ANNUAL REVIEWS

Annual Review of Control, Robotics, and Autonomous Systems

Inferring Human Intent and Predicting Human Action in Human–Robot Collaboration

Guy Hoffman,[1] Tapomayukh Bhattacharjee,[2] and Stefanos Nikolaidis[3]

[1]Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, New York, USA; email: hoffman@cornell.edu

[2]Department of Computer Science, Cornell University, Ithaca, New York, USA

[3]Thomas Lord Department of Computer Science, University of Southern California, Los Angeles, California, USA

Annu. Rev. Control Robot. Auton. Syst. 2024. 7:73–95

First published as a Review in Advance on November 29, 2023

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at control.annualreviews.org

https://doi.org/10.1146/annurev-control-071223-105834

Keywords

human–robot collaboration, intention inference, motion prediction, probabilistic methods, human–robot interaction

Abstract

Researchers in human–robot collaboration have extensively studied methods for inferring human intentions and predicting their actions, as this is an important precursor for robots to provide useful assistance. We review contemporary methods for intention inference and human activity prediction. Our survey finds that intentions and goals are often inferred via Bayesian posterior estimation and Markov decision processes that model internal human states as unobserved variables or represent both agents in a shared probabilistic framework. An alternative approach is to use neural networks and other supervised learning approaches to directly map observable outcomes to intentions and to make predictions about future human activity based on past observations. That said, due to the complexity of human intentions, existing work usually reasons about limited domains, makes unrealistic simplifications about intentions, and is mostly constrained to short-term predictions. This state of the art provides opportunity for future research that could include more nuanced models of intents, reason over longer horizons, and account for the human tendency to adapt.

1. INTRODUCTION

For robots to be useful collaborators in human–robot teams, they should be able to choose the correct action at the right time. A key capability that can help with the machine's optimal on-time activity is to understand the human's intentions and goals beyond their immediate physical action and be able to predict future actions (1, 2). Take, for instance, a manufacturing scenario, where a human places automobile parts for a robot to weld together. A collaborative robot could wait for a human command that explicitly requests a welding action, or it could wait for the parts to be placed and then act based on their positions and features. However, if a robot could instead reason about the human, understanding their larger goals, such as what assembly stage the team should currently be working on, or predicting the human's actions, such as where the human is likely to place the next part, the robot's collaborative capability would be greatly improved.

For that reason, researchers studying human–robot collaboration have dedicated significant effort to developing methods for human intention inference and action prediction. In early work, Hoffman & Breazeal (3) proposed that anticipating human actions based on a probabilistic inference method about human plans could lead to higher fluency and user satisfaction in a collaborative robot task. This observation has been replicated by other researchers (e.g., 4–6).

In this review, we map the prevalent techniques used in the literature to understand and anticipate human intentions, goals, and actions. We divide the literature on this topic into three sections, as depicted in **Figure 1**: inferring the human's intentions and goals, inferring specific collaborative features of the human, and predicting the human's future movement in space.

Section 2, Inferring Intentions and Goals, covers work that reasons about the unobservable human goal-oriented constructs that underlie their behavior. This includes inferring what overall strategy the team should be using, or it could be about a more specific activity goal, like what part of an assembly the human is trying to attach to an object at the current stage. Even more specific intention inference occurs when the robot tries to understand what point in space or object the human is trying to reach or wants the robot to reach.

In many cases, the focus is not on inferring intentions or goals, but rather on specific features of the human that are important to collaboration, such as their level of fatigue or stress, their capabilities, or their trust in the robot. Correctly inferring these features can help the robot predict future human behavior and help it choose its own best course of action. We cover research dealing with this inference challenge in Section 3, Inferring Collaborative Features.

Finally, Section 4, Predicting Human Motion, deals with methods that reason about the observable physical movement of the human. These works are concerned with the motion of the human, without necessarily inferring any internal state on the human's part. We review works that

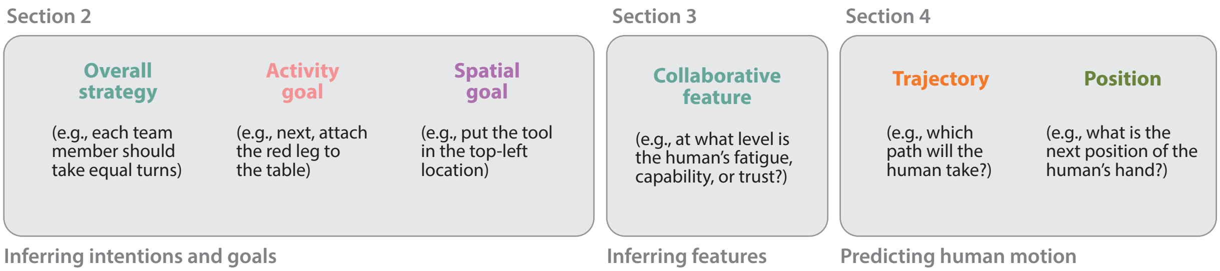

Figure 1

This review maps the recent literature on intention inference and human activity prediction across three categories. Intention and goal inference are presented in Section 2, collaborative feature inference in Section 3, and spatial motion prediction in Section 4.

are trying to infer or predict whole trajectories, as well as research focused on the instantaneous modeling of the next human position.

These categories are not mutually exclusive. Intent inference can happen with or without motion prediction, and action prediction can happen with or without higher-level models of intentions and goals. We further discuss this overlap in Section 1.1 but still find it useful to separate our presentation into the abovementioned categories for clarity.

Of course, a single article could not do justice to the full extent of the work in human–robot collaborative intention inference and activity prediction. Our review therefore focuses on a subset of methods and applications. It generally covers only research published in the last three years, with references to past work where they are pertinent. We discuss mostly work where humans and robots directly collaborate or interact on a shared task. This excludes, for example, a large literature concerned with intention inference in exoskeleton control, as well as the excellent work on intention recognition for autonomous driving. Methodologically, we also do not discuss the subfield of machine learning that deals with human preference learning and learning from demonstration, even though there is a clear aspect of intention inference to these methods. Finally, we prioritize research that includes implementations on physical robots, omitting many excellent projects implemented in simulated environments.

1.1. Intention Inference Versus Motion Prediction: A Blurry Boundary

Despite the categorization delineated in the previous section, the demarcation between intention inference and motion prediction is not always distinct. On the one hand, researchers try to map observations to intentions, often using Bayesian inference, Markov decision processes (MDPs), and supervised learning techniques. On the other, there is motion prediction, which focuses on forecasting the future movements or actions of humans. This involves estimating trajectories, paths, or motion patterns to anticipate the future positions and velocities of the entities involved. These works often use time series prediction methods or make use of neural networks.

These categories are not always exclusive, and there can be significant overlap between them. This overlap could be in the methods used or in the combination of goal inference with motion prediction in a single system. For example, Le et al. (7) combined inverse reinforcement learning (IRL) and a goal-conditioned recurrent neural network (RNN) to learn both discrete goals and continuous movements. Their method first infers the high-level goal of the user using IRL and then passes the inferred goal as input to the goal-conditioned RNN, which outputs a low-level trajectory. Cheng & Tomizuka (8) followed the inverse approach: Their method first estimates the motion type (e.g., reaching or installing) using a long short-term memory (LSTM) network and then uses the motion type to infer the intended goal. Another way to combine goal inference and motion estimation is to compare the current human motion with an expected motion given a target goal. This can be achieved by leveraging encoded demonstrations, such as dynamic movement primitives (DMPs) (9), as a reference for comparison, as proposed by Qiao et al. (10). Using this comparison, the robot can infer the intended goal behind the observed motion.

1.2. Collaborative Contexts

In the human–robot collaboration literature, intention inference and activity prediction are studied across a variety of contexts: Robots try to predict handover goal positions, reason over human plans in collaborative workspaces, infer operator intentions, and predict human motion for shared navigation. To provide a sense of where and how intention inference and action prediction are used, we describe some popular application contexts below, along with specific needs and challenges they present.

1.2.1. Handovers. Many collaborative tasks involve robot-to-human object handovers, human-to-robot handovers, or both (11). These can be direct object transfers, as in an IKEA coassembly study (12), or indirect ones involving an intermediate step, such as hand-to-table and table-to-pickup maneuvers (13). When coordinating a handover, a robot can benefit from estimating where the handover will occur. This requires the robot to infer the goal position of the human's movement or predict which object the human would like to hand over or receive. These inference methods usually take into account past human actions, features of the objects to be handed over, or a combination of both (14).

1.2.2. Shared workspace collaboration. In a related collaborative context, a human and a robot might work together on a shared workspace. This can include assembly tasks, where the human and robot take turns operating on an object being assembled (3). In other scenarios, the robot holds an object for a human to work on (15) or provides tools for the human to complete a task. In these contexts, the robot can benefit from inferring the human's planned sequence of actions, in order to predict the next action and formulate the best collaborative plan. The robot might also be required to infer what the human would want the robot to do. In some cases, the human provides corrections to the robot's activity, which then leaves the robot to extrapolate intention from corrective cues (16, 17).

A particular kind of shared workspace collaboration is direct comanipulation of an object during a collaborative task, sometimes called physical human–robot interaction. In this context, the robot may have to infer the goal position of an object (18). Alternatively, the robot might have to predict the human's movement in space or the applied forces and torques based on a model of human comanipulation (19). In some cases, comanipulation features can be used to infer higher-level states and intents (20).

1.2.3. Shared autonomy. A third context for human–robot collaboration is shared autonomy, where the human indicates an action for the robot to perform. This can be in the context of an assistive technology, in which case the robot performs an action based on limited control input from the human being assisted, or in the context of hybrid teleoperation. In both cases, partial or imperfect human input has to result in accurate robot activity.

These contexts necessitate methods for the robot to infer goal locations from human input or make inferences about higher-level meanings of human controls. In one scenario, a robot is tasked with the challenge of choosing one of several goal positions from partial or imperfect human instruction (6, 21, 22). In other cases, the robot does not have a set of possible goals or intents, but builds a probability distribution over its full operational space (10, 23).

Since the shared autonomy scenario generally assumes imperfect input, and since the human input is usually explicitly robot directed, it rarely poses a direct motion prediction problem, and more often is concerned with the inference of unobservable and underlying intents.

1.2.4. Social navigation. Finally, in many collaborative contexts, a human and a robot move around a shared space. As a result, a subset of human–robot interaction research deals with the problem of social navigation, roughly defined as the challenge of a robot navigating around humans. Robots moving in human spaces have to find ways to safely cross a human's path, lead a human, follow a human, or walk alongside them (e.g., 24, 25).

To avoid collisions and provide people with a sense of safety around the robot, the control algorithm often has to predict where the human will go. Researchers use a wide variety of computational models, such as grid-based probability distributions (26), hidden Markov models (HMMs) (27), MDPs (28), or sets of Gaussian distributions (29, 30). Kollmitz et al. (30), for example, used Gaussian distributions to build a social cost map for the robot, whereas Bennewitz et al. (29) clustered human trajectories into a number of movement patterns or classes and used these to both track and predict future human motion. These predictions and cost maps then allow the robot to plan accordingly.

The goal of the robot performing social navigation is usually physical in nature, such as avoiding humans or staying close to them. Therefore, high-level intent inference is rare in this context, and most works are concerned with momentary probabilistic predictions of human motion.

1.3. Reasoning About Humans Is Difficult

As described in the previous section, the survey presented here covers a variety of inference challenges, across a range of application areas, using different computational methods. The common factor to all the works surveyed in this article is that they deal with reasoning about human behavior. Understanding and predicting human behavior pose unique challenges for robotics and AI research, beyond the baseline problems of handling uncertainty in the robot's operation and its environment.

First, humans are complex systems with rich internal representations, cultural backgrounds, and social contexts. The full complexity of the psychological and cognitive processes that underlie human physical behavior is beyond the reach of existing computational models and methods. As a result, any reasoning about human intentions and goals is bound to rely on oversimplification and therefore will only achieve partial success in its inference and prediction.

Second, humans are highly adaptive creatures. Any change in the environment or any new or continued robot behavior is likely to affect the future cognition and behavior of the human who is interacting with the robot. While obviously beneficial to humans, their adaptiveness makes any inference expire quickly and any prediction imprecise. Robots would benefit from taking a similarly adaptive approach to reasoning about humans, but in reality, most computational systems do not.

Finally, the very notion of intention is complicated and ill defined. In the work we survey in this article, the word intention or goal can mean vastly different things, from strategies, to object choices, to preferences. The fact that intentions are unobservable also adds to the vagueness of their use in the AI and robotics literature.

When reading this survey, it is worth keeping in mind these challenges and limitations of the human–robot intention inference literature. We offer a longer discussion of these topics, along with some ways that researchers try to tackle them, in Section 5.

2. INFERRING INTENTIONS AND GOALS

We now turn to the challenge of inferring a human collaborator's intention or goal. This could be an attempt to predict a spatial goal position, i.e., where the human's action will end up and where the robot might meet the human's action. It could also be in the form of trying to understand a collaborative strategy from imperfect data, for example, whether the human wants to take strict and equal turns or not. Alternatively, a robot might face a selection problem: choosing one of a set of possible goals that the human is trying to achieve, or deciding which of a number of strategies the human is currently following.

There are several key approaches that researchers have used to make these inferences, which we survey in this section. A robot could use direct Bayesian inference to determine the posterior probability of each possible goal or intent; this inference could be done in a one-shot manner or could take the form of a dynamic probabilistic process, such as a Markov model or MDP. A different route is to classify the intent using supervised learning techniques, including neural network models. Supervised learning methods, including neural networks, attempt to directly map observations to probabilities over intents using function approximations based on training data.

Each of these families of methods has its advantages and shortcomings. Bayesian methods can be computationally efficient, and their meaning is straightforwardly interpretable, because these models explicitly encode the conditional probabilities between observations and the objects of inference. This explicit statement of probabilities is also a potential issue with Bayesian inference, as this approach is highly dependent on the modeling decisions of the researcher and is sensitive to assumed prior distributions. Most simple Bayesian methods also make inferences only about instantaneous constructs, a restriction overcome by sequential methods like HMMs and MDPs. The latter are able to encode temporal relationships and the evolving dynamics of collaboration. They also naturally integrate the robot's reasoning about the human's intentions and goals with its choice of action. HMMs and MDPs, however, also suffer from the explicit modeling constraints of Bayesian inference methods. Furthermore, in their basic forms they are less accurate than neural network models.

Supervised learning and neural network approaches make less explicit assumptions about the probabilistic relationships between variables and allow for the modeling of more complex and nonlinear mapping functions. These characteristics have made this route popular over the last few years. That said, any supervised learning method relies on the quality of the available training data for accuracy. Traditional supervised learning approaches, like regression models, are also sensitive to the choice of features, a problem overcome by the more feature-agnostic neural network models. However, the more complex a neural network model is, the more data it needs to perform well, and good training data are difficult to collect about humans, especially in specialized contexts like human–robot collaboration. Large models also result in difficult- or impossible-to-interpret models, which have negative effects for reproducibility and are linked to ethical concerns about the fairness and controllability of the output from these models.

2.1. Bayesian Methods

Most simply stated, intent or goal inference in human–robot collaboration can be represented as the confidence in an intended goal given a set of observations. Many systems operate by computing a probability distribution over the space of possible goals within a Bayesian framework. In general, Bayesian inference uses conditional probabilities to infer the value of unknown variables by combining a priori probability distributions with the likelihood of observed evidence. This likelihood is computed based on a model that relates the variable of interest to the evidence.

A common example scenario is shared autonomy in the assistive domain, where a robot tries to assist a human who indicates their intentions through imperfect control interfaces. The robot can use a Bayesian measure of confidence in the inferred human goal or intent to provide the assistance that best corresponds to it. For example, Jain & Argall (21) modeled the uncertainty of a user's intended goal during assistive tasks using a Bayesian filtering framework. The belief $b_t(g)$ in the intended goal g at time t is represented as a posterior $P(g_t|\theta_{0:t}) \propto P(\theta_t|g_t, \theta_{0:t-1})P(g_t|\theta_{0:t-1})$, where θ is an observation source. Their approach considers multiple observation sources, such as goal proximity and nonverbal human action, and uses the maximum a posteriori decision method (31) to predict the most likely goal.

Iregui et al. (22) used a similar posterior maximization approach to infer the intended goal position of an assistive robotic arm. Specifically, they used the maximal posterior probability $P(O_{\text{pred}}|x_{\text{hmi}})$ to choose a predicted goal object O_{pred} given a user interface input position x_{hmi}, derived from either eye tracking or a touchscreen. The "known" probabilities $P(x_{\text{hmi}}|O_{\text{pred}})$ are modeled as multivariate Gaussian distributions around the known object positions in the user interface.

Jonnavittula & Losey (23) used assistive manipulation of a coffee cup as an example domain to learn human objectives from user demonstrations by considering a user's limitations in providing good demonstrations due to their disability. Using repeated user demonstrations of trajectories $\mathcal{D} = \{\xi_1 \ldots \xi_N\}$ from an assumed choice set Ξ of all possible demonstrations, the robot models the human intent via an underlying reward function r. They used Bayes' theorem to model the belief $b(r) \propto P(r) \prod_{\xi \in \mathcal{D}} P(\xi | r, \Xi)$, where $P(r)$ is the robot's prior over human reward and $P(\xi | r, \Xi)$ is the likelihood of observing a trajectory ξ given that the human has a reward r and choice set Ξ. This likelihood is commonly calculated using a Boltzmann-rational model (32, 33) that assumes humans are noisily optimal and predicts that a human will select a trajectory with a probability proportional to the exponentiated return on the trajectory. Other ways of modeling human behavior are also possible (see, e.g., 34).

Researchers have also used Bayesian approaches for intent or goal inference in other human–robot collaboration scenarios. For example, in a human–robot collaborative pick-and-place operation, Felip et al. (6) inferred reach target locations using approximate Bayesian computation, a sampling-based inference method that originated in the field of genetics. This method obtains samples from an unknown posterior distribution by first generating synthetic data from a prior distribution and then using a similarity metric to reject samples that are far from the observations. In Felip et al.'s (6) method, a human upper-left body kinematics model serves to generate the likelihood of a hand trajectory given a target goal location, which in turn is sampled from a Gaussian centered on the gaze–table intersection. When given actual human hand trajectories as observations, their approach uses the mean squared error as a similarity metric to determine how close the generated synthetic data are to the observed data points.

A similar method to infer reach target locations was proposed by Zanchettin & Rocco (35), who used a Bayesian approach in a human–robot assembly task. Their method uses a model-based trajectory generator and compares it with new observations to update the probabilities associated with each target location. The model used for trajectory generation minimizes the overall curvature of the path. This model is used to generate a trajectory for a given goal location and compare the angle between the tangent to the modeled path and the tangent to the actually observed path. Their process results in a Gaussian distribution that represents the conditional probability of a new observation given a goal location and the previous observation.

2.2. Markov Models and Markov Decision Processes

In addition to one-shot Bayesian inference, researchers also use sequential probabilistic decision models to make intention inferences. Sequential models can be a good fit for collaborative assembly tasks, for example, where a robot has to infer the step-by-step assembly plan that the human intends to use. Knowing the human-intended sequence of actions can help the robot assist with the correct action at the right time. Note that even in assemblies with a small number of parts, the space of possible subassemblies and operation orderings scales rapidly (3).

To address these situations, researchers often use Markov models, MDPs, and partially observable MDPs (POMDPs) to reason about the human's next steps. MDPs are mathematical models that consist of a set of states and actions, where the system moves from one state to another based on the chosen action, according to a transition probability distribution. A key assumption in MDPs is the Markov property, stating that the history of the system can be entirely modeled by its current state, and subsequent states depend only on the current state and the chosen action. POMDPs extend MDPs to situations where the agent does not directly observe the underlying state but instead receives partial and noisy observations about the state. Both MDPs and POMDPs are highly applicable in human–robot collaboration, as they provide a principled way to represent the uncertainty that the robot faces regarding the human's state and actions.

Using MDPs and POMDPs, a robot's objective is usually to find an action policy that maximizes the expected cumulative reward over time. The success of MDP models is therefore highly dependent on the accuracy or usefulness of the reward function, making manually specified reward functions suboptimal. To overcome this issue, reward functions are sometimes learned from explicit human feedback or are inferred from human demonstrations, a method known as IRL.

In human–robot collaboration, the robot's inference about the human's sequence preferences, strategies, or intentions using Markov models or MDPs is often coupled with the robot's policy optimization process. Hoffman & Breazeal (3), for example, used a simple first-order Markov process to model the state of a shared workbench, in order to anticipate the next human action and act accordingly. More recently, Cramer et al. (36) modeled the human-desired assembly sequence as a path through an assembly state graph. They proposed a POMDP formulation that has a hidden state variable representing the currently assumed human assembly path and then uses human object selections as observations. They implemented the SARSOP (successive approximations of the reachable space under optimal policies) algorithm (37) to find a robot policy given this formulation.

Zhao et al. (38) used a combined HMM and MDP approach to infer different high-level strategies in a joint human–robot cooking scenario. Starting with data from humans collaborating on the task, their method uses expectation–maximization to learn the parameters for an HMM representing the relationship between hidden strategy states and observed action sequences. Clustering the hidden state sequences gives the robot a set of high-level strategies, which are used to find an optimal policy for the two-agent MDP as follows: During collaboration, the robot samples its actions from a weighted model of possible strategies, where each weight is determined by the likelihood of the human's actions under the corresponding strategy.

Nikolaidis & Shah (15) used an MDP model in a cross-training approach, where a human and a robot switch roles in the collaboration to learn a shared model of the task. Cross-training has two phases: a forward phase, where the human observes the learned behavior, and a rotation phase, where the human switches roles with the robot. This allows the robot to learn both the reward and transition functions of the MDP. In the forward phase, the robot learns a predictive model of the human's actions, and thus the state transitions. In the rotation phase, the robot learns the MDP's reward function by learning human preferences about the task from demonstrations.

Hadfield-Menell et al. (39) proposed a cooperative IRL approach. Similar to cross-training, the human and robot switch between a learning phase, where a human demonstrates their preference to the robot, and a deployment phase, where the robot exhibits its learned behavior. Cooperative IRL uses a game-theoretic model that allows improved learning of the reward function through implicit active teaching behavior.

Providing demonstrations can be tedious and time-consuming. An alternative approach is to let the robot execute a trajectory and only provide informative corrections that enable the robot to infer the human intent. Losey et al. (40) modeled the human's true intention about the robot's trajectory by employing a POMDP framework that views the human's intention as a hidden parameter. They then solved for an approximation of the ideal policy while taking the human's corrections into account as observations. Li et al. (41) also modeled the robot's behavior in an MDP context but made use of human corrections while taking into account the temporal connection between sequential corrections, rather than just considering each correction individually. To avoid time-consuming demonstrations on a complex task, Nemlekar et al. (42) proposed learning a human preference model by collecting demonstrations on a short, canonical task. They then used this preference model as a prior in the more complex task, further refining the prior via corrections by the human.

Another way to combine the human's intention inference with the robot's policy is to model them as a mixed observability MDP (MOMDP). In this approach, exemplified by Bandyopadhyay et al. (28), the human's behavior is modeled as optimizing an MDP with a reward structure that is dependent on their goal. The robot then solves for a larger MDP model, which includes unobservable beliefs over the human's goal. The MOMDP can be manually specified or learned from human demonstrations (43).

2.3. Supervised Learning and Regression Models

Going back to goal position or object inference, an alternative method is to directly learn a mapping between human activity data and actually selected goals, without having explicit conditional probability models. This can be done via any number of supervised learning methods. Tsitos et al. (44) presented such an approach, collecting 400 examples of reaching for one of four object positions while tracking the human's wrist using the OpenPose algorithm (45). In a comparative study, they found that linear regression and support vector machines outperform naive Bayes and decision tree classifiers.

A more informed way to learn a mapping between observations and goals is in the form of DMPs. DMPs represent the modeled motion as coming from a spring–damper system with basis functions that encode the shape of the movement. These models are then used to predict the user's goal by assessing the agreement between the user's input and the learned DMPs. Qiao et al. (10), for example, have successfully utilized this approach in a range of shared control teleoperation tasks (**Figure 2**), whereas Sidiropoulos et al. (46) used a similar DMP approach combined with extended Kalman filtering to predict end positions for human trajectories in a handover task.

While DMPs capture the mean motion of the demonstrations, they may lack the expressiveness necessary to capture the variance in the demonstrated motions. Probabilistic movement primitives (ProMPs), on the other hand, address this limitation by representing distributions over trajectories (47). Ly et al. (48) showed how ProMPs can be learned from demonstrations in shared control teleoperation scenarios.

Taking a different approach, researchers have also employed Gaussian process (GP) regression to infer human intent. In a collaborative manipulation task, Haninger et al. (18) trained separate GP models for each goal and employed Bayesian inference to compute a belief over the human's goal based on the applied human force. This belief was then integrated into a model predictive controller to compute a trajectory for the robot.

Figure 2

Dynamic movement primitives can be used to predict human goals in assisted teleoperation, for example in reaching, grasping, and pouring tasks. Figure adapted with permission from Reference 10.

Supervised learning models have also been used to infer intent from human corrective actions. In the work by Jin et al. (16), the human only needs to give a rough directional correction throughout the robot's operation. The robot's trajectory is modeled as optimizing a cost function specified by a weight vector over features. Each directional correction defines a hyperplane constraining the robot's weight vector search space, which the robot uses to adjust its cost function for future trajectory calculations.

2.4. Neural Network Models

In recent years, many researchers have turned to varieties of neural network architectures for the task of mapping observations to goals and intents. While neural networks can be categorized as a form of supervised learning, the high prevalence of these models in recent years merits a section of their own.

Neural network models can be used to predict locations, as in the work of Choi et al. (13), which deals with human-to-robot object handover through a shared workspace. The authors trained an RNN by taking as input the human's arm pose and the intersection of the human's gaze with the workspace. The network's output was a probabilistic map across the workspace of the object's placement location (**Figure 3**). When the uncertainty around the goal drops below a threshold, the robot finalizes a definitive action toward the inferred handover position. Hu et al. (49) also used neural networks to infer pointing gestures. They additionally classified different types of hand configurations to distinguish between grasp intents for the robot.

Instead of inferring the goal from the user's trajectory, Urkmez & Bozma (50) used pointing gestures for goal inference. They developed a learning model that utilizes convolutional neural networks to detect hands from RGB data and classify observations into pointing gestures. Their method uses hand geometry information to estimate the pointing direction toward the intended goal object, offering an intuitive means of goal inference.

In another example of using neural networks, Schrum et al. (17) interpreted intents from corrective actions. They started by assuming that humans have personal idiosyncrasies that drive the

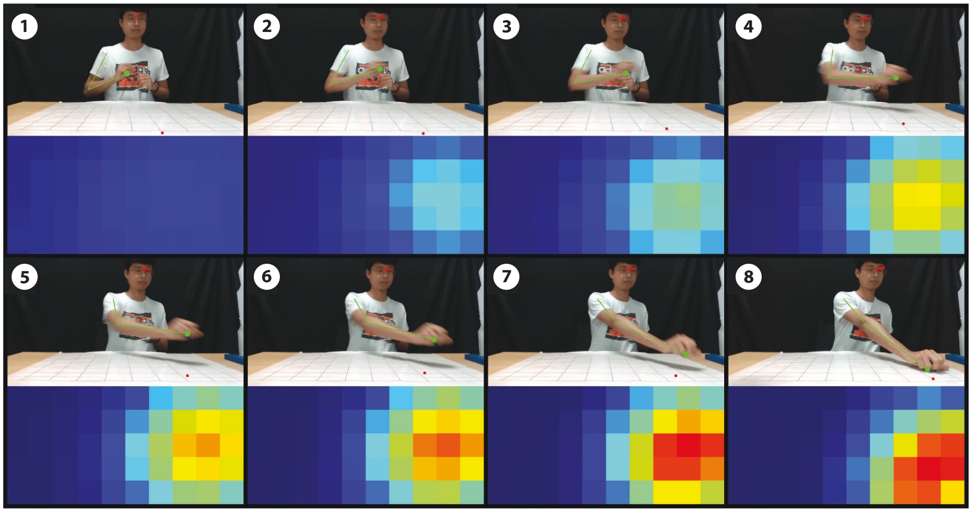

Figure 3

A recurrent neural network learns the mapping from the human's arm pose and gaze–table intersection to the predicted object placement position. Figure adapted with permission from Reference 13.

form of their feedback and suggested learning the mapping between the individual feedback style (e.g., anticipatory or delayed) and the intended corrective feedback by learning a personalized embedding in a bidirectional LSTM model.

There are many examples of using existing neural network architectures for visual or auditory recognition and transferring them into the human–robot collaboration domain. For example, Zhuang et al. (12) used an inflated 3D convolutional network architecture (51) for human action prediction in a collaborative IKEA assembly task. Cui et al. (52), in turn, proposed to convert natural language utterances (e.g., "pick up the book") and corrections (e.g., "no, to the left!") into demonstrated trajectories in a shared autonomy domain. They developed a mapping from states, joystick inputs, and language utterances to robot actions. To account for variations in natural language utterances, they used a frozen bidirectional encoder representations from transformers (BERT)–type model, which maps utterances to the nearest training examples.

In summary, there are three major approaches used for robots to infer a human's intent or goal. Bayesian methods calculate posterior probabilities for different intents given observations; these methods rely on assumed prior and conditional probability distributions that tie intentions to observations. Dynamic Bayesian processes, most commonly variations on POMDPs, model the uncertainty about the human's intentions as a simultaneous belief inference and policy selection problem; using this approach, robots can infer human goals and strategies and use these inferences to plan optimal actions. Finally, supervised learning methods, including conventional parametric and nonparametric methods as well as neural network models, attempt to directly map observations to intentions using training data. Each of these methods can also be used to interpret corrective actions rather than directly infer a specific intent.

3. INFERRING COLLABORATIVE FEATURES

So far, we have discussed ways in which a robot can make inferences about the human collaborator's intention or goal. In many cases, however, researchers study how to infer or predict a specific feature of the human's internal state to improve the collaboration. For example, a system could try to model a human's trust in the robot, the human's capabilities, or the human's level of fatigue. Each of these could then be used to predict future human activity or guide the robot's actions toward a more successful collaboration. This section discusses examples of work that deals with inferring such collaborative features.

3.1. Trust in the Robot

The trust a human has in a collaborative robot is key to the success of the human–robot team, as it can reduce interventions in the robot's operation, reversals by the human, and delays. This has led many researchers to propose ways to model human trust in robots and to use it in subsequent robot action selection.

Early on, Lee & Moray (53) found that an autoregressive moving-average vector model can predict almost 80% of the variance of experimentally obtained trust-in-machine ratings, using machine failures and performance rates as predictors. In recent years, trust in the robot has usually been estimated as a temporally evolving metric through a dynamic probabilistic model. In their seminal work, Xu & Dudek (54) tested the relationship between the trust an operator has in an unmanned aerial vehicle and several predictors, like human intervention events and robot errors, leading them to propose a dynamic Bayesian network that encodes these relationships and their evolution over time. The model's parameters were learned via expectation–maximization. Based on these findings, many researchers today rely on robot errors and human interventions as key variables to track when estimating human trust in a machine.

Once a robot has a model of trust evolution, it can use this inference to either engage in trust-repair behavior or utilize trust in other ways to improve human–robot collaboration. We refer the reader to Zahedi et al.'s (55) excellent review of the trust-aware planning literature.

In many cases, researchers use a POMDP formulation to integrate the evolution of human trust with the choice of the robot's trust-based policy. For example, Chen et al. (56) described such a POMDP that allows the robot to maintain a belief state about the human trust based on observable variables such as robot performance and human actions. The trust variable is modeled as a probabilistic dynamic model dependent on previous trust and robot performance. Guo et al. (57) also used a POMDP but proposed a more nuanced model of trust evolution via a beta distribution with parameters that change based on the robot's performance. Azevedo-Sa et al. (58) extended the simple scalar probabilistic models of trust to separately track trust across a set of agent capabilities, rather than just overall trust in the robot's performance.

While much of the trust-related human inference work is concerned with estimating the trust the human has in the robot, Wang et al. (59) conversely described a human trustworthiness measure that is calculated from a combination of human features, including their position, velocity, acceleration, and change in applied force.

3.2. Capabilities

Whereas trust in the robot is linked mostly to the robot's performance, researchers have also worked on modeling various aspects of the human's collaborative capabilities, such as their expertise and availability. This can help predict human behavior and aid in robot decision-making.

Recently, Carreno-Medrano et al. (60) proposed a framework where human expertise is represented by a parameter in a maximum entropy policy. This policy guides the human's trajectory by maximizing the likelihood of actions chosen based on the internal objective. Through Bayesian inference and analysis of demonstrated trajectories, the researchers estimate this human expertise parameter.

Since human capabilities can vary over time, Liu et al. (61) investigated the learning curve of individuals as they acquire proficiency in a given task. They proposed a model that describes the learning curve using an exponential function of the form $y = c + ke^{-\beta i}$, where i is the number of repetitions and c, k, and β are individual-specific parameters. These parameters are inferred online using an extended Kalman filter, capturing the learning rate and the potential for improvement. By tracking the learning curve, this method constructs robust schedules and optimizes task allocations based on an individual's proficiency level.

Some research takes a less agnostic view of human capabilities and instead relates them to specific cognitive traits. Kolb et al. (62) pretested participants on two traits: retaining situational awareness and modeling network structure. They demonstrated that the latter can effectively predict performance in tasks related to multirobot network management.

However, mere capability might not be enough to be a useful collaborator. The human also has to have a willingness to collaborate. Nanavati et al. (63) therefore attempted to infer how helpful a bystander would be when a robot needs to ask for help. Their method looks at inherent helpfulness, as well as the human's instantaneous busyness and previous interactions with the human, and uses a general linear mixed model to predict helpfulness. The robot then uses a POMDP formulation to plan under the uncertainty of this inference.

3.3. Workload, Fatigue, and Frustration

On the converse side of human capabilities, researchers have also explored the estimation of negative features of humans collaborating with robots, including workload, fatigue, and frustration. These phenomena can be related to physical factors, which may increase the risk of injury, or to cognitive factors, which may lead to lower technology acceptance. Modeling and countering these features can help increase the effectiveness of a collaborative task.

Estimating physical workload is often done using either musculoskeletal models or ergonomic posture models. Messeri et al. (64) used a deep neural network to learn a mapping from human motion to muscle activation in order to predict how muscles activate during task execution. Their method uses this mapping to estimate the physical fatigue accumulated by a worker. This model is then used to dynamically allocate tasks between a human and a robot to minimize physical fatigue for efficient human–robot collaboration. El Makrini et al. (19) estimated physical workload using virtual spring systems that generate torques corresponding to the deviation of the human skeleton kinematic chain from standard human ergonomic postures, then used hierarchical finite-state machines for role allocation based on this physical workload estimation.

For cognitive workload estimation, researchers have generally used nonverbal communication such as gaze, body language, and physiological signals. In an example of the latter, Lagomarsino et al. (65) used electrocardiograms during human–robot collaboration. They monitored the mean interval between two consecutive heartbeats over a time window along with the intervals' variation over time to estimate the cognitive workload. This estimate was then used to adapt the execution time and smoothness of robot trajectories. Kalatzis et al. (66) also used heart-rate variability as an indicator of cognitive workload or fatigue. In their work, machine learning models, including support vector machines, *k*-nearest neighbors, logistic regression classifiers, AdaBoost, and random forests, map heart rate to fatigue during a human–robot collaborative surface-polishing task.

Finally, Mohamed et al. (67) showed that incorporating thermal imaging can improve inference about a user's frustration over just using RGB facial images. They used a *k*-nearest-neighbors approach as a supervised learning inference method.

In summary, inferring specific features about the human can be important for human–robot collaboration. In many cases, features such as trust in the robot and human capabilities are inferred and simultaneously used in planning via dynamic probabilistic models such as POMDPs. When inferring negative features, such as workload or frustration, researchers tend to employ supervised learning methods that often take into account biological models and use biometric measures as input variables.

4. PREDICTING HUMAN MOTION

So far, we have discussed work that tries to infer something about the human's unobservable internal state. In many cases, however, researchers in human–robot collaboration are interested in detecting and predicting the human's physical motion through space. While this process of inference and prediction is sometimes tied to a model of intentions or goals, in other cases it is treated as a direct prediction problem, which does not take an intentional stance.

A variety of approaches have been proposed to tackle this prediction problem. We divide the discussion into two categories. First, deterministic methods make specific future predictions from past positional data, mostly using supervised learning. In recent years, these function approximation approaches have generally converged on neural network models. The second category, probabilistic predictions, includes research that explicitly models the uncertainty inherent to the motion prediction task. These works study methods to calculate probability distributions of possible movement predictions for the robot to take into account.

In some cases, the methods used in this section predict momentary motions of the human. In other works, the prediction is made over longer horizons and full trajectories.

Figure 4

A robot traces a person's arm in a bathing task using a neural network model that maps measurements from the capacitive sensor mounted on the robot to the human's pose. Figure adapted with permission from Reference 69.

4.1. Deterministic Predictions

Predicting motion can be done using simple regression methods. Nguyen & Xie (68), for example, predicted fingertip trajectories during human arm reaching movements using regression learning on a five-parameter logistic model. This proved to be more effective in representing natural fingertip trajectories than a commonly used minimum-jerk model. MDPs are also sometimes used to model motion trajectories. For instance, Bandyopadhyay et al. (28) combined intention inference, motion prediction, and planning in an MOMDP formulation.

Today, however, researchers are increasingly using neural networks to estimate the human's position and predict their motion. For example, Erickson et al. (69) used a fully connected neural network to infer the pose of a human limb based on a window of measurements from a robot-mounted capacitive sensor. This method offers a noninvasive and real-time solution for understanding human limb movements and can allow a robot to trace a person's arm in assistive tasks, such as bathing a patient (**Figure 4**). Ondras et al. (70) have used a multichannel deep convolutional neural network to predict mouth-opening timing for taking a bite in robot-assisted feeding in a social-dining scenario. They used features such as gaze, head pose, arm gestures, and speaking status from other diners to predict the bite timing of the user being fed.

In comanipulation tasks, researchers have combined information on the human skeleton with information on the manipulated objects. Wan et al. (71) trained a graph convolutional neural network that uses features of human motion and of manipulated objects to predict human motions. The network was trained from videos of human actors manipulating objects of different properties. The benefit of using a graph convolutional neural network is that it can incorporate contextual information about the human, objects, and interactions between the two, leading to more accurate predictions. In related work, Laplaza et al. (14) used a multihead attention architecture that combines the human motion, the robot end effector, and the position of obstacles to predict human motion in a handover task.

In addition to the above methods, RNNs have also been used to predict human motion. An example is the social LSTM approach by Alahi et al. (72), which pools together several LSTM networks to account for multiple humans moving in the same space.

Researchers have also integrated neural networks that structure low-level human motions into hierarchical approaches. Le et al. (7) combined low-level and high-level intention prediction by first learning a high-level goal policy using IRL and then implementing a goal-conditioned RNN to generate a series of complete body trajectories that correspond to the high-level goals. This approach allowed them to predict longer-trajectory sequences that represent sequential collaborative behavior.

4.2. Probabilistic Predictions

Whereas the above inference and prediction methods try to find a best-guess outcome of the model, a robot can also explicitly model the uncertainty inherent in the prediction. Chang et al. (73), for example, observed that users approached a robot's end effector in a linear motion and applied a first-order linear system to model the approach dynamics; a Kalman filter then provides continuous estimates of the human pose while representing uncertainty in the estimate through a covariance matrix of a Gaussian distribution. Thompson et al. (26) used a grid-quantized probabilistic map to predict human motion in a social navigation setting, whereas Kollmitz et al. (30) devised a Gaussian prediction model for social navigation and translated it into a social cost map for the robot to follow.

To model more complex motion dynamics, one approach used by researchers is the utilization of encoder–decoder architectures. Yasar & Iqbal (74) employed an encoder–decoder architecture that takes into account information on the human's position, velocity, and acceleration. The encoder captures the input information, while the autoregressive decoder generates predictions via the latent representation summarizing the observation history. This approach not only enables the model to make longer-horizon predictions but also includes an interpretable latent distribution.

While encoder–decoder architectures can be useful for predicting motion dynamics and representing lower-dimensional information in a probabilistic manner, it is equally important to consider the uncertainty associated with future predictions, especially when human safety is a concern. GPs have emerged as a popular model for probabilistic prediction due to their capacity to represent uncertainty and make predictions even with limited data.

Li et al. (75) proposed such a GP-based approach, assuming that human motion is smooth with respect to a reproducing kernel in Hilbert space. By integrating motion prediction into a model predictive controller in an assistive dressing scenario, their method ensured safety while assisting the human. Similarly, Jin et al. (76) employed GPs to model human hand velocities in a comanipulation task. Their approach integrated the predicted velocities into a controller that combines a proactive policy, optimizing task execution using human velocities, with a conservative policy focused on safety. The balance between these policies was achieved by incorporating uncertainty estimates encoded as a covariance matrix from the learned GP model.

In addition to predicting hand velocities, the estimation of human posture is crucial in comanipulation tasks for ergonomic considerations. Vianello et al. (77) tackled this problem by formally defining it as the modeling of the distribution of the null space of the Jacobian of a human skeleton with articulated joints. They utilized a GP to capture the distribution, resulting in a probabilistic estimation of future postures that satisfy the kinematic constraints imposed by the manipulated object operated by both the human and the robot.

In summary, in human–robot collaboration, many motion prediction approaches that do not explicitly model the internal state, intent, or goal of the human take a supervised learning approach. This is in contrast to traditional motion prediction tasks that do not involve humans, which often resort to simpler physics-based or kinematic models. Researchers in human–robot collaboration appreciate that a human's future trajectory is not easily captured by a physics-only approach. Today, most of these supervised learning methods converge on neural-network models.

While many of the works mentioned above make a best-guess prediction about the human's future trajectory, some also acknowledge the need to model the underlying uncertainty about these predictions. This is most often achieved by incorporating uncertainty-capturing representations in the model, for example, via GPs, covariance matrices, or latent representations in a neural network.

5. CHALLENGES WHEN REASONING ABOUT HUMAN INTENTIONS

In the sections above, we have focused exclusively on inferences and predictions about humans. That said, the computational methods that were discussed build on a large literature of robot estimation and planning methods outside the human domain, specifically those that deal with uncertainty. In the general robotics literature, this uncertainty is related to the robot's imperfect sensors and actuators, as well as to the uncertainty about the environment. In this section, we discuss how and why reasoning about humans, particularly about their intents and future actions, adds unique challenges on top of those generally considered in robotics.

5.1. Humans Are Complex and Hard to Predict

While any attempt to infer future states and to plan for them entails uncertainty, humans are—as the human–robot interaction adage goes—a special kind of environment. Humans are autonomous agents, driven by a host of internal, societal, cultural, and idiosyncratic factors, all of which are difficult to observe or model.

By and large, the human–robot collaboration community has used and extended existing probabilistic inference, prediction, and planning models to work with and around humans. These models rely on simplifications, which are less applicable to humans than to nonliving processes. To integrate human-specific aspects of internal or predictive processes, researchers have often proposed the ad hoc integration of social or psychological constructs into their computational model. For instance, MDP models can assume that humans have idiosyncratic preferences and decide their actions by maximizing a reward function based on those preferences (15, 39). Social constructs such as competition, cooperation, or coercion can also be built into an MDP model (78). Temporal changes in human behavior can be explained using learning curves (61), assumed cognitive states such as trust (56), or affective states such as frustration (67).

Still, the complexity underlying human unpredictability means that inferences often can be made only in highly limited domains, such as choosing one of a number of possible goals, and predictions generally operate on a short time horizon. These limitations arise in part from the inherent disparities between computational modeling frameworks and real human behavior. A notable example is the inability to succinctly represent human sensory experiences within a Markov process state (79). Human perception often involves the recall of past experiences triggered by stimuli, invalidating the Markov assumption. Additionally, the reward function mechanism often makes assumptions about rationality, without acknowledging that rationality and expertise can be highly subjective and contextual. Existing models also fail to account for habitual or impulsive behaviors that humans exhibit, further highlighting the divergence between existing computational models and the complexity of human behavior.

5.2. Humans Are Adaptive

A further complication in the challenge of reasoning about humans is that humans tend to adapt to a changing situation. As a result, any action on a collaborative robot's part is likely to be taken into account by a human counterpart. Changes in the environment that are external to humans may also affect future human states and actions. Any model or method that does not take into account the effect of the robot's action or changes in the environment on the human's internal state is unlikely to function correctly when the human adapts.

One approach for modeling the adaptation of human behavior in human–robot collaboration is through models that capture how human intent changes over time. For instance, Nikolaidis et al. (80) modeled the human as a bounded-memory agent who infers the strategy being followed by the robot based on a window of past observations. Subsequently, the model allows for the human to adapt and switch to another strategy with a certain probability. Similarly, a method devised by Parekh et al. (81) learns from a history of past interactions how the human's strategy evolves over time. By dynamically adapting to the changing human behavior, the robot can optimize its own actions to enhance the overall performance of the collaboration.

While these adaptive models can capture the observed changes in human states and actions, they do not explain why such changes occur and thus have limited predictive power across different tasks. An alternative route is to attempt to model why human actions change as a function of a changing human internal state. For instance, we have discussed previously how Chen et al. (56) modeled the dynamics of human trust, as the human observes the robot succeeding or failing in actions of different risk. Liu et al. (61) modeled the human's learning of a task using an exponential function and leveraged this information to schedule team activities. Tian et al. (82) modeled the human's learning of the robot's dynamics as a set of time-varying features. By accounting for the evolving human perception, the robot can, in turn, influence the human's learning process.

5.3. Intentions Are Complex and Ill Defined

An additional challenge lies in the lack of a clear definition of what intents are, especially compared with more straightforward action or motion prediction. If a computational model can accurately predict future spatial trajectories, one can say that it is successful in its task. The question of inferring intents, in contrast, confronts a deeper challenge, as the very notion of intent is complex and not well defined.

It is useful to consider the work reviewed here in light of the earlier belief–desire–intention (BDI) models popular in the 1990s (83, 84). These models owe their formulations to philosophical discussions on the nature of intent cognition in the 1980s (85). BDI models attempted to build systems of computation that could model the relationships between events and internal states reasoning about these events—namely, beliefs, desires, and intentions. While the systems built under this framework had clear limitations (for a discussion, see 86), they still strove for a theoretically complex and nuanced understanding of intentionality. The models included, for example, false beliefs, long time horizons, and conditional intents.

Today's human–robot collaboration literature takes a decidedly more pragmatic approach, and most of the historical lineage of intention modeling has been replaced by strictly probabilistic models, making use of the vastly increased sensor capability and computational power available today. At the same time, the computational notion of intent as understood in human–robot collaboration does not rise to the complexity that human intent encompasses. In most of the work reviewed here, intents are reduced to a selection of one of several goals, a spatial position, or a selection of one of a limited number of strategies. In reality, intents can span a much richer space, including plans that can diverge into different alternatives given specific outcomes, or weak preferences over actions and results. The research community would be well served by a more thoughtful and nuanced definition of intents in the context of human–robot collaboration.

6. FUTURE DIRECTIONS

In light of the above, research in intent inference and action prediction in human–robot collaboration can expand its boundaries in several ways. Robots could learn to reason about intents over longer time horizons, probably necessitating more complex models and higher-level symbolic representations about how human intent is structured. More research is also necessary to build models and systems that can account for the mutual adaptation of human and robot when they collaborate. Finally, the community would benefit from better definitions of intent and broadly agreed-upon benchmarks of success.

6.1. Longer Time Horizons

Most of the intent-related work in human–robot collaboration deals with short time horizons, mostly reasoning about the immediate next action or activity. The BDI architectures of the 1990s had higher ambitions, trying to model and reason about enduring intents and events further in the future than the work we surveyed here.

Perhaps the high reliance on probabilistic models that reason over high-frequency sensor data contributes to the short-term nature of much of the current work in intention inference. That said, some recent work uses computational architectures that are able to reason about states across longer timescales, such as the work by Le et al. (7) discussed in previous sections, which combines IRL with goal-conditioned RNNs. Another example is recent work by Patel & Chernova (87) that uses a graph representation of object relations along with a generative neural network to predict object state configurations up to several hours into the future. These predictions can be used by a robot to proactively assist a human in their daily tasks. The dataset used to train these models also includes human activity patterns over days-long periods of time and could be used for longer-horizon intention inference and action prediction in future work. Further developing inference methods based on such models that make room for longer-term predictions is a worthwhile direction for collaborative robotics.

6.2. Modeling Mutual Adaptation

A robot cannot assume that the intents of a human are static. It is more likely that the human will adapt to the robot's behavior and to changes in the environment. A promising future research direction would therefore be to develop models and systems that take the human adaptation into account and can provide for robots to adapt their inferences, predictions, and plans accordingly.

Tabrez et al. (5) presented a survey of mental modeling techniques for human–robot teaming in which second- and higher-order models address some of this recursive reasoning. Some researchers have taken this path using multiagent MDPs, as they can explicitly model the plans and reward functions of both agents (see Section 5.2). Others have used POMDP formulations that include hidden state features that can change over time and be used for the robot to find optimal strategies for adapting to humans. At least part of the solution could be in designing expressive, legible, and explainable robot actions that would guide the human to adapt in ways that are more predictable and productive.

At any rate, robust models for intention inference and action prediction will have to consider the dyadic adaptive relationship unfolding during the human–robot collaboration. To do so, it may be necessary to consider human models that go beyond the traditional assumptions of bounded rationality and noisy rationality to represent a wider range of real-world human–robot interaction behaviors.

6.3. Better Definitions and Benchmarks

Finally, the complex and ill-defined nature of human intentions necessitates more discussion around what researchers mean when they presume to reason about intents. In contrast to early work, which attempted to formalize intentions through calculus relations over events, beliefs, goals, and actions, the current human–robot collaboration literature takes a looser view toward defining intentions. In many cases, intentions are synonymous with strategies, goal locations, policies, or future trajectories. Having a clearer understanding of what kind of intention is being modeled and inferred would clarify the discourse in this literature.

Clear definitions could also lead to benchmarks for intention inference that can help the community compare approaches and measure progress. Usually, researchers use task success or other metrics related to collaborative outcomes as proxies for successful reasoning about intentions. While it is difficult—and perhaps impossible—to measure a human's internal state directly, it would still be useful to have some robust way to assess the accuracy of intention inferences.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The authors would like to thank Reuben Aronson, Yuchen Cui, and Maithili Patel for helpful comments on an earlier draft. G.H. was supported in part by the Defense Advanced Research Projects Agency (DARPA) under contract number W911NF2010004. T.B. was supported in part by the National Science Foundation under grant numbers 2132846 and 2238792. S.N. was supported in part by an Agilent Early Career Award and the National Science Foundation under grant numbers 2024936 and 2145077.

LITERATURE CITED

1. Grosz BJ. 1996. Collaborative systems (AAAI-94 Presidential Address). *AI Mag.* 17(2):67–85
2. Hoffman G, Breazeal C. 2004. Collaboration in human-robot teams. In *AIAA 1st Intelligent Systems Technical Conference*, pap. 2004-6434. Palo Alto, CA: AAAI
3. Hoffman G, Breazeal C. 2007. Cost-based anticipatory action selection for human–robot fluency. *IEEE Trans. Robot.* 23(5):952–61
4. Chang ML, Gutierrez RA, Khante P, Short ES, Thomaz AL. 2018. Effects of integrated intent recognition and communication on human-robot collaboration. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3381–86. Piscataway, NJ: IEEE
5. Tabrez A, Luebbers MB, Hayes B. 2020. A survey of mental modeling techniques in human–robot teaming. *Curr. Robot. Rep.* 1(4):259–67
6. Felip J, Gonzalez-Aguirre D, Nachman L. 2022. Intuitive & efficient human-robot collaboration via real-time approximate Bayesian inference. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3093–99. Piscataway, NJ: IEEE
7. Le AT, Kratzer P, Hagenmayer S, Toussaint M, Mainprice J. 2021. Hierarchical human-motion prediction and logic-geometric programming for minimal interference human-robot tasks. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication*, pp. 7–14. Piscataway, NJ: IEEE
8. Cheng Y, Tomizuka M. 2022. Long-term trajectory prediction of the human hand and duration estimation of the human action. *IEEE Robot. Autom. Lett.* 7(1):247–54
9. Schaal S. 2006. Dynamic movement primitives—a framework for motor control in humans and humanoid robotics. In *Adaptive Motion of Animals and Machines*, ed. H Kimura, K Tsuchiya, A Ishiguro, H Witte, pp. 261–80. Tokyo: Springer
10. Qiao CZ, Sakr M, Muelling K, Admoni H. 2021. Learning from demonstration for real-time user goal prediction and shared assistive control. In *2021 IEEE International Conference on Robotics and Automation*, pp. 3270–75. Piscataway, NJ: IEEE
11. Ortenzi V, Cosgun A, Pardi T, Chan WP, Croft E, Kulic D. 2021. Object handovers: a review for robotics. *IEEE Trans. Robot.* 37(6):1855–73

12. Zhuang Z, Ben-Shabat Y, Zhang J, Gould S, Mahony R. 2022. GoferBot: a visual guided human-robot collaborative assembly system. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 8910–17. Piscataway, NJ: IEEE
13. Choi A, Jawed MK, Joo J. 2022. Preemptive motion planning for human-to-robot indirect placement handovers. In *2022 International Conference on Robotics and Automation*, pp. 4743–49. Piscataway, NJ: IEEE
14. Laplaza J, Garrell A, Moreno-Noguer F, Sanfeliu A. 2022. Context and intention for 3D human motion prediction: experimentation and user study in handover tasks. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication*, pp. 630–35. Piscataway, NJ: IEEE
15. Nikolaidis S, Shah J. 2013. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *HRI '13: Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 33–40. Piscataway, NJ: IEEE
16. Jin W, Murphey TD, Lu Z, Mou S. 2023. Learning from human directional corrections. *IEEE Trans. Robot.* 39(1):625–44
17. Schrum ML, Hedlund-Botti E, Moorman N, Gombolay MC. 2022. MIND MELD: personalized meta-learning for robot-centric imitation learning. In *HRI '22: Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 157–65. Piscataway, NJ: IEEE
18. Haninger K, Hegeler C, Peternel L. 2022. Model predictive control with Gaussian processes for flexible multi-modal physical human robot interaction. In *2022 International Conference on Robotics and Automation*, pp. 6948–55. Piscataway, NJ: IEEE
19. El Makrini I, Omidi M, Fusaro F, Lamon E, Ajoudani A, Vanderborght B. 2022. A hierarchical finite-state machine-based task allocation framework for human-robot collaborative assembly tasks. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 10238–44. Piscataway, NJ: IEEE
20. Al-Saadi Z, Hamad YM, Aydin Y, Kucukyilmaz A, Basdogan C. 2023. Resolving conflicts during human-robot co-manipulation. In *HRI '23: Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 243–51. New York: ACM
21. Jain S, Argall B. 2019. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Trans. Human-Robot Interact.* 9(1):2
22. Iregui S, Schutter JD, Aertbelien E. 2021. Reconfigurable constraint-based reactive framework for assistive robotics with adaptable levels of autonomy. *IEEE Robot. Autom. Lett.* 6(4):7397–405
23. Jonnavittula A, Losey DP. 2021. I know what you meant: learning human objectives by (under)estimating their choice set. In *2021 IEEE International Conference on Robotics and Automation*, pp. 2747–53. Piscataway, NJ: IEEE
24. Repiso E, Garrell A, Sanfeliu A. 2020. People's adaptive side-by-side model evolved to accompany groups of people by social robots. *IEEE Robot. Autom. Lett.* 5(2):2387–94
25. Hu Y, Ryu J, Gundana D, Petersen KH, Kress-Gazit H, Hoffman G. 2023. Nudging or waiting? Automatically synthesized robot strategies for evacuating noncompliant users in an emergency situation. In *HRI '23: Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 603–11. New York: ACM
26. Thompson S, Horiuchi T, Kagami S. 2009. A probabilistic model of human motion and navigation intent for mobile robot path planning. In *2009 4th International Conference on Autonomous Robots and Agents*, pp. 663–68. Piscataway, NJ: IEEE
27. Bui HH, Venkatesh S, West G. 2001. Tracking and surveillance in wide-area spatial environments using the abstract hidden Markov model. *Int. J. Pattern Recognit. Artif. Intell.* 15(1):177–96
28. Bandyopadhyay T, Won KS, Frazzoli E, Hsu D, Lee WS, Rus D. 2013. Intention-aware motion planning. In *Algorithmic Foundations of Robotics X*, ed. E Frazzoli, T Lozano-Pérez, N Roy, D Rus, pp. 475–91. Berlin: Springer
29. Bennewitz M, Burgard W, Thrun S. 2003. Adapting navigation strategies using motions patterns of people. In *2003 IEEE International Conference on Robotics and Automation*, Vol. 2, pp. 2000–5. Piscataway, NJ: IEEE
30. Kollmitz M, Hsiao K, Gaa J, Burgard W. 2015. Time dependent planning on a layered social cost map for human-aware robot navigation. In *2015 European Conference on Mobile Robots*. Piscataway, NJ: IEEE. **https://doi.org/10.1109/ECMR.2015.7324184**

31. Van Trees HL. 2002. *Optimum Array Processing*. New York: Wiley-Intersci.
32. Baker CL, Saxe R, Tenenbaum JB. 2009. Action understanding as inverse planning. *Cognition* 113(3):329–49
33. Ziebart BD, Maas A, Bagnell JA, Dey AK. 2008. Maximum entropy inverse reinforcement learning. In *AAAI '08: Proceedings of the 23rd National Conference on Artificial Intelligence*, Vol. 3, pp. 1433–38. Palo Alto, CA: AAAI Press
34. Bobu A, Scobee DRR, Fisac JF, Sastry SS, Dragan AD. 2020. LESS is more: rethinking probabilistic models of human behavior. In *HRI '20: Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 429–37. New York: ACM
35. Zanchettin AM, Rocco P. 2017. Probabilistic inference of human arm reaching target for effective human-robot collaboration. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6595–600. Piscataway, NJ: IEEE
36. Cramer M, Kellens K, Demeester E. 2021. Probabilistic decision model for adaptive task planning in human-robot collaborative assembly based on designer and operator intents. *IEEE Robot. Autom. Lett.* 6(4):7325–32
37. Kurniawati H, Hsu D, Lee WS. 2009. SARSOP: efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems IV*, ed. O Brock, J Trinkle, F Ramos, pp. 65–72. Cambridge, MA: MIT Press
38. Zhao M, Simmons R, Admoni H. 2022. Coordination with humans via strategy matching. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 9116–23. Piscataway, NJ: IEEE
39. Hadfield-Menell D, Dragan A, Abbeel P, Russell S. 2016. Cooperative inverse reinforcement learning. In *Advances in Neural Information Processing Systems 29*, ed. D Lee, M Sugiyama, U Luxburg, I Guyon, R Garnett, pp. 3916–24. Red Hook, NY: Curran
40. Losey DP, Bajcsy A, O'Malley MK, Dragan AD. 2021. Physical interaction as communication: learning robot objectives online from human corrections. *Int. J. Robot. Res.* 41(1):20–44
41. Li M, Canberk A, Losey DP, Sadigh D. 2021. Learning human objectives from sequences of physical corrections. In *2021 IEEE International Conference on Robotics and Automation*, pp. 2877–83. Piscataway, NJ: IEEE
42. Nemlekar H, Dhanaraj N, Guan A, Gupta SK, Nikolaidis S. 2023. Transfer learning of human preferences for proactive robot assistance in assembly tasks. In *HRI '23: Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 575–83. New York: ACM
43. Nikolaidis S, Ramakrishnan R, Gu K, Shah J. 2015. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *HRI '15: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pp. 189–96. New York: ACM
44. Tsitos AC, Dagioglou M, Giannakopoulos T. 2022. Real-time feasibility of a human intention method evaluated through a competitive human-robot reaching game. In *HRI '22: Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 1080–84. Piscataway, NJ: IEEE
45. Cao Z, Hidalgo G, Simon T, Wei SE, Sheikh Y. 2021. OpenPose: realtime multi-person 2D pose estimation using part affinity fields. *IEEE Trans. Pattern Anal. Mach. Intell.* 43(1):172–86
46. Sidiropoulos A, Karayiannidis Y, Doulgeri Z. 2021. Human-robot collaborative object transfer using human motion prediction based on Cartesian pose dynamic movement primitives. In *2021 IEEE International Conference on Robotics and Automation*, pp. 3758–64. Piscataway, NJ: IEEE
47. Paraschos A, Daniel C, Peters J, Neumann G. 2013. Probabilistic movement primitives. In *Neural Information Processing Systems 25*, pp. 2616–24. Red Hook, NY: Curran
48. Ly KT, Poozhiyil M, Pandya H, Neumann G, Kucukyilmaz A. 2021. Intent-aware predictive haptic guidance and its application to shared control teleoperation. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication*, pp. 565–72. Piscataway, NJ: IEEE
49. Hu Z, Xu Y, Lin W, Wang Z, Sun Z. 2022. Augmented pointing gesture estimation for human-robot interaction. In *2022 International Conference on Robotics and Automation*, pp. 6416–22. Piscataway, NJ: IEEE

50. Urkmez M, Bozma HI. 2022. Detecting 3D hand pointing direction from RGB-D data in wide-ranging HRI scenarios. In *HRI '22: Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 441–50. Piscataway, NJ: IEEE
51. Carreira J, Zisserman A. 2017. Quo vadis, action recognition? A new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4724–33. Piscataway, NJ: IEEE
52. Cui Y, Karamcheti S, Palleti R, Shivakumar N, Liang P, Sadigh D. 2023. No, to the right: online language corrections for robotic manipulation via shared autonomy. In *HRI '23: Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 93–101. New York: ACM
53. Lee J, Moray N. 1992. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics* 35(10):1243–70
54. Xu A, Dudek G. 2015. OPTIMo: Online Probabilistic Trust Inference Model for asymmetric human-robot collaborations. In *HRI '15: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pp. 221–28. New York: ACM
55. Zahedi Z, Verma M, Sreedharan S, Kambhampati S. 2023. Trust-aware planning: modeling trust evolution in iterated human-robot interaction. In *HRI '23: Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 281–89. New York: ACM
56. Chen M, Nikolaidis S, Soh H, Hsu D, Srinivasa S. 2020. Trust-aware decision making for human-robot collaboration. *ACM Trans. Human-Robot Interact.* 9(2):9
57. Guo Y, Shi C, Yang XJ. 2021. Reverse psychology in trust-aware human-robot interaction. *IEEE Robot. Autom. Lett.* 6(3):4851–58
58. Azevedo-Sa H, Yang XJ, Robert LP, Tilbury DM. 2021. A unified bi-directional model for natural and artificial trust in human–robot collaboration. *IEEE Robot. Autom. Lett.* 6(3):5913–20
59. Wang Q, Liu D, Carmichael MG, Aldini S, Lin CT. 2022. Computational model of robot trust in human co-worker for physical human-robot collaboration. *IEEE Robot. Autom. Lett.* 7(2):3146–53
60. Carreno-Medrano P, Smith SL, Kulic D. 2023. Joint estimation of expertise and reward preferences from human demonstrations. *IEEE Trans. Robot.* 39(1):681–98
61. Liu R, Natarajan M, Gombolay MC. 2021. Coordinating human-robot teams with dynamic and stochastic task proficiencies. *ACM Trans. Human-Robot Interact.* 11(1):5
62. Kolb J, Kishore M, Shaw K, Ravichandar H, Chernova S. 2021. Predicting individual human performance in human-robot teaming. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication*, pp. 45–50. Piscataway, NJ: IEEE
63. Nanavati A, Mavrogiannis C, Weatherwax K, Takayama L, Cakmak M, Srinivasa S. 2021. Modeling human helpfulness with individual and contextual factors for robot planning. In *Robotics: Science and Systems XVII*, ed. D Shell, M Toussaint, MA Hsieh, pap. 16. N.p.: Robot. Sci. Syst. Found.
64. Messeri C, Bicchi A, Zanchettin AM, Rocco P. 2022. A dynamic task allocation strategy to mitigate the human physical fatigue in collaborative robotics. *IEEE Robot. Autom. Lett.* 7(2):2178–85
65. Lagomarsino M, Lorenzini M, Momi ED, Ajoudani A. 2022. Robot trajectory adaptation to optimise the trade-off between human cognitive ergonomics and workplace productivity in collaborative tasks. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 663–69. Piscataway, NJ: IEEE
66. Kalatzis A, Hopko S, Mehta RK, Stanley L, Wittie MP. 2022. Sex parity in cognitive fatigue model development for effective human-robot collaboration. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 10951–58. Piscataway, NJ: IEEE
67. Mohamed Y, Ballardini G, Parreira MT, Lemaignan S, Leite I. 2022. Automatic frustration detection using thermal imaging. In *HRI '22: Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 451–59. Piscataway, NJ: IEEE
68. Nguyen A, Xie B. 2021. Human arm motion prediction in reaching movements. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication*, pp. 1117–23. Piscataway, NJ: IEEE
69. Erickson Z, Clever HM, Gangaram V, Xing E, Turk G, et al. 2023. Characterizing multidimensional capacitive servoing for physical human–robot interaction. *IEEE Trans. Robot.* 39(1):357–72
70. Ondras J, Anwar A, Wu T, Bu F, Jung M, et al. 2022. Human-robot commensality: bite timing prediction for robot-assisted feeding in groups. In *Proceedings of the 6th Conference on Robot Learning*, ed. K Liu, D Kulic, J Ichnowski, pp. 921–33. Proc. Mach. Learn. Res. 205. N.p.: PMLR

71. Wan W, Yang L, Liu L, Zhang Z, Jia R, et al. 2022. Learn to predict how humans manipulate large-sized objects from interactive motions. *IEEE Robot. Autom. Lett.* 7(2):4702–9
72. Alahi A, Goel K, Ramanathan V, Robicquet A, Fei-Fei L, Savarese S. 2016. Social LSTM: human trajectory prediction in crowded spaces. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 961–71. Piscataway, NJ: IEEE
73. Chang P, Luo R, Dorostian M, Padr T. 2021. A shared control method for collaborative human-robot plug task. *IEEE Robot. Autom. Lett.* 6(4):7429–36
74. Yasar MS, Iqbal T. 2021. A scalable approach to predict multi-agent motion for human-robot collaboration. *IEEE Robot. Autom. Lett.* 6(2):1686–93
75. Li S, Figueroa N, Shah AJ, Shah JA. 2021. Provably safe and efficient motion planning with uncertain human dynamics. In *Robotics: Science and Systems XVII*, ed. D Shell, M Toussaint, MA Hsieh, pap. 50. N.p.: Robot. Sci. Syst. Found.
76. Jin Z, Liu A, Zhang WA, Yu L, Su CY. 2023. A learning based hierarchical control framework for human–robot collaboration. *IEEE Trans. Autom. Sci. Eng.* 20(1):506–17
77. Vianello L, Mouret JB, Dalin E, Aubry A, Ivaldi S. 2021. Human posture prediction during physical human-robot interaction. *IEEE Robot. Autom. Lett.* 6(3):6046–53
78. Tejwani R, Kuo YL, Shu T, Stankovits B, Gutfreund D, et al. 2022. Incorporating rich social interactions into MDPs. In *2022 International Conference on Robotics and Automation*, pp. 7395–401. Piscataway, NJ: IEEE
79. Gershman SJ, Daw ND. 2017. Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annu. Rev. Psychol.* 68:101–28
80. Nikolaidis S, Kuznetsov A, Hsu D, Srinivasa S. 2016. Formalizing human-robot mutual adaptation: a bounded memory model. In *HRI '16: The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, pp. 75–82. Piscataway, NJ: IEEE
81. Parekh S, Habibian S, Losey DP. 2022. RILI: robustly influencing latent intent. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2135–42. Piscataway, NJ: IEEE
82. Tian R, Tomizuka M, Dragan AD, Bajcsy A. 2023. Towards modeling and influencing the dynamics of human learning. In *HRI '23: Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 350–58. New York: ACM
83. Cohen PR, Levesque HJ. 1990. Intention is choice with commitment. *Artif. Intell.* 42(2–3):213–61
84. Rao AS, Georgeff ML. 1995. BDI agents: from theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pp. 312–19. Palo Alto, CA: AAAI Press
85. Bratman M. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard Univ. Press
86. Georgeff M, Pell B, Pollack M, Tambe M, Wooldridge M. 1999. The belief-desire-intention model of agency. In *Intelligent Agents V: Agents Theories, Architectures, and Languages*, ed. JP Müller, AS Rao, MP Singh, pp. 1–10. Berlin: Springer
87. Patel M, Chernova S. 2023. Proactive robot assistance via spatio-temporal object modeling. In *Proceedings of the 6th Conference on Robot Learning*, ed. K Liu, D Kulic, J Ichnowski, pp. 881–91. Proc. Mach. Learn. Res. 205. N.p.: PMLR