Influencing Behavioral Attributions to Robot Motion During Task Execution

Nick Walker, Christoforos Mavrogiannis, Siddhartha Srinivasa, Maya Cakmak

***Abstract*— Recent literature has proposed algorithms for autonomous generation of robot motion that communicates functional attributes of a robot's state such as intent or incapability. However, less is known about how to automate the generation of motion for communicating higher-level behavioral attributes such as curiosity or competence. In this work, we consider a coverage task resembling robot vacuum cleaning in a household. Through a virtual interface, we collect a dataset of human attributions to robot trajectories during task execution and extract a probabilistic model that maps robot trajectories to human attributions. We then incorporate this model into an trajectory generation mechanism that balances between task completion and communication of a desired behavioral attribution. Through an online user study on a different household layout, we find that our prediction model accurately captures human attribution for coverage tasks. Further, our generation mechanism produces trajectories that are thematically consistent, but more research is required to understand how to balance attribution and task performance.**

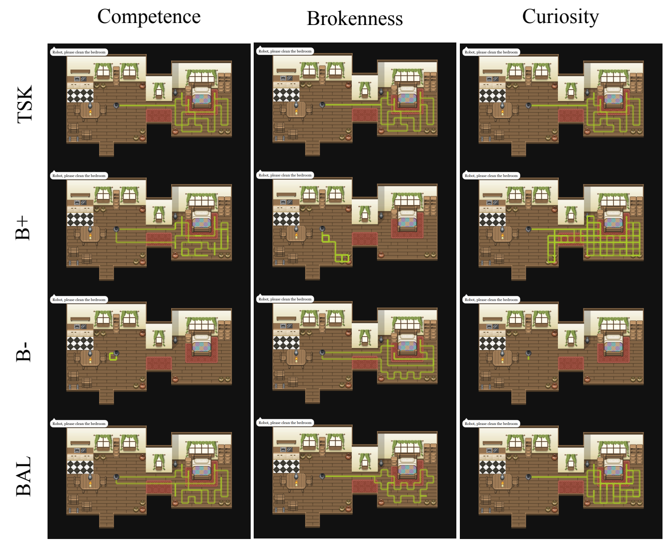

Fig. 1: Traces of trajectories in our virtual household coverage domain. Labels describe their usage in Experiment I.

I. Introduction

As robots enter households and public spaces, it is increasingly important to account for human perceptions of their behavior [26, 25, 17, 6]. A common setting involves a robot performing a task while a human bystander observes a portion of its behavior. While the robot's actions might be driven by unambiguous internal objectives, solely optimizing such criteria might result in robot behavior that is difficult to interpret or disruptive to surrounding humans. For example, a highly articulated robot may follow a non-humanlike trajectory that makes observers uncomfortable [29] or a home robot such as a robot vacuum cleaner can exhibit unpredictable motion that interrupts home activity.

Accounting for high-level attributions to robot behavior is a complex problem relying on the mechanisms underlying human inference and behavior generation. In psychology, there is a long history of work on understanding human attribution for behavior explanation or inference of behavior traits [10, 20]. Humans are highly attuned to how their actions are perceived and adapt their behavior to elicit a desired impression from others or adhere to social norms, a concept that is known as *presentation of self* [8]. The tendency for humans to attribute even situational behaviors to deeper character traits is so pervasive that it is known as "the fundamental attribution error" [23].

Inspired by these theories, we envision that robots in human environments can leverage an understanding of humans' attribution mechanisms to generate behaviors that elicit desired human impressions. This will enable robots to more seamlessly integrate in human spaces and increase their acceptance. In this work, we use a virtual interface to collect a dataset of human attributions to a vacuum cleaning robot's coverage trajectories. Based on this dataset, we extract a set of dominant attribution dimensions—`competence`, `brokenness` and `curiosity`—and a low-dimensional trajectory representation that captures salient features of the robot's motion. We learn a probabilistic model that maps a robot trajectory to the expected distribution of human ratings. We then incorporate this model into a trajectory generation framework that balances between task-related and attribution elicitation objectives. We conduct an online user study to verify that our generation framework produces trajectories that are thematically consistent both within the same domain and also within a domain variant. Our experiments highlight the need for further work to understand how to balance attribution and task-related specifications.

Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, 98195, USA. email: {nswalker,cmavro,siddh,mcakmak}@cs.uw.edu

II. Related Work

Several studies have illustrated the value of robot motion as a communicative modality [7, 15, 5, 18, 22, 11]. Some works propose algorithms for legible robot motion generation, which have been shown to enable effective human-robot collaboration in manipulation tasks [7], or smooth robot navigation in close proximity to humans [5, 22].

Other works focus on conveying higher-level information such as the robot's objective function [11] or the source of failure [18] in failure cases. Animation principles [27] or movement analysis [16] are often employed to inform the design of expressive robot behaviors. Finally, related graphics research focuses on the generation of stylistically distinct but functionally equivalent motion primitives for walking and other activities [3, 9].

The complex interplay of embodiment and communicative motion has motivated research on understanding human perceptions of robot behavior. For instance, early work looked at the effect of robot gaze on human impressions [13]. Sung et al. [26] study human attitudes towards robot vacuum cleaners and propose design principles aimed at enhancing the acceptance of robots in domestic environments. [25] report a relation between robot motion and perceived affect. Lo et al. [19] and [22] investigate human perceptions of different robot navigation strategies whereas Walker et al. [28] study human perceptions of robot actions that deviate from the robot's assigned task.

Our work draws inspiration from recent work on design and methods for extracting human perceptions and attributions for robot motion [26, 26, 25]. However, it moves beyond the problem of understanding and analyzing human perceptions, and focuses on the problem of *synthesizing* implicitly communicative motion. Our work is closely related to past work on the generation of legible robot motion [7, 15, 21, 18] in that we also incorporate a model of human inference into the robot's motion generation pipeline. However, unlike these works which emphasize the communication of task-related attributes, our focus is instead on communicating high-level, behavioral attributes through robot motion.

III. A Framework for Behavioral Attribution

We consider a robot performing a task $\mathcal{T}$ in a human environment. We denote by $s \in \mathcal{S}$ the robot state where $\mathcal{S}$ is a state space and define a robot trajectory as a sequence of states $\xi = (s_0, \ldots, s_t)$ where indices correspond to timesteps following a fixed time parametrization of step size dt. Let us define the task as a tuple $\mathcal{T} = (\Xi, \mathcal{A}, \mathcal{P}, C)$ where Ξ is a space of robot trajectories, $\mathcal{A}$ denotes the robot action space, $\mathcal{P} : \Xi \times \mathcal{A} \rightarrow \Xi$ represents a deterministic state transition model, and $\mathcal{C} : \Xi \rightarrow \mathbb{R}$ is a trajectory cost.

We assume that the robot starts from an initial state s_0 and reaches a terminal state s_T at the end of the task execution (time T) by executing a trajectory $\xi = (s_0, \ldots, s_T)$. We assume that this trajectory ξ is fully observable by a human observer who is aware of the task specification $\mathcal{T}$. The observer makes an inference of the form $\mathcal{I}_B : \Xi \times \mathcal{T} \rightarrow \mathcal{B}$, where Ξ is a space of trajectories and $\mathcal{B}$ is a space of behavioral attributions. The form of $\mathcal{B}$ will vary, but should be selected to capture the range, combinations, and intensities of attributions that the robot should be sensitive to.

Conversely, given a behavioral attribution $\boldsymbol{b}$ from a space of behavioral attributions $\mathcal{B}$ and a task $\mathcal{T}$, the observer expects to see a trajectory $\xi_b \in \Xi$, corresponding to an inference of the form $\mathcal{I}_\xi : \mathcal{B} \times \mathcal{T} \rightarrow \Xi$. In other words, we assume that there is a "way" that a curious—or any other attribution—robot would be expected to execute a particular task. Equivalently, we assume that there is a set of trajectories that the observer could more often imagine to be consistent with some attribution.

In this paper, we aim to provide a general framework for modeling inferences of the form $\mathcal{I}_B$, and $\mathcal{I}_\xi$. Our goal is to enable robots to understand and account for the communicative effects of their motion on human observers.

IV. Modeling and Influencing Attributions

We consider a scenario in which a mobile robot performs a coverage task in a two-dimensional discrete workspace while a human is observing from a top-down view. We employ a virtual environment[1] that resembles a house and stylize the agent as a robot vacuum cleaner (see Fig.1) since the general population is already somewhat familiar[2] with such robots [26, 12], making it easier for participants to develop mental models about their motion than, for example, that of a manipulator.

In this scenario, the robot state space is the complete home workspace and Ξ is the space of all possible trajectories of any length that could be followed in the space. The robot action space $\mathcal{A}$ consists of the cardinal directions and actions are deterministic. The cost of a state transition from a state s_t to a state s_{t+1} after having followed a trajectory ξ_t is defined as 0 if the state hasn't been visited before, -5 if the state is a small traversable obstacle (e.g. a vase), and -1 otherwise. A penalty proportional to the number of unvisited goal states is applied on termination.

A. Understanding Behavioral Attribution for Coverage Tasks

Through exploratory studies on Amazon Mechanical Turk, we sought to extract domain knowledge for attributions to robot motion within coverage tasks. Using the home layout of Fig.1, we generated a set of trajectories exhibiting qualitatively distinct ways the robot could respond to the prompt to "clean the bedroom," ranging from a near optimal coverage plan to a trajectory that barely visited the target room. Each participant viewed videos of a random selection of three of these trajectories. After each video, participants were asked: a) to provide three words to describe the robot's behavior; b) to rate their agreement that "the robot is ______" for a range of adjectives drawn from relevant literature on human attributions [1, 4, 28]; c) to "explain what factors contributed to their strongest ratings." In addition to attributions, participants were asked to use an interactive interface to demonstrate how they would "clean the bedroom in a way that makes the robot look ______" where the blank was filled at random with an adjective from the attribution rating items. Across all exploratory studies, we collected 375 sets of attribution ratings from 115 participants (73 male, 41 female) aged 21-70 (M = 38.3, SD = 10.7) covering 63 trajectories and a total of 193 demonstrations.

[1]The environment is built in the Phaser game engine (https://phaser.io/) and uses art by Bonsaiheldin under a CC-BY-SA license.

[2]A recent presentation by iRobot [12] reports that about 14M households had robot vacuum cleaners in the U.S. in 2019.

TABLE I: Low-dimensional trajectory representation.

Feature	Description
Coverage (%)	Goal region states visited at least once.
Redundant coverage (%)	Goal region states visited more than once.
Overlap (%)	Plan states visited more than once.
Length (%)	Normalized plan length[3].
Hook template (%)	Frequency of "U" shape patterns in plan.
Straight template (%)	Frequency of action repetition in plan.
Start-stop template (%)	Frequency of idle-move-idle patterns in plan.
Idleness (%)	Frequency of idle actions in plan.
Map coverage (%)	Ratio of map states visited at least once.
Collision (%)	Ratio of obstacle states from O in plan.
Goal deviation (%)	Ratio of of plan before first goal state.

1) Extracting the Space of Attributions: To understand the inter-correlation of participant adjective ratings, we conducted an exploratory factor analysis. We selected a three-factor model (promax rotation) which explained 74% of the observed variance due to its parsimony and coherence. The first factor, which we call "competence" for its similarity to the relevant factor described by Carpinella et al. [4] consists of six items (responsible, competent, efficient, reliable, intelligent, focused) centered on the capability and diligence of the robot. The second consists of four items (lost, clumsy, confused, broken) alluding to a temporary or extended negative state, for which we title the factor "brokeness". The third contains two items (curious, investigative) and matches the curiosity factor examined by Walker et al. [28].

The extracted model enables the computation of standardized factor scores roughly in the range $[-3, 3]$ which summarize how much a participant's ratings for the items deviate from the mean along each factor. Reflecting the format of the component items, a high or low factor score denotes agreement or disagreement that a trajectory expresses an attribution, respectively. Based on this model, we represent the attribution for a trajectory ξ as a tuple $\boldsymbol{b} = (b_{competent}, b_{broken}, b_{curious}) \in \mathcal{B}$ where the space of attributions is the set $\mathcal{B} = [-3, 3]^3$.

2) Low-dimensional Trajectory Representation: The space of possible trajectories in this domain is too large to explore directly, so we constructed a low-dimensional space Φ based on features relevant to the formation of attribution ratings. This allows us to describe a trajectory ξ as a vector $\phi_\xi = \phi(\xi) \in \Phi$. The feature space was inspired by relevant literature on human behavioral attribution to robot motion and enriched with features appearing in participants' explanations. The final set of 11 features used in further experiments is listed in Table I.

B. Mapping Trajectories to Attribution Scores

Given a trajectory ξ, an observer's inference $\mathcal{I}_B$ of behavioral attribution can be expected to vary from individual to individual and as a result of measurement error. For this reason, we model $\mathcal{I}_B$ as a conditional probability density $f_{\mathcal{B}|\Xi}(\boldsymbol{b}|\phi_\xi) : \mathcal{B} \to \mathbb{R}$. We observed multimodality in the distribution of factor scores for some trajectories, so we use a Mixture Density Network (MDN) [2] to approximate each conditional density as a mixture distribution $f_{\mathcal{B}|\Xi}(\boldsymbol{b}|\phi_\xi) = \sum_{i=1}^{C} \alpha_i(\phi_\xi) k_i(\boldsymbol{b}|\phi_\xi)$ where α_i, $i = 1, \ldots, C$, is a mixing coefficient, and k_i is a multivariate Gaussian kernel function with mean $\boldsymbol{\mu}_i$ and covariance Σ_i. Note that the mixing coefficients α_i and the Gaussian parameters $\boldsymbol{\mu}_i$ and Σ_i are functions of the featurized trajectory ϕ_ξ. In our models, these functions are implemented as linear transformations of features produced by a shared multi-layer perceptron.

[3]Scaled so that 1 corresponds to three times the size of V.

TABLE II: Test performance of models.

Model	Parameters	Average NLL	SD
Uniform	-	5.38	0.00
MDN, C=1	120	$3.13 \pm .05$	$1.35 \pm .09$
MDN, C=4	300	$2.66 \pm .08$	$1.57 \pm .05$
MDN Ensemble, C=4 N=8	2400	$2.53 \pm .06$	$1.38 \pm .04$

To make efficient use of scarce data, we created ensembles of MDNs using bootstrap aggregation, i.e., we trained N models with different data splits and uniformly weight their predictions: $f^{ens}_{\mathcal{B}|\Xi}(\boldsymbol{b}|\phi_\xi) = \frac{1}{N}\sum_{i=1}^{N} f^{i}_{\mathcal{B}|\Xi}(\boldsymbol{b}|\phi_\xi)$.

We studied three different model configurations; single and four component MDNs, i.e., $C = 1$ and $C = 4$ and an ensemble of 8 MDNs each with four components, i.e., $C = 4$, $N = 8$. We trained all models using an average negative log likelihood (NLL) loss function, the Adam optimizer [14], noise regularization [24], and early stopping. We configured the input MLP to use a single hidden layer with 5 units and a hyperbolic tangent activation. The dataset used was an expanded version of the set collected in our exploratory studies containing 126 trajectories with 671 attribution ratings. Table II compares the NLL of the models over held-out data. The mean indicates the typical quality of the prediction and the standard deviation indicates the degree to which this varied from sample to sample. Both quantities are averaged over 16 random folds and reported with bootstrapped 95% confidence intervals. All models compare favorably to a uniform baseline, which simply assigns equal probability to all outcomes. The ensemble model performs best and is used in further experiments in the remainder of the paper.

C. Generating Trajectories that Elicit Desired Attributions

We represent the behavior specification as a one-dimensional Gaussian $b^* \sim \mathcal{N}(\mu_b, \sigma_b^2)$ centered on a desired rating $\mu_b \in [-3, 3]$ for a single attribution dimension where σ_b is a variance representing a tolerance parameter modeling the acceptable distance from the desired behavioral rating. We use a density representation as it more closely matches the output of our model for $\mathcal{I}_B$.

Together with the task requirements as described by the cost function $\mathcal{C}$, we realize the inference $\mathcal{I}_\xi$ as an optimization scheme of the form:

$$\xi^* = \arg\min_{\xi \in \Xi} \Big(\mathcal{C}(\xi) + w D_{\mathrm{KL}}(f_{\boldsymbol{B}_i} || \mathcal{N}(\mu_b, \sigma_b^2)) \Big), \qquad (1)$$

where D_{KL} denotes the KL divergence, $f_{\boldsymbol{B}_i}$ is the density $f_{\boldsymbol{B}|\Xi}|\phi(\xi)$ marginalized across dimensions other than i, and w is a weight representing the balance between the task and

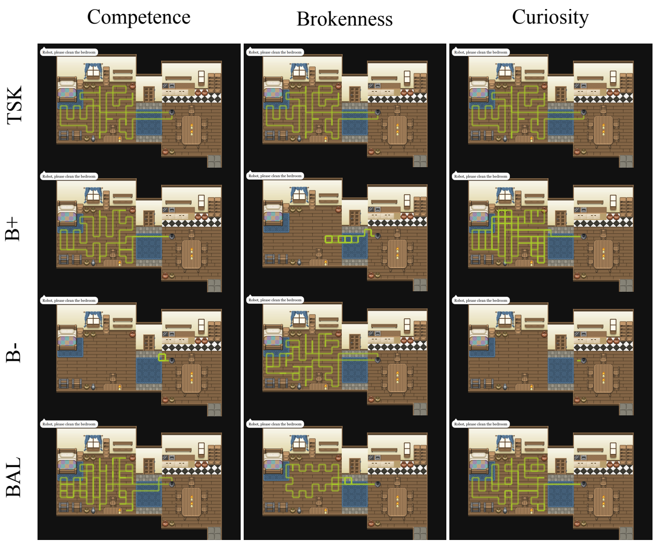

Fig. 2: Traces of the trajectories used in Experiment II.

attribution specifications. We implement this optimization using a hill-descending search in the space of trajectories.

V. Evaluation

We conduct a user study to evaluate the efficacy of the framework as a means of producing trajectories that elicit desired attributions. In Experiment I, participants observe and rate trajectories in the same home layout used for data collection, while in Experiment II, trajectories are generated in a new, arbitrarily modified home layout. In both experiments, we systematically control the trajectories selected to span a range of attributions.

a) Experiment Design: Our experiments are within-subjects, video-based user studies, both instantiated in three parallel sets corresponding to the three attribution dimensions considered. For each dimension, we consider four distinctly produced robot trajectories: the task-optimal trajectory (**TSK**), the trajectories optimized to alternately maximize and minimize a particular attribution using a goal of $\mathcal{N}(\pm 1, 0.3)$ (**B+** and **B-**), and a trajectory that optimizes for a balance between the task cost and the attribution goal following eq. (1) (**BAL**). The full set of trajectories is shown in Fig.1 and Fig.2. In all experiments, participants rate and describe each trajectory using the same items and questions used in the exploratory studies of Sec IV. After watching all trajectories in a randomly assigned order, they also respond to additional comparative questions: "which robot seemed the most ______" and "which robot seemed the least ______", where the blanks are filled with the adjective corresponding to the dimension of attribution studied. Both comparisons are accompanied with an open-ended question asking for a brief explanation of the choice.

b) Participants: A total of 144 participants (75 male, 69 female) aged 20-72 (M = 37.0, SD = 12.0) were recruited via Amazon Mechanical Turk. 16 had taken part in our earlier exploratory studies. Participants were equally distributed amongst the six total sets of conditions. Condition orderings were fully counterbalanced.

c) Results: Average negative log likelihoods were calculated as 2.51 (SD = 1.17) for Experiment I and 2.93 (SD = 1.40) for Experiment II. B+ was selected as the "most ______" trajectory for brokenness and curiosity by a majority of participants across both experiments. "Most" selection for competence was mixed with strong modes on B+ and TSK. B- was selected as the "least ______" trajectory for competence and curiosity by a majority of participants across both experiments. "Least" selection for brokenness was mixed between B- and TSK.

VI. Discussion

The average NLL of the attributions observed across Experiment I was lower than a uniform model, indicating that the model was able to meaningfully predict attributions in the layout it was trained in. The same held in Experiment II where the environment layout was perturbed. The variance indicates that the model's performance is not even across conditions, but the variation is not significantly different than that seen in the model selection experiments (see Table II).

When optimizing for curiosity, the model emphasized over-coverage of the goal region and visiting penalized states depicted with vases (*Exp. I&II-Curiosity B+*), leading to mixed results, with a strong effect in Experiment I (Fig.1) and little effect in Experiment II (Fig.2). Some participants highlight the extra coverage as a kind of exploration; "It seemed to check the same places multiple times as if it was discovering". Others attribute curiosity to the task-optimal trajectory for avoiding repetition: "because it visited all the areas it hadn't instead of going back over the same spots."

Trajectories optimized to look "not broken" tend to over-cover the goal region (*Exp.I&II-Brokenness B-*), but are seen to be indistinguishable from the task-optimal trajectories (see Fig.1: *Exp.I&II TSK*). The model similarly emphasizes redundant goal coverage to an extreme when attempting to appear competent (*Exp. I&II-Competence B+*), but this is not successful in either of the experiments. While some participants appreciated the extra coverage, noting that it "may take longer but would clean the best," many felt it went too far and preferred the tempered BAL condition saying it "did the best job of hitting all the areas without a lot of backtracking and confusion."

Only the "brokeness" factor was demonstrated to be tunable, with BAL falling below "B+" and above "B-" in both experiments, likely because the self-overlap and length features that are relied upon have a linear relation with task cost. Balanced trajectories for curiosity retain their tendency to over-cover the goal, but lose the propensity to make contact with the vases due to the high associated task-cost penalty.

We are encouraged by the predictive performance of the model. Our future work includes exploring improved methods for leveraging the model in generation and expanding to more realistic domains.

Acknowledgments

This work was (partially) funded by the National Science Foundation IIS (#2007011), National Science Foundation DMS (#1839371), the Office of Naval Research, US Army Research Laboratory CCDC, Amazon, and Honda Research Institute USA. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant #1762114.

References

[1] C. Bartneck, E. Croft, and D. Kulic. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1): 71–81, 2009.

[2] C. Bishop. Mixture density networks. Technical Report NCRG/94/004, Aston University, 1994.

[3] M. Brand and A. Hertzmann. Style machines. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, page 183–192. ACM Press/Addison-Wesley Publishing Co., 2000.

[4] C. M. Carpinella, A. B. Wyman, M. A. Perez, and S. J. Stroessner. The robotic social attributes scale (rosas): Development and validation. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI*, pages 254–262, 2017.

[5] D. Carton, W. Olszowy, and D. Wollherr. Measuring the effectiveness of readability for mobile robot locomotion. *International Journal of Social Robotics*, 8(5): 721–741, 2016.

[6] E. Cha, A. D. Dragan, and S. S. Srinivasa. Perceived robot capability. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 541–548, 2015.

[7] A. Dragan and S. Srinivasa. Integrating human observer inferences into robot motion planning. *Autonomous Robots*, 37(4):351–368, 2014.

[8] E. Goffman. *The Presentation of Self in Everyday Life*. Anchor, 1959.

[9] S. J. Guy, S. Kim, M. C. Lin, and D. Manocha. Simulating heterogeneous crowd behaviors using personality trait theory. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 43–52, 2011.

[10] F. Heider. *The psychology of interpersonal relations*. John Wiley & Sons Inc., 1958.

[11] S. H. Huang, D. Held, P. Abbeel, and A. D. Dragan. Enabling robots to communicate their objectives. *Autonomous Robots*, 43(2):309–326, 2019.

[12] iRobot. *Investor Presentation - August 2019*. August 2019. URL https://investor.irobot.com/static-files/cd40f151-aede-4047-a24f-7e14d317723d. Accessed: 2021-03-01.

[13] T. Kanda, H. Ishiguro, and T. Ishida. Psychological analysis on human-robot interaction. In *Proceedings IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 4166–4173 vol.4, 2001.

[14] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.

[15] R. A. Knepper, C. Mavrogiannis, J. Proft, and C. Liang. Implicit communication in a joint action. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 283–292, 2017.

[16] H. Knight and R. Simmons. Expressive motion with x, y and theta: Laban effort features for mobile robots. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 267–273, 2014.

[17] M. Kwon, M. F. Jung, and R. A. Knepper. Human expectations of social robots. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 463–464, 2016.

[18] M. Kwon, S. H. Huang, and A. D. Dragan. Expressing robot incapability. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 87–95, 2018.

[19] S. Lo, K. Yamane, and K. Sugiyama. Perception of pedestrian avoidance strategies of a self-balancing mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1243–1250, 2019.

[20] B. F. Malle. Attribution theories: How people make sense of behavior. In D. Chadee, editor, *Theories in social psychology*, page 72–95. Wiley Blackwell, 2011.

[21] C. Mavrogiannis, W. B. Thomason, and R. A. Knepper. Social momentum: A framework for legible navigation in dynamic multi-agent environments. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 361–369, 2018.

[22] C. Mavrogiannis, A. M. Hutchinson, J. Macdonald, P. Alves-Oliveira, and R. A. Knepper. Effects of distinct robotic navigation strategies on human behavior in a crowded environment. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 421–430, 2019.

[23] L. Ross. The intuitive psychologist and his shortcomings: Distortions in the attribution process. volume 10 of *Advances in Experimental Social Psychology*, pages 173–220. Academic Press, 1977.

[24] J. Rothfuss, F. Ferreira, S. Boehm, S. Walther, M. Ulrich, T. Asfour, and A. Krause. Noise regularization for conditional density estimation. *arXiv:1907.08982*, 2019.

[25] M. Saerbeck and C. Bartneck. Perception of affect elicited by robot motion. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 53–60, 2010.

[26] J.-Y. Sung, L. Guo, R. E. Grinter, and H. I. Christensen. "my roomba is rambo": Intimate home appliances. In J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang, editors, *UbiComp 2007: Ubiquitous Computing*, pages

145–162. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[27] L. Takayama, D. Dooley, and W. Ju. Expressing thought: Improving robot readability with animation principles. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 69–76, 2011.

[28] N. Walker, K. Weatherwax, J. Alchin, L. Takayama, and M. Cakmak. Human perceptions of a curious robot that performs off-task actions. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 529–538, 2020.

[29] F. Zacharias, C. Schlette, F. Schmidt, C. Borst, J. Rossmann, and G. Hirzinger. Making planned paths look more human-like in humanoid robot manipulation planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1192–1198, 2011.