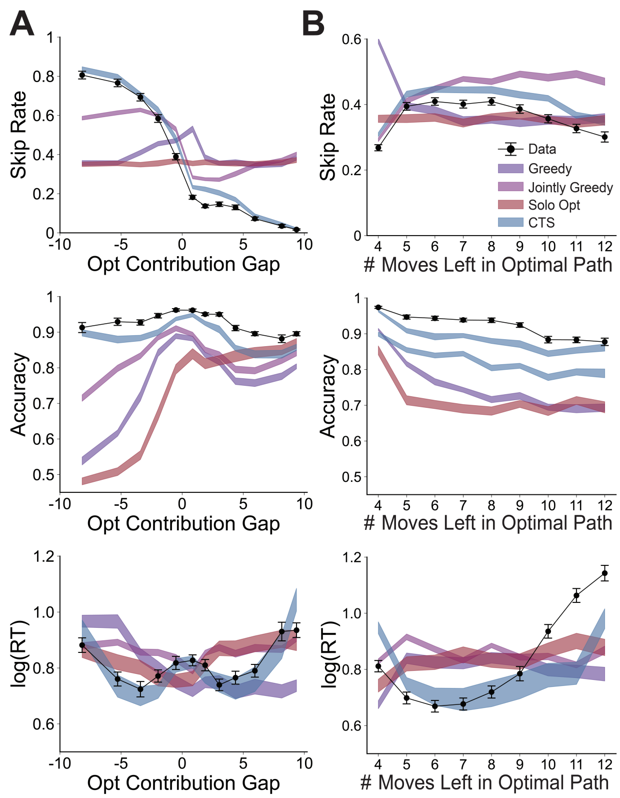

Figure 2: A and B.Data and model simulation summary statistics.

$$V_{\text{self}}(a) = b + \sigma \cdot \eta \quad , \quad V_{\text{Other}}(a) = b + \sigma \cdot \eta$$

Solo Optimal Modell does a full tree search of the gameplays considering only its own future moves, then assigns value as the maximum number of collectible cities (n_{max}) plus noise to each available action. It then either chooses to take the highest value action or skips with a uniform probability or makes a random move based on a lapse rate.

$$V(a) = n_{\text{max}} + \sigma \cdot \eta$$

Collaborative Heuristic Tree Search Model (CTS) uses a value function in conjunction with tree search to evaluate actions for both self and the other (Fig. 3A). Here the value function is defined as the linear combination of the number of cities connected (n_c), square root of the number of cities within reach (n_r), remaining budget and noise.

$$V(a) = n_c + w_r \cdot \sqrt{n_r} + w_b \cdot b + \sigma \cdot \eta$$

The tree search algorithm gradually improves the accuracy of such value estimates by expanding nodes of a decision tree and recursively backpropagating the maximal value of the successor nodes to the predecessor nodes. This model predicts the future action of the "other" using the same value function, weights and tree search mechanism as their own but applied to the current state of the collaborator. The model then evaluates the actions available to "self", including skipping which is valued as the maximum value action available to the "other".

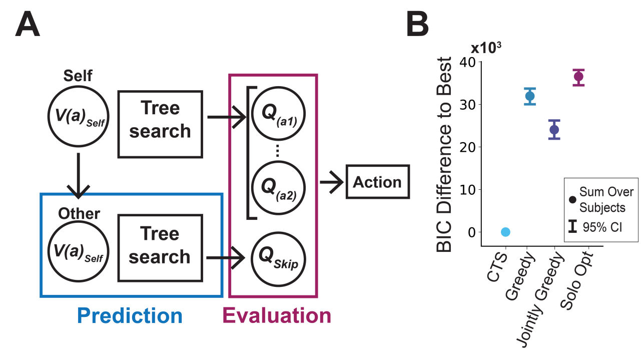

Figure 3: A.Schematic of the CTS model. B. BIC difference between each model and the CTS model, summed across subjects, with 95 percent confidence interval

Model Fitting. We fit the models to the move-by-move data from each participant separately using maximum-likelihood estimation (van Opheusden, Acerbi, & Ma, 2020) and Bayesian adaptive direct search (Acerbi & Ma, 2017).

Model Comparison. The CTS model performs the best with both BIC and AIC lower than the competing models(Fig. 3B). We simulated data for all encountered states using models with best-fitted parameters for each participant. We estimated response times for the models as the value gap (difference between the two available actions of highest value), scaled to match the actual response times average (Fig. 2A and B). We saw a better performance in predicting data trends for the CTS model compared to the competing models. Performance was specifically better for skipping trends and accuracy, at lower optimal contribution gaps, suggesting the combined effect of planning and prediction in the decision process.

Conclusion and Future Direction

This task and the modeling approach here have allowed us to examine the intricacy of the algorithmic mechanisms underlying the intertwined nature of planning and action prediction during collaboration. We revealed empirical evidence for the incorporation of action prediction into the collaborative planning process. We showed that a collaborative heuristic search model performs better than competing models in this context.

Our CTS model assumed the same value function as itself for the prediction of the other. As a next step, we will examine models in which the agent assumes a different value function for the prediction of the future moves of the collaborator. We will then compare the performance of models that use their own selves to predict others (similar to simulation theory) to the ones that use an inference about the other (similar to theory of mind) during their planning process.

References

Acerbi, L., & Ma, W. J. (2017). Practical bayesian optimization for model fitting with bayesian adaptive direct search. *Advances in neural information processing systems*, *30*.

Baker, C., & Saxe, R. (2011). Bayesian theory of mind: Modeling joint belief-desire attribution [Journal Article]. *Proceedings of the Thirty-Third Annual Conference of the Cognitive Science Society*.

Ho, M. K., Saxe, R., & Cushman, F. (2022). Planning with theory of mind. *Trends in Cognitive Sciences*, *26*(11), 959–971.

Joiner, J., Piva, M., Turrin, C., & Chang, S. W. (2017). Social learning through prediction error in the brain. *NPJ science of learning*, *2*(1), 8.

Khalvati, K., Park, S. A., Mirbagheri, S., Philippe, R., Sestito, M., Dreher, J.-C., & Rao, R. P. (2019). Modeling other minds: Bayesian inference explains human choices in group decision-making. *Science advances*, *5*(11), eaax8783.

Strachan, J. W. A., & Török, G. (2020). Efficiency is prioritised over fairness when distributing joint actions [Journal Article]. *Acta Psychologica*, *210*, 103158. Retrieved from

van Opheusden, B., Acerbi, L., & Ma, W. J. (2020). Unbiased and efficient log-likelihood estimation with inverse binomial sampling. *PLoS computational biology*, *16*(12), e1008483.