



 Latest updates: <https://dl.acm.org/doi/10.1145/3610977.3634999>

RESEARCH-ARTICLE

Generative Expressive Robot Behaviors using Large Language Models

KARTHIK MAHADEVAN, University of Toronto, Toronto, ON, Canada

JONATHAN CHIEN, Google LLC, Mountain View, CA, United States

NOAH BROWN, Google LLC, Mountain View, CA, United States

ZHUO XU, Google LLC, Mountain View, CA, United States

CAROLINA PARADA, Google LLC, Mountain View, CA, United States

FEI XIA, Google LLC, Mountain View, CA, United States

[View all](#)

Open Access Support provided by:

Google LLC

University of Toronto



PDF Download
3610977.3634999.pdf
20 December 2025
Total Citations: 27
Total Downloads: 3583



Published: 11 March 2024

[Citation in BibTeX format](#)

HRI '24: ACM/IEEE International Conference on Human-Robot Interaction
March 11 - 15, 2024
CO, Boulder, USA

Conference Sponsors:

SIGCHI

SIGAI

Generative Expressive Robot Behaviors using Large Language Models

Karthik Mahadevan
University of Toronto
karthikm@dgp.toronto.edu

Zhuo Xu
Google Deepmind
zhuoxu@google.com

Andy Zeng
Google Deepmind
andyzeng@google.com

Jonathan Chien
Google Deepmind
chienj@google.com

Carolina Parada
Google Deepmind
carolinap@google.com

Leila Takayama
Hoku Labs
takayama@hokulabs.com

Noah Brown
Google Deepmind
noahbrown@google.com

Fei Xia
Google Deepmind
xafei@google.com

Dorsa Sadigh
Google Deepmind
dorsas@google.com

ABSTRACT

People employ expressive behaviors to effectively communicate and coordinate their actions with others, such as nodding to acknowledge a person glancing at them or saying “excuse me” to pass people in a busy corridor. We would like robots to also demonstrate expressive behaviors in human-robot interaction. Prior work proposes rule-based methods that struggle to scale to new communication modalities or social situations, while data-driven methods require specialized datasets for each social situation the robot is used in. We propose to leverage the rich social context available from large language models (LLMs) and their ability to generate motion based on instructions or user preferences, to generate *expressive robot motion* that is *adaptable* and *composable*, building upon each other. Our approach utilizes few-shot chain-of-thought prompting to translate human language instructions into parametrized control code using the robot’s available and learned skills. Through user studies and simulation experiments, we demonstrate that our approach produces behaviors that users found to be competent and easy to understand. Supplementary material can be found at <https://generative-expressive-motion.github.io/>.

CCS CONCEPTS

- Computing methodologies → *Online learning settings*.

KEYWORDS

Generative expressive robot behaviors, in-context learning, language corrections

ACM Reference Format:

Karthik Mahadevan, Jonathan Chien, Noah Brown, Zhuo Xu, Carolina Parada, Fei Xia, Andy Zeng, Leila Takayama, and Dorsa Sadigh. 2024. Generative Expressive Robot Behaviors using Large Language Models. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24)*, March 11–14, 2024, Boulder, CO, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3610977.3634999>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HRI '24, March 11–14, 2024, Boulder, CO, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0322-5/24/03.

<https://doi.org/10.1145/3610977.3634999>

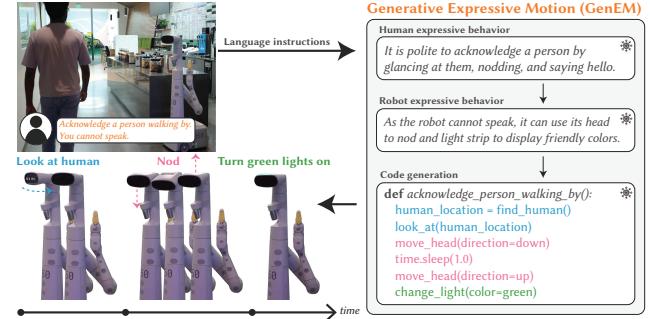


Figure 1: We present Generative Expressive Motion (GenEM), a new approach to autonomously generate expressive robot behaviors. GenEM takes a desired expressive behavior (or a social context) as language instructions, reasons about human social norms, and generates control code for a robot using pre-existing robot skills and learned expressive behaviors. Iterative feedback can quickly modify the behavior according to user preferences. Here, the * symbols denote frozen large language models.

Interaction (HRI '24), March 11–14, 2024, Boulder, CO, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3610977.3634999>

1 INTRODUCTION

People employ a wide range of expressive behaviors to effectively interact with others on a daily basis. For instance, a person walking by an acquaintance may briefly glance at them and nod to acknowledge their presence. A person might apologetically say, “excuse me!” to squeeze through a tight hallway, where a group of people are conversing. In much the same manner, we would like robots to also demonstrate expressive behaviors when interacting with people. Robots that don’t have expressive capabilities will need to re-plan their paths to avoid the crowded hallway. On the other hand, robots that have expressive capabilities might actually be able to persuade the group of people to make room for them to squeeze by, thereby improving the robot’s efficiency in getting its job done.

Prior work has demonstrated the value of expressive robot behaviors, and explored approaches for generating behaviors for various

purposes and contexts, including general-purpose use [8], manipulation settings, where transparency is important [21], and everyday scenarios where social norms must be observed (such as interacting with a receptionist) [36]. Approaches can be rule- or template-based [2, 7, 33], which often rely on a rigid template or a set of rules to generate behaviors. This often leads to robot behaviors that can be expressive, but do not scale to new modalities or variations of human preferences. On the other hand, data-driven techniques offer the promise of flexibility and the ability to adapt to variations. Prior work have studied data-driven techniques that generate expressive motion [42], but these methods also have their shortcomings as they often need specialized datasets for each social interaction where a particular behavior is used (e.g., for affective robot movements [41, 42]).

Our goal is to enable robots to generate expressive behavior that is flexible: behaviors that can *adapt* to different human preferences, and be *composed* of simpler behaviors. Recent work show that large language models (LLMs) can synthesize code to control virtual [44] and embodied agents [25, 39], help design reward functions [22, 48], enable social and common-sense reasoning [20], or perform control and sequential decision making tasks through in-context learning [10, 29, 30] by providing a sequence of desirable inputs, and outputs in the prompt. Our key insight is to tap into the rich social context available from LLMs to generate adaptable and composable expressive behavior. For instance, an LLM has enough context to realize that it is polite to make an eye contact when greeting someone. In addition, LLMs enable the use of corrective language such as “bend your arm a bit more!” and the ability to generate motion in response to such instructions. This makes LLMs a useful framework for autonomously generating expressive behavior that flexibly respond to and learn from human feedback in human-robot interaction settings.

Leveraging the power and flexibility provided by LLMs, we propose a new approach, Generative Expressive Motion (GenEM), for autonomously generating expressive robot behaviors. GenEM uses few-shot prompting and takes a desired expressive behavior (or a social context) as language instructions, performs social reasoning (akin to chain-of-thought [45]), and finally generates control code for a robot using available robot APIs. GenEM can produce multimodal behaviors that utilize the robot’s available affordances (e.g., speech, body movement, and other visual features such as light strips) to effectively express the robot’s intent. One of the key benefits of GenEM is that it responds to live human feedback – adapting to iterative corrections and generating new expressive behaviors by composing the existing ones.

In a set of online user studies, we compared behaviors generated on a mobile robot using two variations of GenEM, with and without user feedback (a non-expert in HRI behavior design), to a set of behaviors designed by a professional character animator (or the *oracle animator*). We show that behaviors generated by GenEM and further adapted with user feedback were positively perceived by users, and in some cases better perceived than the oracle behaviors.

In additional experiments with the mobile robot and a simulated quadruped, we show that GenEM: (1) performs better than a version where language instructions are directly translated into code, (2) allows for the generation of behaviors that are agnostic to embodiment, (3) allows for the generation of composable behaviors

that build on simpler expressive behaviors, and finally, (4) adapt to different types of user feedback.

2 RELATED WORK

Expressive Behavior Generation. Researchers have made significant efforts towards generating socially acceptable behavior for both robots and virtual humans. These can largely categorized into *rule-based*, *template-based*, and *data-driven* [33] behavior generation approaches. We define rule-based approaches as those that require a formalized set of rules and operations (typically provided by a person) which are used to generate subsequent robot behavior.

Rule-based approaches enable behavior generation through formalized sets of rules and operations [2]. Some methods include interfaces that lets users manually specify interaction rules and logic [4, 5, 23, 24, 35]. Other methods work by observing and modelling humans [3, 13, 14, 18]. Despite their use, rule-based approaches face several issues, including limited expressivity in the generated behavior due to the requirement of formal rules, and the reduced ability to produce multimodal behaviors as the number of modalities increases [33]. Template-based methods formulate generic templates for interaction by learning from traces of interaction data [7, 11]. Templates can translate few examples of human traces into reusable programs through program synthesis [19, 34]. Traces can be collected by observing humans interacting [34, 36], or through approaches such as sketching [37] or tangibles on a tabletop [38]. Overall, prior rule- and template-based methods enforce strong constraints to enable behavior generation but are limited in their expressivity. In contrast, GenEM enables increased expressivity in the initial behavior generation as well as iterative improvements through live user feedback.

On the other hand, data-driven approaches produce behaviors using models trained on data. Some methods learn *interaction logic* through data and use this to produce multimodal behaviors via classical machine learning methods [9, 15, 27]. Other methods train on hand-crafted examples through generative models [28, 42]. For instance, predicting when to use backchanneling behaviors (i.e., providing feedback during conversation such as by nodding) has been learned through batch reinforcement learning [17] and recurrent neural networks [31]. Lastly, recent work has investigated how to learn cost functions for a target emotion from user feedback [49], or even learn an emotive latent space to model many emotions [40]. However, these approaches are data inefficient and require specialized datasets per behavior to be generated, while GenEM is able to produce a variety of expressive behaviors with a few examples through in-context learning.

LLMs for Robot Planning and Control. Recent work has achieved great success by leveraging LLMs in downstream robotics tasks specifically by providing sequences of desirable input-output pairs in context [10, 29, 30]. In addition, LLMs have been used for long-horizon task planning [1, 26], and can react to environmental and human feedback [16]. LLMs have been leveraged for designing reward functions for training reinforcement learning agents [22, 48]. Research has also shown that LLMs can enable social and common-sense reasoning [20] as well as infer user preferences by summarizing interactions with humans [47]. Most relevant to our approach are prior work where LLMs synthesize code to control virtual [44]

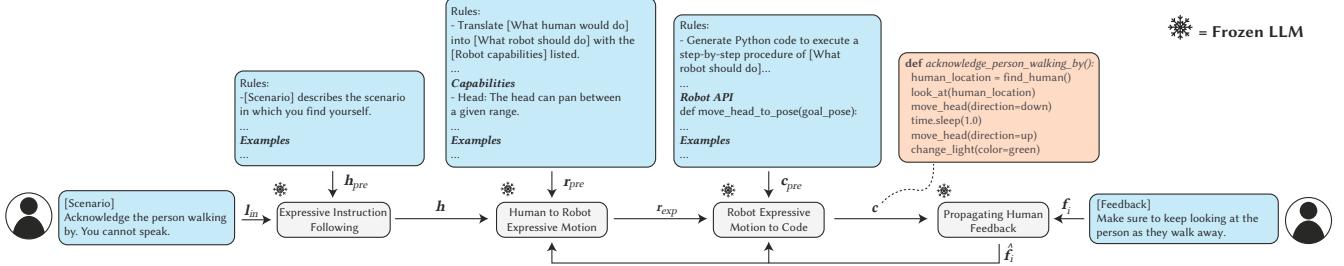


Figure 2: Generative Expressive Motion. Given a language instruction l_{in} , the *Expressive Instruction Following* module reasons about the social norms and outputs how a human might express this behavior (h). This is translated into a procedure for robot expressive behavior using a prompt describing the robot’s pre-existing capabilities (r_{pre}) and any learned expressive behaviors. Then, the procedure is used to generate parametrized robot code c that can be executed. The user can provide iterative feedback f_i on the behavior which is processed to determine whether to re-run the robot behavior module first followed by the code generation module or just the code generation module. Note: * shown on top of all the gray modules denotes them as frozen LLMs.

and robotic agents [25, 39] by using existing APIs to compose more complex robot behavior as programs. We are also encouraged by work demonstrating that language can be used to correct robot manipulation behaviors online [6]. Taken together, we propose to leverage the rich social context available from LLMs, and their ability to adapt to user instructions, to generate expressive robot behaviors. To our knowledge, LLMs have not previously been used to generate expressive robot behaviors that adapt to user feedback.

3 GENERATIVE EXPRESSIVE MOTION

Problem Statement. We aim to tackle the problem of expressive behavior generation that is both adaptive to user feedback and composable so that more complex behaviors can build on simpler behaviors. Formally, we define being *expressive* as the distance between some expert expressive trajectory that could be generated by an animator (or demonstrated) τ_{expert} and a robot trajectory τ . $dist(\tau, \tau_{expert})$ can be any desirable distance metric between the two trajectories, e.g., dynamic time warping (DTW). GenEM aims to minimize this distance $d^* = \min dist(\tau, \tau_{expert})$.

Our approach (Figure 2) uses several LLMs in a modular fashion so that each *LLM agent* plays a distinct role. Later, we demonstrate through experiments that a modular approach yields better quality of behaviors compared to an end-to-end approach. GenEM takes user language instructions $l_{in} \in L$ as input and outputs a robot policy π_θ , which is in the form of a parameterized code. Human iterative feedback $f_i \in L$ can be used to update the policy π_θ . The policy parameters get updated one step at a time given the feedback f_i , where $i \in \{1, \dots, K\}$. The policy can be instantiated from some initial state $s_0 \in S$ to produce trajectories $\tau = \{s_0, a_0, \dots, a_{N-1}, s_N\}$ or instantiations of expressive robot behavior. Below we describe one sample iteration with human feedback f_i . Please refer to **Appendix A** for full prompts.

Expressive Instruction Following. The input to our approach is a language instruction $l_{in} \in L$, which can either be a description of a social context where the robot needs to perform an expressive behavior by following social norms (e.g., "A person walking by waves at you.") or an instruction that describing an expressive behavior to be generated (e.g., "Nod your head"). The input prompt is of the form $u = [h_{pre}, l_{in}]$ where h_{pre} is the prompt prefix that adds

context about the role of the LLM and includes few-shot examples. The output of the LLM call is a string of the form $h = [h_{cot}, h_{exp}]$ consisting of Chain-of-Thought reasoning h_{cot} [45] and the human expressive motion h_{exp} in response to the instruction. For example, for $l_{in} = "Acknowledge a person walking by. You cannot speak."$, the *Expressive Instruction Following* module would output $h_{exp} = Make eye contact with the person. Smile or nod to acknowledge their presence.$ Examples of h_{cot} could be: *"The person is passing by and it's polite to acknowledge their presence. Since I cannot speak, I need to use non-verbal communication. A nod or a smile is a universal sign of acknowledgement."*

From Human Expressive Motion to Robot Expressive Motion.

In the next step, we use an LLM to translate human expressive motion h to robot expressive motion r . The prompt takes the form $u = [r_{pre}, l_{in}, h, r_{i-1, opt}, f_{i-1, opt}]$ where r_{pre} is the prompt prefix setting context for the LLM, contains few-shot examples, and describes the robot’s capabilities some of which are pre-defined (e.g., the ability to speak or move its head) and others which are learned from previous interactions (e.g., nodding or approaching a person). Optionally, the prompt can include the response from a previous step r_{i-1} and response to user iterative feedback from a previous step f_{i-1} . The output is of the form $r = [r_{cot}, r_{exp}]$ consisting of the LLM’s reasoning and the procedure to create expressive robot motion. An example response r_{exp} could include: *"1) Use the head’s pan and tilt capabilities to face the person who is walking by. 2) Use the light strip to display a pre-programmed pattern that mimics a smile or nod."*. An example of r_{cot} could be: *"The robot can use its head’s pan and tilt capabilities to make "eye contact" with the person. The robot can use its light strip to mimic a smile or nod."*

Translating Robot Expressive Motion to Code. In the following step, we use an LLM to translate the step-by-step procedure of how to produce expressive robot motion into executable code. We propose a skill library in a similar fashion to that of Voyager [44] containing existing robot skill primitives, and parametrized robot code π_θ representing previously learned expressive motions. To facilitate this, the prompt encourages modular code generation by providing examples where small, reusable functions with docstrings and named arguments are used to generate more complex functions that describe an expressive behavior. To generate code, the prompt to

the LLM takes the form $u = [c_{pre}, l_{in}, h_{exp}, r_{exp,i-1_{opt}}, c_{i-1_{opt}}, \hat{f}_{i-1}, r_{exp}]$. Here, c_{pre} provides context about its role as a code generating agent to the LLM, includes the robot's current skill library, and contains few-shot examples. Optionally, the expressive robot motion $r_{exp,i-1}$, and code c_{i-1} from a previous step can be provided as well as LLM output \hat{f}_{i-1} responding to the user feedback f_{i-1} . The output c is parametrized robot code representing the policy π_θ for the expressive behavior (see Figure 2 for sample output). Later, the generated code can be incorporated into the robot's skill library to utilize in future expressive behavior generations.

Propagating Human Feedback. In the final (optional) step, we use an LLM to update the generated expressive behavior in response to human feedback f_i if the user is not satisfied with the generated behavior. The prompt is of the form $u = [f_{pre}, l_{in}, r_{exp}, c, f_i]$, where f_{pre} provides context to LLM, and includes both the procedure for expressive robot motion r_{exp} and the generated code c . The output is of the form $f = [f_{cot}, \hat{f}_i]$ and includes the LLM's reasoning and the changes \hat{f}_i needed to improve the current expressive motion based on human feedback. The output also classifies whether the changes require an iterative call to modify the procedure for generating the robot's expressive behavior r and then translating it to code c , or just modifying the generated code c .

For example, the user could state $f_i = "When you first see the person, nod at them."$, and the output \hat{f}_i could be: "*[Change: What robot should do]...As soon as the robot sees the person, it should nod at them. After nodding, the robot can use its light strip to display a pre-programmed pattern that mimics a smile or nod...*". As an example, f_{cot} could state: "*The feedback suggests that the robot's action of acknowledging the person was not correct. This implies that the robot should nod at the person when it first sees them.*"

4 USER STUDIES

We conducted two user studies to assess whether our approach, GenEM, can be used to generate expressive behaviors that are perceivable by people. We generated two versions of behaviors: *GenEM*, and *GenEM* with iterative Feedback (or *GenEM++*). In both studies, all comparisons were made against behaviors designed by a professional animator and implemented by a software developer, which we term the *oracle animator*. In the *first study*, our goal was to assess whether behaviors that are generated using *GenEM* and *GenEM++* would be perceived similarly to the behaviors created using the *oracle animator*. In the *second study*, we attempted to generate behaviors using *GenEM* and *GenEM++* that were similar to the behaviors created using the *oracle animator*. Both studies aim to demonstrate that our approach is *adaptable* to human feedback.

Behaviors. All behaviors were generated on a mobile robot platform (please see website ¹ for full clips). The robot has several capabilities that can be used to generate behaviors through existing APIs, including a head that can pan and tilt, a base that can translate, rotate, and navigate from point to point, a light strip that can display different colors and patterns, and finally, a speech module that can generate utterances and nonverbal effects. To enable the comparison of behaviors produced in the three conditions – *oracle animator*, *GenEM*, and *GenEM++*, we recorded video clips of each

behavior (see Figure 3). To ensure consistency across conditions, behaviors in each condition were recorded in the same physical locations under similar lighting conditions. The *GenEM* and *GenEM++* behaviors were generated by sampling OpenAI's GPT-4 APIs for text completion [32] (gpt-4-0613) with the temperature set to 0.

Study Procedure. After providing informed consent, participants completed an online survey to evaluate the robot's expressive behaviors in both studies. The survey is divided into three sections (one per behavior condition) and clips within each condition randomly appeared. To minimize ordering effects, a Balanced Latin Square design (3×3) was used. For each behavior in each condition, participants watched an unlabeled video clip ¹, and then answered questions. All participants received remuneration after the study.

Measures. In both studies, participants completed a survey to assess each behavior, answering three 7-point Likert scale questions assessing their confidence on their understanding of the behavior, the difficulty in understanding what the robot is doing, and the competency of the robot's behavior. Participants also provided an open-ended response describing what behavior they *believed* the robot was attempting to express.

Analysis. One-way repeated-measures ANOVA were performed on the data with post-hoc pairwise comparisons where there were significant differences with Bonferroni corrections applied. When reporting comparisons between conditions, we define *instances* as pairwise significant conditions for at least one of the three Likert-scale questions asked about a behavior.

4.1 Study 1: Benchmarking Generative Expressive Motion

To determine whether our approach produces expressive behaviors that people can perceive, we conducted a within-subjects user study with thirty participants (16 women, 14 men), aged 18 to 60 (18-25: 3, 26-30: 9, 31-40: 9, 41-50: 7, 51-60: 2). One participant did not complete the entire survey and their data was omitted.

Behaviors. We generated ten expressive behaviors (see Figure 3) ranging in complexity: *Nod*, *shake head (Shake)*, *wake up (Wake)*, *excuse me (Excuse)*, *recoverable mistake (Recoverable)*, *unrecoverable mistake (Unrecoverable)*, *acknowledge person walking by (Acknowledge)*, *follow person (Follow)*, *approach person (Approach)* and *pay attention to person (Attention)*. The input included a one-line instruction (e.g., *Respond to a person saying, "Come here. You cannot speak."*).

Conditions. The *oracle animator* condition consisted of professionally animated behaviors that were implemented on the robot through scripting. To create the *GenEM* behaviors, we sampled our approach five times to generate five versions of each behavior. Since the behaviors were sampled with a temperature of 0, they shared significant overlap with small variations amongst them (due to non-determinism in GPT-4 output; please see **Appendix C** for samples generated using the same prompt). Then, six participants experienced in working with the robot were asked to rank them. The best variation for each behavior was included as part of the *GenEM* behaviors. To generate the *GenEM++* behaviors, we recruited one participant experienced in using the robot (but inexperienced in HRI behavior design) and asked them to provide feedback on the best rated version of each behavior. Feedback was used to iteratively

¹<https://generative-expressive-motion.github.io/>

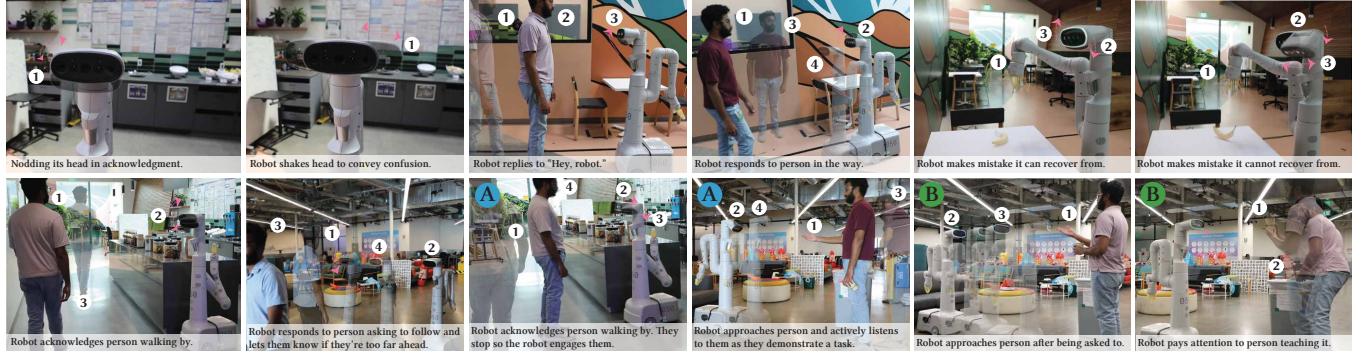


Figure 3: Behaviors tested in the two user studies where the behaviors labelled in green denote those unique to the first study and behaviors labelled in blue denote those unique to the second study. The remaining behaviors (8) were common among the two studies.

modify the expressive behavior until the participant was satisfied with the result, or upon reaching the maximum number of feedback rounds ($n = 10$). We note that although participants rated the behaviors in the studies, the behavior generation is personalized to the user who provided the initial feedback, which may not reflect the preferences of all potential users (e.g., study participants).

Hypotheses. We hypothesized that the perception of the *GenEM++* behaviors would not differ significantly from the oracle animator behaviors (**H1**). We also hypothesized that the *GenEM* behaviors would be less well-received compared to the *GenEM++* and the oracle animator behaviors (**H2**).

Quantitative Findings. Figure 4 summarizes participants' responses to the survey questions for each behavior. The results show that the *GenEM++* behaviors were worse than the oracle animator behaviors in 2/10 instances (*Shake* and *Follow*). In contrast, the *GenEM++* behaviors received higher scores than the oracle animator behaviors in 2/10 instances (*Excuse* and *Approach*). Hence, **H1** is supported by our data – the *GenEM++* behaviors were well received and the oracle animator behaviors were not significantly better received than the *GenEM++* behaviors.

The *GenEM* behaviors were worse received compared to the oracle animator behaviors in 2/10 instances (*Acknowledge Walk* and *Follow*) whereas the *GenEM* behaviors were better received than the oracle animator behaviors in 2/10 instances (*Excuse* and *Approach*). This was surprising because user feedback was not incorporated into the behavior generation in this condition. Besides 1/10 instances (*Shake*), there were no significant differences in the perceptions of the *GenEM* and *GenEM++* behaviors. Hence, we did not find support for **H2**. We performed equivalence tests (equivalence bound: ± 0.5 Likert points) but did not find any sets of behaviors to be equivalent. Overall, the results support the finding that *GenEM* (even with an untrained user providing feedback) produces expressive robot behaviors that users found to be competent and easy to understand.

4.2 Study 2: Mimicking the Oracle Animator

We conducted an additional within-subjects user study with twenty-four participants (21 men, 2 women, 1 prefer not to say), aged 18-60 (18-25: 4, 26-30: 3, 31-40: 12, 41-50: 4, 51-60: 1) to assess

whether using *GenEM* to generate behaviors that resembled the oracle animator would be perceived differently. One participant did not complete the entire survey and their data was omitted.

Behaviors. We generated ten expressive behaviors ranging in complexity, with eight overlapping² behaviors from the first study (see Figure 3): nod (*Nod*), shake head (*Shake*), wake up (*Wake*), excuse me (*Excuse*), recoverable mistake (*Recoverable*), unrecoverable mistake (*Unrecoverable*), acknowledge person walking by (*Acknowledge Walking*), acknowledge person stopping by (*Acknowledge Stop*), follow person (*Follow*), and teaching session (*Teach*). Behaviors that were different from the first study were chosen to add further complexity – e.g., longer single-turn interactions such as *teaching*, that started with a person walking up a robot, teaching it a lesson, and lastly the robot acknowledging that it understood the person's instructions. Unlike in the first study, the prompts were more varied and sometimes included additional descriptions such as for the more complex behaviors (see **Appendix B** for full prompts for each behavior). To generate each *GenEM* behavior, we sampled our approach ten times after which an experimenter selected the version that appeared most similar to the equivalent oracle animator behavior when deployed on the robot. To create each *GenEM++* behavior, an experimenter refined the *GenEM* behavior through iterative feedback until it appeared similar to the equivalent oracle animator behavior or after exceeding the maximum number of feedback rounds ($n = 10$)¹.

Hypotheses. We hypothesized that user perceptions of the *GenEM++* behaviors would not significantly differ when compared to the oracle animator behaviors (**H3**). We also suppose that the behaviors in the *GenEM* condition would be perceived as worse than the *GenEM++* and oracle animator behaviors (**H4**).

Quantitative Findings. The results of the study are summarized in Figure 5. They show that the *GenEM++* behaviors were worse received than the oracle animator behaviors in 2/10 instances (*Acknowledge Walk* and *Follow*) whereas the *GenEM++* behaviors were more positively received than the oracle animator in 2/10 instances (*Excuse* and *Teach*). Hence, our hypothesis is supported by the data

²Some behaviors in the second study differ from the first study as they are too complex to express as a single line instruction which we maintained for consistency in the first study. Instead, in the first study, these complex behaviors were broken down into simpler behaviors (e.g., teaching is equivalent to approaching and paying attention).

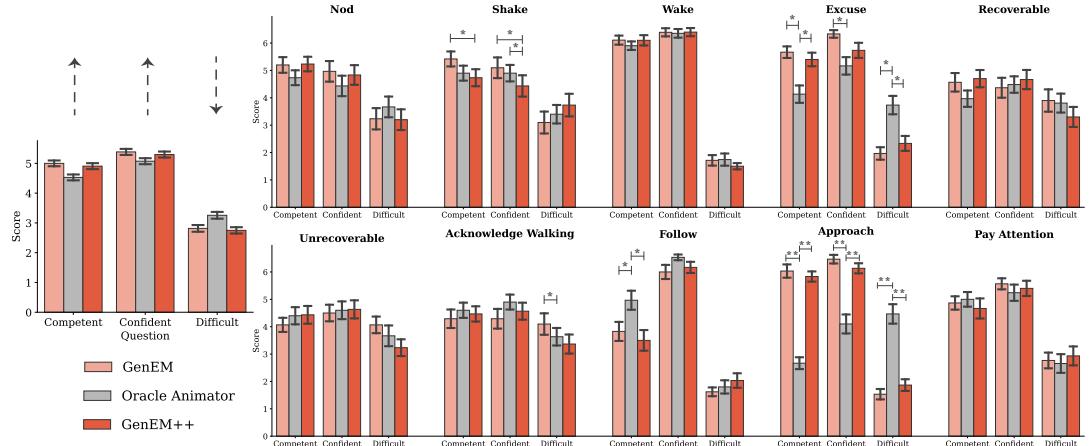


Figure 4: Plots showing participants' survey responses to three questions about each behavior (of 10) in each condition (of 3) in the 1st user study. Bars at the top denote significant differences, where (*) denotes $p < .05$ and (**) denotes $p < .001$. Error bars represent standard error. The first plot shows the average score for each question across conditions. The arrows reflect the direction in which better scores lie.

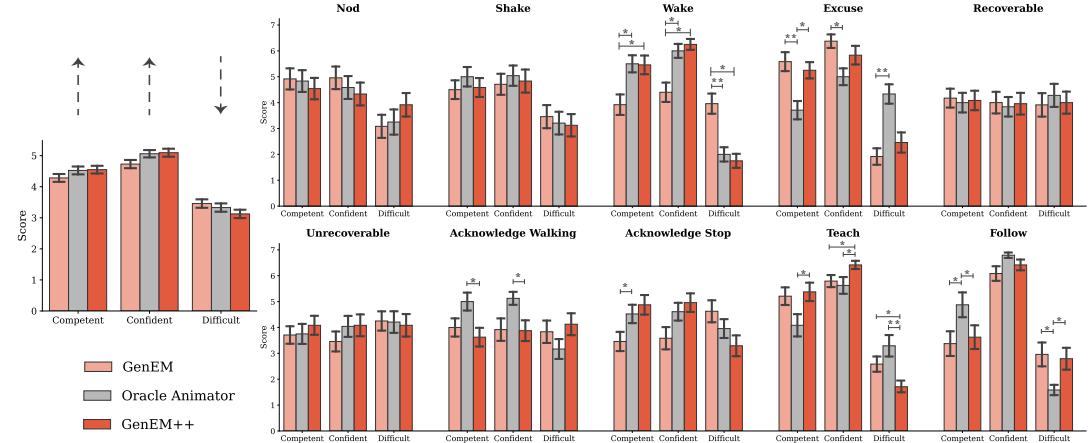


Figure 5: Plots showing participants' survey responses to three questions about each behavior (of 10) in each condition (of 3) in the 2nd user study. Bars at the top denote significant differences, where (*) denotes $p < .05$ and (**) denotes $p < .001$. Error bars represent standard error. The first plot shows the average score for each question across conditions. The arrows reflect the direction in which better scores lie.

(H3) – the GenEM++ behaviors well received and the oracle animator behaviors were not significantly better perceived. When comparing the oracle animator behaviors and GenEM behaviors, there were 4/10 instances where the GenEM behaviors were worse received (Wake, Acknowledge Walk, Acknowledge Stop, and Follow), and 1/10 instances where the GenEM behaviors were more positively rated (Excuse). As with the first study, it is somewhat surprising that the GenEM behaviors were better received than the baselines in one instance; although they resemble them, they do not capture all the nuances present in the oracle animator behaviors since user feedback is not provided. Lastly, the GenEM behaviors were rated worse than the GenEM++ behaviors in 2/10 instances (Wake and Teach) whereas there were 0/10 instances where the reverse was true. Hence, we did not find support for the last hypothesis **(H4)**. Upon performing equivalence tests (equivalence

bound: ± 0.5 Likert points), we did not find any sets of behaviors to be equivalent. Overall, the findings suggest that expressive robot behaviors produced using our approach (with user feedback) were found competent and easy to understand by users.

5 EXPERIMENTS

We conducted a set of experiments to carefully study different aspects of GenEM. This includes ablations to understand the impact of our prompting structure and the modular calls to different LLMs versus an end-to-end approach. Further, through an experiment, we demonstrate that GenEM can produce modular and composable behaviors, i.e., behaviors that build on top of each other. The behaviors were generated by sampling OpenAI's GPT-4 APIs for text

	GenEM		Ablated	
	Execution	Norms	Execution	Norms
Nod	5	0	5	2
Shake	5	0	5	2
Wake	4	2	3	0
Excuse	5	3	0	-
Recoverable	3	0	5	1
Unrecoverable	5	0	5	0
Acknowledge	5	1	5	0
Follow	3	1	0	-
Approach	5	1	5	3
Attention	4	0	1	0

Table 1: Ablations on the mobile robot platform showing the successful attempts of behavior generation when sampling each prompt five times to compare our approach (without feedback) against a variation without the *Expressive Instruction Following* module and subsequently the module translating human expressive motion to robot expressive motion. The *Execution* column indicates the number of successful attempts (/5). The *Norms* column indicates the number of attempts where social norms were not appropriately followed (coded by the experimenter).

	Execution	Norms
Nod	5	0
Shake	5	0
Wake	5	0
Excuse	3	0
Recoverable	5	2
Unrecoverable	4	0
Acknowledge	4	1
Follow	2	2
Approach	5	5
Attention	1	0

Table 2: Behaviors generated on the quadruped in simulation showing successful attempts of behavior generation when sampling each prompt five times. The *Execution* column indicates the number of successful attempts (/5). The *Norms* column indicates the number of attempts where social norms were not properly observed (coded by the experimenter).

completion [32] (gpt-4-0613) with the temperature set to 0. In addition to our user study and experiments on the mobile manipulator, we conducted further experiments using a quadruped simulated in Gazebo/Unity via ROS (see Figure 6).

Ablations. We performed ablations to compare GenEM to an end-to-end approach that takes language instructions and makes one call to an LLM to generate an expressive behavior. The ablations were performed using existing APIs for the mobile robot. The behaviors examined were identical to the first user study along with the prompts. Each prompt was sampled five times to generate behaviors and executed on the robot to verify *correctness*. Further, an experimenter examined the code to check whether the behavior code incorporated reasoning to account for human social norms. The results for code correctness and social norm appropriateness

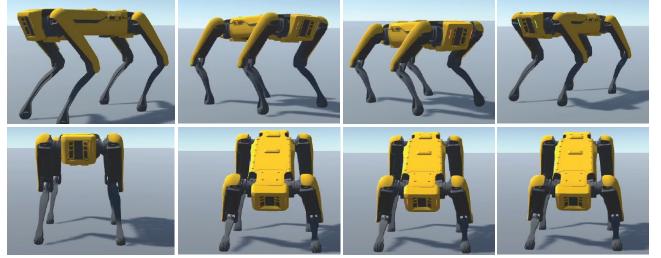


Figure 6: Quadruped simulated in Gazebo performing the *Recoverable mistake* behavior (top) and *Unrecoverable mistake* (bottom) generated by GenEM prior to feedback. After making a recoverable mistake, the robot demonstrates it made a mistake by turning away, lowering its legs, and flashing red lights to convey regret but then returns to its initial position and flashes a green light. In contrast, an unrecoverable mistake causes the robot to lower its height, display red lights for a brief period, and bow forwards and maintains this pose.

are shown in Table 1. Overall, our approach produced higher success rates compared to the ablated variation where no successful runs were generated for 2 behaviors – *Excuse* and *Follow*. For the *Excuse* behavior, the robot must check the user’s distance and signal to a person that they are in its way. However, for the ablated variation, the distance was never checked in the attempts. For the *Follow* behavior, the code called functions that were not previously defined, and used the wrong input parameter type when calling robot APIs, resulting in zero successful attempts. Further, nearly all generated functions were missing docstrings and named arguments, which could make it difficult to use them in a modular fashion for more complex behaviors (despite providing few-shot code examples).

We qualitatively observed that behaviors generated by GenEM reflected social norms, particularly for more complex behaviors, and looked similar for simpler behaviors. For instance, the *Excuse* behavior generated by GenEM used the speech module to say, “*Excuse me*”. For the *Attention* behavior, the ablated variations looked at the person, turned on the light strip, and then turned it off, whereas the GenEM variations also incorporated periodic nodding to mimic “active listening”. For the *Approach* behavior, the GenEM variations always incorporated a nod before moving towards the person while the ablated variations never used nodding; instead lights were used in two instances.

Cross-Embodiment Behavior Generation. We sampled the same prompts in the first user study five times per behavior using API for a simulated Spot robot. The results, summarized in Table 2, show that we were able to generate most expressive behaviors using the same prompts using a different robot platform with its own affordances and APIs. However, some generated behaviors such as *Approach* included variations where the robot navigated to the human’s location instead of a safe distance near them, which would be considered a social norm mismatch (possibly due to the lack of a distance threshold parameter in the translate API), while some did not account for the human (e.g., the robot rotating an arbitrary angle instead of towards the human for *Attention*). Overall, the success rates hint at the generality of our approach to differing robot embodiments.

	Eye contact	Blinking lights	Look around	Shake head	Nod head
Acknowledge Walk	5	-	-	-	5
Approach	4	5	-	-	0
Confusion	-	4	1	5	-

Table 3: Number of times (out of 5 attempts) where previously-learned behaviors (columns) are used when composing new behaviors (rows) using GenEM. Dashes indicate that the given learned behavior API is not provided when prompting the creation of the new behavior.

	Insert actions	Swap actions	Loop actions	Remove capability
Excuse	4	5	5	5
Approach	4	5	5	3
Acknowledge Stop	5	5	4	3

Table 4: Success rates (out of 5 attempts) when providing different types of feedback to behaviors generated using GenEM, where: *Insert actions* request a new action be added ahead of other actions, *Swap actions* request to swap the order of existing actions, *Loop actions* request to add loops to repeat actions, and *Remove capability* requests to swap an existing action with an alternate one.

Composing Complex Expressive Behaviors. In the user studies, all behaviors were generated from scratch using few-shot examples and existing robot APIs. We attempted to generate more complex behaviors using a set of learned expressive behaviors from previous interactions — these skills (represented as functions with docstrings) were appended to the prompts describing the robot’s capabilities (step 2 of our approach) as well as the robot’s API (step 3 of our approach). The learned behaviors used in the prompt were: *nodding*, *making eye contact*, *blinking the light strip*, *looking around*, and *shaking*. We prompted GenEM to generate three behaviors, varying in complexity: *Acknowledge Walk*, *Approach*, and expressing confusion (*Confusion*). All of these behaviors were generated on the quadruped without providing feedback, using instructions that contained a single line description of the desired behavior. We sampled GenEM five times to assess the frequency with which learned behaviors would be included in the outputted program. To assess success, an experimenter checked whether the generated code utilized a combination of robot APIs and learned APIs (see Table 3). For the approach behavior, it was surprising to note that the *nod head* behavior was never utilized whereas blinking lights were always used. For expressing confusion, it was surprising that 4/5 instances generated code for looking around, but only 1/5 instances used the existing *looking around* behavior.

Adaptability to Human Feedback. In the user studies, feedback had some effect on the perception of the generated behaviors. Further, we qualitatively observed that feedback could steer the behavior generation in different ways. We studied this in an experiment where we generated three behaviors from the two prior studies:

Excuse, *Approach*, and *Acknowledge Stop*. Each behavior was generated using a single-line description as before, and without any learned robot APIs. We attempted to modify the generated behavior through four types of feedback: (1) adding an action and enforcing that it must occur before another action, (2) swapping the order of the actions, (3) making a behavior repeat itself (loops), and (4) removing an existing capability without providing an alternative (e.g., removing the light strip as a capability after producing a behavior that uses the light strip). Overall, the results (see Table 4) suggest that it is possible to modify the behavior according to the type of feedback provided, though removing capabilities lead to calling undefined functions more often.

6 DISCUSSION

Summary. In this work, we proposed an approach, GenEM, to generate and modify expressive robot motions using large language models by translating user language instructions to robot code. Through user studies and experiments, we have shown that our framework can quickly produce expressive behaviors by way of in-context learning and few-shot prompting. This reduces the need for curated datasets to generate specific robot behaviors or carefully crafted rules as in prior work. In the user studies, we demonstrated that participants found the behaviors generated using GenEM with user feedback competent and easy to understand, and in some cases perceived significantly more positively than the behaviors created by an expert animator. We have also shown that our approach is *adaptable* to varying types of user feedback, and that more complex behaviors can be *composed* by combining simpler, learned behaviors. Together, they form the basis for the rapid creation of expressive robot behaviors conditioned on human preferences.

Limitations and Future Work. Despite the promise of our approach, there are a few shortcomings. Our user studies were conducted online through recorded video clips, and although this is a valid methodology [12, 43], it may not reflect how participants would react when in the physical proximity of the robot [46]. Hence, further studies involving interactions with the robot should be pursued. Some inherent limitations of current LLMs should be noted, including small context windows and the necessity for text input.

In our work, we only evaluate single-turn behaviors (e.g., acknowledging a passerby), but there are opportunities to generate behaviors that are multi-turn and involve back-and-forth interaction between the human and the robot. Future work should also explore generating motion with a larger action space such as by including the manipulator and gripper. Although we have shown that our approach can adapt to user feedback and their preferences, there is currently no mechanism to learn user preferences over a longer period. In reality, we expect that users will exhibit individual differences in their preferences about the behaviors they expect robots to demonstrate in a given situation. Hence, learning preferences in-context [47] may be a powerful mechanism to refine expressive behaviors.

Despite these limitations, we believe our approach presents a flexible framework for generating adaptable and composable expressive motion through the power of large language models. We hope that this inspires future efforts towards expressive behavior generation for robots to more effectively interact with people.

REFERENCES

[1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. 2023. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. In *Conference on Robot Learning*. PMLR, 287–318.

[2] Amir Aly and Adriana Tapus. 2013. A model for synthesizing a combined verbal and nonverbal behavior based on personality traits in human-robot interaction. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 325–332.

[3] Niklas Bergstrom, Takayuki Kanda, Takahiro Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. 2008. Modeling of natural human-robot encounters. In *2008 ieee/rsj international conference on intelligent robots and systems*. IEEE, 2623–2629.

[4] Nina Buchina, Sherin Kamel, and Emilia Barakova. 2016. Design and evaluation of an end-user friendly tool for robot programming. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 185–191.

[5] Michael Jae-Yoon Chung, Justin Huang, Leila Takayama, Tessa Lau, and Maya Cakmak. 2016. Iterative design of a system for programming socially interactive service robots. In *Social Robotics: 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1–3, 2016 Proceedings 8*. Springer, 919–929.

[6] Yuchen Cui, Siddharth Karamchetti, Raj Palleti, Nidhya Shivakumar, Percy Liang, and Dorsa Sadigh. 2023. No, to the Right: Online Language Corrections for Robotic Manipulation via Shared Autonomy. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 93–101.

[7] Porfirio David, Maya Cakmak, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2022. Interaction Templates: A Data-Driven Approach for Authoring Robot Programs. In *PLATEAU: 12th Annual Workshop at theintersection of PL and HCI*.

[8] Ruta Desai, Fraser Anderson, Justin Matejka, Stelian Coros, James McCann, George Fitzmaurice, and Tovi Grossman. 2019. Geppetto: Enabling semantic design of expressive robot behaviors. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.

[9] Malcolm Doering, Dylan F Glas, and Hiroshi Ishiguro. 2019. Modeling interaction structure for robot imitation learning of human social behavior. *IEEE Transactions on Human-Machine Systems* 49, 3 (2019), 219–231.

[10] Qingxiao Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234* (2022).

[11] Paola Ferrari, María T Lázaro, and Luca Iocchi. 2018. Design of robot teaching assistants through multi-modal human-robot interactions. In *Robotics in Education: Latest Results and Developments*. Springer, 274–286.

[12] Guy Hoffman and Wendy Ju. 2014. Designing robots with movement in mind. *Journal of Human-Robot Interaction* 3, 1 (2014), 91–122.

[13] Chien-Ming Huang and Bilge Mutlu. 2012. Robot behavior toolkit: generating effective social behaviors for robots. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. 25–32.

[14] Chien-Ming Huang and Bilge Mutlu. 2013. The repertoire of robot behavior: Enabling robots to achieve interaction goals through social behavior. *Journal of Human-Robot Interaction* 2, 2 (2013), 80–102.

[15] Chien-Ming Huang and Bilge Mutlu. 2014. Learning-based modeling of multimodal behaviors for humanlike robots. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. 57–64.

[16] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. 2023. Inner Monologue: Embodied Reasoning through Planning with Language Models. In *Conference on Robot Learning*. PMLR, 1769–1782.

[17] Nusrat Hussain, Engin Erzin, T Metin Sezgin, and Yücel Yemez. 2022. Training socially engaging robots: modeling backchannel behaviors with batch reinforcement learning. *IEEE Transactions on Affective Computing* 13, 4 (2022), 1840–1853.

[18] Yusuke Kato, Takayuki Kanda, and Hiroshi Ishiguro. 2015. May I help you? Design of human-like polite approaching behavior. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. 35–42.

[19] Alyssa Kubota, Emma IC Peterson, Vaishali Rajendren, Hadas Kress-Gazit, and Laurel D Riek. 2020. Jessie: Synthesizing social robot behaviors for personalized neurorehabilitation and beyond. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 121–130.

[20] Minae Kwon, Hengyuan Hu, Vivek Myers, Siddharth Karamchetti, Anca Dragan, and Dorsa Sadigh. 2023. Toward Grounded Social Reasoning. *arXiv preprint arXiv:2306.08651* (2023).

[21] Minae Kwon, Sandy H Huang, and Anca D Dragan. 2018. Expressing robot incapability. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 87–95.

[22] Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. 2023. Reward Design with Language Models. In *International Conference on Learning Representations (ICLR)*.

[23] Nicola Leonardi, Marco Manca, Fabio Paternò, and Carmen Santoro. 2019. Trigger-action programming for personalising humanoid robot behaviour. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.

[24] Zhongyu Li, Christine Cummings, and Koushil Sreenath. 2020. Animated cassie: A dynamic relatable robotic character. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3739–3746.

[25] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. 2023. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 9493–9500.

[26] Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. 2023. Text2Motion: From Natural Language Instructions to Feasible Plans. *Auton. Robots* 47, 8 (Nov 2023), 1345–1365. <https://doi.org/10.1007/s10514-023-10131-7>

[27] Phoebe Liu, Dylan F Glas, Takayuki Kanda, and Hiroshi Ishiguro. 2016. Data-driven HRI: Learning social behaviors by example from human–human interaction. *IEEE Transactions on Robotics* 32, 4 (2016), 988–1008.

[28] Mina Marmpena, Angelica Lim, Torbjørn S Dahl, and Nikolas Hemion. 2019. Generating robotic emotional body language with variational autoencoders. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 545–551.

[29] Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Dylan Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837* (2022).

[30] Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao, Dorsa Sadigh, and Andy Zeng. 2023. Large Language Models as General Pattern Machines. In *Proceedings of the 7th Conference on Robot Learning (CoRL)*.

[31] Michael Murray, Nick Walker, Amal Nanavati, Patricia Alves-Oliveira, Nikita Filippov, Allison Sauppé, Bilge Mutlu, and Maya Cakmak. 2022. Learning backchanneling behaviors for a social robot via data augmentation from human–human conversations. In *Conference on robot learning*. PMLR, 513–525.

[32] OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774 [cs.CL]*

[33] Nurziya Oralbayeva, Amir Aly, Anara Sandygulova, and Tony Belpaeme. 2023. Data-Driven Communicative Behaviour Generation: A Survey. *ACM Transactions on Human-Robot Interaction* (2023).

[34] David Porfirio, Evan Fisher, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2019. Bodystorming human-robot interactions. In *proceedings of the 32nd annual ACM symposium on user Interface software and technology*. 479–491.

[35] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2018. Authoring and verifying human-robot interactions. In *Proceedings of the 31st annual acm symposium on user interface software and technology*. 75–86.

[36] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2020. Transforming robot programs based on social context. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–12.

[37] David Porfirio, Laura Stegner, Maya Cakmak, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2023. Sketching Robot Programs On the Fly. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 584–593.

[38] David J Porfirio, Laura Stegner, Maya Cakmak, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2021. Figaro: A tabletop authoring environment for human-robot interaction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.

[39] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. 2023. Progprompt: Generating situated robot task plans using large language models. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 11523–11530.

[40] Arjun Sripathy, Andreea Bobu, Zhongyu Li, Koushil Sreenath, Daniel S Brown, and Anca D Dragan. 2022. Teaching robots to span the space of functional expressive motion. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 13406–13413.

[41] Michael Sugitan, Mason Bretan, and Guy Hoffman. 2019. Affective robot movement generation using cyclegans. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 534–535.

[42] Michael Sugitan, Randy Gomez, and Guy Hoffman. 2020. MoveAE: modifying affective robot movements using classifying variational autoencoders. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 481–489.

[43] Leila Takayama, Doug Dooley, and Wendy Ju. 2011. Expressing thought: improving robot readability with animation principles. In *Proceedings of the 6th international conference on Human-robot interaction*. 69–76.

[44] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291* (2023).

[45] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.

[46] Sarah Woods, Michael Walters, Kheng Lee Koay, and Kerstin Dautenhahn. 2006. Comparing human robot interaction scenarios using live and video based methods: towards a novel methodological approach. In *9th IEEE International Workshop on Advanced Motion Control*, 2006. IEEE, 750–755.

[47] Jimmy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. 2023. TidyBot: Personalized Robot Assistance with Large Language Models. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3546–3553. <https://doi.org/10.1109/IROS55552.2023.10341577>

[48] Wenhao Yu, Nimrod Gileadi, Chuyuan Fu, Sean Kirmani, Kuang-Huei Lee, Montse Gonzalez Arenas, Hao-Tien Lewis Chiang, Tom Erez, Leonard Hasenclever, Jan Humplik, Brian Ichter, Ted Xiao, Peng Xu, Andy Zeng, Tingnan Zhang, Nicolas Heess, Dorsa Sadigh, Jie Tan, Yuval Tassa, and Fei Xia. 2023. Language to Rewards for Robotic Skill Synthesis. In *Proceedings of the 7th Conference on Robot Learning (CoRL)*.

[49] Allan Zhou and Anca D Dragan. 2018. Cost functions for robot motion style. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3632–3639.