



OPEN ACCESS

EDITED BY

Nutan Chen,
Volkswagen Group, Germany

REVIEWED BY

Izidor Mlakar,
University of Maribor, Slovenia
Xingyuan Zhang,
ARGMAX.AI Volkswagen Group Machine
Learning Research Lab, Germany

*CORRESPONDENCE

Anna Deichler,
✉ deichler@kth.se

SPECIALTY SECTION

This article was submitted to Robot
Learning and Evolution, a section of the
journal Frontiers in Robotics and AI

RECEIVED 28 November 2022

ACCEPTED 16 March 2023

PUBLISHED 30 March 2023

CITATION

Deichler A, Wang S, Alexanderson S and
Beskow J (2023), Learning to generate
pointing gestures in situated embodied
conversational agents.
Front. Robot. AI 10:1110534.
doi: 10.3389/frobt.2023.1110534

COPYRIGHT

© 2023 Deichler, Wang, Alexanderson
and Beskow. This is an open-access
article distributed under the terms of the
[Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is
permitted, provided the original author(s)
and the copyright owner(s) are credited
and that the original publication in this
journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Learning to generate pointing gestures in situated embodied conversational agents

Anna Deichler*, Siyang Wang, Simon Alexanderson and
Jonas Beskow

Division of Speech, Music and Hearing, KTH Royal Institute of Technology, Stockholm, Sweden

One of the main goals of robotics and intelligent agent research is to enable them to communicate with humans in physically situated settings. Human communication consists of both verbal and non-verbal modes. Recent studies in enabling communication for intelligent agents have focused on verbal modes, i.e., language and speech. However, in a situated setting the non-verbal mode is crucial for an agent to adapt flexible communication strategies. In this work, we focus on learning to generate non-verbal communicative expressions in situated embodied interactive agents. Specifically, we show that an agent can learn pointing gestures in a physically simulated environment through a combination of imitation and reinforcement learning that achieves high motion naturalness and high referential accuracy. We compared our proposed system against several baselines in both subjective and objective evaluations. The subjective evaluation is done in a virtual reality setting where an embodied referential game is played between the user and the agent in a shared 3D space, a setup that fully assesses the communicative capabilities of the generated gestures. The evaluations show that our model achieves a higher level of referential accuracy and motion naturalness compared to a state-of-the-art supervised learning motion synthesis model, showing the promise of our proposed system that combines imitation and reinforcement learning for generating communicative gestures. Additionally, our system is robust in a physically-simulated environment thus has the potential of being applied to robots.

KEYWORDS

reinforcement learning, imitation learning, non-verbal communication, embodied interactive agents, gesture generation, physics-aware machine learning

1 Introduction

1.1 Overview

Humans rely on both verbal and non-verbal modes of communications in physically situated conversational settings. In these settings, non-verbal expression, such as face and hand gestures, often contain information which is not present in the speech. A prominent example is how people point to an object instead of describing it with words. This complementary function makes communication more efficient and robust than speech alone. In order for embodied agents to interact with humans more effectively, they need to adapt similar strategies, i.e., they need to both comprehend and generate non-verbal communicative expressions.

Regarding comprehension, one focus has been on gesture recognition in order to accomplish multimodal reference resolution and establishing common ground between human and agent (Abidi et al. (2013); Häring et al. (2012); Wu et al. (2021)). Regarding generation, the main focus has been on generating co-speech gestures that typically uses supervised learning techniques to map text or speech audio to motion [see Liu et al. (2021) for an overview]. While these approaches can generate natural looking gesticulation, they only model beat gestures, a redundant aspect of non-verbal communication. Moreover, these supervised methods require large labeled datasets tied to a specific embodiment and are not physics-aware, posing a major issue for transferring the results to robots.

The focus of our study is the generation of pointing gestures. According to McNeill's influential classification, pointing gestures are among the four primary types of gestures, alongside with iconic, metaphorical and beat gestures (McNeill (1992)). Iconic gestures depict concrete objects or actions with the hands. Metaphoric gestures express abstract concepts or ideas with the hands. Beat gestures are simple movements that emphasize or segment speech. Deictic gestures point to objects, locations, or persons in space. Pointing is the most specialized of all gestures in terms of its association with particular verbal expressions, specifically reference and referring expressions (Kibrik (2011)). Linguistic reference is the act of using language to refer to entities or concepts. Linguistic elements that perform a mention of a referent are called referring expressions or referential devices. People use reference in discourse to draw the listener's attention to a referent or target, which can be in either the speech-external (deictic reference) or speech-internal (anaphoric reference) environment. The anaphoric referent is an element of the current discourse, while the deictic referent is outside the discourse in the spatio-temporal surroundings (Talmy (2020)). Deictic expressions are used to indicate a location or point in time relative to the deictic centre, which is the center of a coordinate system that underlies the conceptualization of the speech situation. They are essential parts of human communication, since they establish a direct referential link between world and language. Demonstratives like "this," "that" or "there" are the simplest form of deictic expressions that focus the interlocutor's attention to concrete entities in the surrounding situation (Peeters et al. (2021)). In everyday conversations people often use non-verbal means (e.g., eyes, head, posture, hands) to indicate the location of the referent alongside with verbal expressions to describe it. Pointing and eye gaze are the most prominent non-verbal means of deictic reference and also play a key role in establishing joint attention in human interactions (Diessel and Coventry (2020)), a prerequisite for coordination in physical spaces. Pointing is ubiquitous in adult interaction across settings and it has been described as "a basic building block" of human communication (Kita (2003)). Demonstratives produced with pointing gestures are more basic than demonstratives produced without pointing (Cooperrider (2020)). Speakers also tend to prefer short deictic descriptions when gesture is available, in contrast to longer non-deictic descriptions when gesture is not available (Bangerter (2004)).

We proposed a new framework based on imitation and Reinforcement Learning (imitation-RL) for this for generating referential pointing gestures in physical situated environment. Our method, adapted from Peng et al. (2021), learns a motor

control policy that imitates examples of pointing motions while pointing accurately. By providing only a few pointing gesture demonstrations, the model learns generalizable accurate pointing through combined naturalness and accuracy rewards. In contrast to supervised learning, our method achieves high pointing accuracy, requires only small amounts of data, learns physically valid motion and has high perceived motion naturalness. It is also considerably more lightweight than compared supervised-learning methods. To implement the method, we first collected a small dataset of motion-captured pointing gesture with accurate pointing positions recorded. We then trained a motor control policy on a humanoid character in a physically simulated environment using our method. The generated gestures were then compared against several baselines to probe both motion naturalness and pointing accuracy. This was done through two subjective user studies in a custom-made virtual reality (VR) environment. In the accuracy test, a full-body virtual character points at one of several objects, and the user is asked to choose the object that the character pointed at, a setup known as referential game (Lazaridou et al. (2018), Steels (2001)). This setup is also motivated by the notion that language use is a triadic behaviour involving the speaker, the hearer, and the entities talked about (Bühler (1934)).

We found that the proposed imitation-RL system indeed learns highly accurate pointing gestures while retaining high level of motion naturalness. It is also shown that imitation-RL is much better at both accuracy and naturalness than a supervised learning baseline. The results suggest that imitation-RL is a promising approach to generate communicative gestures.

1.2 Related work

Our work concerns learning communication in embodied agents through reinforcement learning, specifically learning pointing gestures in a physically simulated environment with imitation-RL.

Below we first review related work in the area of non-verbal deictic expression and gesture generation for embodied agents, followed by an overview of two reinforcement learning fields related to our work: (a) reinforcement learning for communication, (b) reinforcement learning for motion control. We believe that our work is the first that learns non-verbal communication with reinforcement learning, which can be seen as bridging the two currently disjunct reinforcement learning fields.

1.2.1 Gesture and non-verbal deictic expression generation

There has been substantial work considering the generation of speech-driven gestures in embodied agents. Early work in synthesis of gestures for virtual agents employed methods based on rules, inverse kinematics and procedural animation to generate co-speech gestures (Cassell et al. (1994), Cassell et al. (2001); Kopp and Wachsmuth (2004); Ng-Thow-Hing et al. (2010); Marsella et al. (2013)). Rule-based approaches, however, often require laborious manual tuning and often struggle to achieve a natural motion quality. More recently, supervised learning systems for co-speech gesture generation have achieved high naturalness in certain settings [Liu et al. (2021)]. Some of these systems take only speech

audio as input [Hasegawa et al. (2018); Kucherenko et al. (2021a); Ginosar et al. (2019); Ferstl et al. (2020); Alexanderson et al. (2020)] but are limited to generating beat gestures aligned with speech. Extending these supervised learning based gesture generation methods to produce iconic, metaphoric or deictic gestures would require more elaborate modelling and a large dataset containing these gestures Ferstl and McDonnell (2018). Adding text as input can potentially generate more semantic gestures but is still limited Kucherenko et al. (2020); Yoon et al. (2020); Ahuja et al. (2020); Korzun et al. (2020). Common for all these methods is that they do not take into account the surrounding environment and thus cannot communicate other than redundant information. On the contrary, works related to deictic expression generation for intelligent agents take into account the physical environment in which the agent is situated. Pointing and gaze generation in general have been extensively studied in virtual agents and robots. Rule-based approaches for deictic expression generation are often used in virtual agents Noma et al. (2000); Rickel and Johnson (1999); Lester et al. (1999), as well as robotics. Fang et al. (2015) and Sauppé and Mutlu (2014) both implemented rule-based pointing gestures in a humanoid robot for referential communication. Holladay et al. (2014) proposes a mathematical optimization approach to make legible pointing, i.e., disambiguation of closely situated objects. In Sugiyama et al. (2007) natural generate rule-based deictic expressions consisting of pointing and a reference term are generated in an interactive setting with a robot. Referential gaze has been studied in virtual agents Bailly et al. (2010), Andrist et al. (2017), as well as robotics Mutlu et al. (2009). Some studies have adopted learning-based approaches to develop pointing gestures generation. Huang and Mutlu (2014) proposes a learning based modeling approach using dynamic Bayesian network to model speech, gaze and gesture behavior in a narration task. Zabala et al. (2022) developed a system of automatic gesture generation for a humanoid robot that combines GAN based beat gesture generation with probabilistic semantic gesture insertion, including pointing gestures. Chao et al. (2014) takes a developmental robotics approach to generating pointing gestures using reinforcement learning for a 2D reaching task. None of the above works consider human-like and physics based pointing motion generation in 3D space.

1.2.2 RL for communication in task-oriented embodied agents

There have been a substantial amount of works in recent years that focused on extending the capabilities of embodied RL agents with language. This field has been reviewed in a recent survey Luketina et al. (2019) which separates out these works into two categories: language-conditioned RL, where language is part of the task formulation and language-assisted RL, where language provides useful information for the agent to solve the task. The simplest form of language-conditioned is instruction following, in which a verbal instruction of the task is provided to the agent. Examples are manipulation Stepputtis et al. (2019) and navigation task Qi et al. (2020). These studies focused more on comprehension of verbal expressions and usually no interaction takes place between the user and the agent. In more complex settings, the agent interacts with humans using natural language. In Lynch et al. (2022) the authors present a RL and imitation based framework, Interactive Language, that is capable of continuously adjusting its behavior

to natural language based instructions in a real-time interactive setting. There have also been a recent wave of datasets and benchmarks created by utilizing 3D household simulators and crowd sourcing tools to collect large-scale task-oriented dialogue aimed at improving the interactive language capabilities of embodied task-oriented agents Padmakumar et al. (2022), Gao et al. (2022), Team et al. (2021). Most of the above mentioned works focus on the verbal mode of communication and largely on the comprehension side (e.g., instruction following). Less work has explored the non-verbal communication in situated embodied agents Wu et al. (2021), especially on the generation side, which we aim to address in our work.

1.2.3 Learning motor control through RL

RL has been extensively applied to learning motor control for both robotics and graphics applications Duan et al. (2016); Heess et al. (2017). The learned motion dynamic for humanoid characters from these methods are usually not very human-like. To improve both the motion dynamics and also to facilitate learning, it is proposed for the learning agent to imitate expert demonstration or motion-capture animation Merel et al. (2017); Ho and Ermon (2016). However, it is not until DeepMimic Peng et al. (2018) that the RL-imitation approach achieved human-like motion. The drawback in DeepMimic is that it uses several manually tuned hyper-parameters in the imitation reward, a problem addressed by a follow-up work Adversarial Motion Prior (AMP) Peng et al. (2021) which replaces the “hard” imitation in DeepMimic with a learned discriminator (similar to GAN). Through a combination of imitation reward and task reward, methods like AMP can learn human-like motion dynamics while completing a manually defined task, for example, heading in a given direction while walking or spin-kick a given target. It has been shown that even with just one motion clip as imitation target, AMP learns generalizable motion control given random task input (e.g., randomly sampled walking direction) while maintaining high motion naturalness.

2 Materials and methods

2.1 Pointing gesture data collection

Our method for synthesizing human-like pointing gestures for the humanoid character requires an appropriate full-body motion capture dataset with diverse and accurate target locations covering the 3D space surrounding the character. Multiple datasets exist of referring expression with pointing as non-verbal modality Schauerte and Fink (2010); Matuszek et al. (2014); Shukla et al. (2015); Chen et al. (2021), but since most of these focus on comprehension, they lack full body motion capture, which is essential for our method. Existing full-body motion capture datasets for gesture generation Ferstl and McDonnell (2018), Ferstl et al. (2020) focus on beat gestures with the aim to train speech-to-gesture generative models and do not contain pointing gestures. A recent dataset Islam et al. (2022) records multi-modal referring expressions, including pointing gestures, but it is restricted to a tabletop setting.

To obtain a collection of ground truth examples for training and evaluation, we recorded a pointing gesture dataset in a optical

motion capture studio equipped with 16 Optitrack Prime 41 cameras. The actor wore a suit with 50 passive markers, a pair of Manus data gloves (hand motion capture) and a head-mounted iPhone 12 (face motion and voice capture). In addition, we also recorded a pointing target, consisting of a rigid structure equipped with 4 markers. The marker data was solved to a representation of the actor's skeleton and the location of the target in 3D space using the systems software (Motive 3.0).

Three different pointing tasks were recorded: single target pointing, two targets selection and two targets moving (point-and-place). In each setting the targets were moved around in order to get a coverage of the surrounding space, while the actor was in a stationary position. Between each pointing task the actor returned to a neutral stance. The beginning and the end of the movement is defined by the pointing hand leaving and returning to the initial, downward position. We used this information to parse the continuous recording of motion capture data into pointing clips. For more details, see [Deichler et al. \(2022\)](#).

We focus on single target pointing in this study. Here, we provide an overview of that data. We first divided the pointing target positions to front and back. We only use the front data in our study. The total sum of single target positions in the dataset is 83, from which 52 are front target pointing movements (25 left-handed and 27 right-handed). The actor's dominant hand is right hand. The target distribution of this subset is visualized in 3D space in [Figure 1](#) (blue balls). The single target dataset is parsed into single arm pointing movement based on peaks in displacement on the sagittal plane (yz). We further analyse the pointing movements in terms of pointing accuracy ([Figure 2](#)) and velocity ([Figure 3](#)) dividing into dominant (right) and non-dominant (left) hand. The pointing accuracy is calculated based on inverse of alignment angle between arm and arm to pointing target (Eq. 2), the more accurate pointing the higher this measure is. As seen in [Figure 2](#), both dominant and non-dominant hands have a bell-shaped accuracy curve. This reflects the pointing motion trajectory: the hand first moves to pointing position as the accuracy rises, it stays there for some time (holding), and then retracts. Moreover, we note that neither hand achieves maximum accuracy. This shows that human pointing is not accurate in terms of alignment angle, which has been shown in previous pointing gesture studies [Lücking et al. \(2015\)](#). The velocity trajectory ([Figure 3](#)) shows clear correspondence with accuracy trajectory. No significant difference between dominant and non-dominant hands was found.

We applied left-right mirroring on the mocap data, which doubled the total amount of data available for model training. This is mainly due to the baselines MoGlow ([Section 2.4.1.1](#)) and GT-nn ([Section 2.4.1.2](#)) require large amount of data to perform well. For fair comparison, we used the same mirrored data for training all systems including our proposed method even though our proposed method does not have the same data quantity requirement as the baselines.

2.2 Method

Our method is based on imitation-RL method AMP [Peng et al. \(2021\)](#), capable of learning complex naturalistic motion on humanoid skeletons and showing good transferrability of

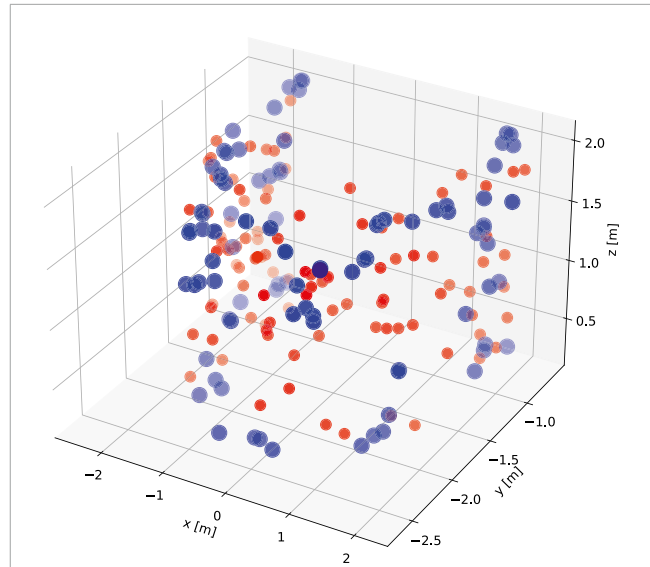


FIGURE 1
Targets from training set (blue) and perceptual test set (red) visualized in 3D space.

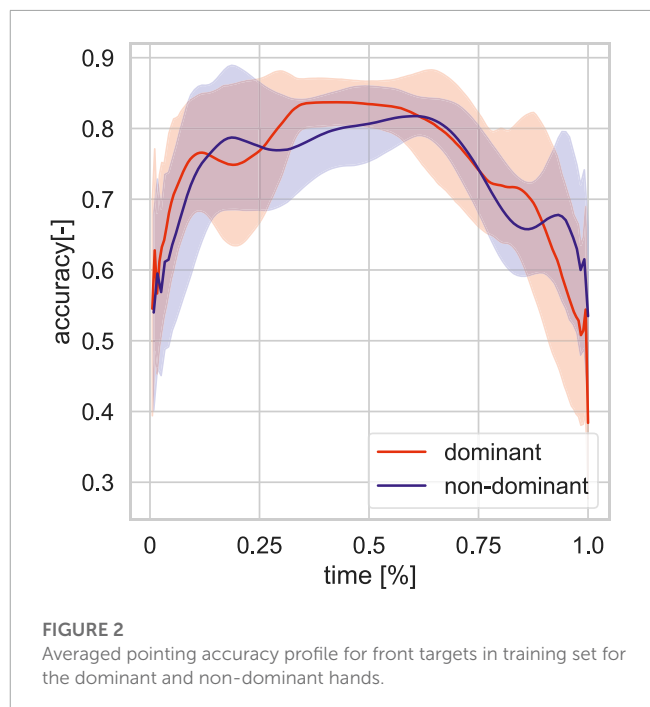
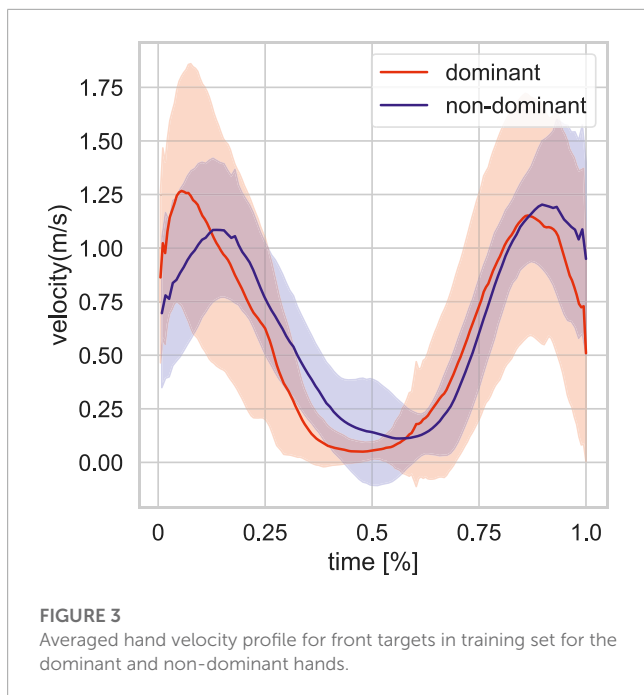


FIGURE 2
Averaged pointing accuracy profile for front targets in training set for the dominant and non-dominant hands.

simulation-learned policy to real-world robots [Escontrela et al. \(2022\)](#), [Vollenweider et al. \(2022\)](#). AMP is a subsequent work of Generative Adversarial Imitation learning [Ho and Ermon \(2016\)](#), that takes one or multiple clips of reference motion and learns a motor control policy π_θ that imitates the motion dynamics of the reference (s) through a discriminator network D_ϕ . An optional task reward can be added for the specific motion task to encourage task completion in the learned policy, for example, heading in a given direction while walking. The full reward function is the combination of rewards from motion imitation r_t^I and the task reward r_t^G , weighted



respectively by ω^I, ω^G ,

$$r_t = \omega^I r_t^I + \omega^G r_t^G. \quad (1)$$

The policies are trained using a combination of GAIL and PPO [Proximal Policy Optimization Schulman et al. (2017)]. The schematic view of the system is presented on Figure 4. We kept $\omega^I = \omega^G = 0.5$ and in all AMP models. We used same action and state space as humanoid motion example presented in original AMP Peng et al. (2021). The state is defined as the character's configuration in the character's local root joint based coordinate system. The state vector includes the relative positions of each link with respect to the root, their rotation and linear and angular velocities. The actions specify the target positions in the PD controllers for the controllable joints.

We designed a reward function r_t^{Pt} for pointing accuracy that takes the angle between the hand-to-target vector $\overrightarrow{V_{HT}}$ and the elbow-to-hand vector $\overrightarrow{V_{EH}}$, and gives higher reward to smaller angle, i.e., more accurate pointing. This is illustrated in Figure 5. The angular distance from the index finger ray has been used as a heuristics for pointing accuracy Lücking et al. (2015), here we use the lower-arm extension, due to the limitations of the simulated agent's morphology. Similarly to previous closely related work in humanoid motion generation with RL and imitation Peng et al. (2021), Peng et al. (2018), we apply exponential scaling to angular distance term to form the task reward function as follows,

$$\hat{\theta} = 1 - \frac{\mathcal{L}(\overrightarrow{V_{HT}}, \overrightarrow{V_{EH}})}{\pi}, \quad r_t^{Pt} = \frac{e^{\hat{\theta}} - 1}{e} \quad (2)$$

Since our task is to point at a given position, the x-y-z coordinates of the pointing position are concatenated to policy network input. Additionally, it is common in methods like AMP

to sample the task input randomly to achieve good generalization Peng et al. (2021); Peng et al. (2018), for example, to randomly sample walking directions when learning directional walking. We also applied this approach by randomly sampling pointing target positions in the training process. The sampling process takes in a ground-truth pointing position (corresponding to the current pointing gesture demonstration) as input and uniformly samples points within a square box of size $20 \times 20 \times 20$ cm centered at that ground-truth position.

Through preliminary experiments, we found that the original AMP algorithm does not learn pointing gesture effectively. We subsequently implemented changes described in the following sections to improve AMP in learning pointing gesture.

2.2.1 Phase input to network: AMP base

In motion imitation it is common to provide as input normalized running time called phase ($p \in [0, 1]$, 0 is beginning of a clip, 1 is the end) to the policy network in order to synchronize the simulated character with the reference motion. This was proven to be effective in imitation of single motion clips in DeepMimic Peng et al. (2018). In AMP this is not needed, since the policies are not trained to explicitly imitate a single motion clip. However, we found that vanilla AMP without phase input struggles to learn pointing gesture motion dynamics, thus we provided the policy network with phase to help the system learn the dynamics of the pointing motion. We call this model AMP base.

2.2.2 Phase-functioned neural network with pointing target clustering: AMP-pfnn

We found that the discriminator network in AMP, which learns what "plausible" motion looks like through adversarial loss and provide imitation reward to the policy network, has difficulty learning pointing gesture dynamics even with phase as input. Specifically, it has trouble distinguishing between the raising phase of pointing gesture and the retraction phase, resulting in the learned motion stuck in the middle between idle and pointing positions. This is unsurprising since both phases of the pointing gesture have similar trajectories (in opposite directions). We tackled this issue by changing discriminator network from a simple MLP (as in original AMP) to a phase-functioned neural network (PFNN) Holden et al. (2017).

Through early experiments, we also found that having a single network to learn pointing in all directions can be challenging. Thus, we grouped the training set pointing positions into clusters in the 3D space. The clustering is done with the MeanShift algorithm, which is based on kernel density estimation (KDE). Using Scott's rule the estimated bandwidth is 0.78 for the univariate KDE, which results in 8 clusters in the target space. The estimated bandwidth is quite wide given the target space limits, but this is expected with a limited amount of data. Based on this clustering, we separated the motion clips into subgroups and trained specialized networks for the different subgroups. During test time, the subgroup which contains the closest training target to the test target in terms of Euclidean distance is chosen to generate the pointing movement. We denote this version of AMP with pfnn and cluster group training as AMP-pfnn.

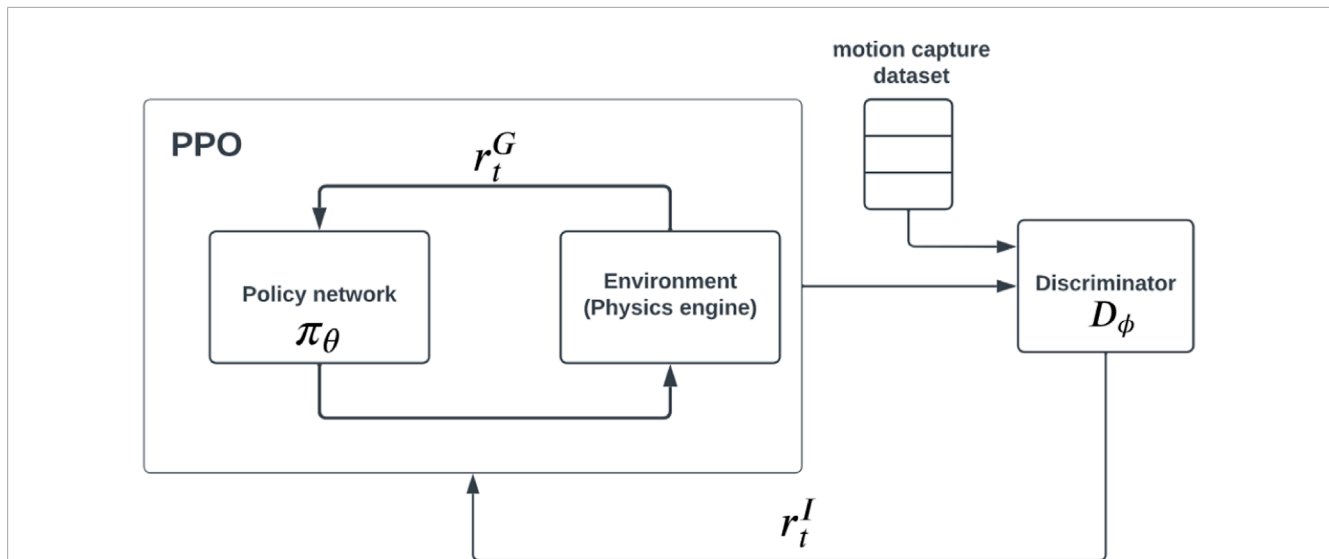


FIGURE 4

Schematic view of the AMP framework, which allows using a composite reward of task r_t^G and imitation r_t^I rewards by combining RL policy gradient methods (PPO) with a GAN-like discriminator.

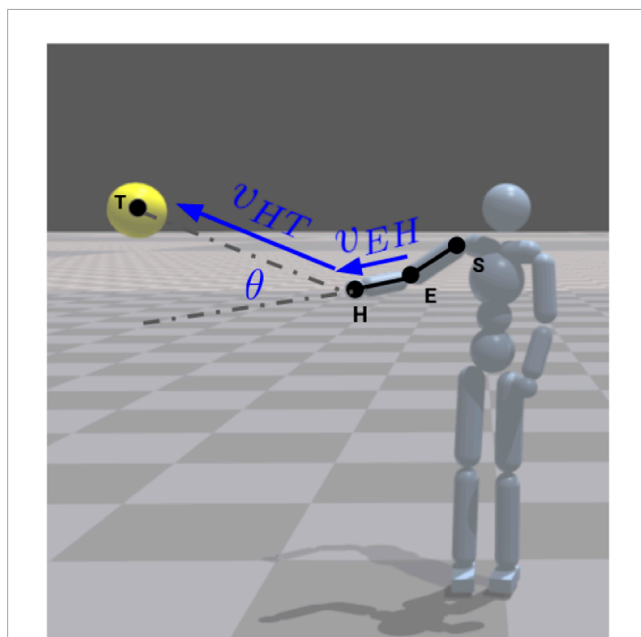


FIGURE 5

Pointing reward based on angle between vector from elbow joint (E) to hand joint (H) and vector from hand joint to target position (T). Agent rendered in the IsaacGym physics engine Makoviychuk et al. (2021).

2.3 Implementation

We use a GPU-based physics simulation software Isaac gym Makoviychuk et al. (2021), a rendered example of which can be seen in Figure 5. For implementation of our proposed method, we modified an existing AMP implementation provided by authors

of Isaac gym¹. The modifications also include enabling learning pointing gesture in humanoid characters and adding our proposed reward function for pointing accuracy.

2.4 Experiment

The goal of our proposed system is to produce natural and accurate pointing gestures in an interactive agent. In order to evaluate both aspects in an embodied interaction, we created a novel virtual reality (VR) based perceptual evaluation using a 3D game engine. This is described with more details in Section 2.4.2. We further probed the specific aspects of our proposed model in an object evaluation. We benchmarked our proposed method with two baselines, a supervised learning motion synthesis model MoGlow, and a simple nearest-neighbour retrieval method.

2.4.1 Baselines

2.4.1.1 MoGlow

We used a state-of-the-art animation synthesis model MoGlow Henter et al. (2020) as supervised learning baseline. This model is based on machine learning framework normalizing flow Dinh et al. (2014) Dinh et al. (2016) which directly estimates data probability distribution with invertible neural network. MoGlow utilizes a popular normalizing flow architecture Glow Kingma and Dhariwal (2018) and is able to synthesize animation in a conditional or unconditional probabilistic manner, and has shown state-of-the-art performance in synthesizing locomotion Henter et al. (2020), dancing Valle-Pérez et al. (2021), and gesture Alexanderson et al. (2020).

¹ <https://github.com/NVIDIA-Omniverse/IsaacGymEnvs>.

We extended the conditional generation mode of MoGlow to achieve pointing gesture synthesis. We provided the model with a 4-dimensional vector as condition, consisting of pointing target position (x-y-z coordinates) and an indicator variable of whether the current frame is pointing or idling. This indicator variable differs from the AMP phase variable, as it takes binary values, whereas the phase variable is the normalized running time. This difference stems from a fundamental difference between the two models. The MoGlow model is based on autoregressive supervised learning, and therefore has a “sense of time,” AMP is RL based and does not get the history as input, therefore the phase is used as an input.

The data preparation consists of first downsampling the mocap data to 20 frames per second, then cutting it into training clips with a sliding window of window length 6 s and step size 0.5 s. Such data preparation is consistent with prior studies using MoGlow for locomotion generation [Henter et al. \(2020\)](#) and gesture generation [Alexanderson et al. \(2020\)](#).

2.4.1.2 Nearest-neighbor queried ground-truth animation (GT-NN)

We also created a simple baseline through nearest-neighbor querying of the dataset. Given an input pointing position, this baseline model queries the closest pointing position in the dataset and plays the corresponding pointing gesture animation. We call this baseline Ground-truth Nearest-Neighbor (GT-NN). This baseline is the topline in motion naturalness since it is playing back motion-captured animation. On the other hand, its performance in referential accuracy would be determined by how close is the input position to positions in the dataset. If the input is close to some pointing position already existing in the dataset, then this method is expected to perform well, otherwise poorly. Thus, comparing a more complex learned model with this simple baseline in referential accuracy should reveal how much the learned model is able to generalize beyond the training data to unseen pointing target positions.

2.4.2 VR perceptual test

The perceptual test is an embodied interaction in VR. A user is put in a shared 3D space with an embodied virtual agent. In this setting, the virtual agent makes pointing gestures and the user is asked to evaluate motion naturalness of the gesture and play a referential game, i.e., guessing which object is the virtual agent pointing at. The two aspects are evaluated in two separate stages as seen in [Figure 6](#).

In the first stage, the naturalness of the pointing gesture is evaluated by presenting a pointing gesture to the user, then asking the user to rate “How natural do you find the animation?” on a 1-5 scale [in [Figure 6A](#)]. The avatar is in a fixed position, facing the user in this stage. We intended for this naturalness test to only evaluate motion dynamics and not pointing accuracy, so we did not show the pointing object at this stage so that the user is not distracted by potentially inaccurate pointing. It is likely that a user would rate a pointing motion as unnatural even though the motion dynamics itself is natural but it does not point accurately at the object. We avoided this by not showing the pointing object at this stage. This evaluation setup is similar to the one in the GENE co-speech generation challenge [Kucherenko et al. \(2021b\)](#). In the GENE challenge the evaluation metrics included

“human-likeness” and “appropriateness” where the first one measures motion quality, like naturalness in our study, and the second one measures task performance - in their case how well motion matches speech, in our case the equivalent metric is pointing accuracy.

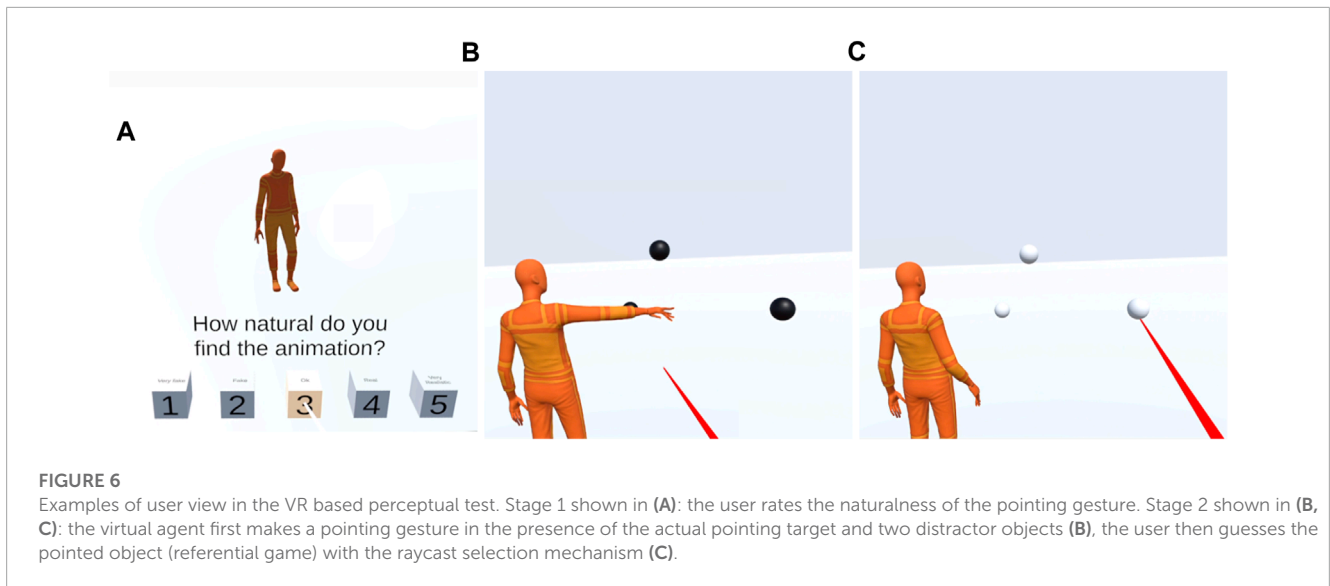
In the second stage, the pointing accuracy is evaluated through a simplified embodied referential game [Steels \(2001\)](#), [Lazaridou et al. \(2018\)](#). The participants were presented with 3 balls in the environment, 1 pointing target and 2 distractors (the sampling process are described in [Section 2.4.3](#)). In this stage, the avatar’s relative position to the user is varied between “across” and “side-by-side” conditions. In case of the “side-by-side” condition, the user stands by the shoulder of the avatar, allowing for a shared perspective on the target. In case of the “across” condition, the avatar is facing the user. After observing the agent’s pointing motion, participants were asked to guess which object the agent was pointing at. The target selection mechanism was implemented using a raycast selection [red ray seen in of [Figures 6B, C](#)]. The raycast selection becomes available to the user after the motion has ended and this is indicated to the user by the objects color changing from black to white. This setup prevents the user from selecting before seeing the entire motion.

Each participant saw 5 samples for each of the 4 models in motion naturalness test (first stage) and 10 samples for each model in referential accuracy test (second stage). Thus each participant sees 20 samples in total for motion naturalness test and 40 in total for referential accuracy test. There are more samples in the referential accuracy test because a model’s pointing accuracy may vary depending on the pointing position thus requires more samples to estimate true mean, while motion naturalness is easier to judge with less samples since it varies little given different pointing positions. We separated the two stages, i.e., the participant would finish motion naturalness test (first stage) before doing referential accuracy test (second stage). We also randomized ordering of the samples within each test to mitigate unwanted factors that could bias resulting statistics, such as the participant might take some time to get used to the VR game setup thus making data in the early part of the test less reliable. Lastly, the pointing positions in both tests (5 for stage 1, 10 for stage 2) were randomly picked from the 100 sampled test positions ([Section 2.4.3](#)) and were shared among models for the same user to ensure fair comparison.

2.4.3 Sampling test pointing targets and distractors

We sampled 100 test pointing target positions within the range of ground-truth positions. The ground-truth positions roughly form a half-cylindrical shape as seen in [Figure 1](#). We thus defined the range of ground-truth position as a parameterized half-cylinder. We used 3 parameters, the height of the cylinder, the arc, and the radius of the cylinder. Given the ranges of these 3 parameters obtained from the ground-truth data, we then sampled new positions by uniformly sampling these 3 parameters within their respective ranges. The sampled 100 test positions are visualized in [Figure 1](#) (red balls are sampled positions, blue balls are ground-truth).

For each sampled test position, we sampled 2 distractors in the following way. For each test position, we first sampled a distractor position the same way as the test positions (within the parameter ranges), we then checked if that distractor is within 20–40 cm away



from the input test position. If not, we resampled. We repeated this process until we got 2 distractors that met the distance condition. This process was done for each of the 100 test pointing positions.

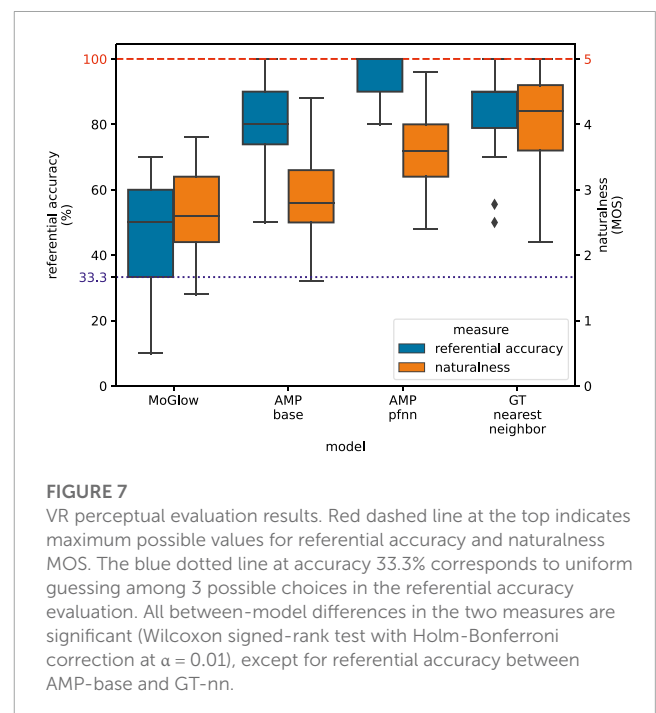
We note that the distance range 20–40 cm effectively determines the difficulty of the distractors. Referential game is more difficult if the distractors are closer to the actual pointing target. Through preliminary experiments, we found that 20–40 cm range is difficult enough that the model would have to point with a high level of accuracy, and at the same time not too difficult that even ground-truth pointing would not be able to distinguish the correct position from distractors. Admittedly, our process is not fully theoretically driven, and future studies could explore varying distractor distance or choose this distance in a more theoretical manner.

3 Results

3.1 VR perceptual test results

We recruited 39 participants (age: min = 20, median = 23, max = 39; female: 16, male: 23). Before conducting statistical analysis on the results, we first averaged each participant's scores for each model in the two measures, referential accuracy and motion naturalness, thus obtaining a participant's average scores for the four models in the two measures. Subsequent analysis is done on this data.

Results from VR perceptual evaluation are shown in **Figure 7**. AMP-pfnn (2.2.2) obtains highest referential accuracy (mean = 95.5%) and highest naturalness MOS (mean = 3.57) among models. GT-nn (2.4.1.2) is second best overall at accuracy (mean = 85.6%) and is the naturalness MOS topline (mean = 4.01) since it is playing back motion-captured animation. AMP-base (2.2.1) obtains comparable referential accuracy (mean = 83.0%) as GT-nn, but has much worse naturalness (mean = 2.93) than either AMP-pfnn or topline GT-nn. MoGlow (2.4.1.1) performs worst in both accuracy (mean = 45.9%) and naturalness (mean = 2.61). All but one



between-sample differences are significant at $\alpha = 0.01$ according to Wilcoxon signed-rank test with Holm-Bonferroni correction. The only insignificant difference is between AMP-base and GT-nn in accuracy measure. These results show that AMP-pfnn is able to point with very high referential accuracy in the presence of challenging distractors, and with high motion naturalness approaching ground-truth motion capture. Analysis on the avatar's relative position to the user in stage two shows that there is no significant between-sample difference for the "across" and "side-by-side" conditions at $\alpha = 0.01$ according to Wilcoxon signed-rank test with Holm-Bonferroni correction for any of the models. Further results on the position analysis can be found in **Table 1**.

TABLE 1 Referential accuracy results (mean and standard deviation) for “across” and “side-by-side” agent positions in VR perceptual evaluation.

Model	Across	Side-by-side
All	0.78 ± 0.41	0.76 ± 0.42
GT nearest neighbor	0.84 ± 0.37	0.88 ± 0.33
MoGlow	0.5 ± 0.5	0.39 ± 0.49
AMP base	0.83 ± 0.38	0.84 ± 0.37
AMP PFNN	0.96 ± 0.20	0.95 ± 0.22

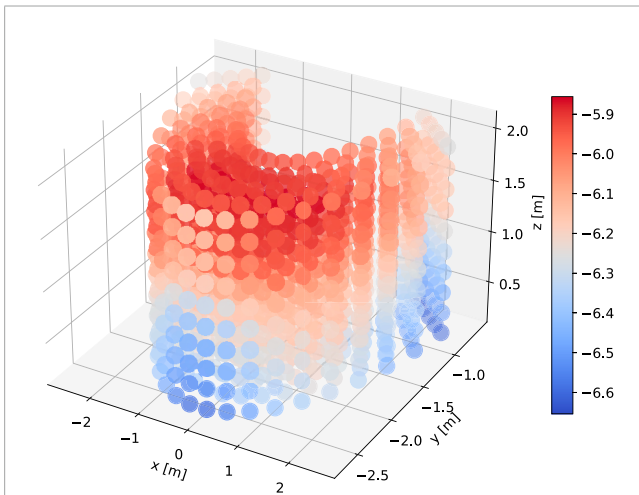


FIGURE 8 Visualization of KDE estimation values over the 3D grid for objective evaluation. The KDE was fitted on the training set targets.

3.2 Objective evaluation results

We also conducted objective evaluations in order to gain further insights into the performance of the different models. This is done by calculating the pointing accuracy (Eq. 2) over a uniform grid in the 3D space. This is also a more principled way to measure the generalization capabilities of the learning systems. In this evaluation, the pointing accuracy was calculated from the pointing reward function (Figure 5). The grid was created by taking the x,y,z limits of the training set targets and generating a spherical grid of 1,000 test points (Figure 8). We defined the accuracy of the pointing movement as the maximum of the reward curve 1, where the velocity is below 0.5 m/s in a window of frames corresponding to 0.15 s around the time of pointing. Apart from the systems in the perceptual evaluation, we also added ablations to examine the effect of run time (phase) input to the AMP policy network (AMP base no runtime, compared to AMP base which has run time input Section 2.2.1), as well as varying the number of clusters in AMP-pfnn training (Section 2.2.2). Cluster variations include 1, 3 and 8, denoted as AMP-pfnn 1, AMP-pfnn 3 and AMP-pfnn 8, where the 8-cluster model is AMP-pfnn model evaluated in the perceptual test.

We visualized the results for the objective evaluation in Figure 9. AMP-pfnn 8 obtains the highest accuracy overall (mean = 0.62),

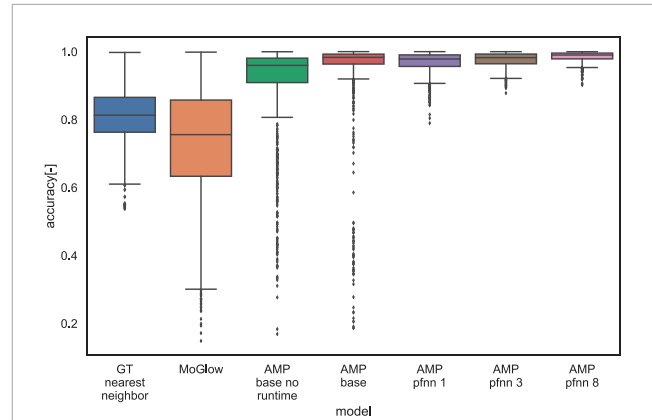


FIGURE 9 Comparison of objective evaluation results of normalized pointing accuracy for the different systems.

TABLE 20 Correlation (Spearman-r) between accuracy and KDE density measure.

Model	Acc corr. KDE	p-value
GT nearest neighbor	-0.118	< 0.001
MoGlow	0.718	< 0.001
AMP base no runtime	-0.020	0.532
AMP base	-0.128	< 0.001
AMP-pfnn 1	-0.043	0.176
AMP-pfnn 3	0.149	< 0.001
AMP-pfnn 8	-0.135	< 0.001

consistent with perceptual evaluation results. The graphs also shows that using 8 clusters slightly outperforms using 1 or 3 clusters and also has less outliers. This further validates the use of clusters (and specifically 8) in AMP-pfnn for learning generalizable accurate pointing. The performance gain from concatenating the run time to the policy network input is clearly visible by comparing “AMP base no runtime” (mean = 0.56) and “AMP base” (mean = 0.59), where the latter has the runtime input. We can also observe that similarly to the perceptual evaluation, MoGlow performs worst amongst the compared systems (mean = 0.46), again showing the limitation of a supervised system for learning communicative gesture. Overall, the results from the objective evaluation show the same performance ranking as the perceptual referential accuracy with actual users, suggesting that our objective evaluation could also be used as a system development tool for pointing gestures in the future.

Since both MoGlow and AMP are systems learned from data, it is interesting to examine how well the models generalize outside training data distribution. We hypothesized that generalization (as measured by pointing accuracy) is correlated with how far away a test pointing position is from training position distribution. We first quantify training position distribution by fitting a multivariate kernel density estimation (KDE) on the training target positions. The bandwidth for KDE is estimated with Scott’s rule and the resulting bandwidth is [0.80,0.36,0.35]. The estimated KDE density function

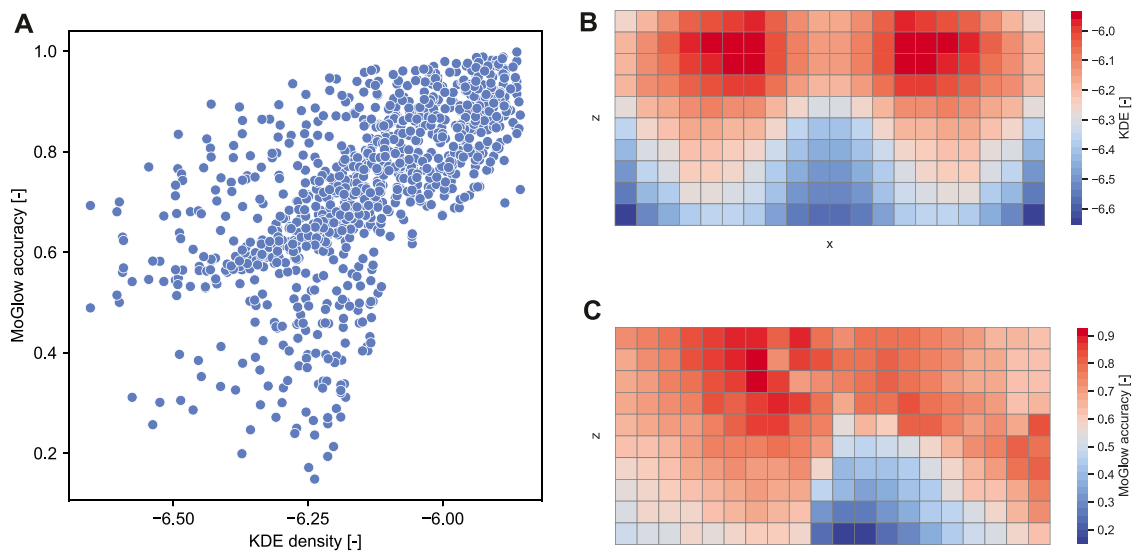


FIGURE 10

Correlation analysis between MoGlow model accuracy and training data density (measured by KDE). **(A)** Shows moderate correlation between the accuracy and the KDE that corresponds to correlation Spearman- $r = 0.718$ ($p < 0.001$) as presented in [Table 2](#). **(B)** and **(C)** show the KDE grid density and MoGlow accuracy respectively condensed to the x - z axis by averaging over y coordinates.

is evaluated over 1,000 test grid points as depicted on [Figure 8](#). The resulting density estimate is then correlated with the corresponding accuracy values for these points in the examined systems; [Table 2](#) shows the results of the correlation analysis. Only MoGlow shows strong correlation between accuracy and data density. We fully visualized the MoGlow data points in [Figure 10A](#), which shows clearly that MoGlow is more accurate in test points with high training data density. We also plotted the KDE in flattened 2D heatmap ([Figure 10B](#)) and MoGlow's accuracy ([Figure 10C](#)), and it is shown that the regions with high training data density (red regions in [Figure 10B](#)) correspond to regions where MoGlow obtained high accuracy (red regions in [Figure 10C](#)). This result is not surprising since MoGlow is a supervised learning system that highly depends on training data distribution. In other systems, GT-*nn*, AMP base, AMP-pfnn 3, and AMP-pfnn 8 have small amount of significant correlation between training position density and model accuracy. Curiously, some of these correlations are negative, meaning that the models are more accurate further away from the training position distribution. This could be due to several factors, such as, KDE may not fully represent data density, or that the pointing position sampling mechanism in AMP models helps generalization outside of training position distribution.

4 Discussion

While perceptual test results show a trend of correlation between motion naturalness and referential accuracy, i.e., models with higher naturalness also obtained higher accuracy, it is still not clear what role does motion dynamics play in referential accuracy. We did not show referential target in motion naturalness test in order to not distract naturalness judgement with some models' inaccurate pointing. But we hypothesize that the motion dynamics itself

contributes to referential accuracy. A factor of motion dynamics that we already noticed could have contributed to referential accuracy is how long was the pointing phase held and how much the hand moved during that phase, i.e., how stable was the pointing phase. One of the baselines, MoGlow, sometimes has long and unstable pointing phase; it points to one position and slowly moves to a nearby position before retraction. Different users reported both choosing the initial position or the ending position as the perceived referent.

However, one user mentioned that this type of unstable and confusing pointing gesture feels quite natural and human-like, as if the avatar is trying to convey that the pointing is not certain. This example suggests a level of communicative capacity in motion style. That is, pointing gesture style, i.e., pointing at the same position with different dynamics, can convey more information to the interlocutor than just what is the referent. It can also convey, for example, degrees of certainty of the made reference. A quick motion with short but steady pointing phase shows certainty. On the other hand, a slow motion with unsteady/constantly moving pointing phase can convey uncertainty. Previous research in cognitive science has shown that speakers tailor the kinematics of their pointing gestures to the communicative needs of the listener, by modulating the speed and duration of the different sub-components of the pointing gesture [Peeters \(2015\)](#) and gesture kinematics has also been exploited in human robot interaction, e.g., to signal hesitation [Moon et al. \(2021\)](#). Future research could explore ways to model pointing styles and uncover what communicative capacity motion style has in pointing gesture.

In our current study we focused solely on the generation of non-verbal mode for referent localization in the 3D space. However, in most cases pointing gestures appear as part of multi-modal expressions, where the non-verbal mode (gesture) and verbal (language, speech) mode carry complementary information. In

these multi-modal expressions pointing gestures are accompanied by verbal descriptors, which can range from simple demonstratives (“this,” “that”) Peeters and Özyürek, (2016) to more complex forms of multi-modal referring expressions describing the referent Clark and Bangerter (2004) (e.g., “that green book on the table, next to the lamp”). Furthermore, since the goal of our study was to create accurate and human-like pointing gestures, we used a simple 3D position target representation as input to the network. In order to make the pointing gesture generation more realistic and flexible, image features could be provided to the agent. This raises further questions about efficient visual representation for pointing and establishing eye-body coordination in embodied intelligent agents Yang et al. (2021).

It should also be noted that the AMP systems achieved higher pointing accuracy (Figure 9) than the ground-truth training data (Figure 2) in terms of the defined pointing reward function Section 2.2. On one hand, this is not surprising, since AMP is trained to maximize the pointing reward, as part of its full reward. On the other hand, it raises the question of how well-suited is the designed reward function for the pointing task. Pointing is an interactive communication process, which can be modeled as a referential game. Referential games are a type of signaling game Lewis (1969), where the sender (gesture producer) sends a signal to the listener (observer), who needs to discover the communicative intent (localize referent) from the message. In recent years language games have been used in emergent communication studies, mostly in text and image based agents Lazaridou et al. (2018), but also in embodied agents Bullard et al. (2020). In these, the sender and the listener jointly learn a communication protocol, often through an RL setup. This could also be relevant for our framework, since modeling the observer as an agent could substitute the current reward function based on simple geometric alignment. However, it is important to note that this kind of reward signal in emergent communication setup is very sparse, compared to our continuously defined dense reward function, which could pose further learning difficulties. Furthermore, current studies within emergent communication have largely focused on the theoretical aspects and the learned communication protocols are not interpretable to humans by design, therefore they have limited use in interactive applications. Extending these emergent communication frameworks with our pointing gesture generation framework is a promising direction in learning more general non-verbal communication in embodied interactive agent.

5 Conclusion

In this paper, we presented the results of learning pointing gestures in a physically simulated embodied agent using imitation and reinforcement learning. We conducted perceptual and objective evaluations of our proposed imitation-RL based method, and showed that it can produce highly natural pointing gestures with high referential accuracy. Moreover, we showed that our approach generalizes better in pointing gesture generation than a state-of-the-art supervised gesture synthesis model. We also presented a novel interactive VR-based method to evaluate pointing gestures in situated embodied interaction.

Data availability statement

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

Ethics statement

Ethical review and approval was not required for the study on human participants in accordance with the local legislation and institutional requirements. Written informed consent for participation was not required for this study in accordance with the national legislation and the institutional requirements.

Author contributions

AD: Data collection, implementations and modeling in AMP, VR experiment design. Contributed to writing. SW: Co-designed/implemented AMP models, MoGlow baseline. Co-designed VR experiment. Contributed to writing. SA Contributed with expert knowledge in data recording and processing, gave feedback to model development and helped out with the VR implementation. He also contributed to the paper writing. JB: Co-PI of the project. Planning and coordination of data collection, strategic discussions on modelling and evaluation experiments, manuscript editing.

Funding

The work is funded by Advanced Adaptive Intelligent Agents project (Digital Futures) and Swedish Research Council, grant no 2018-05409.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frobt.2022.1110534/full#supplementary-material>

References

- Abidi, S., Williams, M., and Johnston, B. (2013). "Human pointing as a robot directive," in 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI), Tokyo, Japan, 03-06 March 2013, 67–68.
- Ahuja, C., Lee, D. W., Ishii, R., and Morency, L.-P. (2020). *No gestures left behind: Learning relationships between spoken language and freeform gestures*. Toronto, Canada: Association for Computational Linguistics, 1884–1895. Proc. EMNLP.
- Alexanderson, S., Henter, G. E., Kucherenko, T., and Beskow, J. (2020). Style-controllable speech-driven gesture synthesis using normalising flows. *Comput. Graph. Forum* 39, 487–496. Wiley Online Library. doi:10.1111/cgf.13946
- Andrist, S., Gleicher, M., and Mutlu, B. (2017). "Looking coordinated: Bidirectional gaze mechanisms for collaborative interaction with virtual characters," in Proceedings of the 2017 CHI conference on human factors in computing systems, Hamburg, Germany, April 23 - 28, 2023, 2571–2582.
- Bailly, G., Raidt, S., and Elisei, F. (2010). Gaze, conversational agents and face-to-face communication. *Speech Commun.* 52, 598–612. doi:10.1016/j.specom.2010.02.015
- Bangerter, A. (2004). Using pointing and describing to achieve joint focus of attention in dialogue. *Psychol. Sci.* 15, 415–419. doi:10.1111/j.0956-7976.2004.00694.x
- Bühler, K. (1934). *Sprachtheorie*. vol. 2 (Jena Fischer).
- Bullard, K., Meier, F., Kiela, D., Pineau, J., and Foerster, J. (2020). *Exploring zero-shot emergent communication in embodied multi-agent populations*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2010.15896.
- Cassell, J., Pelachaud, C., Badler, N., Steedman, M., Achorn, B., Becket, T., et al. (1994). *Animated conversation: Rule-based generation of facial expression, gesture & spoken intonation for multiple conversational agents*. Philadelphia, PA: University of Pennsylvania ScholarlyCommons, 413–420. Proc. SIGGRAPH.
- Cassell, J., Vilhjálmsdóttir, H. H., and Bickmore, T. (2001). *Beat: The behavior expression animation toolkit*. Cambridge, MA: Massachusetts Institute of Technology, 477–486. Proc. SIGGRAPH.
- Chao, F., Wang, Z., Shang, C., Meng, Q., Jiang, M., Zhou, C., et al. (2014). A developmental approach to robotic pointing via human-robot interaction. *Inf. Sci.* 283, 288–303. doi:10.1016/j.ins.2014.03.104
- Chen, Y., Li, Q., Kong, D., Kei, Y. L., Zhu, S.-C., Gao, T., et al. (2021). "Yourefit: Embodied reference understanding with language and gesture," in Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, QC, Canada, 10-17 October 2021, 1385–1395.
- Clark, H. H., and Bangerter, A. (2004). "Changing ideas about reference," in *Experimental pragmatics* (Berlin, Germany: Springer), 25–49.
- Cooperrider, K. (2020). *Fifteen ways of looking at a pointing gesture*. Berlin, Germany: Research Gate.
- Deichler, A., Wang, S., Alexanderson, S., and Beskow, J. (2022). "Towards context-aware human-like pointing gestures with r1 motion imitation," in Context-Awareness in Human-Robot Interaction: Approaches and Challenges, workshop at 2022 ACM/IEEE International Conference on Human-Robot Interaction.
- Diessel, H., and Coventry, K. R. (2020). Demonstratives in spatial language and social interaction: An interdisciplinary review. *Front. Psychol.* 11, 555265. doi:10.3389/fpsyg.2020.555265
- Dinh, L., Krueger, D., and Bengio, Y. (2014). *Nice: Non-linear independent components estimation*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1410.8516.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. (2016). *Density estimation using real nvp*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1605.08803.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., and Abbeel, P. (2016). "Benchmarking deep reinforcement learning for continuous control," in International conference on machine learning (PMLR), 1329–1338.
- Escontrela, A., Peng, X. B., Yu, W., Zhang, T., Iscen, A., Goldberg, K., et al. (2022). *Adversarial motion priors make good substitutes for complex reward functions*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2203.15103.
- Fang, R., Doering, M., and Chai, J. Y. (2015). "Embodied collaborative referring expression generation in situated human-robot interaction," in Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, Portland, OR, USA, 02-05 March 2015, 271–278.
- Ferstl, Y., and McDonnell, R. (2018). "Investigating the use of recurrent motion modelling for speech gesture generation," in Proceedings of the 18th International Conference on Intelligent Virtual Agents, NSW, Sydney, Australia, November 5 - 8, 2018, 93–98.
- Ferstl, Y., Neff, M., and McDonnell, R. (2020). Adversarial gesture generation with realistic gesture phasing. *Comput. Graph.* 89, 117–130. doi:10.1016/j.cag.2020.04.007
- Gao, X., Gao, Q., Gong, R., Lin, K., Thattai, G., and Sukhatme, G. S. (2022). *Dialfred: Dialogue-enabled agents for embodied instruction following*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2202.13330.
- Ginosar, S., Bar, A., Kohavi, G., Chan, C., Owens, A., and Malik, J. (2019). *Learning individual styles of conversational gesture*. Washington D.C: The Computer Vision Foundation. Proc. CVPR.
- Häring, M., Eichberg, J., and André, E. (2012). "Studies on grounding with gaze and pointing gestures in human-robot-interaction," in *International conference on social robotics* (Berlin, Heidelberg: Springer), 378–387.
- Hasegawa, D., Kaneko, N., Shirakawa, S., Sakuta, H., and Sumi, K. (2018). *Evaluation of speech-to-gesture generation using bi-directional LSTM network in IVA '18: International Conference on Intelligent Virtual Agents*, Sydney NSW Australia, November 5 - 8, 2018.
- Heess, N., Tb, D., Sriram, S., Lemmon, J., Merel, J., Wayne, G., et al. (2017). *Emergence of locomotion behaviours in rich environments*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1707.02286.
- Henter, G. E., Alexanderson, S., and Beskow, J. (2020). Moglow: Probabilistic and controllable motion synthesis using normalising flows. *ACM Trans. Graph. (TOG)* 39, 1–14. doi:10.1145/3414685.3417836
- Ho, J., and Ermon, S. (2016). "Generative adversarial imitation learning," in *Advances in neural information processing systems 29* (British Columbia, Canada: NeurIPS Proceedings).
- Holden, D., Komura, T., and Saito, J. (2017). Phase-functioned neural networks for character control. *ACM Trans. Graph. (TOG)* 36, 1–13. doi:10.1145/3072959.3073663
- Holladay, R. M., Dragan, A. D., and Srinivasa, S. S. (2014). "Legible robot pointing," in The 23rd IEEE International Symposium on robot and human interactive communication (IEEE), 217–223.
- Huang, C.-M., and Mutlu, B. (2014). "Learning-based modeling of multimodal behaviors for humanlike robots," in Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction, Bielefeld, Germany, 03-06 March 2014, 57–64.
- Islam, M. M., Mirzaiee, R. M., Gladstone, A., Green, H. N., and Iqbal, T. (2022). "Caesar: An embodied simulator for generating multimodal referring expression datasets," in Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track.
- Kibrik, A. A. (2011). *Reference in discourse*. Oxford, United Kingdom: Oxford University Press. doi:10.1093/acprof:oso/9780199215805.001.0001
- Kingma, D. P., and Dhariwal, P. (2018). "Glow: Generative flow with invertible 1x1 convolutions," in *Advances in neural information processing systems 31* (British Columbia, Canada: NeurIPS Proceedings).
- Kita, S. (2003). "Pointing: A foundational building block of human communication," in *Pointing* (London, England: Psychology Press), 9–16.
- Kopp, S., and Wachsmuth, I. (2004). Synthesizing multimodal utterances for conversational agents. *Comput. Animat. Virtual Worlds* 15, 39–52. doi:10.1002/cav.6
- Korzun, V., Dimov, I., and Zharkov, A. (2020). "The finemoto entry to the genea challenge 2020," in *Proc. GENE workshop* (Glasgow, UK: Zenodo). doi:10.5281/zenodo.4088609
- Kucherenko, T., Hasegawa, D., Kaneko, N., Henter, G. E., and Kjellström, H. (2021a). Moving fast and slow: Analysis of representations and post-processing in speech-driven automatic gesture generation. *Int. J. Human-Computer Interact.* 37, 1300–1316. doi:10.1080/10447318.2021.1883883
- Kucherenko, T., Jonell, P., van Waveren, S., Henter, G. E., Alexanderson, S., Leite, I., et al. (2020). "Gesticulator: A framework for semantically-aware speech-driven gesture generation," in Proceedings of the 2020 International Conference on Multimodal Interaction, Virtual Event, Netherlands, October 25 - 29, 2020, 242–250.
- Kucherenko, T., Jonell, P., Yoon, Y., Wolfert, P., and Henter, G. E. (2021b). "A large, crowdsourced evaluation of gesture generation systems on common data: The genea challenge 2020," in 26th international conference on intelligent user interfaces, TX, College Station, USA, April 14 - 17, 2021, 11–21.
- Lazaridou, A., Hermann, K. M., Tuyls, K., and Clark, S. (2018). "Emergence of linguistic communication from referential games with symbolic and pixel input," in *Iclr* (Ithaca, NY: Cornell University).
- Lester, J. C., Voerman, J. L., Towns, S. G., and Callaway, C. B. (1999). Deictic believability: Coordinated gesture, locomotion, and speech in lifelike pedagogical agents. *Appl. Artif. Intell.* 13, 383–414. doi:10.1080/088395199117324
- Lewis, D. K. (1969). *Convention: A philosophical study*. Cambridge, MA, USA: Wiley-Blackwell.
- Liu, Y., Mohammadi, G., Song, Y., and Johal, W. (2021). "Speech-based gesture generation for robots and embodied agents: A scoping review," in Proceedings of the 9th International Conference on Human-Agent Interaction, Virtual Event, Japan, November 9 - 11, 2021, 31–38.
- Lücking, A., Pfeiffer, T., and Rieser, H. (2015). Pointing and reference reconsidered. *J. Pragmat.* 77, 56–79. doi:10.1016/j.pragma.2014.12.013
- Luketina, J., Nardelli, N., Farquhar, G., Foerster, J., Andreas, J., Grefenstette, E., et al. (2019). *A survey of reinforcement learning informed by natural language*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1906.03926.

- Lynch, C., Wahid, A., Tompson, J., Ding, T., Betker, J., Baruch, R., et al. (2022). *Interactive language: Talking to robots in real time*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2210.06407.
- Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., et al. (2021). *Isaac gym: High performance gpu-based physics simulation for robot learning*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2108.10470.
- Marsella, S., Xu, Y., Lhommet, M., Feng, A., Scherer, S., and Shapiro, A. (2013). *Virtual character performance from speech*. Eindhoven: Eurographics Association, 25–35. Proc. SCA.
- Matuszek, C., Bo, L., Zettlemoyer, L., and Fox, D. (2014). Learning from unscripted deictic gesture and language for human-robot interactions. *Proc. AAAI Conf. Artif. Intell.* 28 (1), 9051. doi:10.1609/aaai.v28i1.9051
- McNeill, D. (1992). *Hand and mind 1*. Chicago: University of Chicago press. *Advances in Visual Semiotics* 351.
- Merel, J., Tassa, Y., Tb, D., Srinivasan, S., Lemmon, J., Wang, Z., et al. (2017). *Learning human behaviors from motion capture by adversarial imitation*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1707.02201.
- Moon, A., Hashmi, M., Loos, H. M. V. D., Croft, E. A., and Billard, A. (2021). Design of hesitation gestures for nonverbal human-robot negotiation of conflicts. *ACM Trans. Human-Robot Interact. (THRI)* 10, 1–25. doi:10.1145/3418302
- Mutlu, B., Yamaoka, F., Kanda, T., Ishiguro, H., and Hagita, N. (2009). “Nonverbal leakage in robots: Communication of intentions through seemingly unintentional behavior,” in Proceedings of the 4th ACM/IEEE international conference on Human robot interaction, La Jolla, CA, USA, 11–13 March 2009, 69–76.
- Ng-Thow-Hing, V., Luo, P., and Okita, S. (2010). “Synchronized gesture and speech production for humanoid robots,” in 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, Taipei, Taiwan, 18–22 October 2010.
- Noma, T., Zhao, L., and Badler, N. I. (2000). Design of a virtual human presenter. *IEEE Comput. Graph. Appl.* 20, 79–85. doi:10.1109/38.851755
- Padmakumar, A., Thomason, J., Shrivastava, A., Lange, P., Narayan-Chen, A., Gella, S., et al. (2022). Teach: Task-driven embodied agents that chat. *Proc. AAAI Conf. Artif. Intell.* 36, 2017–2025. doi:10.1609/aaai.v36i2.20097
- Peeters, D. (2015). *A social and neurobiological approach to pointing in speech and gesture*. Nijmegen, Netherlands: Radboud University Nijmegen. Ph.D. thesis.
- Peeters, D., Krahmer, E., and Maes, A. (2021). A conceptual framework for the study of demonstrative reference. *Psychonomic Bull. Rev.* 28, 409–433. doi:10.3758/s13423-020-01822-8
- Peeters, D., and Özyürek, A. (2016). This and that revisited: A social and multimodal approach to spatial demonstratives. *Front. Psychol.* 7, 222. doi:10.3389/fpsyg.2016.00222
- Peng, X. B., Abbeel, P., Levine, S., and Van de Panne, M. (2018). Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph. (TOG)* 37, 1–14. doi:10.1145/3197517.3201311
- Peng, X. B., Ma, Z., Abbeel, P., Levine, S., and Kanazawa, A. (2021). Amp: Adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph. (TOG)* 40, 1–20. doi:10.1145/3476576.3476723
- Qi, Y., Wu, Q., Anderson, P., Wang, X., Wang, W. Y., Shen, C., et al. (2020). “Reverie: Remote embodied visual referring expression in real indoor environments,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020, 9982–9991.
- Rickel, J., and Johnson, W. L. (1999). Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Appl. Artif. Intell.* 13, 343–382. doi:10.1080/088395199117315
- Sauppe, A., and Mutlu, B. (2014). “Robot deictics: How gesture and context shape referential communication,” in 2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI), Bielefeld, Germany, 03–06 March 2014, 342–349.
- Schauerte, B., and Fink, G. A. (2010). “Focusing computational visual attention in multi-modal human-robot interaction,” in International conference on multimodal interfaces and the workshop on machine learning for multimodal interaction, Beijing, China, November 8 – 10, 2010, 1–8.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). *Proximal policy optimization algorithms*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1707.06347.
- Shukla, D., Erkent, O., and Piater, J. (2015). “Probabilistic detection of pointing directions for human-robot interaction,” in 2015 international conference on digital image computing: techniques and applications (DICTA, Adelaide, SA, Australia, 23–25 November. 1–8.
- Steels, L. (2001). Language games for autonomous robots. *IEEE Intell. Syst.* 16, 16–22. doi:10.1109/mis.2001.956077
- Stepputtis, S., Campbell, J., Phielipp, M., Baral, C., and Amor, H. B. (2019). *Imitation learning of robot policies by combining language, vision and demonstration*. Ithaca, NY: Cornell University. arXiv preprint arXiv:1911.11744.
- Sugiyama, O., Kanda, T., Imai, M., Ishiguro, H., and Hagita, N. (2007). “Natural deictic communication with humanoid robots,” in 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems, San Diego, CA, USA, 29 October 2007 – 02 November 2007, 1441–1448.
- Talmy, L. (2020). Targeting in language: Unifying deixis and anaphora. *Front. Psychol.* 11, 2016. doi:10.3389/fpsyg.2020.02016
- Team, D. I. A., Abramson, J., Ahuja, A., Brussee, A., Carnevale, F., Cassin, M., et al. (2021). *Creating multimodal interactive agents with imitation and self-supervised learning*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2112.03763.
- Valle-Pérez, G., Henter, G. E., Beskow, J., Holzapfel, A., Oudeyer, P.-Y., and Alexanderson, S. (2021). Transflower: Probabilistic autoregressive dance generation with multimodal attention. *ACM Trans. Graph. (TOG)* 40, 1–14. doi:10.1145/3478513.3480570
- Vollenweider, E., Bjelonic, M., Klemm, V., Rudin, N., Lee, J., and Hutter, M. (2022). *Advanced skills through multiple adversarial motion priors in reinforcement learning*. Ithaca, NY: Cornell University. arXiv preprint arXiv:2203.14912.
- Wu, Q., Wu, C.-J., Zhu, Y., and Joo, J. (2021). “Communicative learning with natural gestures for embodied navigation agents with human-in-the-scene,” in 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Prague, Czech, 27 September 2021 – 01 October 2021, 4095–4102.
- Yang, R., Zhang, M., Hansen, N., Xu, H., and Wang, X. (2021). “Learning vision-guided quadrupedal locomotion end-to-end with cross-modal transformers,” in International Conference on Learning Representations.
- Yoon, Y., Cha, B., Lee, J.-H., Jang, M., Lee, J., Kim, J., et al. (2020). Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Trans. Graph.* 39, 1–16. doi:10.1145/3414685.3417838
- Zabala, U., Rodriguez, I., and Lazkano, E. (2022). “Towards an automatic generation of natural gestures for a storyteller robot,” in In 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), Napoli, Italy, 29 August 2022 – 02 September 2022, 1209–1215.