Inferring user intent with Bayesian inverse planning: Making sense of multi-UAS mission management

Brian Riordan
Sylvain Bruni
Nathan Schurr
Jared Freeman
Gabriel Ganberg
Aptima, Inc.
12 Gill Street, Suite 1400
Woburn, MA 01801
{briordan, sbruni, nschurr, freeman, gganberg}@aptima.com

Nancy J. Cooke
Noel Rima
Cognitive Engineering Research Institute
5810 South Sossaman Road
Mesa, AZ 85212
Nancy.Cooke@asu.edu

Keywords:
Bayesian inference, intent, inverse planning, Markov Decision Processes, Unmanned Aerial Systems

ABSTRACT: *The goal of intent inference is to use observed behavior to predict underlying mental states and causal processes that are likely to have generated the behavior. A potentially powerful technique for inferring intent uses Bayesian inference in structured generative models for planning. We describe our adaptation of the Bayesian inverse planning framework to the multi-unmanned systems mission planning domain. We describe three experiments that elucidate the space of planning priorities in the domain, infer users' goals and priorities given their actions with a planning user interface, and predict users' next planning actions using inferences about their goals and priorities.*

1. Introduction

Inferring the intent of a volitional agent is an ill-posed inductive problem, which necessitates working backward from observations of the agent's actions – which are often fragmentary or incomplete – to a representation of the agent's goals, preferences, and beliefs.

If we view any sequence of actions as a kind of plan, we can apply the framework of Bayesian inverse planning to the problem of intent inference (Baker, Saxe, & Tenenbaum, 2009). We specify a generative model of action sequences that incorporates structured prior knowledge about agents' possible goals and preferences. Using Bayesian inference, we invert this model to capture a representation of the mental model underlying agents' behavior.

In this work we focus on intent inference for planning involving Unmanned Aerial Systems (UAS), as part of the MIMIC project, a Phase II SBIR funded by the Office of Naval Research. UAS are now a core U.S. military capability for intelligence-gathering and other tasks, yet the intelligence that drives planning and decision-making often remains uncodified and firmly lodged in the heads of expert operators.

Here we describe our work using Bayesian inverse planning to learn about the goals and priorities of users during the planning of UAS intelligence, surveillance and reconnaissance (ISR) missions. The ability to infer users' dynamically evolving goals and priorities based on their interactions with user interfaces for planning has a variety of potential applications, including designing decision aids for novice operators or developing autonomous agents for training that mimic and/or complete the actions of expert operators.

In the next section, we describe the Bayesian inverse planning framework and our extensions for capturing priorities. We then present the results of three experiments. In Experiment 1 we elucidated the space of priorities and planning behaviors of naïve subjects in a variety of UAS mission planning scenarios. The results of this experiment were used to parameterize our model. In Experiment 2, we validated the model's ability to account for participants' reported priorities. Finally, in Experiment 3, we explored the model's predictions of the next actions of participants given their observed action sequences and the model's inferences about their goals and priorities.

2. Bayesian inverse planning

Inverse planning is a computational framework for inferring the underlying intent of a volitional agent that is observed performing actions in an environment. In a Bayesian inverse planning framework, the input consists of a hypothesis space of agent priorities and goals and an observed sequence of actions. By specifying a generative model of action sequences, our aim is to infer the *goals* that are most likely to have generated the agent's actions (Baker et al, 2009). Here, we extend this framework to include the *priorities* or preferences of the agent, each goal being a concrete realization of a particular priority Figure 1.

Priorities are high-level outcomes that agents try to optimize by taking actions in pursuit of goals in the environment. We assume that priorities give rise to agents' concrete goals. Goals represent different ways that a priority could be implemented. Goals are shaped both by the agent's priorities and the environment. The environment and goals combine to generate planning actions.

The causal relation between an agent's goals, the environment, and planning actions is specified as a Markov Decision Process (MDP) (Puterman, 2005) based on rational probabilistic planning. Each goal is encoded in a separate MDP. Bayesian inference is used to invert the causal relation between goals, environment, and actions. Information from an agent's observed action sequences is combined with prior knowledge about the agent's goals to infer the agent's most probable goals and priorities at a given point in time.

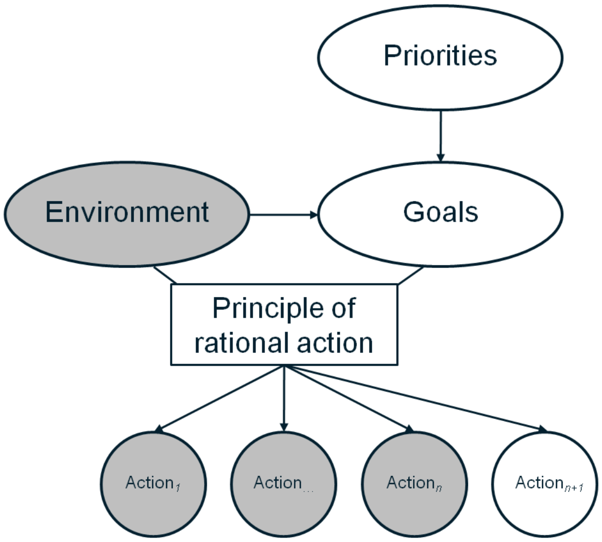

Figure 1. The hypothesized causal structure of goal- and priority-based planning. Shaded nodes represent observed variables; unshaded nodes are latent variable whose values must be inferred.

In our work, "agents" are human participants in experiments performing UAS planning missions. We wish to account for participants' goals and priorities based on our prior knowledge of the UAS planning domain and the observed actions of participants carrying out planning missions.

The core properties of the Bayesian inverse planning framework are:

- Agents can be effectively modeled as approximately rational planners. In attempting to capture the priorities and behaviors of agents, we assume that they will choose the actions that most efficiently lead them to the accomplishment of their goals while maintaining their overarching priorities.
- Inverse planning can be accomplished by integrating bottom-up information from observed data and top-down constraints from a hypothesis space of possible goals. This approach allows inference of an agent's latent goals and preferences, as well as prediction of the agent's future actions.

At the core of the Bayesian inverse planning approach is the model of how actions are generated given the configuration of the environment and the agent's current goals. This is modeled with MDPs that incorporate a mechanism for rational probabilistic planning (Baker et al, 2009). MDPs are a framework for representing an agent's interaction with its environment and computing the optimal action for an agent at any given state. MDPs include a representation of the state of the environment and the agent, the actions that can be taken in the environment, and the rewards or costs for taking an action and transitioning to a new state.

We assume that agents act rationally in making planning decisions, choosing actions that are likely to bring them closer to achieving their goals. However, we assume that agents take actions only probabilistically, so that at any given time we have a probability distribution over possible actions,

$$P(\mathrm{A} \mid \mathrm{G}, \mathrm{E}) \qquad (1)$$

where A is an action sequence, G is a goal, and E is an environment. Actions that are likely to lead to greater rewards are given higher probabilities than less desirable ones. This represents the likelihood of an observed action sequence given our current beliefs about an agent's goals and the state of the environment. A parameter, β, controls the degree to which a participant using the policy chooses actions according to the actions' expected values. High values for β will fit agents whose actions follow the shortest progression through states to the goal encoded in an MDP. Low values for β assume that the agent is noisier in action selection, and hence fit agents whose action sequences deviate from the optimal state sequence.

At each time step, the model infers the posterior probability of each goal. To perform this inference, it compares all possible goal hypotheses (each encoded in an MDP) against each other in terms of how well they explain the agent's observed action sequence. The posterior probability of a goal hypothesis, given observed an Action sequence and the Environment, is:

$$P(\mathrm{G} \mid \mathrm{A,E}) \propto P(\mathrm{A} \mid \mathrm{G,E}) P(\mathrm{G} \mid \mathrm{E}) \qquad (2)$$

$P(\mathrm{G} \mid \mathrm{E})$ is the prior probability of a Goal given the Environment. The structure of the goal hypothesis space and the prior probabilities of goals can be structured in a variety of ways (see Baker et al, 2009). In our work, we make the simplifying assumption that agents have a single optimal state, or goal, that they attempt to achieve throughout a given scenario. Figure 2 is a graphical model that depicts this goal hypothesis and the dependencies between goals, the sequence of observed states, and actions.

Using this generative model of goal-based actions, we can also predict the probability of future actions, given the current observed action sequence and the environment:

$$P(\mathrm{A}' \mid \mathrm{A,E}) = \sum_{\mathrm{G}} P(\mathrm{A}' \mid \mathrm{G,E}) P(\mathrm{G} \mid \mathrm{A,E}) \qquad (3)$$

This is the posterior predictive distribution over actions. The probability of each action a_t is obtained by marginalizing over the posterior probabilities of g given the observed state sequences (cf. Figure 2).

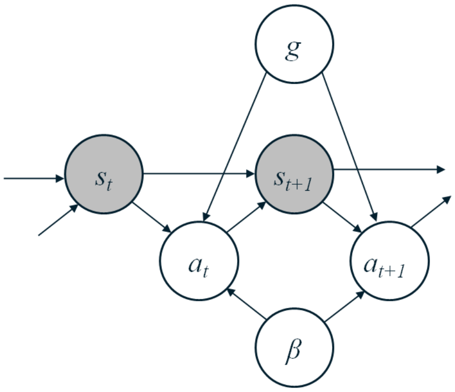

Figure 2. Graphical model of MDP-based probabilistic planning, assuming a single invariant goal. s_t is a state at time t, a_t is an action taken in state s_t, and g is a goal.

We extend the framework of Bayesian inverse planning of Baker et al (2009) by explicitly including priorities, which are updated after each observed action, and by adopting a representation for handling multiple simultaneous goals and priorities.

Our model uses an MDP state representation that allows us to account for multiple simultaneous goals and priorities. The state variable is a feature vector that, along with features describing the configuration of the environment and the agent's state in the environment, includes features that score how well an agent maintains the priorities of interest with each successive observed action. Scores for priorities are computed independently. An agent that maintains an optimal score for one of the priorities indicates that the agent is most focused on a single priority, while an agent that maintains an optimal score for all priorities suggests that the agent's planning is based on equally weighting each of the priorities.

In our model each goal state, encoded in an individual MDP, is associated with a particular priority. The goal states associated with a priority are those states where the score for that priority is optimal. The posterior probability of a priority P is computed by summing across goals associated with a priority and normalizing:

$$P(\mathrm{P}) \propto \sum_{\mathrm{G}} P(\mathrm{G} \mid \mathrm{A,E}) \qquad (4)$$

To model an individual experiment participant, prior to each experimental scenario, we encode each possible goal in a separate MDP. For each MDP, we compute the optimal policy to attain its associated goal from each state using value iteration (Puterman, 2005). As we observe the participant performing actions, the MDP state is updated. Each MDP gives a value for being in the current state if the participant's true goal was the goal encoded in that MDP. The observed action sequence at each time step is combined with the model's current beliefs to infer the most probable goals and priorities for the participant.

3. Experiment 1: Exploring the space of priorities

Experiment 1 was part of a detailed analysis of the domain of multi-UAS air tasking order planning. Experiment 1 was designed to elucidate and record common mission planning priorities, strategies, and action sequences. The data we collected on priorities allowed us to create priority and goal representations in the inverse planning model that covered much of the space of possible priorities and goals in the planning environment.

3.1 Method

Participants. Forty-two undergraduate and graduate students, aged 18-69 years (average 30 years ± 12.5), participated. Some participants had familiarity with the UAS domain, but none were experts in mission planning.

Stimuli. The multi-UAS simulation environment was generated using Aptima's Dynamic, Distributed, Decision-Making (DDD) software. Participants created plans using the interface shown in Figure 3. In the planning tool, a plan consists of a sequence of plan items, each of which is an assignment of a UAS to a Target with a particular set of operating characteristics, including Priority, Speed, Launch Time, and Endpoint Action. The Priority element determines the order in which plan items referring to the same UAS are to be executed. Using the map display, participants can add waypoints to create specific routes – for example, to fly around regions of bad weather (red polygons).

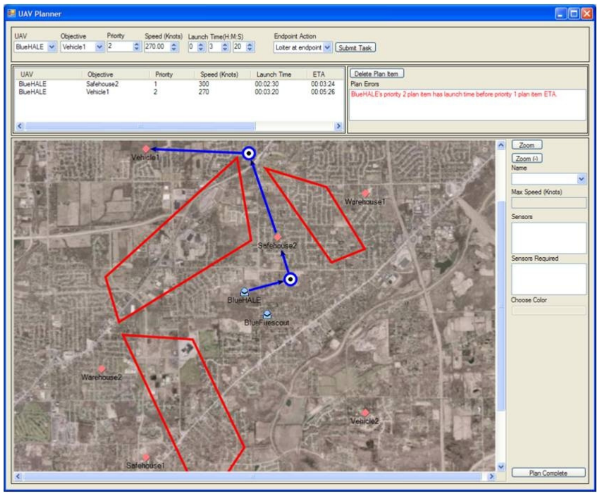
UAV Planner

UAV: BlueHALE
Objective: Vehicle1
Priority: 2
Speed (Knots): 270.00
Launch Time(H:M:S): 0 3 20
Endpoint Action: Loiter at endpoint

UAV	Objective	Priority	Speed (Knots)	Launch Time	ETA
BlueHALE	Safehouse2	1	300	00:02:30	00:03:24
BlueHALE	Vehicle1	2	270	00:03:20	00:05:26

Plan Errors
BlueHALE's priority 2 plan item has launch time before priority 1 plan item ETA.

Figure 3. UAS mission planning interface used in all three experiments.

Procedure. The experiment duration averaged three hours, and included three phases: training, data-gathering, and debriefing. In the training phase, participants underwent a 45-minute training session, and completed two practice scenarios. In the data-gathering phase, participants underwent multiple simulation scenarios in which they were tasked with creating a mission plan for two UAS in a hostile environment, based on a set of mission objectives. Participants were responsible for assigning UAS to targets, setting operating characteristics and constructing the flight routes. Debriefing involved a cognitive walkthrough during a video replay of the last scenario performed by the participant. The participant described out loud what they did during said scenario, what they tried to accomplish, what their objectives were, and what difficulties they encountered.

Design. This experiment included two independent variables: planning time (that is, how long the participant was allowed to plan the missions of the UAS) and scenario complexity. There were two levels of planning time: short (3 minutes) and long (6 minutes). There were three levels of scenario complexity: low (10 targets and few Rules of Engagement (ROE) constraints), medium (20 targets and few ROE constraints), and high (20 targets and many ROE constraints). ROE constraints typically included the number of weather zones to avoid, the maximum total flight time, fuel allowance, and the number and duration of actions to perform at targets. A repeated-measures design was implemented, in which all participants saw three replicates of the six conditions (two planning times by three scenario complexities), yielding a total of eighteen trials per participant. Blocking of the randomized replications was counterbalanced across participants.

3.2 Results

We collected a suite of subjective and objective measures designed to reveal users' strategies and priorities in multi-UAS mission planning under various contexts. Here we report on the results concerning participant priorities elicited following the cognitive walkthrough. Participants were asked to mention and rank-order what they considered to be their planning priorities. Fourteen subjective planning priorities were described by participants. Figure 4 displays the priorities, the number of participants who mentioned each priority, and each priority's total score based on the rankings provided by participants. The scoring metric is a weighted function where the score is increased by 5, 4, 3, 2 or 1 point(s) when the priority is ranked respectively 1st, 2nd, 3rd, 4th or 5th. *Use closest targets* was, by far, the prevalent priority, measured by frequency or score. *Use adequate covertness* and *Avoid weather areas* came in second and third position respectively. All other priorities were rarely cited.

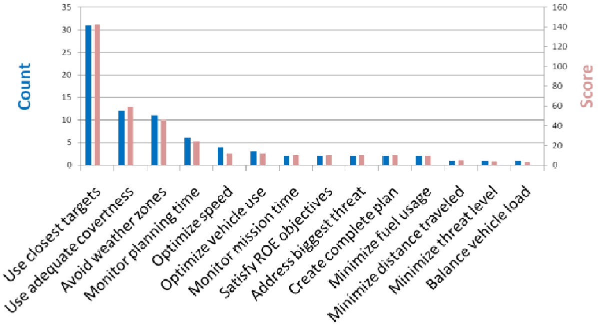

Figure 4. Count and score for each priority mentioned by participants.

3.3 Implementing the inverse planning model for the UAS mission planning domain

The experimental data indicated that a relatively small number of priorities and constraints dominated the strategies of the participants. Three priorities were selected among them, to be represented in the Bayesian inverse planner: (1) Avoiding weather zones (avoiding weather); (2) planning based on proximity of UAS and target (proximity); and (3) assigning UAS to targets following the covertness requirements for targets specified in a scenario's ROE (covertness).

We assumed that each action taken by a participant was an effort to maintain a priority and achieve a goal state. Based on how the participant's actions change the state of the environment, the model classifies each action into one of the following categories:

- *New plan: avoid weather.* A UAS is assigned to an available target such that the route avoids weather zones.
- *New plan: proximity.* A UAS is assigned to an available target that is closer to it than to the other UAS.
- *New plan: maintain covertness.* A UAS is assigned to an available target such that the UAS chosen meets the covertness requirements specified in the scenario.
- *Plan modification: avoid weather.* The user alters an existing plan so that the trajectory of the UAS avoids weather zones.
- *Plan modification: proximity.* The user implements modifications based on proximity. These modifications include both the objective (i.e., change the target to a different, unassigned target to minimize distance traveled) and the assigned UAS (i.e., assign a different UAS to fly to the target to minimize distance travelled).
- *Plan modification: maintain covertness.* The user changes the target of an existing plan so that a UAS is assigned to a target based on the covertness requirements specified in the scenario.

4. Experiment 2: Priority inference

Experiment 2 explored the inverse planning model's ability to account for people's high-level priorities during UAS mission planning. After completing a planning scenario, participants were asked to report their priorities for the scenario. We tested whether the model's inferences about participants' priorities – based on a set of common priorities for mission planning derived from Experiment 1 – were able to predict the participants' reported priorities.

4.1 Method

Participants and stimuli. 15 participants, aged 18-47 years (average 27 years ± 8.3), carried out planning actions in same scenarios as were used in Experiment 1. Some participants had familiarity with the UAS domain, but none were experts in mission planning. The stimuli and the conditions that participants experienced were the same as Experiment 1.

Procedure. Following each planning scenario, each participant filled out a priority rating grid. The grid listed six priorities from Experiment 1, including the top four priorities from Figure 4. For each of the six priorities, participants were asked to give ratings of the priority's importance during planning and the time they spent attending to carrying out the priority. To reduce the possibility of participants giving the same ratings to priorities after each scenario: 1) the vertical ordering of the priorities was randomized for each scenario; 2) the participants' ratings for the previous scenario were obscured. This experiment focused on modeling people's ratings in response to the question: How important was this priority when planning this mission? The rating scale ranged from "1 = not important at all" to "5 = very important".

Modeling. Our inverse planning model accounts for a subset of the priorities included in the ratings questionnaire: *covertness, proximity,* and *avoiding weather.* We focused on modeling these priorities in this experiment; modeling the other priorities for which we collected data is future work. Model predictions were probability distributions over these priorities.

The model was run on the planning sequences of each participant in each scenario. After each new plan or plan modification by a participant, the model inferred a probability for each of the three major planning priorities: maximizing covertness, using proximity, and avoiding weather zones. This probability distribution models the subjective importance attached to each planning priority by the current participant at a particular point in the planning process. We compared the correlation between the human priority ratings with the inferred distribution over priorities at the end of the scenario (i.e. after the last action performed by the participant). In order to convert the participants' ratings to the same scale as the model predictions, ratings were normalized to [0, 1]. For example, a rating of 5 (very important) became a 1.0, while a rating of 0 (not important at all) became a 0.0.

4.2 Results

For our analysis, we first explored the parameter space of the β parameter to find the best-fitting model for the human priority ratings data in terms of the correlation between human priority ratings and model priority predictions.

We explored five β parameter settings from 0.5 to 2.5. Table 1 shows the correlations between the human priority ratings and the model's inferred priority distributions for each parameter setting. The highest correlation, 0.315, was obtained for $\beta = 0.5$.

Table 1. Correlations r between human and model priorities across the β parameter space.

β	0.5	1.0	1.5	2.0	2.5
r	0.315	0.283	0.264	0.256	0.248

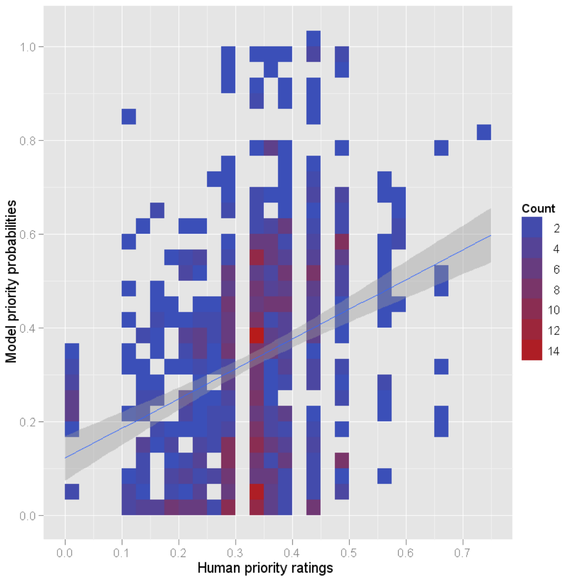

Figure 5. Scatterplot/heatmap of *human priority ratings* and *model priority predictions* for all participants and all scenarios. Larger numbers of points are shown in red. The best-fitting linear regression line is also shown (blue line), with 95% confidence interval (grey cone).

Figure 5 displays a scatterplot of the human priority ratings and model predictions for the best-fitting model. The best-fitting linear regression line clearly increases from left to right, reflecting the positive correlation of human data and model predictions.

Table 2 shows the correlations of human ratings and model predictions for each priority. The model predicted human ratings for the *covertness* priority with the most accuracy. Ratings for *proximity* and *avoiding weather* were more difficult to predict. The reason for this pattern may be a discrepancy between participants' priorities and what the model observed about the participants' behavior using the planning interface. For example, participants may not have always closely monitored when their planning actions corresponded to carrying out the *proximity* or *avoiding weather* priorities. The model, on the other hand, keeps close track of all actions, noting when any priority has been violated. As a result, at the end of a scenario, participants' recollections of the actions they took, as reflected in their priority ratings, and the model's history of those actions may sometimes be at odds.

Table 2. Correlations r of human priority ratings and model priority predictions for each priority.

Priority	*Covertness*	*Proximity*	*Avoiding weather zones*
r	.562	.365	.359

We also analyzed the variability in the correlations of human and model data across scenarios and participants. Table 3 displays the ranges, means, and standard deviations of the correlations for scenarios and participants. There was a much greater variability between participants compared with scenarios. This may suggest that participants' attention to the rating task varied.

Table 3. Ranges, means, and standard deviations of the correlations between human priority ratings and model priority predictions.

	Scenarios	*Participants*
***r* range**	.09 - .48	-.12 - .72
***r* mean (st. dev.)**	.31 (.11)	.32 (.23)

These results address important questions about the model's parameters and assumptions in the context of the task of predicting latent priorities. First, to what extent do the model results depend on the setting of the β parameter? We saw that a low setting of the β parameter gave the best fit to the human data. Still, the magnitude of the range of correlations achieved by the different settings of β (.067) was not large, indicating that the model is not overly sensitive to the choice of β. Second, what does the best-fitting β parameter (0.5) say about the fit of the model's assumptions of rationality to the human data? For this dataset, the best-fitting β value was low. As Figure 5 showed, model predictions often did not match human ratings well. The low β setting allows the model to account for the human data better by adjusting the model's evolving posterior probability distribution over goals and priorities more conservatively after each action, essentially integrating observed human actions into the model more slowly over time.

4.3 Discussion

In this experiment we demonstrated that the MIMIC model of inverse planning for the UAS planning domain can predict human planners' high-level priorities. These results extend the work of Baker et al (2009), showing that a Bayesian inverse planning framework can not only account for concrete human goals in an environment, but also the latent preferences that guide the selection of goals to pursue.

The space of possible high-level priorities and concrete goals in this domain is very large, and there are potentially many ways of encoding these factors in an inverse planning model. Here, we showed that our simple inverse planning model, coupled with a straightforward representation of multiple simultaneous high-level priorities, was sufficient to achieve significant correlations with planners' ratings of their priorities.

Although Experiment 2 showed the ability of our approach to account for human priorities in a planning domain, applications such as behavior modeling or decision aiding may not have as their primary interest the inference of human priorities, but rather the accurate prediction of future human behavior. In the next experiment, we tested the ability of the MIMIC model to use inferred priorities and goals to make predictions about the subsequent actions of humans.

5. Experiment 3: Action prediction

Our third experiment sought to use the MIMIC model to account for the next actions of planners. Participants' actions in the planning interface were classified by human coders as particular high-level actions as in the MIMIC model. The question that was addressed in the current experiment was, then, given the model's inferences about a planner's priorities and goals at a particular time – and the model's assumption of approximately rational planning – could the model accurately predict the planner's next planning action?

5.1 Method

Participants, stimuli, design, procedure. 8 participants, aged 18-54 (average 26 years ± 9.3), were exposed to the same scenarios as in the previous experiments. Some participants had familiarity with the UAS domain, but none were experts in mission planning. All aspects of the stimuli, the conditions that participants were exposed to, and the experimental procedure were identical to Experiment 1.

Modeling. Each action of each participant was coded by human coders. Coders assigned a probability that the human action fell into each of the six action types represented in the MIMIC model (e.g. *New plan: Avoid weather*; see Section 3.3 for a list of action types). Model predictions took the form of probability distributions over the six actions encoded in the model. Correlations between the human and model were analyzed across participants and actions.

5.2 Results

First, we explored the parameter space of the β parameter to find the best-fitting model for the human data. In contrast with Experiment 2, we searched for the best fit by comparing the human action ratings with the model action predictions. As in Experiment 2, we explored five β parameter settings from 0.5 to 2.5. The correlations between the human action ratings and the model action predictions are shown in Table 4. The highest correlation was .385, when β = 1.0. The magnitude of the range of correlations was similar to that found in Experiment 2; a broad range of β parameter settings produced similarly good fits to the human data.

Table 4. Correlations between human action ratings and model action predictions across the β parameter space.

$\boldsymbol{\beta}$	0.5	1.0	1.5	2.0	2.5
$\boldsymbol{r}$	.372	.385	.373	.356	.338

Using the best-fitting model, for each time step, we compared the human's action distribution to the model's predicted next action distribution from the previous time step. Figure 6 depicts the positive correlation between human and model data. Note that the human action ratings are very often either 1.0 or 0.0 – meaning the coders thought that actions belonged only to a single action category – while the model predictions are more graded. The best-fitting model's β parameter was low because the model had to explain this bimodal distribution as noise, with many actions deviating from a deterministic path through hyperspace to a set of goals.

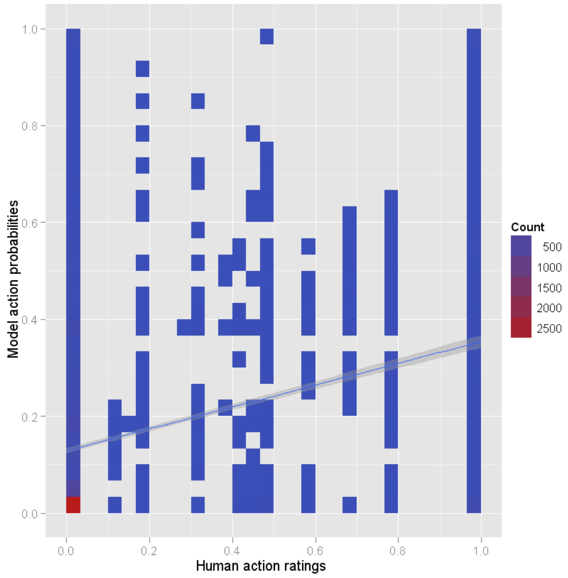

Figure 6. Scatterplot/heatmap of human action ratings and model action predictions for all participants and all scenarios. Larger numbers of points are shown in red. The best-fitting linear regression line is also shown (blue line), with 95% confidence interval (grey).

Preliminary analysis indicates that participants often planned one action and worked to refine it immediately ("sequential" planning). The model, however, often predicted *New plan* actions ("batch" planning), since these actions would quickly bring the participants closer to the goal of assigning UAS to each target. These results indicate that the rationality encoded in the model may better fit certain planning styles.

5.3 Discussion

Experiment 3 demonstrated that the MIMIC model can predict many of the high-level actions of human planners in the UAS mission planning domain. We have begun to identify areas where the model makes errors in its predictions regarding next actions. We expect that further analysis will clarify the source of these errors and allow the model to be modified to produce even better fits to the human action data.

6. Conclusion

We presented three experiments from the multi-UAS mission planning domain that point to the promise of the Bayesian inverse planning framework for using human user actions to infer their likely goals, priorities, and future actions. We demonstrated that a Bayesian inverse planning framework can be used to account for both concrete goals and latent priorities that underlie these goals.

In this work, we implemented a simple inverse planning model with structured prior knowledge of the domain. This model was limited to using users' actions to infer a single invariant goal representation across each scenario. In ongoing work, we are exploring more richly structured goal priors. We are also exploring how the current model and extensions can benefit from prior probabilities for goals and priorities learned from expert UAS mission planners. A third direction we are pursuing is the development of more detailed representations for the UAS mission planning domain that can afford more fine-grained inferences about goals, priorities, and future actions, using factored and partially-observable MDPs.

7. Acknowledgements

We thank Marc Steinberg of the Office of Naval Research (ONR) for his support of this research. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the view of ONR or the U.S. government.

8. References

Baker, C. L., Saxe, R., and Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition,* 113, 129-349.

Puterman, M. L. (2005). Markov decision processes: Discrete stochastic dynamic programming. New York: John Wiley and Sons.

Author Biographies

DR. BRIAN RIORDAN is Research Scientist at Aptima, Inc. He specializes in machine learning, statistical natural language processing, and cognitive modeling. His recent research includes machine learning approaches to inferring user intent and predicting human performance, large-scale text-analytics, and pattern recognition in multi-source intelligence data.

DR. SYLVAIN BRUNI is Human Systems Engineer at Aptima, Inc., where he provides expertise in human-automation collaboration, interface and information design, and the statistical design of experiments. His research focuses on designing and testing collaborative decision-support systems used in military command and control for multiple unmanned systems.

DR. NATHAN SCHURR is Senior Research Scientist and leads the Artificial Intelligence Team at Aptima, Inc. His interests center on allowing humans and artificially intelligent systems to collaborate and interact for superior mission performance.

DR. JARED FREEMAN is Chief Research Officer at Aptima, Inc. Dr. Freeman investigates human problem solving and decision making in real-world settings, and defines methods of monitoring and managing these processes using modeling and training applications.

GABRIEL GANBERG is Senior Software Engineer at Aptima, Inc., and the architect of the DDD®, Aptima's user-configurable simulation for team decision-making research and training.

DR. NANCY J. COOKE is Professor of Cognitive Science and Engineering at Arizona State University Polytechnic and is Science Director of the Cognitive Engineering Research Institute. Her research focuses on the assessment of team cognition in a number of applied settings including ground control of remotely piloted aircraft.

NOEL RIMA is a Research Assistant at the Cognitive Engineering Research Institute in Mesa, AZ.