Where Do You Think You're Going?: Characterizing Spatial Mental Models from Planned Routes

BRANDON S. PERELMAN, ARTHUR W. EVANS III, and KRISTIN E. SCHAEFER,
CCDC Army Research Laboratory, Aberdeen Proving Ground, MD, USA

Route planning is a critical behavior for human-intelligent agent (H-IA) team mobility. The scientific community has made major advances in improving route planner optimality and speed. However, human factors, such as the ability to predict and understand teammates' actions and goals, are necessary for trust development in H-IA teams. Trust is especially critical when agents' behaviors do not match human team members' expectations, or the human cannot understand the agent's underlying reasoning process. To address this issue, the artificial intelligence community has pushed toward creating *human-like* agent behaviors using machine learning. The problem with this approach is that we do not yet have a clear understanding of what constitutes human-like behavior across the breadth of tasks that H-IA teams undertake. This article describes an investigation and comparison of human and agent route planning behaviors, the interplay between humans and agents in collaborative planning, and the role of trust in this collaborative process. Finally, we propose a data-driven methodology for characterizing and visualizing differences among routes planned by humans and agents. This methodology provides a means to advance compatible mental model metrics and theory by informing targeted transparency manipulations, thereby improving the speed and quality of routes produced by H-IA teams.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; • **Computing methodologies** → **Planning and scheduling**; • **Computer systems organization** → **Embedded and cyber-physical systems**; **Robotics**;

Additional Key Words and Phrases: Human-autonomy teaming, human-robot interaction, mental models, spatial planning, spatial cognition, spatial navigation, route planning, GPS systems, navigation systems, autonomous systems

ACM Reference format:

Brandon S. Perelman, Arthur W. Evans III, and Kristin E. Schaefer. 2020. Where Do You Think You're Going?: Characterizing Spatial Mental Models from Planned Routes. *ACM Trans. Hum.-Robot Interact.* 9, 4, Article 23 (May 2020), 37 pages.
<https://doi.org/10.1145/3385008>

Research was sponsored by the CCDC Army Research Laboratory and supported by Oak Ridge Associated Universities (ORAU) through Cooperative Agreement W911NF-16-2-0008. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the CCDC Army Research Laboratory or the US Government. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Authors' addresses: B. S. Perelman, A. W. Evans III, and K. E. Schaefer, Army Research Laboratory, 7101 Mulberry Point Rd, Bldg. 459, Aberdeen Proving Ground, MD 21005-5425; emails: {brandon.s.perelman.civ, Arthur.w.evans20.civ, Kristin.e.schaefer-lay.civ}@mail.mil.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-9522/2020/05-ART23

<https://doi.org/10.1145/3385008>

1 INTRODUCTION

Route planning is a critical research focus task for future human-intelligent agent (H-IA) teaming. Historically, route planning is treated as an optimization problem where the goal is to find the best solution according to a given criterion (e.g., the shortest path) [1], however the underlying mechanisms used by humans and intelligent agents (IAs) to optimize routes are typically studied in isolation. More specifically, while research exists on human acceptance and rejection of IA-generated routes [2–6], there is comparatively little information regarding the characteristics that determine the extent to which an IA-generated route will be acceptable to, and predictable by, humans. This is important to future teaming initiatives, because research has shown that humans and optimal route planning algorithms often produce different solutions to a given spatial problem [7]. This incompatibility between the human’s preferences and the algorithm’s behavior has implications for the predictability and acceptance of autonomous systems’ route planning behaviors. The rising prevalence of interdependent human-autonomy teams (in which the IA is either a computer agent or embodied robotic agent) will necessitate IAs that can exhibit spatial planning behaviors that are predictable and trustworthy to human teammates. Research in autonomy has acknowledged that human behaviors tend to converge based on generally accepted sets of rules or mechanisms (that is, the field has recognized the need to produce behaviors that are accepted by humans, *generally*). Here, we demonstrate that this convergence is not unanimous, and we argue that algorithms intended for H-IA teams should be sensitive to potential discordance among teammates.

Mental model theory [8] (see also preferred mental model theory [9]) proposes that humans execute spatial reasoning tasks in phases. During the first phase (comprehension), humans construct a mental model from the information available (e.g., a map or description). As an example, during this phase the environment might be represented as, “point A is closest to my starting point.” Next, during the description phase, they inspect this model to derive novel information not readily apparent; for example, “the obvious route to point A traverses a main road.” During the final phase, validation, the human explores alternative models. For example, “I plan to embark during rush hour. There is also a back road that will lead me to point A and may be faster.” Consistent with the definitions of mental model theory and related works [7, 10–13], we define a spatial mental model here as a mental representation of the environment (i.e., spatial problem) and extend the definition to include all pertinent information and problem solving mechanisms as perspectives applied during the planning process. This therefore includes the human’s understanding of the physical problem space, availability of information about the environment, bottom-up (sensory) and top-down (cognitive) planning mechanisms including heuristics, and the mission criteria and their weightings.

To permit more *human-like* behavior that is compatible with users’ spatial mental models, and to allow autonomous systems to be retrained by humans in the field, researchers have begun to consider machine learning techniques that allow the human to teach the autonomy, e.g., Reference [14]. However, there are potential challenges to the use of these machine learning approaches for effective H-IA teaming. The first challenge is human variability, such that a group of humans typically produces a wide variety of routes during planning and execution for an autonomous system based on individual differences (e.g., information access and information weighting) and task-specific factors such as changing mission considerations [7, 10, 15], as well as team composition [16]. We take these behavioral differences as evidence of differences in humans’ underlying spatial mental models. From the engineering perspective, this serves to obfuscate the nature of *human-like* behavior and makes it difficult to create route planning behaviors that generalize well across human teammates across contexts. A second challenge lies in human-autonomy interaction. Human mental models of robot capabilities are frequently inaccurate and are often influenced by

superficial morphological characteristics or irrelevant prior experiences (for a review, see Reference [17]). Acquiring mental models of a robot's capabilities, such as its information processing mechanisms, mobility, and decision-making capacity, is critical for building appropriate trust in autonomous teammates [18] and employing those assets appropriately. To summarize, if the human team members' expectations of the IA's behaviors do not match the IA's actual behaviors or intent (even if the intended actions are optimal), the opportunity for trust degradation and potential disuse of the IA can develop [19, 20].

These challenges may be mitigated to some extent via communication and experience. In human teams, teammates may develop shared mental models¹ that can facilitate prediction of team member actions, support adapting decisions based on team demands, and coordinate actions especially in circumstances where communication is limited [21]. Trust and transparency research suggests that trust can be retained if team members understand underlying intent [22–24]. The difficulty in understanding and communicating intent is that humans and IAs trained by machine learning may not have declarative access to the reasons they choose a particular route or action. Therefore, new processes are needed for communication between humans and IAs that characterize, to the extent possible, the underlying reasoning processes and criteria considered during route planning. The ability to characterize a human's spatial mental model, or an IA's underlying reasoning process, can be used to communicate that information among team members to better understand and predict teammates' behaviors. To address these current limitations, our current research investigates human and IA approaches to solving route planning problems to (1) quantify the similarity among human and IA planned routes, (2) test the dynamics of replanning in a mixed-initiative system, (3) understand the role of trust in those dynamics, and (4) develop a data-driven methodology for inferring the characteristics of spatial mental models from the routes.

This research is important, because the ability to predict differences in route planning among members of human-autonomy teams can assist in the identification of situations where potential breakdowns in trust could occur. Furthermore, understanding where the teammates' route planning expectations and priorities differ can provide a means to mitigate these potential breakdowns through bidirectional communication and system transparency. This work aims to remove subjective bias in descriptions of an individual's mental models by providing evidence for a new process for quantifying and comparing general characteristics of different routes, which we assume are a product of distinct spatial mental models. This process results in unbiased characterization of the criterion that both humans and IA team members may use when planning a specific route. By knowing precisely how teammates' conceptualizations of ideal routes differ, it is possible to develop ways to communicate intent reasoning so that team members can understand one another's spatial decision-making process, thereby mitigating that degradation in trust. The core of the problem lies in differences in the ways that humans and IAs approach the route planning problem.

2 BACKGROUND

In a simple form of human-IA interaction, such as supervisory control, the human makes the majority of the spatial decisions while the IA handles mobility policy and perhaps some local decisions related to tasks like obstacle avoidance. In cases of limited autonomy, the simple behaviors are likely to match human expectations due to the low level of uncertainty and restricted set of

¹Shared mental models are "knowledge structures held by members of a team that enable them to form accurate explanations and expectations for the task, and, in turn, to coordinate their actions and adapt their behavior to the demands of the task and other team members" [21]. We use the term *compatible mental models* to indicate spatial mental models that lead humans to produce behaviors similar to those produced by a given autonomous system, as well as human spatial mental models that lead humans to accept behaviors by a given autonomous system.

future actions that an agent could potentially take on its own [25]. However, with increasing levels of agent intelligence, we expect a wider range of potential agent behaviors, thereby increasing uncertainty in the agent's future actions. This variance sets the stage for potential conflicts between the agent's behavior and the human's expectations, and as such increases the opportunities for degradations in trust to appear. These potential conflicts arise from differences in the mechanisms by which humans and IAs plan routes and the extent to which they are aware of and consider mission-specific criterion (i.e., information) to make spatial decisions, which can lead the human to prefer or expect routes that are different from those suggested or executed by the IA.

2.1 Route Planning Mechanisms

Human spatial navigation and route planning are often studied in the context of naturalistic tasks. These include ground vehicle routing to multiple destinations [26]; UAV mission planning for reconnaissance [27, 28], force protection [29], and search and rescue operations [10, 30, 31]; navigation through cities by pedestrians [32]; orienteering [33]; and even vacation planning [34]. Typically, these studies focus on the heuristics employed during navigation that are necessary to permit humans to approximate near-optimal solutions to spatial problems. As such, the primary goal of most of these domains is to understand human spatial decision-making, while a secondary goal is to inform the development of algorithms that approximate rapid high-quality solutions where brute force (i.e., exhaustive) calculation is too computationally intensive. Evidence for this near-optimal performance comes from systematic investigations of human performance in a simple but NP-hard task, the traveling salesman problem (TSP). In the visual form of this task, the subject attempts to draw the shortest tour that visits a number of nodes, representing cities, and returns to the starting location. Despite the computational complexity of TSP, human solutions are typically less than 10% longer than the optimal route, and solution times increase only linearly with increasing problem size [35–37]. The available literature suggests that humans use a combination of top-down spatial heuristics and bottom-up mechanisms dependent upon the structure of biological systems (the eye and the hippocampus, specifically) to solve these problems, making this process for route planning inherently different from the algorithms that IAs use to plan routes.

Many mechanisms are currently available for IA route planning. While an exhaustive discussion of all route or path planning algorithms is beyond the scope of this article, we focus instead on mechanisms that are specifically prevalent in robotics as a base for highlighting the difference between these mechanisms and those employed by humans. Optimization algorithms (many based on Dijkstra's algorithm [38]) were among the first means of finding least-cost routes in robotics (e.g., A^* [1]), and persist in domains in which a great deal is known about both the environment and the IA. One such example lies in the development of robotic planetary rovers. Specifically, optimization algorithms have been proposed to maximize the robot's time spent in sunlight [39, 40], promote the robot's exploration of novel locations [41], and minimize the risk of mechanical failures due to terrain [42]. Contemporary modifications to these optimization algorithms allow the IA to plan routes optimally to goal locations (e.g., Reference [43]) or decrease run times through the use of heuristics [44]. While these algorithms are capable of producing least-cost solutions, the inherent problem with these types of optimization algorithms is that solutions can be counter-intuitive to human users, who may prefer simpler routes to more complex ones that offer trivial advantages in the optimization criterion (e.g., travel time).

For complex applications requiring close coordination with human users, other mechanisms provide an attractive alternative to optimization in terms of the IA's performance in that context. Where robots are expected to conform to human users' intentions and expectations regarding tradeoffs among task-relevant criteria, more sophisticated algorithms based on homotopy class or pareto-optimality can provide reasonable solutions. For example, a series of algorithms have

been created to solve complex problems with human input, including an approximation algorithm designed to maximize information gathering [45], homotopy-aware algorithms [46], and sampling-based algorithms designed to find pareto-optimal solutions based on human input [47]. Other laboratories have explored similar concepts for allowing autonomous systems to function appropriately in complex environments, including indoors [48], under emergency conditions [49], and under dynamic conditions [50, 51]. Reviews of this body of literature are available [52, 53].

Another approach altogether is to allow the human to train the IA to behave independently from demonstration using reinforcement and inverse reinforcement learning [54, 55]. One suggested benefit of these types of approaches is that they can teach robots to behave in ways that are generally acceptable, and perhaps predictable, to most humans. For example, inverse reinforcement learning permits robots to plan paths in human-populated areas that are socially adaptive with respect to human movements and concepts of personal space [56, 57]. Another advantage of learning through demonstration is that teammates do not need to understand the robot's underlying reasoning process to retrain it in the field, due to changes in mission or context, much the same way that humans adapt to these changes [14]. Importantly, since the robot learns its control policy from a human demonstrator (i.e., a *human way* of completing the task), there is the potential for individual differences in strategy, informational weighting, or understanding of context or priorities to produce different demonstrations (and thus alter the robot's control policy).

Contrary to the goal of producing acceptable, trustworthy, and predictable behaviors, these differences between human demonstrators could manifest in robot behaviors that do not generalize well across all human users (especially with respect to culture, training, and demographics). One major contribution of this work is to explore human variance in a simple and generalizable task, planning a route for errands, in attempt to answer the question, "is there a *human way* to complete this task?" If humans generally complete a given task the same way, then a robot trained by one should exhibit behaviors compatible with the others' expectations. However, if a group of humans do not exhibit consensus in their solutions, then a robot trained by one to complete the task may behave in ways that are incompatible with other humans' expectations.

While machine learning may produce more *human-like* solutions generally, its underlying processes are perhaps the most difficult to explain to human teammates (its exact mechanisms are, in fact, usually opaque even to the designers). In summary, the mechanisms that IAs use to plan routes are likely to produce different solutions than human teammates, and the IA's underlying reasoning processes may be difficult to explain, which poses a potential source of conflict for H-IA mixed-initiative teams. For mixed-initiative teams, the interactive nature of the planning process is dependent on both the capabilities of the team members and the task or mission needs. In general, human intervention in a control loop for route planning can improve the quality of the IA-generated plan [3]. This dynamic assumes that humans and IAs have different but valuable and often complementary capabilities. For example, while humans are capable of generating reasonable solutions quickly, there are cognitive and physical workload burdens that can degrade performance in complex dynamic control situations, ameliorable to some extent by constraining the process in an automated fashion [4, 5]. Thus, the effectiveness of a human-IA team can be viewed as the synergistic product of both contributors.

2.2 Route Planning Criteria

In addition to differences in the mechanisms that IAs and humans use to plan routes, these agents also differ in regards to the types of information and context they consider during route planning and the weighting by which they consider that information for different missions. For example, applied to planetary rovers, the algorithms treat route planning as an optimization problem and, as such, minimize cost functions related to system performance [58] or the rover's ability to maintain

its batteries using solar power [39, 40]. In these cases, optimal solutions to the problems exist. By comparison, real-world human decision-making is often about rapidly generating solutions that are good enough to accomplish a task (i.e., *satisficing* [59]). As such, human route planning is flexible and robust to changing task requirements [10] but may also be influenced by other criteria of which the human is not aware. Therefore, it can be posited that a human's spatial mental model for a specific problem space is a complex mixture of the strategies they use to solve the problem, the mission criterion, and cognitive or visual biases that may impact their preferred solutions. In an early publication on the topic of preferences in route planning, Golledge [60] identified several criteria that humans use for planning paths. Some closely match the optimization criteria of artificial intelligence (AI) algorithms (e.g., shortest distance and least time). Others, however, are difficult to quantify (e.g., a preference for the "most scenic" route), while some are even contradictory (some participants listed their preferred path as having the longest leg first; others, the shortest first).

While any number of criteria could be considered during the spatial decision-making process, some criteria are specific to certain domains. Take, for example, the task of planning a route for a vacation. Popular GPS navigation systems integrate distance and traffic data to find routes minimizing travel duration. By comparison, human route planning incorporates criteria and biases to which they have a varying ability to report or quantify. Some of these criteria are relatively abstract but could be quantified and implemented in an IA's route planning algorithm, such as the requirement to visit a particular location at a specific time to watch a sunset [34] or a preference to avoid a high crime area. Other qualities of human solutions to spatial problems can arise from biases that the human may not even realize they have. For example, given two routes of equal length, humans generally prefer the southern-most route to the northern route [61]. In addition, humans also choose routes consisting of fewer turns and longer straightaways [62]. An example of bias during free-form route planning comes from the TSP literature; the TSP only requires that the solver produce the shortest tour that visits all of the nodes. However, prior research shows that humans prefer routes that reduce the distance to first contact with a goal [37]. This preference is adaptive for solving real-world tasks, such as search and orienteering [63], which prioritize distance-discounted reward. Reaching goal locations early in a route is important in these problems, since it may not be possible to execute a proposed route exhaustively.

2.3 Route Planning Criteria for Self Versus an IA

In certain contexts, humans may plan routes that will be executed by themselves, other humans, IAs, or a heterogeneous group of humans and agents. Prior work directly comparing routes planned for oneself versus other humans versus IAs is notably sparse, but there are reasons to believe that route planning behaviors will differ for each case. For example, in a military context, humans plan more risk-averse routes for other humans relative to robots under their control [16]. In the context of driving in the civilian world, however, the prospect of risk is notably less salient, but other factors may be relevant. Related areas of research, human-to-human route planning and direction-giving, have been topics of study for decades (for review, see Reference [64]). For example, humans provide less-complex directions to other humans compared to routes that they plan for themselves [65]. The lion's share of IA route planning research has focused on autonomy and supervisory control for IAs without a direct comparison to humans (especially for UAVs; for review of trends, see Reference [66]). Therefore, it is difficult to predict how exactly these contexts will alter the criteria that humans consider when planning routes for themselves, other humans, or IAs.

2.4 Current Work

At present, IAs and humans use different mechanisms to plan routes, and the underlying reasoning processes are not easy to communicate bidirectionally. Therefore, the overarching goal of the

current work is the development of a quantitative process for identifying and characterizing distinct spatial mental models of humans planning routes for themselves and a driverless vehicle, in collaboration with IAs. This is important for mixed initiative H-IA teams, because this process allows users to systematically identify and operationalize the criteria employed during route planning to classify and communicate intent. The present study investigates spatial mental models using four different analytical approaches. The first two approaches investigate consensus and shifts in spatial mental models. Next, we investigate the relationship between a separate variable, trust in the IA, and those shifts in spatial mental models. Finally, we describe a process of objectively characterizing spatial mental models (i.e., in a data-driven fashion).² The core experimental task requires participants to draw routes under different conditions both before (Subtrial 1) and after (Subtrial 2) viewing a route suggested by the IA. In this way, we can capture participants' initial spatial mental model of the ideal solution to the route planning problem prior to a recommended route, then test the interaction between participant and IA in a mixed-initiative system in which the IA proposes a route by how the participant replans the route. Finally, this research attempts to objectively characterize the criteria that produce different route planning solutions to identify where they differ from one another. Each drawn route is expected to reflect different prioritization of mission-relevant criteria. This quantitative method of analysis will be used to explore the underlying reasoning process for route planning that could be leveraged in future research to develop more-effective bidirectional communication for mixed-initiative teams.

3 METHOD

3.1 Hypotheses

3.1.1 Spatial Mental Model Consensus. To assess spatial mental model consensus, participants will draw a path on the roadways from a set start location, visit every stop location, and return to the start location, similar to the TSP. In this task, however, participants' routes are constrained to a road network. A given sample of participants might exhibit consensus among their routes for each spatial problem (that is, they all draw the same route). Alternatively, participants' routes may converge to several different points of consensus, each corresponding to a reasonable route given differences in the weighting and prioritization of criteria. We take convergence in route planning behaviors to indicate similar spatial mental models among the participants who produced them. Prior work in robotics [45–47] and with the TSP [35–37, 63] indicates the need to account for human variability in spatial planning, between users and under different contexts. These aforementioned studies of human variability did not explicitly examine convergence, however. Based on prior work using similar methodology (References [7, 10], which found such effects when humans planned routes for themselves or a UAV under supervisory control, respectively), we expect that participants will exhibit this convergence when planning a route for themselves, as well as when planning a route for the driverless vehicle.

H1a: Humans will not exhibit a single group consensus in routes that they plan for themselves. Rather, their routes will converge to multiple groups of similar solutions when planning a route for themselves. The hypothesis is best reflected as $1 < k < n$, where k is the resultant number of groups and n is the number of participants.

²This will focus on criteria that are sufficiently general to be applicable across different domains and contexts, including (1) tour length, (2) complexity of the route, (3) distance-discounted reward associated with reaching each of the goal locations, (4) distance to first contact with a goal location, (5) distance required to return home after visiting all of the goals, and (6) overlap in the route. Note, these criteria are known to be important in spatial decision-making, but differences between individuals' spatial mental models may cause them to minimize a particular criterion value for different reasons and in different ways [60].

H1b: Similarly, humans will not all plan the same route for a driverless vehicle, nor will all of their routes be different. Instead, routes planned for a driverless vehicle will converge to multiple groups of similar solutions (i.e., $1 < k < n$).

3.1.2 Shifts in Spatial Mental Models. When participants view a route suggested by an IA, they have the opportunity to evaluate the spatial problem from a perspective that uses different route planning mechanisms and has different informational access and weighting. To understand the impact of this on a participant's spatial mental model, participants will have the opportunity to revise or retain their previously developed route plan after observing the route recommended by one of two IAs (real-world GPS navigation system algorithms³). The experiment employed two different IAs, each exhibiting different behavior on each route planning problem, to explore the effect of different suggestions on participants' behaviors. Using two different measures of route agreement (i.e., a continuous measure of similarity and categorical measure of cluster membership) provides insight into this effect at two different operational definitions and grain sizes. Here, similarity is quantified using an area-based measure (see Section 3.5.1), while cluster membership is obtained from the results of the analyses used to test H1a and H1b (see Section 3.1.1). These hypotheses reflect the effect of viewing the IA-suggested route on participants' spatial mental models and resultant routes produced after viewing the IA recommended route relative to the initial planned route.

H2a – IA Influence on Route by Similarity: Participants' routes drawn after viewing the IA's suggestion (Subtrial 2) will be more similar to the IA's route than those drawn during Subtrial 1.

H2b – IA Influence on Route by Cluster Membership: More participants' routes will cluster with the IA's route suggestion *after* they have viewed that route.

3.1.3 The Role of Trust in Shifting Mental Models. While there are many factors that could influence why humans would change their route plans, this research looks at the relationship between trust and shifts in spatial mental model compatibility with the IA. This relationship is important to explore because of the potential role of trust in participants' willingness to accept route suggestions from IAs.

H3a: Because the IAs in this experiment generated different routes, participants will have different trust ratings for each IA⁴ based on their perception of the IA's abilities.

H3b: Higher trust in the IA will positively correlate with the proportion of trials on which participants produced solutions congruent with those routes after viewing them.

3.2 Experimental Design

In this study, participants planned routes using five different maps (see Appendix A.1), starting and ending in a single location and visiting four goal locations. On each trial, participants drew a route on the map, before (Subtrial 1) and after (Subtrial 2) viewing a route suggested by one of two IAs (for this article, denoted as IA_A and IA_B), for both themselves and for a driverless vehicle.

The experiment used a within-subjects 2 Traveler (Self vs. Driverless Vehicle) * 2 IA (IA_A vs. IA_B) * 5 Map (1 through 5) factorial design, blocked across the first two factors (see Table 1). Each of these trials contained two Subtrials, as participants planned routes before (Subtrial 1) and after (Subtrial 2) receiving the IA recommendations. Across all experimental conditions, participants generated 40 routes in 20 trials. Levels of each factor (i.e., Traveler, IA, and Map) were presented randomly to each participant to mitigate potential order effects.

³Note, we do not specifically name the two algorithms here, as we have not obtained legal permission to do so and do not want their inclusion to be viewed as endorsement. The algorithms used in this study were selected due to their popularity and the fact that they generally produced different solutions to each of the route planning problems used in the experiment.

⁴Both GPS navigation systems use proprietary software, potentially different algorithms with access to different types of information; therefore, it is not possible to propose that trust will be higher in one IA versus the other.

Table 1. Experiment design

Independent Variable	Experimental Block			
Traveler: Self vs. Driverless Vehicle ¹	A		A'	
Intelligent Agent: A vs. B ²	N	N'	N	N'

¹A and A' refer to the two levels of the factor Traveler (Self vs. Driverless Vehicle). Each block contains 10 trials.

²N and N' refer to the two levels of the factor IA (A vs. B). Each block contains five trials.

3.3 Participants

Thirty participants were recruited for the present study from the CCDC Army Research Laboratory participant pool. Three participants failed to complete the study. Data from the remaining 27 (16 men, 11 women), aged 19–63 years of age ($M = 35.67$ years, $SD = 10.59$ years), were retained for analysis. All data were collected per research protocol ARL 16-109, approved by the CCDC U.S. Army Research Laboratory's IRB, authorizing data collection from human research participants.

3.4 Apparatus

All experiment materials were coded in the Psychology Experiment Building Language v. 0.14 [67] and presented to participants on a laptop computer equipped with a keyboard and mouse.

3.4.1 Route Planning Task. Participants viewed instructions at the beginning of each Traveler experimental block. At the beginning of the block, instructions told participants that the task was to draw the *best* route that would allow the Traveler to pick up several packages, consistent with shopping errands (i.e., visit each goal location), that begins and ends in the starting location, and visits all of the goal locations. For the driverless vehicle, participants were instructed that the vehicle could receive the deliveries autonomously, and that the time required to do so would be similar to them doing it themselves. Participants received no further instruction regarding the criteria that would constitute *best*, with the goal of encouraging them to construct routes based on their own subjective evaluations of route quality. In addition, instructions told participants that after drawing their route, they would be shown a route suggested by one of two IAs, GPS navigation system algorithms, and that they could retain their initially drawn route or draw a new one. Participants were provided no information about the IAs (so participants could base their decisions on the behaviors of these IAs alone, and to avoid the specific brands biasing participants' responses) or *why* the IA suggested this particular route, but they were simply informed that the IAs were "popular GPS navigation systems." More detailed information regarding this procedure follows below.

On each trial, participants viewed a map (Appendix A.1) containing a starting location (green star) and four target locations (red circles). These maps were generated using Google Maps images of real-world locations, modified using a contrast feature to improve visibility. The maps depicted suburban neighborhoods in the eastern United States. Map locations were selected to represent a suburban road network environment with multiple possible routes among the start and goal locations. The locations shown in the maps were selected for their relative complexity (i.e., they had more than one clear route among all of the locations) and interpretability (i.e., all roads could be seen easily). Participants indicated that they were not familiar with the locations shown in the maps. The experimental software otherwise presented these maps as generated to preserve the naturalistic quality of the stimuli, and it is therefore possible that street names or landmarks indicated on the maps influenced participants' route planning strategies.

Participants drew their routes using a computer mouse. Each click on the map drew a waypoint connected by line segment to the previous point. Participants could delete the last drawn waypoint

Fig. 1. Workflow in the main route planning task. In Subtrial 1, the participant plans a route independently, beginning at the starting location (green star), visiting all four target locations (red circles), then returning to the starting location. In Subtrial 2, the participant views the IA's suggested route (blue, center, and right panels; numbers indicate serial order of visitation). The right panel shows the state of the route planning task during Subtrial 2, as the participant is given an opportunity to replan their route. Note, in this example image the participant has already begun constructing a route (red segments); the participant has chosen to visit the first two targets in the same order as the suggested route, but has omitted the portion of the suggested route in the top-right corner of the map.

and segment by clicking the *BACK UP* button. Once satisfied with the route, the participant clicked the *DONE* button to end the trial.

After initially drawing the route unaided (i.e., Subtrial 1), participants viewed the same map overlaid with route suggestion from the appropriate IA. The IA's direction of travel was indicated by numbers positioned next to each of the goal locations, which indicated the serial order of visitation. In this Subtrial 2, participants were informed that they could keep the route they had previously drawn (in which case, they were instructed to simply click *DONE*), or they could revise (i.e., re-draw) their route. This process is depicted in Figure 1.

Importantly, when participants elected to replan their route, they were shown the IA's suggestion, but not their previously planned route. This decision was made to prevent the two routes from occluding one another, and because presenting both routes on separate maps required downsizing the images such that they were no longer easily legible. This presentation potentially biased participants electing to replan their routes, and we acknowledge this as a limitation of the study.

3.4.2 Questionnaires. In addition to the primary experimental route planning task, participants also completed a number of questionnaires during the course of the experiment. Before the route planning task, participants completed demographic questionnaires to establish participants' gender, age, and video-game-playing habits (results not presented here). Participants completed a trust scale after each IA block and each Traveler block (for four total observations; Appendix A.2). The items included intelligence, fluidity of information, intended use, route safety, and trust on a seven-point semantic differential trustworthiness (adapted from Reference [68] and previously used in Reference [69] to assess trust in driverless vehicles). The items were designed to measure aspects of trust related to route planning.

3.5 Spatial Mental Model Analysis

A goal of the present study was to characterize the variability among participants' spatial decision-making behaviors before and after viewing a route suggested by an IA, given a simple multi-goal route planning task. For each of the 20 spatial problems, the similarity among participants' routes was calculated using the procedure first described by Perelman & Mueller [10]; see Sections 3.5.1 and 3.5.2, below). The data from the experimental trials resulted in a total of 1,080 routes: Each of the 27 participants generated routes for themselves and a driverless vehicle (i.e., *Traveler*), with assistance from both IAs, before and after viewing the IA-suggested route, on five maps per condition. The data processing and visualization techniques described herein are intended to quantify and visualize the similarity among participants' routes spatially, and infer consensus among them, which we operationalize as evidence of similar spatial mental models.

3.5.1 Quantifying Route Similarity. To calculate and visualize the similarity among participants' routes drawn for each spatial problem, the data were aggregated by map, IA, and Traveler, yielding 20 separate data subsets, each containing participants' Subtrial 1 and 2 (i.e., before and after assistance) routes as well as the IA's suggested route. The Algorithm for Finding the Least Cost Areal Mapping between Paths (ALCAMP [70]) was used to calculate pairwise divergence values between all of the routes in each data subset (i.e., for each spatial problem). ALCAMP generates a polygon from two routes, the area of which is a measure of the dissimilarity between them. We refer to the area of the polygon between the two routes under comparison as the divergence value (measured in pixels). The theoretical bounds of this metric, the minimum and maximum areas of such a polygon in the present study, are therefore 0 and the area of the map in pixels, respectively. Two routes producing a divergence value of 0 are identical, whereas the magnitude of a divergence value above 0 indicates the extent to which those routes are dissimilar. The divergence metric does not provide a meaningful comparison of routes drawn between maps due to differences in map size. While it is possible to scale the divergence values by the upper bound (maximum area of the polygon) for each map to normalize for map size, this would not provide a meaningful comparison between maps, since the road network and targets were distributed differently on each. This process is shown in Appendix A.3.

3.5.2 Testing for Consensus in Spatial Mental Models. To check consensus among participants' spatial mental models, multidimensional scaling (MDS) was used to visualize the similarity among participants' routes. MDS requires a distance matrix as input. We used the divergence values described in Section 3.5.1, which reflect the pairwise dissimilarity between routes, to populate these distance matrices. In these distance matrices, each cell contained the divergence value associated with that pair of routes, and the distance matrix as a whole contained all possible pairwise comparisons for the given spatial problem. With the inclusion of the IA-suggested route, each distance matrix had the dimension 55*55 (i.e., two routes from each participant, one generated before and the other after viewing the IA's suggested route, and the route generated by the IA itself). This procedure was repeated for every spatial problem in all experimental conditions (see Table 1), resulting in 20 distance matrices.

Using each of the distance matrices described above as input, Kruskal's non-metric MDS algorithm (using the isoMDS function of the MASS package [71]) provided a means of visually representing the similarity relationships among routes. Non-metric MDS was selected over classical MDS because, while the distance matrices exhibit symmetry and minimality, they do not necessarily satisfy the other property of a Euclidean space, triangle inequality (i.e., three cases can be found where $d_{ab} + d_{bc} < d_{ac}$). Applied to each of the 20 subsets, the gradients of Scree plots testing two through $n-1$ dimensions indicated that two dimensions were sufficient in every

case; additional dimensions did not meaningfully reduce the MDS solution stress values. This is consistent with prior work using a similar approach [10]. In summary, the procedure resulted in 2D visual mappings of the similarity relationships among participants' routes, both before and after viewing the IA-suggested routes, along with that IA-suggested route itself. MDS solutions use Euclidean distance to represent the *similarity* among data points, thus the orientation of data points should not be taken to convey any information. Cluster membership was provided using the procedure described below.

To test for the presence and number of spatial mental models among participants, operationalized as consensus among participants' routes, Flexible Mixture Modeling (FMM [72]) clustered each of the 2D MDS solutions. FMM is a model-based clustering technique that fits a range (user-specified, in this case 2 to 7) of bivariate Gaussian models to the data points to generate a clustered dataset (i.e., aggregate model). The stepwise FMM process employed a custom mixture of Gaussians driver (designed by Perelman & Mueller [10]; see Appendix A.4 for the driver code and R package [flexmix] for more information on the use of clustering drivers) that allowed each model flexibility in terms of volume and orientation. This flexibility is critical, as FMM must fit every data point to a model, and the aggregate models must be sufficiently flexible as to allow one or more models to capture routes that did not cluster well with others. Gaussian models within each MDS solution were iteratively fit using expectation-maximization, and the aggregate model with the lowest Bayes Information Criterion value indicated the best fit within the specified range. This procedure allows for aggregate model selection in a way that rewards goodness of fit but penalizes model complexity, and it allows inference of similar spatial mental models without *a priori* hypotheses regarding the number of distinct clusters in the sample.

4 RESULTS AND DISCUSSION

This section is organized by distinct analyses aimed at testing the experimental hypotheses. The first analysis aims to establish consensus among participants' routes, which we take to indicate similar spatial mental models, using clustering methods described above. The second analysis tests the influence of each of the IA's suggested routes on participants' route planning behaviors using continuous and categorical metrics of route agreement. The third analysis aims to determine whether trust plays a role in participants' willingness to alter their route so it is more congruent with the IA's suggestion. The final analysis quantifies characteristics of the routes in each cluster to determine where they differ; the findings of this analysis hold direct implications for the bidirectional communication of criteria and informational weighting that H-IA teams may use during collaborative route planning. Where appropriate, hypotheses were tested using the appropriate null hypothesis significance test (primarily factorial ANOVA). Effect sizes are reported in terms of *Cohen's d* (two groups), η_p^2 (3+ groups), or ϕ (categorical variable tests).

4.1 Analysis 1: Consensus in Spatial Mental Models

The first hypothesis for this experiment states that humans will exhibit consensus among several similar spatial mental models, though we do not have *a priori* hypotheses regarding the number of distinct spatial mental models. For that reason, a data-driven approach was used to characterize the variance among participants' routes drawn in each trial. Stepwise FMM, performed on each MDS solution, converged at 3 to 6 Gaussian models (i.e., clusters; see Table 2) for all the experimental maps⁵ and for both traveler conditions. Using Figure 2 as an example FMM converged to four

⁵Note that this clustering approach must account for *every* point in the MDS solution; therefore, not all clusters indicate equal convergence to a shared spatial mental model.

Table 2. Number of Clusters Identified by FMM in Each Experimental Condition

Driverless Vehicle	Self					
		Map 1	Map 2	Map 3	Map 4	Map 5
	IA _B	3	4	3	4	4
	IA _A	4	5	5	2	4
Vehicle	IA _B	4	4	6	6	6
	IA _A	4	4	5	5	4

Fig. 2. Sample clustered MDS solutions containing all experimental routes, drawn on Map 5 for a driverless vehicle, prior to viewing the route suggested by an IA. Each point indicates a route, color- and shape-coded according to cluster membership. Crosshair indicates the IA's route. Note, the coloring scheme provides a means of visualizing cluster membership, and the color and shape of each cluster are arbitrarily assigned with the exception of the cluster containing the IA's route, which is always shown in black.

clusters, and only four solutions were initially clustered with the IA's route (represented by the black crosshairs). All solutions are provided in Appendix A.5.

The results described here support H1a and H1b, as participants' routes generated for a given trial converged to a small number of clusters when planning routes for themselves or a driverless vehicle, respectively. The practical implication of this result is that there is not a single *human way* of solving this problem, but human route planning behavior is not random and unpredictable. Instead, there appears to be several ways that a group of humans will solve a given route planning problem, each of which might be reasonable according to a different spatial mental model.

This finding is important to future human-autonomy teams, because disagreement among team members can increase decision-making time due to deliberation, especially when the consequences of any decision are great (e.g., Reference [49]). In addition, lack of consensus for the ideal course of action can create situations in which teammates are unable to predict one another's future actions in the absence of communication. In mixed-initiative systems, disagreement between the human and the IA manifests as replanning [2]. Replanning frequency generally produces inconsistent

effects on mission performance metrics. Some studies [4, 5] have found that increased replanning frequency degrades performance, while others [3] have concluded that frequent replanning improves performance. These effects are typically attributed to changes in workload, and they interact with the replanning *duration* (i.e., the amount of time the human spends replanning the route during each trial) such that replanning improves performance but only if done rapidly. Importantly, metrics of replanning can provide insight into the H-IA dynamic but do not describe the underlying reasons that replanning occurs outside of general disagreement.

In spatial planning, differences leading to disagreement arise from a multitude of factors, some of which are easy to describe and others that arise from subconscious biases that may be more difficult to communicate. For example, a teammate would have no problem understanding and describing that one route is longer than another or reaches more goal locations early in the route. In fact, prior research has shown that humans can tailor their solutions based on specific instructions to produce routes that solve for these criteria [10]. However, teammates may have trouble communicating the reasons they prefer a route that is more convex and, consequently, overlaps less. The next section describes route planning dynamics in the context of a collaborative system to understand how viewing an IA's suggested route can cause shifts in a human user's spatial mental model.

4.2 Analysis 2: Shifts in Spatial Mental Models

The intent of this section is to answer the research question, how does access to an IA's route suggestion change a person's spatial mental model, resulting in different route planning behavior? To reach this end, after developing an initial route plan, participants viewed a route suggested by the IA and were given the opportunity to revise their route. Shifts in participants' spatial mental models were examined at two levels of analysis: (1) continuously, as measured by the divergence value changes between the IA's suggested route and their own routes generated during each subtrial; and (2) categorically, as measured using the cluster membership of each route. The first level of analysis examines the relative influence of the IA's suggestion at a high level. The second level of analysis more specifically examines whether viewing the IA's suggestion was sufficient to change participants' spatial mental models and whether those changes made the resulting solutions congruent with the IA's route.

4.2.1 Route Influence Measured by Similarity. To quantify changes in participants' route planning behaviors in response to viewing the IAs' suggestions, a broad approach was used to test the difference in similarity (i.e., ALCAMP-generated divergence values; see Section 3.5.1) between participants' paths produced before and after viewing those suggestions. These divergence values provide a continuous measure of the similarity between routes produced by the IAs and humans under different experimental conditions. A 2 (IA: A vs. B) * 2 (Traveler: Self vs. Driverless Vehicle) * 2 (Subtrial: 1 vs. 2) * 5 (Map: 1 through 5) factorial ANOVA was used to test differences in divergence values between routes produced by participants and the IAs. The ANOVA revealed an expected main effect of Map, $F(1, 1071) = 7.96, p = 0.004, \eta_p^2 = .007$, given that the maps were different sizes. More importantly, the ANOVA also revealed main effects of IA, $F(1, 1071) = 5.25, p = 0.022, d = 0.14$, and Subtrial, $F(1, 1071) = 64.82, p < 0.001, d = 0.49$. Respectively, participants' solutions were overall more similar to IA_A than IA_B, and their routes were significantly more consistent with both IAs' routes after viewing them (i.e., in Subtrial 2). The analysis did not find a main effect of Traveler ($p > 0.999$) or any interaction effects for the independent variables. These findings, taken together, support hypothesis H2a and indicate that both IAs influenced participants' routes and that IA_A's routes were more similar to participants' routes than IA_B. The lack of interaction effect indicates that neither IA was more *influential* as quantified by this continuous measure of similarity.

Fig. 3. MDS solutions for the routes that participants drew on Map 5, for a driverless vehicle, before (left panel; Subtrial 1) and after (right panel; Subtrial 2) viewing the route suggested by IA_B . Each point indicates a route, color- and shape-coded according to cluster membership. Crosshair indicates the IA 's route. As in Figure 2, the coloring scheme merely provides a means of visualizing cluster membership, and the color and shape of each cluster are arbitrarily assigned with the exception of the cluster containing the IA 's route, which is shown in black. The panel on the right shows participants' routes drawn both before and after viewing IA_B 's suggestion (i.e., on both Subtrials). The arrows connect participants' first drawn route to the route they drew after viewing IA_B 's suggestion.

The main effect of Subtrial provides evidence for H2a and demonstrates, on a coarse scale, that viewing a route suggested by an IA significantly impacts the human's spatial mental model, resulting in changes in route planning behaviors. Taken in this context, the main effect of IA shows that IA_A produced routes that were more similar to those produced by humans than IA_B . Finally, the null effect of Traveler should be taken in light of the experimental design. In the introduction to each experimental block, participants were provided only with the detail that either they or a driverless vehicle would travel the route, but they were not given any additional information about the driverless vehicle's capabilities. In the absence of knowledge regarding specific capabilities of an autonomous teammate, participants may have defaulted to a perceived ideal solution that would generalize across teammates. With additional information about the capabilities and limitations of an autonomous system, we might expect that participants' route planning behaviors would be impacted differently by viewing alternative route suggestions.

4.2.2 Spatial Mental Model Shifts by Cluster Membership. A second goal of the experiment was to quantify these shifts in participants' spatial mental models in a categorical fashion to understand whether participants would accept or reject the IA 's suggestions (H2b). The previous analysis indicated that the IA 's route suggestions impacted participants' spatial mental models, but it does not characterize the nature of that influence. Was that effect caused by all participants slightly modifying their routes, or was the effect produced by a subset of participants fully adopting the IA -suggested routes whereas another rejected them?

To answer these research questions, we examined the cluster membership (from Analysis 1; see Section 4.1) of participants' Subtrial 1 and 2 solutions to determine if and how their route planning behaviors shifted in response to viewing the IA 's suggestion (see arrows in Figure 3 connecting Subtrial 1 and 2 solutions). Figure 3 shows an example of these shifts. In some cases, participants' routes in both Subtrials clustered together. In other cases, viewing the IA 's suggested route caused

Table 3. Spatial Mental Model Shift Types by IA and Traveler

		C-C	I-C	C-I	I-I	I-IA
Self	IA _B	15.11% (21.18%)	20.86% (24.87%)	7.19% (3.92%)	46.76% (19.80%)	10.07% (14.12%)
	IA _A	21.54% (18.70%)	26.92% (25.89%)	1.54% (5.43%)	43.08% (28.67%)	6.92% (9.70%)
Driverless Vehicle	IA _B	12.59% (33.30%)	31.11% (46.47%)	0% (0%)	37.78% (48.66%)	18.52% (38.99%)
	IA _A	16.30% (37.07%)	33.33% (47.32%)	0% (0%)	34.07% (45.57%)	16.30% (37.07%)

Note: Results reported as M (SD). Columns C-C and I-C both indicate cases where participants' final (i.e., Subtrial 2) route was congruent with the IA's suggestion.

the participant to modify their route in such a way that it changed clusters. For example, on the lower left region of the right panel, one participant's initial (Subtrial 1) route clustered with the blue crosses cluster (i.e., it differed from the route suggested by the IA). However, after viewing the IA's suggestion (Subtrial 2), this participant produced a route that was very similar to the IA's and therefore clustered with it. Finally, this visualization shows cases where participants changed their routes between Subtrials but did not fully adopt the IA's suggestion (for example, one participant in the red circles cluster modified his route so it clustered with the blue crosses cluster).

To quantitatively analyze these shifts in participants' routes using the clustered MDS solutions, we labeled each participant's Subtrial 1 and Subtrial 2 routes based on whether they were Congruent (C) or Incongruent (I) with the IA's suggested route (i.e., whether they fell into the same cluster). These descriptions resulted in one of the following labels for each route pair drawn before (Subtrial 1) and after (Subtrial 2) viewing the IA's suggested route (Table 3):

- C-C: Both Subtrial 1 and 2 routes are congruent with the IA's suggestion.
- I-C: The Subtrial 1 route is incongruent with the IA's suggestion, but the Subtrial 2 route clustered with the IA-suggested route.
- I-I: The Subtrial 1 route was incongruent with the IA's suggestion, and the participant retained that route in Subtrial 2.
- I-IA: The Subtrial 1 route was incongruent with the IA's suggestion, and the participant revised the route in Subtrial 2, but the revision was not congruent with the IA's suggestion.
- C-I: The Subtrial 1 route clustered with the IA's suggestion, but the participant revised their solution such that their Subtrial 2 route was no longer congruent with the IA's suggestion.

Hypothesis H2b pertained to participants' categorical responses to viewing the IA's suggestion. The tests described herein are analogous to that presented above in Section 4.2.1, only in this case our dependent variable is categorical (did participants accept the IA's suggestion or not?) rather than continuous (how similar was the participant's route to the IA's suggestion?). Across all conditions, only 18.55% of participants' initial solutions were congruent with the IA, versus 44.34% of their revised solutions. A chi-square test of independence found that more participant solutions clustered with those provided by the IAs after viewing the suggestion versus before, $X^2 = 81.95$, $p < 0.001$, $\phi = 0.276$, providing support for H2b. To further explore this phenomenon, we investigated whether shifts differed by IA or Traveler. To accomplish this, pooled frequency counts were used to generate the below table by either IA or Traveler, respectively, for trials in which participants' Subtrial 2 route was congruent with the IA's suggestion.

This test found a significant effect of IA, $X^2 = 4.3$, $p = 0.038$, $\phi = 0.090$, indicating that more participants adopted IA_A's route than the one proposed by IA_B. A second test found no effect of Traveler, $X^2 = 1.00$, $p = 0.316$, $\phi = 0.043$, indicating that participants' decision to adopt the IA's route did not depend upon whether they were planning a route for themselves or the driverless vehicle.

This analysis provides additional context for the findings presented in Section 4.2.1. First, it is concordant with the null effect of Traveler, such that participants' route planning behaviors did not differ significantly depending on the agent for whom they planned the route. Second, while the analysis presented in 4.2.1 indicated that viewing either IA's suggested route did not have a greater influence on participants' route planning behavior (when measured using a continuous metric of similarity), this analysis shows that the influence of each IA differed. A larger subset of participants altered their routes between Subtrials to conform to IA_A's route. These results, juxtaposed with the results described in Section 4.2.1, show that coarse analysis of continuous measures of spatial problem solving may be insufficient to capture the detailed dynamics of collaborative route planning between humans and IAs, since they capture the magnitude of change in aggregate without respect to the behavior of different subgroups. In effect, the sample as a whole did not exhibit a similar behavior to varying degrees; rather, some subgroups within the sample exhibited the behavior (i.e., altered their routes) whereas others did not. Here, we show that a more detailed approach examining different resulting solution types can provide the fidelity necessary to interpret those shifts.

4.3 Analysis 3: The Role of Trust in Spatial Mental Model Shifts

Trust may have played a role in participants' collaborative route planning behaviors described above. To address this possibility, this section describes differences in participant trust of each IA and a relationship between participants' rated trust in each IA and their willingness to accept their proposed routes. First, we describe the differences in trust in each of the Traveler and IA experimental conditions. Second, we correlate these trust scores with individual participants' frequencies of adopting an IA-congruent route to determine if there is a relationship between trust and route acceptance and under which conditions this holds. Note, participants did not have any information about each of the IAs beyond the routes that they generated, and the routes were presented to participants in identical format.

A 2 Traveler (Self vs. Driverless Vehicle) * 2 IA (A vs. B) ANOVA found no significant main effect of Traveler, $F(1, 104) = 0.91$, $p = 0.341$, $d = 0.17$, but did find a main effect of IA, $F(1, 104) = 5.11$, $p = 0.026$, $d = 0.44$, in which participants rated IA_A higher on the trust scale ($M = 5.26$, $SD = 0.83$) than IA_B ($M = 4.87$, $SD = 0.93$). Follow-up paired-samples t -tests were used to explore effects between the two IAs on each dimension of the trust scale. IA_A was rated as more intelligent, $t(53) = 2.55$, $p = 0.014$, $d = 0.52$, fluid, $t(53) = 3.04$, $p = 0.003$, $d = 0.54$, usable, $t(53) = 5.22$, $p < 0.001$, $d = 0.72$, and trustworthy, $t(53) = 3.18$, $p = 0.002$, $d = 0.58$ (see Figure 4). Participants' self-reported preference for IA_A on the trust scale is concordant with the slightly greater frequency with which participants accepted that IA's proposed routes and supports hypothesis H3a. The null effect of Traveler indicates that embodiment, in this design, did not alter participants' trust in the IA. Given that participants had no additional information about each IA, this suggests that the between-IA effects are due to the morphology of the routes themselves, and that there may be a connection between the perceived quality of an IA's suggestion and human teammates' trust in that IA.

To test whether trust played a role in participants' decisions to adopt a route during Subtrial 2 that was congruent with the IA's suggestions, we correlated individual trust scores with

Fig. 4. Trust scale subscale scores for the two IAs. Error bars indicate SE. * indicates $p < 0.05$, ** indicates $p < 0.01$, and *** indicates $p < 0.001$.

the participant-wise frequency count data described in Section 4.2.2 generated during route planning either for themselves or the driverless vehicle. Pearson's correlation tests showed a significant but weak positive relationship between trust and participants' frequency of adopting a route congruent with that IA's suggestion when participants planned a route for themselves, $r = .335$, $n = 54$, $p = .013$, $r^2 = .112$, but not when they planned a route for the driverless vehicle, $r = -0.06$, $n = 54$, $p = .653$, $r^2 = .004$. These results, taken together, provide partial support for hypothesis H3b and indicate that trust played a role in participants' willingness to accept a route proposed by the IAs, but only when planning a route for themselves.

The direct conclusion is that the role of trust in mixed-initiative team performance may depend upon the Traveler that will ultimately execute the route. This finding has implications for these teams and for future research investigating H-IA teams. Specifically, trust may be important when the human teammate will execute the route, since they will also bear the burden of consequence associated with the route. However, when a non-human teammate will execute the route, the human teammate may be more willing to defer to the IA's spatial decisions. We hesitate to extrapolate grand implications from a single weak positive correlation, but based on the present data, we acknowledge the possibility that trust modulates route acceptance or rejection differently, depending on the agent who will actually execute the planned route.

In summary, the analyses described above show that, during the course of solving a relatively free-form multiple-goal route planning task, a group of humans and IAs tends to converge to several different types of solutions. Each of these routes might be reasonable according to a different set of criteria or spatial mental model. The above results also demonstrated that in these types of collaborative mixed-initiative route planning tasks, IAs will not produce solutions that are congruent with the spatial mental models of *all* teammates. This has implications for real operational H-IA teams. In real-world route planning tasks, human teammates may not understand why the IA is making particular route recommendations simply because that information is not made explicit or is of sufficiently high dimensionality that it is difficult to parse and communicate to the human in a meaningful way. In the case of more recent machine learning approaches to IA behavior, the IA's reasoning process may be particularly opaque and difficult to communicate. This mismatch

in spatial decision-making behaviors may be a potential source of trust degradation among team members. One way to counteract this degradation in trust is to develop a shared understanding by effectively communicating the ways in which teammates' solutions differ. In the following section, we propose a process for quantifying relevant characteristics of routes, produced by humans and IAs in the team, with the hope that characterizing the differences between team members will allow them to collaborate more effectively during spatial decision-making.

5 CHARACTERIZING SPATIAL MENTAL MODELS

A route could conceivably be described quantitatively using any number of different cost functions to determine the most-useful information to share with teammates. This section describes a way to infer the characteristics of routes, which can offer insight into the spatial mental models or underlying reasoning processes of the agent that produced them⁶.

Determining how these spatial mental models and processes differ allows us to identify the information that can be leveraged to develop appropriate bidirectional communication practices, using either verbal descriptions or visualizations, for enabling smoother, more effective performance in H-IA teams.

5.1 Rationale for Proposed Methodology

In this exploratory analysis, we focus on characterizing routes drawn by the participants and IAs to provide a basis for comparison among them. In this analysis method, rather than comparing among all routes generated for each spatial problem, we compare among the clusters of routes identified in Section 4.1 (Analysis 1: Consensus in Spatial Mental Models). We take these clusters to indicate points of consensus in spatial mental models among the participants in this study. The differences among these clusters tell us where they differ and give insight into what information should be communicated to resolve conflicts in route planning.

5.2 Criteria for Route Planning

In Section 2.1 (above), we proposed the question, "is there a *human way* to complete this task?" The results presented thus far indicate that there may in fact be not one, but several human ways to solve these problems. The present study focuses on six criteria for characterizing a route, which may be useful for understanding and differentiating those ways that humans solve these problems: tour length, route complexity, distance-discounted reward, distance to first contact, return home distance, and route overlap. Note, while all of these criteria are known to be important in human spatial decision-making, differences between individuals' spatial mental models may cause them to minimize or maximize a particular criterion value in contradictory ways. It is worth noting that the number of criteria that could be considered is practically limitless, and specific domains would benefit from quantifying routes based on mission-specific criteria. These criteria were specifically selected, as they are generalizable across environments and contexts for any route planning task among multiple goal locations. To quantify a route according to these criteria, you must know only the route and goal locations; environmental information is not required.

5.2.1 Tour Length. Tour length refers to the overall length of the route and is the basis for optimization problems such as the TSP and a known prioritization criterion in human route planning

⁶Note, while prior research has explored human models with other pre-defined features (e.g., References [16, 45–47]), the present model differs in that it describes agent behavior using criteria that can be inferred from the paths themselves without knowledge of the environment. While this model will fail to capture variance associated with mission-specific criteria, the criteria used as examples herein are generalizable across the types of tasks that H-IA teams are expected to undertake.

[62]. In the absence of information regarding the agent's speed or aspects of the terrain, the fastest route through a space is the shortest and is generally preferred by humans [73]. Humans are well adapted to solve for this criterion; in the TSP, humans produce near-optimal solutions in times that increase linearly with increasing problem size [35–37, 74] despite the computational complexity of the problem. Here, we operationalize tour length in terms of the drawn route's length in pixels.

5.2.2 Route Complexity. Simpler routes are easier to remember and require the traveler to reference a map or localize itself less frequently. For this reason, some humans aim to reduce route complexity during planning [34]. Reducing route complexity may be an adaptive strategy when planning a route for another traveler to follow [65]. Similarly, given two routes of equal length, humans tend to prefer the route that is less complex (i.e., has fewer turns [62]). The AI community has also addressed this criterion with route planning software specifically designed to reduce route complexity [75, 76]. We define route complexity in the following analyses as the number of segments in the drawn route.

5.2.3 Distance-discounted Reward. Distance-discounted reward refers to a route evaluation structure in which the reward is discounted by the length of the route required to reach it. Reward can also be discounted by other similar route factors, such as travel time. This type of structure forms the basis of multiple objective navigation and search problems, such as the Orienteering Problem. In this problem, orienteers are required to visit a set of nodes, each given a score, with the goal of maximizing the cumulative sum of the collected rewards without exceeding a given tour length [77]. The problem is formally analogous to a TSP with discounted rewards, as each reward's value degrades with time or distance along the route required to reach it [63]. We define discounted reward by simply dividing the reverse n th reward (i.e., a value of 4 for the first reward, 3 for the second, 2 for the third, and 1 for the fourth) obtained by the length of the path (in pixels) required to obtain it and summing these discounted rewards to obtain the discounted reward for that trial.

5.2.4 Distance to First Contact. Distance to first contact (DFC) refers to the distance required to reach the first goal location. DFC is related to greedy local decision-making search strategies (such as the nearest neighbor strategy, according to which the traveler simply navigates to the closest unvisited goal location at each iteration), as well as distance-discounted reward strategies (i.e., reducing DFC increases the reward obtained at that location). Golledge [60] showed that humans prioritize the first leg of a route, indicated by a preference to make the first leg either the shortest or longest in the route. Prior research on the TSP shows that humans are biased toward solution forms that minimize DFC [37]. Furthermore, prior work shows that humans revert to a nearest neighbor strategy for route planning when no information is provided about potential targets in the environment [15]. DFC is defined here as the length of the drawn route, in pixels, between the home (starting) location and the first target location.

5.2.5 Return Home Distance. The return home distance (RHD) criterion is the measurement from the last target location to the final "home" location. After visiting all nodes within the environment, both laboratory tasks such as TSP and naturalistic problems (e.g., search using an autonomous system) require the traveler to return home. Perelman and Mueller [37] noticed that participants tend to neglect the cost associated with this return home, producing tour legs between the final target and the starting location that are far longer than the other legs of the tour. Here, we take a long RHD to indicate a local decision-making strategy that fails to account for the cost associated with the return. Note that in naturalistic search tasks, this is potentially adaptive, since more globally optimal strategies designed to maximize distance-discounted reward will allow the traveler to find a hidden target among candidate locations more quickly, at the expense of

having a longer return home that is essentially irrelevant to the mission. This criterion is defined as the drawn route's distance in pixels between the final goal and home locations.

5.2.6 Route Overlap. The route overlap criterion measures the percentage of the route that overlaps another portion of the route. Evidence from TSP suggests that, during spatial planning, computational models suggest that humans use a heuristic that eliminates inefficient self-intersections [78, 79], resulting in solutions that are very convex with few indentations (e.g., Reference [36]). From a top-down perspective, in real-world route planning problems, humans may have mission-specific reasons to minimize or maximize overlap. For example, in an analysis of route planning behaviors for vacation planning, Tenbrink and Seifert [34] noted that participants expressed a preference for “round, clockwise” tours that would “avoid traveling the same route twice.” In other cases, maximizing overlap may be an adaptive strategy. When traveling to a single destination and returning, one main road with a high speed limit may be the ideal route for traveling both directions. Similarly, it may be easier to remember a single route to and from a destination in an unfamiliar area. A rough estimate of overlap is computed here by discretizing the drawn route into 10-pixel segments, then calculating the percentage of nodes that are closer than 10 pixels to another node.

5.3 Process for Characterizing Spatial Mental Models

The procedure described herein is intended to quantify differences between route clusters on specific criteria to better understand those routes and provide insight into what features lead to consensus in a spatial mental model of the agents who created them. The basis of this approach is to process the MDS solutions to isolate the well-converged clusters, which indicate strong points of consensus among participants, then to apply cost functions to the routes within these clusters, based on criterion known to be important in human spatial decision-making, to determine the characteristics that differentiate each route cluster from the others.

5.3.1 Data Processing for Cluster Analysis. The analysis presented in Section 4.1 identified several points of consensus for each MDS solution. However, the FMM analysis must account for every data point; since we are interested in comparing the relative weightings of route planning criteria, we want to specifically identify and isolate well-converged clusters that represent points of strong consensus among participants' routes. To accomplish this, we implemented a simple rules-based approach for identifying and isolating well-converged clusters of routes (see Figure 5).

For this approach (see Figure 5), the dimensions of each MDS solution are scaled to 0:100 to correspond to percentages of the MDS solution space. Then, for each cluster in the MDS solution, we identify and exclude points lying outside of a distance, d , from that cluster's medoid. In this case, we set $d = 10$ to correspond to 10% of the problem space. This step eliminated routes that did not cluster well with others in its assigned cluster. Second, we excluded clusters containing five or fewer routes (n), since these clusters represented a very small minority (<10%) of the routes produced under those conditions. A complete list of the MDS solutions resulting from this data processing step are shown in Appendix A.5 (Figure A.5.3). This data processing step necessarily resulted in some degree of data loss; data retention information is presented in Table 4. An analysis of percent data retention based on a range of values for both d and n is available in Appendix A.6.

5.3.2 Characterizing Spatial Mental Models. The next step in quantifying route plans was to investigate each individual subset of solutions that reached consensus to understand why individuals made different decisions in this task. Given the relatively open-ended format of the task, there was no single *correct* solution (i.e., route) but rather a set of solutions, each of which would be correct according to different prioritization of mission-relevant criteria. Multiple Criteria Theory

Fig. 5. Data processing flow for eliminating outliers and poorly converged clusters from each MDS solution. Each point indicates a route, with the color and shape corresponding to cluster membership. Crosshair indicates IA's route. Red circled points in the middle panel indicate routes removed as a result of being either outliers or poorly converged clusters. In this particular case, routes belonging to the cluster indicated by blue crosshairs did not converge and were removed from the solution. This step may be generally necessary, since FMM must fit *every* data point, and some points may not cluster well with others.

Table 4. Number of Routes and Percent Data Retention by Condition

		Map 1	Map 2	Map 3	Map 4	Map 5
Self	I _A _B	40 (74.07%)	42 (77.78%)	38 (70.37%)	37 (68.52%)	41 (75.93%)
	I _A _A	43 (79.63%)	36 (66.67%)	39 (72.22%)	22 (40.74%)	40 (74.07%)
Driverless Vehicle	I _A _B	45 (83.33%)	31 (57.41%)	38 (70.37%)	36 (66.67%)	35 (64.82%)
	I _A _A	44 (81.48%)	35 (64.82%)	40 (74.07%)	32 (59.26%)	42 (77.78%)

Note: Participants produced 54 solutions in each of the experimental conditions. The IAs' solutions have been omitted from the percentage calculation.

[80] was used as a framework for investigating these differences in prioritization. According to this theory, spatial plans reflect tradeoffs among multiple mission-relevant criteria. Cost functions were calculated for each route based upon six criteria that have been determined to be important in human spatial decision-making (see Section 5.2, above). The result of this analysis was a distribution of values for each cluster, for each criterion, on each map.⁷ These resultant distributions were compared using MANOVA to determine if each cluster's routes differed meaningfully on these criteria. For each of the 20 clustered MDS solutions, the MANOVA found significant multivariate effects indicating significant differences in the spatial mental models with respect to the aforementioned criteria (all $p < 0.001$). Discussing the post hoc tests for all 20 of these solutions would be unwieldy; therefore, we have selected two MDS solutions to discuss in greater detail below—participants' routes drawn for the driverless vehicle in collaboration with I_A on maps 1 and 4, each with different implications for human-robot interaction (HRI). Concordant with the notion of post hoc explanations in the field of explainable AI [81], these examples highlight cases

⁷While the present application of these visualizations uses data aggregated by cluster, the visualizations could be created just as easily using single observations. This approach would be useful for communicating differences between individual agents' route plans in real-world settings.

Fig. 6. Cluster analysis plots indicate criterion values for each of the route clusters, drawn for a driverless vehicle in collaboration with IA_A on Map 1. The left panel indicates the criterion values for each of the three route clusters. Each row of the plot indicates a different criterion, and X-axis describes the average Z score on that criterion for each cluster. Each cluster's symbol corresponds to its symbol on the solution presented in its MDS solution (see Appendix A.5). The panels on the right show the routes corresponding to each of the clusters, color-coded by cluster membership. The numbers indicate the order of visitation. Routes in the black and green clusters use the same roads in the network but travel in opposite directions on the tour.

in which the agent can communicate to the human where and how its solutions differ from human expectations.

5.3.2.1 Example 1: Selecting Relevant Criteria to Communicate to Human Teammates. One benefit of this type of analysis is that it is useful for indicating the criteria on which clusters of routes differ (see example of IA_A on Map 1; Figure 6). A deeper way to interpret this figure is that, in addition to indicating the relative criterion values of each cluster, it shows where clusters agree and disagree on each of the experimental criterion values. This information can be used to more effectively calibrate spatial mental models in H-IA teams by communicating the criteria upon which paths differ the most.

In this example, routes congruent with the IA's route are shown as black squares, whereas two other clusters are shown as green circles and red diamonds, respectively. This type of depiction clearly shows variations between patterns in response in terms of the paths themselves and their

characteristics. For example, it is clear to see in Figure 6 that the black and green clusters completed the same path just in opposite directions. This produced very similar tour length, overlap, and number of segments between these two clusters. However, routes in the IA-congruent cluster produced substantially lower discounted reward values, a longer first leg, and shorter final leg of the tour. Routes in the red cluster (Figure 6, bottom right panel) followed a longer, more circuitous route with minimal overlap, which was similar to routes in the green cluster in all of the other criterion values. Taken together, these results indicate which criterion differences would be most important to communicate between team members. In a teaming situation, if the goal is to succinctly communicate criterion differences between alternative courses of action, the information requirements of participants in each of these IA-incongruent clusters could differ substantially. For example, if the user's spatial mental model preferred a route that clustered with others indicated by the green circles, they would benefit from knowing that reversing their direction of travel, as suggested by the IA, would reduce the return home distance but also produce a substantially lower discounted reward. However, if the user had produced a route that clustered with those indicated by red diamonds, the IA should communicate that its route was substantially shorter, though its route overlapped more and produced a lower discounted reward value.

5.3.2.2 Example 2: Visualizations for Communicating Differences among Different Plans. Our second example highlights a case in which each cluster differed substantially on most or all of the experimental criteria. Contrary to the example above, where the differences between clusters could be distilled down to a sentence or so (i.e., the IA communicates to the human, "my route is shorter, but it overlaps more and produces less reward"), in this case the differences are sufficiently complex as to warrant a visualization. Participants' routes produced in collaboration with IA_A for the driverless vehicle on Map 4 converged to three markedly different clusters (see Figure 7). Routes clustering with IA_A's route (black square cluster) traveled in a wide circular path and produced criterion values closest to the mean, with the exception of having the lowest amount of overlap. Routes in the blue cluster tended to be the least well-converged and were defined by the extensive use of shortcuts, which allowed them to produce the shortest overall tours, the highest discounted reward values, and the shortest first (DFC) and last RHD legs of the route. In terms of information transfer among teammates, the IA may need to employ a visualization that frames the problem and potential outcomes, allowing the human to compare potential routes or courses of action [82]. Such visualizations have shown promise in HRI for the evaluation of plans for autonomous vehicles [83, 84], and the present methodology presents a means for generating the values necessary to create these visualizations.

The results presented herein demonstrate that the different clusters in each MDS solution reflected distinct groups of solutions to route planning problems, and that these groups differed in terms of the experimental criteria tested. Some participants preferred shortest-route solutions, whereas others attempted to visit locations as early in the route as possible, reflecting a distance-discounted reward strategy that is consistent with search problems [10, 37, 63]. This methodology provides a means for identifying the ways in which routes differ and rapidly visualizing them to sub-select the important differences, which in turn can inform the communication of these differences among routes generated by team members in H-IA teams. When team members understand how and why their proposed solutions conflict, we expect that they will be better able to fairly evaluate alternative choices without being over-burdened by numerical information.

6 JOINT DISCUSSION & CONCLUSIONS

In summary, we conducted the present study to develop a data-driven process for differentiating and characterizing participants' and IAs' solutions to route planning problems. We had several

Fig. 7. Cluster analysis plots indicate criterion values for each of the route clusters, drawn for a driverless vehicle in collaboration with IA_A , on Map 4. For a description dimension, see Figure 6 caption. For this particular map and Traveler, participants produced three qualitatively different solutions, each corresponding to marked difference in the criterion values studied.

goals herein: (1) to show that there is not one “human-like” way of solving these problems, but rather that humans tend to converge around a small number of points of consensus, which we interpret as similar spatial mental models; (2) that viewing an IA’s suggested route will impact some participants’ spatial mental models; and (3) that trust plays a role in the acceptance or rejection of an IA’s proposed route. Finally, we proposed an exploratory quantitative method aimed at characterizing the differences among these points of consensus identified in the first objective. Within the wide body of literature devoted to HRI, using the human model classification types described by Reference [85], this work could be classified as a conventional approach to machine learning using a mixture of Gaussians to model consensus in a group of humans, as well as a simple probabilistic model that describes the behavior of those subgroups using cost functions.

The results of this study have several implications for H-IA teaming and collaborative tasks with an IA, whether that IA be a computer agent or an embodied robotic system. First, there is more than a single human way of solving relatively open-ended route planning tasks, so approaches to AI intended for teamwork applications with humans should be sensitive to individual differences in informational weighting, preferences, and biases. A second related point is that collaborative route planning dynamics are particularly complex especially when the IA produces a

route that does not match the human's spatial mental model. Third, the relationship between trust and route acceptance highlights the importance of trust in the performance of teams operating in mixed-initiative systems. Finally, the method proposed herein for communicating the differences between team members' routes offers a way forward to down-select the types of information that must be communicated among them to improve trust among teammates and mixed-initiative team performance.

One goal for AI is to produce behaviors that are "human-like" [86] or predictable [87, 88]. Achieving this goal is more difficult when there are, in fact, multiple solutions that might be considered human-like. Likewise, a particular solution that might be predictable to one person may not generalize well across a larger sample of individuals. Contemporary approaches to AI employ machine learning techniques in which humans teach the IA "what right looks like," but these approaches suffer from the same problem—the perception of correct behavior differs among individuals for a variety of reasons. The results of this study suggest that AI developers should be cognizant to the potential for individual differences in mental models generally, even in relatively simple tasks; we expect that the manifest differences in behavior arising from these mental models will increase as a function of task complexity.

In the type of mixed-initiative system used in this study, the IA proposes a route to the human that the human can accept or revise. The present study showed that, when the IA produces a route that is similar to the human's preference (i.e., their Subtrial 1 route), they will usually accept the IA's suggestion. Therefore, having the IA produce user-compatible routes can hold benefits in these systems in terms of decreasing operator workload associated with re-planning. However, when the human's spatial mental model is incongruent with that suggestion, the results become more complicated. In some cases, these participants accepted the IA's suggestion. In other cases, they rejected it and retained their initial route. Finally, in the remainder of cases, participants produced a Subtrial 2 route that was in fact different from their initial route and the IA's route. In a way, this fulfills the promise of a mixed-initiative system in which both members of the H-IA team contribute their own knowledge, and their combined product represents the advantages of man and machine.

In a mixed-initiative system, the decision to accept or reject an IA's proposed route may depend upon a complex interplay of IA's behavior, individual differences, and task context. In this study, participants rated one IA as significantly more trustworthy on the trust scale, and they were significantly more likely to accept route suggestions from that same IA. However, participants exhibited trial-to-trial variability in this regard; every individual participant produced at least one Subtrial 2 route that was congruent with an IA's suggestion, and none of the participants accepted the IA-suggested routes on all 20 problems. In terms of task context (i.e., for each of the Travelers), there was no significant difference in either the continuous or categorical dependent measures of IA route acceptance. However, we found a significant positive relationship between the frequency with which participants submitted a Subtrial 2 route congruent with the IA's suggestion and their trust scores, but only when planning a route for themselves. We hesitate to draw strong conclusions from a single weak positive correlation, but we welcome the possibility that the relationship between individual differences in trust and behavior in a mixed-initiative system can vary as a function of the task context (in this case, the agent who would travel the route). This finding is substantiated by recent research on route planning in a military domain, in which Perelman et al. [16] found that participants' route planning behaviors varied as a function of the agent traveling the route as well as the composition of the H-IA team. We expect that IAs capable of delivering individually tailored routes with respect to the human's spatial mental model would be more trustworthy and predictable, thereby reducing human team members' workload associated with

replanning (this is in line with more general understanding of human-robot team trust research [89]).

Finally, we proposed a methodology for identifying differences among solutions to spatial planning problems with the ultimate goal of identifying the types of information (i.e., criterion) that would facilitate bidirectional communication between team members. Transparency in reasoning and decision-making can improve situation awareness among members in a team [24, 90, 91] but only when that information is relevant and accessible. Route planning IAs typically employ algorithms that are somewhat difficult to explain to users, which makes directly communicating the IA's reasoning process difficult. Furthermore, modern machine learning approaches to route planning (i.e., reinforcement learning or deep learning approaches) use mechanisms and information that may be opaque to the designers. The goal of this methodology is to quantify and visualize differences among routes based on characteristics of the routes themselves, with the hope that these visualizations might ease bidirectional communication and the identification of conflict points. The analyses presented in Section 5.3.2 provide examples of these conflict points, and the subsections therein describe two potential ways of resolving them. After all, how the IA arrives at a particular spatial decision may be less important to human teammates than how its decision differs from the teammates' expectations. By visualizing the critical differences among candidate-proposed routes, human teammates can discriminate among the proposed alternates more quickly and accurately, improving overall team performance. Future research will explore users' trust in route planning solutions provided by IAs that are specifically calibrated to individual users' attitudes toward mission-specific criteria.

A APPENDICES

The five appendices compose: a visual depiction of the five maps used in the experimental design (A.1); the trust scale items (A.2); ALCAMP algorithm output and distance matrix example (A.3); the code for the custom mixture-of-Gaussians driver in R (A.4); the clustered solutions for all conditions (A.5.1), the initial planned routes (A.5.2), and the MDS solutions after the data filtering process (A.5.3); and the analysis of data loss by filter parameter value (A.6).

A.1 Experimental Maps

Fig. A.1. The five experiment maps used in the study. Green stars indicate starting locations for each map. Each route was required to begin at and return to this location. Red circles indicate points that the participant's route was required to visit. Participants were instructed to restrict their drawn routes to only the roads, indicated by contrast in each of the images.

A.2 Trust Scale Items

1. The system is...	Not Intelligent	Very Much	Somewhat	Neither/Both	Somewhat	Very Much	Intelligent
2. The system uses information that is...	Static, accounting only for basic information	Very Much	Somewhat	Neither/Both	Somewhat	Very Much	Fluid, accounting for advanced information
3. I will...	Not use this system	Very Much	Somewhat	Neither/Both	Somewhat	Very Much	Use this system often
4. The system...	Does not provide safe routes	Very Much	Somewhat	Neither/Both	Somewhat	Very Much	Provides safe routes
5. I have...	No trust in this agent	Very Much	Somewhat	Neither/Both	Somewhat	Very Much	Complete Trust in this agent

A.3 Sample ALCAMP Output and Distance Matrix Example

Fig. A.3. Data processing flow. ALCAMP compares all paths drawn for a single problem in pairwise fashion. Each comparison produces a divergence value, measured in pixels. The resulting divergence values populate a distance matrix, the input for MDS.

A.4 Custom Mixture-of-Gaussians Driver

This appendix contains the Flexible Mixture Modeling driver code used in the present study. It is an alternative code source to the one that comes with the flexmix package. Here, the R function StepFlexMix in R package [flexmix] requires a driver for the Model argument. In practical terms, the driver specifies the parameters of the Gaussian models that the expectation-maximization algorithm fits to the data. This particular driver permits the Gaussian models to vary in terms of orientation and volume.

```
CustomDriver <- function (formula = . ~ ., diagonal = TRUE)
{
  z <- new("FLXMC", weighted = TRUE, formula = formula, dist = "mvnorm",
    name = "model-based Gaussian clustering")
  z@defineComponent <- expression({
    logLik <- function(x, y){ll <- mvtnorm::dmvnorm(y, mean = center,
      sigma = diag(cov,length(center)), log = TRUE)
    return(ll)}

  predict <- function(x, ...) matrix(center, nrow = nrow(x),
    ncol = length(center), byrow = TRUE)
    new("FLXcomponent", parameters = list(center = center,
      cov = cov), df = df, logLik = logLik, predict = predict)
  })
  z@fit <- function(x, y, w, ...)
  {
    center <- colSums(y * w)/sum(w)
    decentered <- (t(t(y*w) - center))

    sdmin <- sd(y)/50
    ww <- round(w)
    if(sum(ww)==0)
    {
      cov2 <- sdmin^2
    } else {
      covwt <- cov.wt(data.frame(as.numeric(decentered)),
        wt = rep(ww,ncol(decentered)))
      cov2 <- max(sdmin^2,covwt$cov)
    }
    para <- list(center=center,cov=cov2,df=ncol(y)+1)
    with(para, eval(z@defineComponent))
  }
  return(z)
}
```

A.5 Supplemental Analysis: Clustered MDS Solutions

Fig. A.5.1. Clustered MDS solutions containing all experimental routes. Top panel depicts the routes that participants drew for themselves, while the bottom panel shows routes drawn for a driverless vehicle. Each column indicates a different map (1:5), and each row indicates a different GPS navigation system algorithm (top: Algorithm B, bottom: Algorithm A). Each point indicates a route, color- and shape-coded according to cluster membership. Crosshair indicates the GPS navigation system algorithm's route. Note, the coloring scheme provides a means of visualizing cluster membership, and the color and shape of each cluster are arbitrarily assigned with the exception of the cluster containing the GPS navigation system's route, which is always shown in black.

Fig. A.5.2. Clustered MDS solutions containing experimental routes produced *before* viewing the GPS navigation system algorithm's suggestion. Top panel depicts the routes that participants drew for themselves, while the bottom panel shows routes drawn for a driverless vehicle. Each column indicates a different map (1:5), and each row indicates a different trial (note, participants would go on to revise their routes after viewing a suggested route from either Algorithm A or Algorithm B; in this case, participants had not yet viewed the route suggested by each algorithm). Each point indicates a route, color- and shape-coded according to cluster membership. Crosshair indicates the GPS navigation system algorithm's route. Note, the coloring scheme provides a means of visualizing cluster membership, and the color and shape of each cluster are arbitrarily assigned with the exception of the cluster containing the GPS navigation system's route, which is always shown in black.

Fig. A.5.3. Processed MDS solutions after data filtering. Top panel depicts the routes that participants drew for themselves, while the bottom panel shows routes drawn for a driverless vehicle. Each column indicates a different map (1:5), and each row indicates a different IA (top: IAB, bottom: IAA). Crosshair indicates IA's route.

A.6 Data Retention Curves for Filter Parameter Values

Fig. A.6. Percent data loss across a range of parameter values for maximum distance filter (left panel) and minimum cluster size (right panel). Distance parameter $d = 10$ falls at the descent of the curve. Data retention increases nearly linearly with increasing n value. The cluster parameter $n = 5$ represents less than 10% of all clusters in the dataset, permitting the analysis to focus strictly on well-converged clusters of routes.

REFERENCES

- [1] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cyber.* 4, 2 (1968), 100–107.
- [2] Andrew S. Clare. 2013. *Modeling Real-time Human-automation Collaborative Scheduling of Unmanned Vehicles* PhD. Dissertation. MIT, Cambridge, MA.
- [3] Andrew S. Clare, Pierre C. P. Maere, and Mary L. Cummings. 2012. Assessing operator strategies for real-time re-planning of multiple unmanned vehicles. *Intell. Decis. Technologies Technol.* 6, 3 (2012), 221–231.
- [4] Mary L. Cummings, Andrew S. Clare, and Christin Hart. 2010. The role of human-automation consensus in multiple unmanned vehicle scheduling. *Hum. Fact.* 52, 1 (2010), 17–27.
- [5] Emily K. Muthard and Christopher D. Wickens. 2003. Factors that mediate flight plan monitoring and errors in plan revision: Planning under automated and high workload conditions. In *Proceedings of the 12th International Symposium on Aviation Psychology*.
- [6] Kip Johnson, Liling Ren, James Kuchar, and Charles Oman. 2002. Interaction of automation and time pressure in a route replanning task. In *Proceedings of the International Conference on Human-Computer Interaction in Aeronautics*. 132–137.
- [7] Brandon S. Perelman, Kristin E. Schaefer, and Arthur W. Evans III. 2017. Mental model consensus and shifts during navigation system-assisted route planning. In *Proceedings of the Human Factors and Ergonomics Society* 61, 1 (2017), 1183–1187.
- [8] Philip N. Johnson-Laird and Ruth M. J. Byrne. 1991. *Deduction*. Erlbaum, Hillsdale, NJ.
- [9] Reinhold Rauh, Cornelius Hagen, Markus Knauff, Thomas Kuss, Christoph Schlieder and Gerhard Strube. 2005. Preferred and alternative mental models in spatial reasoning. *Spatial Cog. Comput.* 5 (2005), 239–269.
- [10] Brandon S. Perelman and Shane T. Mueller. 2015. Identifying mental models of search in a simulated flight task using a pathmapping approach. In *Proceedings of the 18th International Symposium on Aviation Psychology*.
- [11] Tad T. Brunyé and Holly A. Taylor. 2008. Extended experience benefits spatial mental model development with route but not survey descriptions. *Acta Psychol.* 127 (2008), 340–354.
- [12] Georg Jahn, Philip N. Johnson-Laird, and Markus Knauff. 2005. Reasoning about consistency with spatial mental models: Hidden and obvious indeterminacy in spatial descriptions. In *Spatial Cognition IV: Reasoning, Action, Interaction*, C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, and T. Barkowsky (Eds.). Springer, Berlin, Germany, 165–180.
- [13] Barbara Tversky. 1991. Spatial mental models. *Psychol. Learn. Motiv.* 27 (1991), 109–145.
- [14] Maggie Wigness, John G. Rogers, and Luis E. Navarro-Serment. 2018. Robot navigation from human demonstration: Learning control behaviors. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '18)*. 1150–1157.
- [15] Kristin E. Schaefer, Brandon S. Perelman, Ralph W. Brewer, Julia L. Wright, Nicholas Roy, and Derya Aksaray. 2018. Quantifying human decision-making: Implications for bidirectional communication in human-robot teams. In *Virtual, Augmented and Mixed Reality: Interaction, Navigation, Visualization, Embodiment, and Simulation*, J. Chen and G. Fraomeni (Eds.). Springer, Cham, 361–379.
- [16] Brandon S. Perelman, Arthur W. Evans III, Kristin E. Schaefer, and Susan G. Hill. 2018. Attitudes toward risk and effort tradeoffs in human-robot heterogeneous team operations. In *Proceedings of the Human Factors and Ergonomics Society Meeting* 62, 1 (2018), 1098–1102. SAGE Publications, Los Angeles, CA.
- [17] Elizabeth Phillips, Scott Ososky, Janna Grove, and Florian Jentsch. 2011. From tools to teammates: Toward the development of appropriate mental models for intelligent robots. In *Proceedings of the Human Factors and Ergonomics Society Meeting* 55, 1 (2011), 1491–1495. SAGE Publications, Los Angeles, CA.
- [18] Scott Ososky, David Schuster, Elizabeth Phillips, and Florian Jentsch. 2013. Building appropriate trust in human-robot teams. In *Proceedings of the AAAI Spring Symposium: Trust and Autonomous Systems*. 60–65.
- [19] Raja Parasuraman and Victor Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Hum. Fact.* 39, 2 (1997) 230–253.
- [20] Kristin E. Schaefer and Edward R. Straub. 2016. Will passengers trust driverless vehicles? Removing the steering wheel and pedals. In *Proceedings of the IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA'16)*. 159–165.
- [21] Janis A. Cannon-Bowers, Eduardo Salas, and Sharolyn A. Converse. 1993. Shared mental models in expert team decision making. In *Current Issues in Individual and Group Decision Making*, N. J. Castellan Jr. (Ed.), 221–246. Erlbaum, Hillsdale, NJ.
- [22] John E. Mathieu, Tonia S. Heffner, Gerald F. Goodwin, Eduardo Salas, and Janis A. Cannon-Bowers. 2000. Influence of shared mental models on team process and performance. *J. Appl. Psychol.* 85, 2 (2000), 273–283.
- [23] Katia Sycara and Gita Sukthankar. 2006. *Literature Review of Teamwork Models*. Report CMU-RI-TR-06-50. Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

- [24] Kristin E. Schaefer, Edward R. Straub, Jessie Y. C. Chen, Joe Putney, and Arthur W. Evans III. 2017. Communicating intent to develop shared situation awareness and engender trust in human-agent teams. *Cog. Syst. Res.* 46, 26–39.
- [25] Victoria Alonso and Paloma de la Puente. 2018. System transparency in shared autonomy: A mini review. *Front. Neurobot.* 12, 83 (2018), DOI:10.3389/fnbot.2018.00083.
- [26] George B. Dantzig and John H. Ramser. 1959. The truck dispatching problem. *Manag. Sci.* 6 (1959), 80–91.
- [27] Lanah Evers, Ana I. Barros, Herman Monsuur, and Albert Wagelmans. 2014. Online stochastic UAV mission planning with time windows and time-sensitive targets. *Europ. J. Oper. Res.* 238, 1 (2014), 348–362.
- [28] Lanah Evers, Twan Dollevoet, Ana I. Barros, and Herman Monsuur. 2012. Robust UAV mission planning. *Ann. Ope. Res.* 222, 1 (2012), 293–315.
- [29] Xu C. Ding, Amir R. Rahmani, and M. Egerstedt. 2010. Multi-UAV convoy protection: An optimal approach to path planning and coordination. *IEEE Trans. Rob.* 26, 2 (2010), 256–268.
- [30] Lanny Lin and Michael A. Goodrich. 2014. Hierarchical heuristic search using a Gaussian mixture model for UAV coverage. *IEEE Trans. Cyber.* 44, 2 (2014), 2532–2544.
- [31] Brandon S. Perelman and Shane T. Mueller. 2013. Examining memory for search using a simulated aerial search and rescue task. In *Proceedings of the 17th International Symposium on Aviation Psychology*.
- [32] Hartwig Hochmair and Andrew U. Frank. 2000. Influence of estimation errors on wayfinding decisions in unknown street networks—analyzing the least-angle strategy. *Spatial Cog. Comput.* 2 (2000), 283–313.
- [33] David W. Eccles, Susanne E. Walsh, and David K. Ingledew. 2002. The use of heuristics during route planning by expert and novice orienteers. *J. Sports Sci.* 20 (2002), 327–337.
- [34] Thora Tenbrink and Inessa Seifert. 2011. Conceptual layers and strategies in tour planning. *Cog. Proc.* 12 (2011), 109–125.
- [35] Matthew Dry, Michael D. Lee, Douglas Vickers, and Peter Hughes. 2006. Human performance on visually presented traveling salesperson problems with varying numbers of nodes. *J. Prob. Solv.* 1, 1 (2006), 20–32.
- [36] James N. MacGregor and Tom Ormerod. 1996. Human performance on the traveling salesman problem. *Percep. Psychophys.* 58, 4 (1996), 527–539.
- [37] Brandon S. Perelman and Shane T. Mueller. 2016. Considerations influencing human TSP solutions and modeling implications. In *Proceedings of the 14th International Conference on Cognitive Modeling*.
- [38] Edsger W. Dijkstra. 1959. A note on two problems in connexion with graphs. *Numer. Math.* 1, 1 (1959), 269–271.
- [39] Nathan D. Otten, Heather L. Jones, David S. Wettergreen, and William L. Whitaker. 2015. Planning routes of continuous illumination and traversable slope using connected component analysis. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '15)*. 3953–3958.
- [40] Nathan D. Otten, David S. Wettergreen, and William L. Whitaker. 2018. Strategic autonomy for reducing risk of sun-synchronous lunar polar exploration. In *Proceedings of the International Conference on Field and Service Robotics*. 465–479. Springer, Cham.
- [41] Lawrence Mandow and José L. P. De La Cruz. 2010. Multiobjective A* search with consistent heuristics. *J. ACM* 57, 5 (2010) DOI:10.1145/1754399.1754400.
- [42] Adam R. Short and Douglas L. Van Bossuyt. 2015. Risk attitude informed route planning in a simulated rover. In *Proceedings of the International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*.
- [43] Dongoo Lee and Jaemyung Ahn. 2018. Integrated optimization of planetary rover layout and exploration routes. *Eng. Optim.* 50, 1 (2018), 164–182.
- [44] Adi Botea, Martin Müller, and Jonathan Schaeffer. 2004. Near optimal hierarchical path-finding. *J. Game Dev.* 1, 1 (2004), 7–28.
- [45] Daqing Yi, Michael A. Goodrich, and Kevin D. Seppi. 2014. Informative path planning with a human path constraint. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*. 1752–1758.
- [46] Daqing Yi, Michael A. Goodrich, and Kevin D. Seppi. 2016. Homotopy-aware RRT*: Toward human-robot topological path-planning. In *Proceedings of the 11th ACM/IEEE International Conference on Human Robot Interaction*. 279–286.
- [47] Daqing Yi, Michael A. Goodrich, and Kevin D. Seppi. 2015. MMORF*: Sampling-based multi-objective motion planning. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*.
- [48] Martin Werner and Sebastian Feld. 2014. Homotopy and alternative routes in indoor navigation scenarios. In *Proceedings of the IEEE International Conference on Indoor Positioning and Indoor Navigation*. 230–238.
- [49] Sanjiban Choudhury, Sebastian Scherer, and Sanjiv Singh. 2013. RRT*-AR: Sampling-based alternate routes planning with applications to autonomous emergency landing of a helicopter. In *Proceedings of the International Conference on Robotics and Automation*, 3947–3952.
- [50] Subhrajit Bhattacharya, Vijay Kumar, and Maxim Likhachev. 2010. Search-based path planning with homotopy class constraints. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*.

- [51] Subhrajit Bhattacharya, Maxim Likhachev, and Vijay Kumar. 2012. Topological constraints in search-based robot path planning *Auton. Rob.* 33, 3 (2012), 273–290.
- [52] Zeyad A. Algfoor, Mohd S. Sunar, and Hoshang Kolivand. 2015. A comprehensive study on pathfinding techniques for robotics and video games. *Int. Comput. Games Technol.* 7 (2015), 1–11. DOI:10.1155/2015/736138.
- [53] Emili Hernández Bes. 2012. *Path Planning with Homotopic Constraints for Autonomous Underwater Vehicles*. PhD. Dissertation. Universitat de Girona, Girona, Spain.
- [54] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Rob. Auton. Syst.* 57, 5 (2009), 469–483.
- [55] Christopher G. Atkeson and Stefan Schaal. 1997. Robot learning from demonstration. In *Proceedings of the 14th International Conference on Machine Learning*.
- [56] Beomjoon Kim and Joelle Pineau. 2015. Socially adaptive path planning in human environments using inverse reinforcement learning. *Int. J. Soc. Rob.* 8, 1 (2015), 51–66.
- [57] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. 2016. Socially compliant mobile robot navigation via inverse reinforcement learning. *Int. J. Rob. Res.* 35, 11 (2016), 1289–1307.
- [58] Giuseppe Della Penna, Benedetto Intrigila, Daniele Magazzeni, and Fabio Mercorio. 2010. Planning for autonomous planetary vehicles. In *Proceedings of the 6th IEEE International Conference on Autonomic and Autonomous Systems*. 131–136.
- [59] Herbert A. Simon. 1979. Rational decision making in business organizations. *Amer. Econ. Rev.* 69, 4 (1979), 493–513.
- [60] Reginald G. Golledge. 1995. Path selection and route preference in human navigation: A progress report. In *Proceedings of the International Conference on Spatial Information Theory*. Springer, Berlin, 207–222.
- [61] Tad T. Brunyé, Caroline R. Mahoney, Aaron L. Gardony, and Holly A. Taylor. 2010. North is up (hill): Route planning heuristics in real-world environments. *Mem. Cog.* 38, 6 (2010), 700–712.
- [62] Edward K. Sadalla and Lorin J. Staplin. 1980. The perception of traversed distance: Intersections. *Envir. Behav.* 12, 2 (1980), 167–182.
- [63] Avrim Blum, Shuchi Chawla, David R. Karger, Terran Lane, Adam Meyerson, and Maria Minkoff. 2007. Approximation algorithms for orienteering and discounted-reward TSP. *SIAM J. Comput.* 37, 2 (2007), 653–670.
- [64] Myke Gluck. 1991. Making sense of human wayfinding: Review of cognitive and linguistic knowledge for personal navigation with a new research direction. In *Cognitive and Linguistic Aspects of Geographic Space*, D. M. Mark and A. U. Frank (Eds.). 63, 117–135. Springer, Dordrecht.
- [65] Christoph Hölscher, Thora Tenbrink, and Jan M. Wiener. 2011. Would you follow your own route description? Cognitive strategies in urban route planning. *Cognition* 121, 2 (2011), 228–247.
- [66] Salim A. Mouloua, James Ferraro, Mustapha Mouloua, and Peter A. Hancock. 2018. Trend analysis of unmanned aerial vehicles (UAVs) research published in the HFES proceedings. In *Proceedings of the Human Factors and Ergonomics Society Meeting* 62, 1 (2018), 1067–1071. SAGE Publications, Los Angeles, CA.
- [67] Shane T. Mueller. 2014. PEBL: The psychology experiment building language (version 0.14) [computer experiment programming language]. Retrieved from <http://pebl.sourceforge.net>.
- [68] Kristin E. Schaefer, Tracy L. Sanders, Ryan E. Yordon, Deborah R. Billings, and Peter A. Hancock. 2012. Classification of robot form: Factors predicting perceived trustworthiness. In *Proceedings of the Human Factors and Ergonomics Society Meeting* 56, 1 (2012), 1548–1552. SAGE Publications, Los Angeles, CA.
- [69] Kristin E. Schaefer, Ashley N. Foots, and Edward R. Straub. 2018. *Applied Robotics and Installations and Base Operations: User Perceptions of a Driverless Vehicle at Fort Bragg*. Technical report: ARL-TR-8265. US Army Research Laboratory, Aberdeen Proving Ground, Maryland.
- [70] Shane T. Mueller, Brandon S. Perelman, and Elizabeth S. Veinott. 2016. An optimization approach for mapping and measuring the divergence and correspondence between paths, *Behav. Res. Meth.* 48, 1 (2016), 53–71.
- [71] William N. Venables and Brian D. Ripley. 2002. *Modern Applied Statistics with S (4th Edition)*. Springer, New York.
- [72] Friedrich Leisch. 2004. FlexMix: A general framework for finite mixture models and latent class regression in R. *J. Stat. Softw.* 11, 8 (2004), 1–18.
- [73] Aly M. Tawfik, Hesham A. Rakha, and Shadeeqa D. Miller. 2010. Driver route choice behavior: Experiences, perceptions, and choices. In *Proceedings of the IEEE Intelligent Vehicles Symposium*. 1195–1200.
- [74] Zygmunt Pizlo, Emil Stefanov, John Saalweachter, Zheng Li, Yil Haxhimusa, and Walter G. Kropatsch. 2006. Traveling salesman problem: A foveating pyramid model. *J. Prob. Solv.* 1, 1 (2006), 83–101.
- [75] Matt Duckham and Lars Kulik. 2003. “Simplest” paths: Automated route selection for navigation. In *Spatial Information Theory: Foundations of Geographic Information Science*, W. Kuhn, M. F. Worboys, and S. Timpf (Eds.). Springer, Berlin, 169–185.
- [76] Kai-Florian Richter and Matt Duckham. 2008. Simplest instructions: Finding easy-to-describe routes for navigation. In *Proceedings of the International Conference on Geographic Information Science*. 274–289. Springer, Berlin.

- [77] Pieter Vansteenwegen, Wouter Souffriau, and Dirk Van Oudheusden. 2011. The orienteering problem: A survey. *Europ. J. Oper. Res.* 209, 1 (2011), 1–10.
- [78] Georges A. Croes. 1958. A method for solving traveling salesman problems. *Oper. Res.* 6, 791–812.
- [79] Iris Van Rooij, Ulrike Stege, and Alissa Schactman. 2003. Convex hull and tour crossings in the Euclidean traveling salesperson problem: Implications for human performance studies. *Mem. Cog.* 31, 2 (2003), 215–220.
- [80] Brandon S. Perelman. 2015. *A Naturalistic Computational Model of Human Behavior in Navigation and Search Tasks*. PhD. dissertation. Michigan Technological University, Houghton, MI.
- [81] Brent Mittelstadt, Chris Russell, and Sandra Wachter. 2019. Explaining explanations in AI. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*. 279–288.
- [82] Edward R. Tufte, Susan R. McKay, Wolfgang Christian, and James R. Matey. 1998. Visual explanations: images and quantities, evidence and narrative. *Comput. Phys.* 12, (1998). DOI: [10.1063/1.168637](https://doi.org/10.1063/1.168637)
- [83] Kyle J. Behymer, Elizabeth M. Mersch, Heath A. Ruff, Gloria L. Calhoun, and Sarah E. Spriggs. 2015. Unmanned vehicle planning comparison visualizations for effective human-autonomy teaming. *Proc. Manuf.* 3 (2015), 1022–1029.
- [84] Kimberly Stowers, Nicholas Kasdaglis, Olivia Newton, Shan Lakhmani, Ryan Wohleber, and Jessie Chen. 2016. Intelligent agent transparency: The design and evaluation of an interface to facilitate human and intelligent agent collaboration. In *Proceedings of the Human Factors and Ergonomics Society Meeting* 60, 1 (2016), 1706–1710. SAGE Publications, Los Angeles, CA.
- [85] Laura M. Hiatt, Cody Narber, Esube Bekele, Sangeet S. Khemlani, and J. Gregory Trafton. 2017. Human modeling for human-robot interaction. *Int. J. Rob. Res.* 36 (5–7) (2017), 580–596.
- [86] Christopher A. Miller and Raja Parasuraman. 2007. Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Hum. Fact.* 49, 1 (2007), 57–75.
- [87] Anca D. Dragan, Kenton C. T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*. 301–308.
- [88] Peter Henry, Christian Vollmer, Brian Ferris, and Dieter Fox. 2010. Learning to navigate through crowded environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 981–986.
- [89] Peter A. Hancock, Deborah R. Billings, Kristin E. Schaefer, Jessie Y. C. Chen, Ewart J. De Visser, and Raja Parasuraman. 2011. A meta-analysis of factors affecting trust in human-robot interaction. *Hum. Fact.* 53, 5 (2011), 517–527.
- [90] Jessie Y. C. Chen and Michael J. Barnes. 2012. Supervisory control of multiple robots: Effects of imperfect automation and individual differences. *Hum. Fact.* 54, 2 (2012), 157–174.
- [91] Scott Ososky, Tracy Sanders, Florian Jentsch, Peter A. Hancock, and Jessie Y. C. Chen. 2014. Determinants of system transparency and its influence on trust in and reliance on unmanned robotic systems. In *Proceedings of the International Society for Optics and Photonics: Unmanned Systems Technology*. 9084. DOI: [10.1117/12.2050622](https://doi.org/10.1117/12.2050622)

Received March 2019; revised November 2019; accepted February 2020