Leveraging Memory and Attention in a Kinematically Aware Robot: An Ideomotor-Inspired Approach to Implicit Command Understanding from IMU Sensor Data*

Lena Guinot¹, Ryutaro Matsumoto¹, Shota Takahashi¹ and Hiroyasu Iwata², Member, IEEE

Abstract-As human-robot interaction (HRI) advances, the nuanced interpretation of implicit commands embedded in human gestures becomes paramount for fostering seamless collaboration. In this context, we present a novel machine learning algorithm designed to endow robots with the ability to decipher implicit commands from Inertial Measurement Unit (IMU) sensor data worn at specific locations on the human body. Our approach integrates memory and attention mechanisms inspired by ideomotor cues, allowing the robot to comprehend both temporal and spatial relationships within the sensor data. The attention mechanism operates bidirectionally, enhancing the system's awareness of the temporal sequence of human movements and the spatial interdependencies between sensor data across different body locations. This unique spatial attention enables the robot to understand the kinematic chain between joints during human motion, accommodating variations in sensor data arising from factors such as height differences and motion range capacity. Drawing on prior research in attention mechanisms, ideomotor cues, and memory augmentation, our algorithm represents a significant advancement in addressing the challenges of implicit command understanding in HRI. The proposed system's adaptability and nuanced comprehension of human gestures make it wellsuited for diverse anatomies and movement patterns. Through comprehensive experiments, we demonstrate the effectiveness of our algorithm, paving the way for more intuitive and adaptable robotic systems in real-world applications.

I. INTRODUCTION

Integrating robots into human environments is a complex challenge spanning technical, social, and cognitive domains. This paper introduces a machine learning algorithm to improve robots' interpretive abilities, focusing on implicit command recognition and user-specific adaptation. Traditional Human-Robot Interaction (HRI) systems rely on explicit commands, lacking the flexibility of human-human interactions. To address this, we leverage the ideomotor principle, linking thought processes with motor actions. Our system deduces user intentions from natural, non-verbal cues, streamlining interactions. We use Inertial Measurement Units (IMUs) to capture human motion nuances, processed through a dual attention mechanism considering temporal and spatial

Our algorithm adapts to individual users through adaptive machine learning and a feedback loop, refining predictive

sensor data variations due to user-specific factors.

machine learning and a feedback loop, refining predictive capabilities for a tailored user experience. This enhances implicit command interpretation and enables personalized HRI experiences. This research has implications for robotics in various settings, from industrial collaboration to assistive technologies. By aligning human intentions with robot actions, we move towards collaborative partnerships enhancing human capabilities and experiences.

dimensions. This allows accurate interpretation even with

II. BACKGROUND

A. Robot Control in Human-Robot Interaction

The domain of HRI has been significantly advanced by the integration of user motion as a medium for robot control. The concept of operating robots through the physical movements of a user has its roots in early teleoperation systems and has evolved with the advent of sophisticated sensors and machine learning algorithms. Pioneering work in this field was conducted by Sheridan, who explored the foundational principles of teleoperation and the potential of human motion to guide robotic systems, laying the groundwork for subsequent research in the field [1], [2]. The utility of IMUs for capturing user motion was further highlighted by Roetenberg et al. [3], who demonstrated the effectiveness of IMUs in fullbody motion tracking, thus opening new avenues for userdriven robot control. Significant advancements were made by researchers such as Fang et al., who developed a system for capturing complex human gestures through a combination of IMUs and computer vision, thus enabling more nuanced interactions between humans and robots [4]-[6]. This work underscored the importance of precise motion capture in translating user intent into robotic action. The concept of "learning from demonstration" (LfD) allows robots to learn control policies directly from human demonstrations, using user motion as the primary input. This paradigm shift facilitated a more intuitive and natural means of robot control [7]–[9].

B. Attention Mechanisms in Machine Learning

Research by Vaswani et al. on attention models in neural networks has inspired subsequent work on attention-based machine learning systems [10], especially in the field of sequence prediction tasks. Bahdanau et al. [11] introduced the concept in the context of neural machine translation,

^{*}This work was supported by JST SPRING, Grant Number JPMJSP2128, Waseda University Global Robot Academia Institute and Waseda University Green Computing Systems Research Organization

¹Lena Guinot, Ryutaro Matsumoto and Takahashi Shota are with the Graduate School of Creative Science and Engineering, Waseda University, Tokyo, Japan guinotl@moegi.waseda.jp

²Hiroyasu Iwata is with the Faculty of Science and Engineering, Waseda University, Tokyo, Japan, jubi@waseda.jp

which has since been adapted to various domains. For instance, Xu et al. [12] successfully applied attention models to image captioning, demonstrating their effectiveness in focusing on relevant features of the input data. More recently, attention has been integrated into reinforcement learning and HRI frameworks to enhance the interaction between agents and their environments. For example, Sorokin et al. [13] introduced deep attention recurrent Q-networks, which use attention-driven memory to improve decision-making in video games.

C. Implicit Command Interpretation

The interpretation of user intention through implicit commands, derived from motion data in line with the ideomotor principle, is a growing area of research in HRI. Wearable sensors now capture subtle human movements that indicate underlying intentions, building on early research by Carpenter, Lotz, James, and Pfungst [14]–[17]. They explored the ideomotor principle, linking thoughts with actions initiated without conscious awareness, laying a psychological foundation for interpreting user intentions from motion data. This principle led to systems inferring intent from physical cues, crucial for seamless HRI but challenging. Studies show ideomotor cues' potential in interface design [18]–[21].

In robotics, understanding gestures, gaze direction, and other non-verbal cues is key. Fiore et al. [22] found that robots interpreting such cues engage more meaningfully. Machine learning, especially deep learning like Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, models these cues, often in time-series data form. The challenge is identifying informative cues and their correlation with user intentions.

Personalizing models for user-specific characteristics is crucial. Ziebart et al. [23] explored personalized models for behavior and preferences, enhancing robots' ability to understand and predict user intentions. As HRI progresses, user specificity will likely increase, creating more personalized and adaptable robots for diverse users and environments. The goal is HRI systems as diverse as their users, delivering personalized experiences that improve collaboration.

This study integrates IMUs, attention mechanisms, and adaptive learning to interpret both explicit and implicit user commands. Including user specificity and ideomotor principles enhances the system's ability to understand and predict human intentions, a significant contribution to HRI.

III. Network

A. Concept Outline

In the proposed methodology, outlined in Fig. 1 the system's learning pipeline is initiated by accumulating motion data during user-executed tasks. Simple task execution is accompanied by corresponding vocal commands, whereas complex tasks incorporate a wider range of user motions and command phrases. This data, enriched with the robot's motor encoder readings, forms a comprehensive dataset.

A machine learning model is trained with this dataset to discern user intent from the motion and command data. The



Fig. 1: Training and testing flow of the designed model



Fig. 2: Outline of the network used

model's interpretative accuracy is subsequently assessed in practical scenarios, where correct intention predictions are extracted and reincorporated into the training set, enhancing the dataset.

This iterative refinement process, illustrated in the same figure, cultivates the model's proficiency in interpreting user intentions, thereby improving the robot's interactive performance through a continuous learning loop.

B. Network

In our network architecture (Fig. 2), we use Convolutional Neural Networks (CNNs) and Bidirectional Long Short-Term Memory (Bi-LSTM) networks to process motion data. Each CNN has five layers for feature extraction from Inertial Measurement Units (IMUs) and motor encoders. IMUs are placed strategically on the user's body to capture motion data, while motor encoders provide angle information from the robot arm's axes. After CNN processing, data goes through two Bi-LSTM layers to recognize temporal dependencies. This data then goes through a stacked dual attention mechanism, focusing on relevant features for intention decoding. The output goes through a softmax function and argmax operation to generate the final command label. The input data is 44dimensional, transformed through the network to produce a discrete command label for the robot system.

C. Dual Attention

Central to our approach in decoding implicit human commands through IMU sensor data is our innovative pose reconstruction and motion prediction strategy, underpinned by a sophisticated dual attention mechanism outlined in Fig.3. This mechanism is intricately designed to process and interpret the sensor data in both spatial and temporal



Fig. 3: Dual Attention Block (top) and benchmarking process used (bottom)

dimensions, enhancing the robot's understanding of human motion and intentions.

• Sensor (Spatial) Attention:

The first component of our dual attention structure is the Sensor (Spatial) Attention module. This element is meticulously engineered to analyze data collected from strategically synthesized IMUs placed across vital body joints. By concentrating on these specific body parts, the spatial attention module can capture the nuances of human movements more accurately. This aspect is crucial for achieving a human-like representation of motion, as it enables the system to focus on the intricate movements of key joints and limbs, which are pivotal in conveying implicit commands.

• Sequential (Temporal) Attention:

The second element, Sequential Attention, addresses the dynamic nature of human motion. It operates on the understanding that movement is a continuum influenced not just by the current physical state but also by historical motion patterns. By analyzing data over key time intervals, this module can discern patterns and trends in the sensor data, allowing it to predict subsequent movements based on historical data. This temporal focus is vital for understanding and anticipating human actions, making the robot's responses more fluid and intuitive.

This dual attention mechanism, with its spatial and temporal components, provides a comprehensive framework for interpreting sensor data. By integrating insights from both the current state and historical patterns of movement, our system achieves a more nuanced understanding of human motions and intentions. This enhanced perception is critical in enabling robots to interact with humans in a more natural, responsive, and predictive manner, thereby improving the efficacy and fluidity of human-robot collaboration.

In the present study, using the arbitrary parameters d_{Seq} and d_{IMU} defined as:

$$d_{Sea} = IMUNum \times Features \tag{1}$$

$$d_{IMU} = S \, equence \times Features \tag{2}$$

and the normalized attention scores, determined by the internal product of the query and key vectors:

$$Score_{Seq} = Q_{Seq}K_{Seq}^T \in \mathbb{R}^{Sequence \times Sequence}$$
 (3)

$$Score_{IMU} = Q_{IMU}K_{IMU}^T \in \mathbb{R}^{IMUNum \times IMUNum}$$
(4)

we defined the attention ratio (AR) as obtained from the Softmax function, applied respectively to each row of the input matrix.

$$AR_{seq} = Softmax\left(\frac{Score_{Seq}}{\sqrt{d_{Seq}}}\right)$$
(5)

$$AR_{IMU} = Softmax\left(\frac{Score_{IMU}}{\sqrt{d_{IMU}}}\right)$$
(6)

IV. PERFORMANCE VERIFICATION

A. Task Setting

The human participant is equipped with IMU sensors at designated points on their body.

- Task 1: Simple Handover Task
 - The participant approaches the robot arm and extends their hand with the object towards a designated handover location. The system uses the motion data to predict the intention and guides the robot arm to take the object from the participant.
- Task 2: Dynamic Collaborative Sorting Task
- A set of objects with varied shapes, sizes, and colors are placed on the table. Designated sorting zones are marked on the table for different categories of objects. The participant picks up an object and gestures towards a sorting zone, indicating where the robot should place the item. The robot must interpret the gesture, pick the correct object, and sort it into the designated zone. The task involves various object handovers and placements, requiring the robot to adjust its grip and trajectory based on the participant's motions and the object's characteristics.

Metrics include sorting accuracy, number of successful handovers, number of corrective actions by the participant, and overall task completion time. These tasks aim to test the system's reliability in basic interaction (Task 1) and its adaptability and decision-making in more complex, collaborative scenarios (Task 2). Success in both tasks would validate the system's potential for practical HRI applications.

B. Attention Benchmarking

The evaluation of our dual attention model was benchmarked against traditional BiRNN and sequential models (Fig.3) to quantify improvements in intention estimation accuracy. As Table I illustrates, the dual attention model demonstrated superior performance across all metrics when compared to the baseline BiRNN and sequential models.

Notably, the dual attention model achieved a significant increase in accuracy, reaching 96.2% with six layers, compared to 92.1% for the sequential model and 85.2% for the

TABLE I: Inter	ntion estimation	error	comparison
----------------	------------------	-------	------------

	BiRNN	S	equenti	al	Dua	al Atten	tion
layers		1	3	6	1	3	6
accuracy (%)	85.2	88.7	90.2	92.1	91.3	93.5	96.2
Standard Deviation	±7	± 4	±3	±3	± 4	±2	± 1
Inference (ms)	2.56	1.28	2.73	5.22	2.45	5.60	9.68

BiRNN. This enhancement in accuracy is attributed to the model's ability to concurrently process spatial and temporal data, providing a more holistic understanding of user intent. The standard deviation of accuracy also saw marked improvement, with the dual attention model exhibiting the lowest variability (± 1) at its optimal configuration of six layers, suggesting robust performance across different trials and user interactions. In terms of inference time, the dual attention model maintained competitive performance, with an acceptable increase to 9.68 ms for the six-layer configuration, which is a trade-off for the substantial gains in accuracy and consistency.

C. Sensor Scarcity

Our study, like others using BiRNNs for motion prediction, found that increasing the number of IMUs does not significantly improve intention estimation accuracy. Even adding sensors from 6 to 12 only slightly reduces joint position errors (as seen in Fig.4), indicating that the BiRNN model alone is not enough for precise posture extrapolation.

Research, including ours, suggests that beyond 12 IMUs, there are diminishing returns on accuracy. Further expanding the sensor array unlikely to substantially improve prediction accuracy. However, when focusing on enhancing cognitive depth by adding dual-attention layers instead of more sensors, our results show a progressive improvement in estimating joint positions. This dual-attention mechanism analyzes complex relationships between sensor inputs in both space and time, crucial fori accurate predictions.

While our goal was to optimize a system with minimal sensors, our findings suggest that adding more sensors could enhance the dual-attention framework. This is because a larger sensor array could provide more spatial-temporal relationships for the framework to analyze. Therefore, combining an expanded IMU setup with our dual-attention model could further reduce joint position estimation errors, showing promise for future research.

V. USABILITY EXPERIMENT

A. Participants

The participant pool for this study comprised 20 individuals, detailed in Table II. This group represented a balanced mix of ages (21-40 years), genders, and a variety of professional backgrounds to ensure broad usability testing of the system. Physical attributes ranged with heights between 158 cm and 182 cm and weights from 52 kg to 90 kg, catering to diverse body types. The participants' previous experience with HRI was split equally between novices and those with experience, providing insights into the system's



Fig. 4: Sensor scarcity versus attention layers performance comparison

accessibility to users of varying familiarity with robotic systems. This heterogeneous sample underpins the robustness and generalizability of the study's findings.

B. Simple Handover Task Performance

The simple handover task aimed to assess the system's capability to accurately interpret and respond to straight-forward user commands. Success rates across participants varied (Fig. 5a, with an average starting rate of 88.45% in the initial trials, peaking at 95.55%. This progression indicates a significant learning curve and adaptation to user behaviors, enhancing task execution over time. Conversely, response times showed a decreasing trend, starting from an average of 2.72 seconds and significantly improving to as low as 0.54 seconds, demonstrating the system's growing efficiency in interpreting user intentions promptly.

C. Dynamic Collaborative Sorting Task Performance

The dynamic collaborative sorting task (Fig. 5b), designed to test the system under more complex and variable conditions, revealed success rates beginning at 80.42%, with notable fluctuations reflecting the task's complexity. The highest success rate observed was 87.15%, underscoring the system's capacity to adapt to and manage intricate task demands. Response times for this task also decreased, highlighting the system's ability to streamline its processing and response capabilities amidst challenging scenarios.

D. Average Task Completion Time

The comparison of average task completion times, shown in Fig. 7 revealed a distinct pattern of improvement across both tasks. For the handover task, completion times decreased from 14 minutes in the initial trial to 8 minutes

Participant ID	Age	Gender	Height (cm)	Weight (kg)	Background	Previous HRI Experience
P01	25	Male	172	70	Engineering	Yes
P02	30	Female	165	55	Science	No
P03	27	Male	178	75	Arts	No
P04	22	Female	160	60	Healthcare	Yes
P05	34	Male	180	85	IT	Yes
P06	29	Female	167	65	Education	No
P07	31	Male	175	80	Business	No
P08	28	Female	162	54	Design	Yes
P09	26	Male	170	72	Engineering	No
P10	32	Female	168	58	Science	Yes
P11	24	Male	182	88	Arts	Yes
P12	33	Female	163	56	Healthcare	No
P13	35	Male	176	77	IT	No
P14	23	Female	158	53	Education	Yes
P15	37	Male	179	82	Business	No
P16	36	Female	164	59	Design	Yes
P17	38	Male	181	90	Engineering	Yes
P18	39	Female	169	57	Science	No
P19	40	Male	174	73	Arts	Yes
P20	21	Female	159	52	Healthcare	No

TABLE II: Participant Pool for Verification Experiment

by the fifth trial. The dynamic collaborative sorting task showed a more pronounced reduction, from 22 minutes initially to 10 minutes, indicating substantial efficiency gains as participants and the system became more acquainted with task requirements.

E. Analysis of Robot Error Distribution in the Sorting Task

Fig. 6 offers insightful revelations about the system's performance nuances. The majority of the actions, constituting 76.28%, were correctly identified and sorted, indicating a high level of proficiency in both recognizing and categorizing objects as intended. This significant portion underscores the system's robustness in accurately interpreting user commands and executing tasks correctly. However, the analysis also revealed areas for improvement. A total of 10.8% of actions were correctly identified yet mistakenly sorted, and 9.71% were mistakenly identified but correctly sorted. These errors highlight specific challenges in the system's decisionmaking process, where either the identification or the sorting phase may not align perfectly with the user's intention. This suggests a potential misalignment in the system's understanding of the task context or its execution strategy. Furthermore, the smallest error category, consisting of 3.21% of actions, involved objects that were both mistakenly identified and sorted. This category, while minor, points to instances where the system's interpretation of user intent significantly deviates from the expected outcome.

F. Corrective Actions and Intention Estimation Accuracy

A critical aspect of our analysis focused on the system's need for corrective actions and its accuracy in intention estimation, reflected in Fig. 8. The simple handover task saw a reduction in corrective actions from 6 in the first trial to just 1 by the fifth, alongside an improvement in intention estimation accuracy from 82% to 95%. The dynamic collaborative sorting task also demonstrated progress, with corrective actions decreasing from 13 to 4 and intention estimation



Fig. 5: Average task completion and response time comparison for each participant on the handover task and collaboprative dynamic sorting task



(a) Overall behavior distribution(b) Error break down per objectFig. 6: Robot arm error in identifying intended object to pick up or sorting location.

accuracy increasing from 73% to 87%. These improvements are indicative of the system's learning capabilities and its potential for minimizing human intervention over time.

VI. DISCUSSION AND CONCLUSION

This paper analyzed a dual attention mechanism in an HRI system, showing its effectiveness in two tasks: a Simple Handover Task and a Dynamic Collaborative Sorting Task. The mechanism improved the system's understanding of user intentions, leading to better success rates, faster responses, and fewer corrective actions. Processing both temporal and



Fig. 7: Evolution of the task completion time over the course of 5 trials



Fig. 8: Corrective actions (human intervention) and intention estimation accuracy evolution comparison

spatial sensor data, it accurately interpreted user commands, addressing variability in human behavior.

Despite positive outcomes, limitations were noted, such as variability in success rates and errors, especially in the dynamic collaborative sorting task, highlighting the complexity of real-world HRI scenarios. Further refinement of decisionmaking algorithms is needed for tasks with high variability and unpredictability.

The study was conducted in a controlled experimental setting, possibly not fully capturing real-world interactions. Evaluation over a short term raises questions about longterm adaptability and learning capabilities. Future research should focus on generalizing the system across a wider array of tasks and user behaviors, possibly with more sophisticated machine learning models. Addressing privacy and ethical concerns related to sensor data collection and user monitoring is crucial for responsible deployment.

In conclusion, this study demonstrates the potential of a dual attention mechanism to enhance HRI systems, indicating a promising direction for future research. Refinement of these systems could lead to seamless and intuitive humanrobot collaborations that adapt to and learn from diverse human behaviors.

ACKNOWLEDGMENT

References

 T. Sheridan, "Telerobotics," *Automatica*, vol. 25, no. 4, pp. 487–507, 1989. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0005109889900939

- [2] T. B. Sheridan, *Telerobotics, automation, and human supervisory control.* MIT press, 1992.
- [3] D. Roetenberg, P. Slycke, and P. Veltink, "Ambulatory position and orientation tracking fusing magnetic and inertial sensing," *Biomedical Engineering, IEEE Transactions on*, vol. 54, pp. 883 – 890, 06 2007.
- [4] B. Fang, F. Sun, H. Liu, and C. Liu, "3d human gesture capturing and recognition by the immu-based data glove," *Neurocomputing*, vol. 277, 08 2017.
- [5] L. Fiorini, F. G. C. Loizzo, A. Sorrentino, J. Kim, E. Rovini, A. Di Nuovo, and F. Cavallo, "Daily gesture recognition during human-robot interaction combining vision and wearable systems," *IEEE Sensors Journal*, vol. 21, no. 20, pp. 23 568–23 577, 2021.
- [6] J. Li, X. Liu, Z. Wang, T. Zhang, S. Qiu, H. Zhao, X. Zhou, H. Cai, R. Ni, and A. Cangelosi, "Real-time hand gesture tracking for human-computer interface based on multi-sensor data fusion," *IEEE Sensors Journal*, vol. 21, no. 23, pp. 26642–26654, 2021.
- [7] J. Yang, J. Zhang, C. Settle, A. Rai, R. Antonova, and J. Bohg, "Learning periodic tasks from human demonstrations," *CoRR*, vol. abs/2109.14078, 2021. [Online]. Available: https://arxiv.org/abs/2109.14078
- [8] S. Liu and H. Asada, "Teaching motion/force skills to robots by human demonstration," *IFAC Proceedings Volumes*, vol. 28, no. 15, pp. 679–685, 1995, 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design and Evaluation of Man-Machine Systems 1995, Cambridge, MA, USA, 27-29 June 1995. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667017453110
- [9] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in 2009 IEEE International Conference on Robotics and Automation, 2009, pp. 763–768.
- [10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2017. [Online]. Available: https://arxiv.org/pdf/1706.03762.pdf
- [11] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *ArXiv*, vol. 1409, 09 2014.
- [12] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," *CoRR*, vol. abs/1502.03044, 2015. [Online]. Available: http://arxiv.org/abs/1502.03044
- [13] I. Sorokin, A. Seleznev, M. Pavlov, A. Fedorov, and A. Ignateva, "Deep attention recurrent q-network," 12 2015.
- [14] O. Pfungst and R. Rosenthal, "Clever hans : the horse of mr. von osten," 1911. [Online]. Available: https://api.semanticscholar.org/CorpusID:142217369
- [15] W. B. Carpenter, "On the influence of suggestion in modifying and directing muscular movement, independently of volition." Royal Institution of Great Britain, 1852.
- [16] R. B. Evans, "William james, "the principles of psychology," and experimental psychology," *The American Journal of Psychology*, vol. 103, no. 4, pp. 433–447, 1990. [Online]. Available: http://www.jstor.org/stable/1423317
- [17] R. Lotze, Medizinische Psychologie. Jazzybee Verlag, 2012.
- [18] D. M. Wolpert, Z. Ghahramani, and J. Flanagan, "Perspectives and problems in motor learning," *Trends in Cognitive Sciences*, vol. 5, no. 11, pp. 487–494, 2001. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364661300017733
- [19] E. Todorov, "Optimality principles in sensorimotor control," *Nature neuroscience*, vol. 7, pp. 907–15, 10 2004.
- [20] M. Butz, O. Herbort, and G. Pezzulo, Anticipatory, Goal-Directed Behavior, 09 2008, pp. 85–113.
- [21] O. Herbort, M. V. Butz, and G. Pedersen, Model for Motor Learning and Control of a Redundant Arm: From Modeling Human Behavior to Applications in Robotics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 85–106.
- [22] S. M. Fiore, T. J. Wiltshire, E. J. C. Lobato, F. G. Jentsch, W. H. Huang, and B. Axelrod, "Toward understanding social cues and signals in human-robot interaction: effects of robot gaze and proxemic behavior," *Frontiers in Psychology*, vol. 4, 2013.
- [23] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. AAAI*, 2008, pp. 1433–1438.