
ROS-LLM: A ROS framework for embodied AI with task feedback and structured reasoning

Christopher E. Mower¹, Yuhui Wan^{1,2}, Hongzhan Yu¹, Antoine Grosnit^{1,3},
Jonas Gonzalez-Billandon¹, Matthieu Zimmer¹, Jinlong Wang⁴, Xinyu Zhang⁴, Yao Zhao⁵,
Anbang Zhai⁵, Puze Liu³, Davide Tateo³, Cesar Cadena⁶, Marco Hutter⁶, Jan Peters³,
Guangjian Tian⁷, Yuzheng Zhuang⁸, Kun Shao¹, Xingyue Quan⁸, Jianye Hao⁸,
Jun Wang⁹, Haitham Bou-Ammar^{1,9}

Abstract

We present a framework for intuitive robot programming by non-experts, leveraging natural language prompts and contextual information from the Robot Operating System (ROS). Our system integrates large language models (LLMs), enabling non-experts to articulate task requirements to the system through a chat interface. Key features of the framework include: integration of ROS with an AI agent connected to a plethora of open-source and commercial LLMs, automatic extraction of a behavior from the LLM output and execution of ROS actions/services, support for three behavior modes (sequence, behavior tree, state machine), imitation learning for adding new robot actions to the library of possible actions, and LLM reflection via human and environment feedback. Extensive experiments validate the framework, showcasing robustness, scalability, and versatility in diverse scenarios, including long-horizon tasks, tabletop rearrangements, and remote supervisory control. To facilitate the adoption of our framework and support the reproduction of our results, we have made our code open-source. You can access it at: [ROS-LLM-Code](#).

1 Introduction

When developing a robotic system, whether it be in industry or research, typically roboticists follow a similar development cycle, that we summarize in Figure 1. Initially, a task is *conceptualized*, for example “make me a coffee”. Formally, we will refer to such a task for the robot as an *action*. In our case, we assume actions from domestic scenarios, however, the action could be inspired by any other application (e.g. healthcare, construction, space).

The conceptualized task, or action, is broken down into several sub-tasks, e.g. “reach”, “pick”, “switch”, etc. Throughout this paper, we refer to these sub-tasks as *atomic actions* since a complete action, such as “make me coffee”, comprises several components within a larger sequence or combination of sub-tasks. For each atomic action, a low-level policy representation can be obtained using common approaches such as: reinforcement learning (RL), imitation learning, or optimal

¹Huawei Noah’s Ark Lab, London, UK

²University of Leeds, Leeds, UK

³Technical University of Darmstadt, Darmstadt, Germany

⁴East China Normal University, Shanghai, China

⁵Huawei Technologies, Hangzhou, China

⁶ETH Zurich, Zurich, Switzerland

⁷Huawei Noah’s Ark Lab, Hong Kong, China

⁸Huawei Noah’s Ark Lab, Shenzhen, China

⁹University College London, London, UK

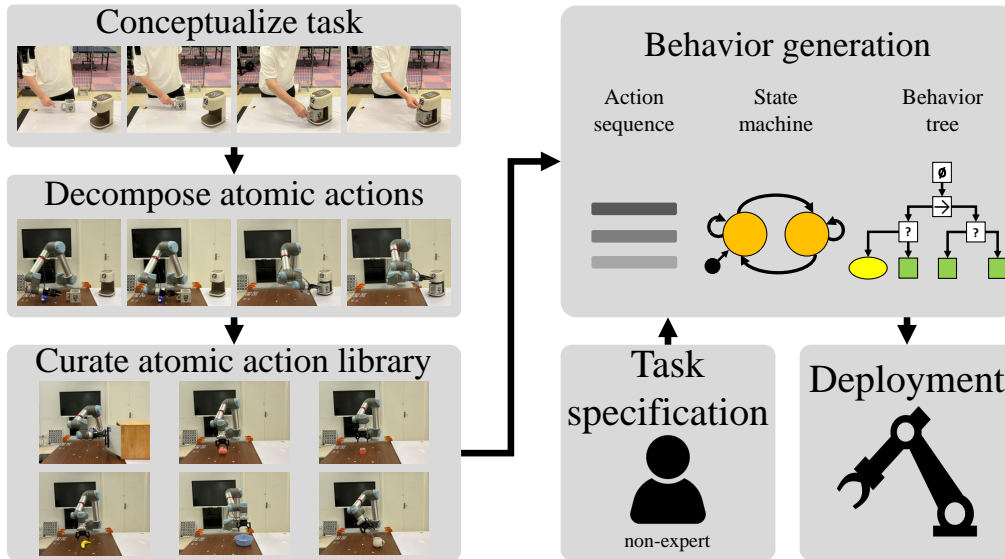


Figure 1: Overview of a typical robotics development workflow.

control (e.g. model-predictive control). Note that, in our case, atomic actions can also include perception sequences (e.g. “locate mug”).

As we conceptualize more tasks, we subsequently develop more and more atomic actions that we can curate in a library. Each atomic action is ultimately a function: an input variable and parameters that are mapped to a return value. The representation of an atomic action could be simple, for instance, “open gripper” that sends an “open” signal to a parallel gripper attached to a robot arm which then returns the final width from the gripper. The representation could take other forms, for example a planner and feedback controller, a neural network policy trained using RL or imitation learning for the end-effector that feeds into an inverse-kinematic controller to compute target joint states. The planner/controller parameters for a given formulation or the neural network weights can be stored in memory and considered as the atomic action representation in the atomic action library. This library of atomic actions can be in the form of a code API or, if the Robot Operating System (ROS) is being used, the library can be a list of ROS actions and services. Textual descriptions for each atomic action are always assumed to be provided, i.e. documentation.

A library of atomic actions is collected and maintained by experts: such as the robot manufacturer (e.g. KUKA, ABB, Clearpath) or companies providing integration services (e.g. Covariant, Berkshire Grey, PickNik Robotics). Interestingly, the robot system can perform unseen tasks by combining some atomic actions. Those novel tasks would typically be specified to the manufacturer/integrator by non-expert customers. Common behavior representations are sequences of atomic actions, state machines, or behavior trees. Given a combination of atomic actions in one of these form, the robot system can then be deployed. We note that, despite this library of atomic actions, teams of expert robot engineers often need to manually develop and deploy the proper combination of atomic actions.

The approach shown in Figure 1 and described above has proven effective in many applications such as manufacturing, logistics, and inspection, e.g. [21, 10]. These applications where robotics have been successfully applied generally involve repetitive situations. Furthermore, these scenarios typically involve task specifications that are unlikely to change frequently. For instance, a vehicle manufacturer shall probably not alter their process once the production line is set up and operational. However, if a new behavior or action is needed for the system, revising or rewriting the system workflow will require the work a team of expert robotics engineers.

The development and adoption of robotics in applications such as domestic use, healthcare, and construction has been slower than in the applications mentioned above. This can be explained by the challenges needed to address such as frequently changing and time-dependent task specifications, interaction and collaboration with non-experts, and the need to extend the range of capabilities of the

system via input from non-experts. For these reasons, it is unreasonable to assume teams of experts (expensive in time and money) will be on-hand to update the capabilities of the system at will and within reasonable time limits.

In general, robot systems are designed in a modular way to allow users to build their own frameworks by easily integrating and modifying existing processes. The most well-known framework used in research and industry is the Robot Operating System (ROS) [41, 32]. Other examples include the Lightweight Communications and Marshaling (LCM) project [16] and the Open Robot Control Software (Orocos) project [5]. The ROS framework provides a well-established ecosystem of packages and libraries that are ready to use and integrated with many common robot systems (e.g. Universal Robots, Robotiq, Clearpath Robotics). Some of most widely used ROS libraries and packages include the TF library [11], MoveIt [9], and the Navigation-stack. The ROS ecosystem of contributors additionally contains many packages for many important requirements such as simulation [26, 34], kinematic modeling [47, 2], and planning and control [19, 36, 8]. Thus, ROS offers many packages providing useful functionalities for both research and commercial applications. These packages include valuable data structures, control interfaces, inverse kinematics (IK) and motion planning tools, perception utilities, and various visualizers. Additionally, with new tools such as the BehaviorTree.ROS library, ROS actions and services enable the generalization of a wide variety of capabilities required by robot systems into a unified execution framework.

In our work, we aim to provide a framework that addresses the needs and limitations described above. We argue that a key bottleneck in the current robot development workflow is the reliance on teams of robot experts. Therefore, our goal is to develop a system that enables *non-experts to program robots*, i.e. provide the robot with new action compositions and extend the capabilities of the system through demonstration. Given its popularity and the general need for modularity, we integrate our framework with ROS.

Many advancements have been made in robotics and machine learning. Notably, the field of AI has made significant advances recently in natural language processing, due primarily to the development of transformer models [53] and subsequently large language models [39, 13, 50]. These language models have been shown to exhibit remarkable performance across a broad spectrum of tasks such as generating code [7], solving advanced mathematics problems [12, 51], generating valid chain-of-thoughts [56, 52], providing search results [24], assisting in medical diagnosis [46], and others.

Language is a natural tool humans use daily to communicate their goals to one another. With recent advancements in language models, there is increasing interest in using these models to map natural language to a sequence of robot actions. By expressing task requirements in natural language and providing contextual information retrieved from the robot’s environment, non-experts can effectively convey their intentions to robotic systems.

Leveraging natural language, our method generates behavior representations tailored to the task specified by a non-expert user, thereby alleviating the burden on expert engineers and enabling rapid prototyping and deployment of robotic systems.

Moreover, since a language model will likely not be capable enough to replace a team of experts, our framework enables continuous learning and adaptation, facilitated by iterative feedback loops. Also, in-order to enable non-experts to update the library of atomic actions via imitation learning, we develop an interface for them to provide several demonstrations. Through teleoperation or kinesthetic teaching, users can augment the system’s atomic action library with skills, enriching its repertoire of behaviors. The human feedback provides the system with the ability to correct its past mistakes and to modify task objectives.

In this paper, we make the following contributions.

- Introduction of a framework that enables non-experts to intuitively program robots using natural language prompts and contextual information retrieved from a ROS environment (i.e. mapping sensor readings to text).
- Integration of an open-source language model with a comprehensive atomic action library, facilitating the composition of action sequences for complex long-horizon tasks that are adaptive to dynamic environments and human feedback.

- Development of a real-world robot setup within a kitchen-like environment and performing various experiments that provide important lessons on prompting strategies to improve the performance of open-source language models.
- Demonstration of the system in handling diverse scenarios, showcasing its robustness and versatility in human-robot collaboration, along with the provision of key insights for using open-source LLMs in robotics applications. Furthermore, we perform a long-distance supervisory control user-study where an operator in Europe controls a robot in Asia to perform a tabletop rearrangement task. Additionally, we describe two use-cases of our system being developed/deployed in systems in the wild.

2 Problem formulation

As described in the introduction, we can think of an atomic action as a single task that the system can perform. Formally, we frame a single task as a Markov Decision Process (MDP) characterized by the tuple $\langle \mathcal{S}, \mathcal{A}, r, \mathcal{P}, \gamma \rangle$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function defined for any state $s \in \mathcal{S}$ and any action $a \in \mathcal{A}$, $\mathcal{P}(s_{t+1}|s_t, a_t)$ is a transition probability distribution, t is a discrete time step, and the scalar $0 < \gamma \leq 1$ is a discount factor.

In contrast, to the standard MDP formulation, we also assume access to a failure flag f that is returned on termination of the MDP, i.e. task completion. The failure flag indicates whether the desired task was completed successfully or not, i.e. $f = 0$ indicates success, $f = 1$ otherwise. For example, if the task is for a robot arm to reach to a target, then at termination $f = 0$ means the target was acquired, and $f = 1$ implies that the robot finished in a configuration far from the target. Thus, our modified MDP is denoted by $\langle \mathcal{S}, \mathcal{A}, r, f, \mathcal{P}, \gamma \rangle$.

In single-task RL, an AI agent generally learns a policy $\mu(a|o)$ choosing an action a given an observation o (e.g. from an image) of the state s . The agent’s objective is to determine the sequence of actions that maximizes the expected return $\mathbb{E}[\sum_t \gamma^t r(s_t, a_t)]$. The specific task and its associated rewards are determined by the reward function r . In our case, we assume access to an atomic action library that corresponds to having a set of N pre-trained or pre-defined policies $\{\hat{\mu}_i\}_{i=1}^N$ ready to use, each based on an underlying modified MDP. Note, we use hat $\hat{\mu}$ to denote that the policy is trained/defined.

Humans exhibit remarkable proficiency in synthesizing complex behaviors by composing various known skills. With regard to a robotic system, this process involves the composition of distinct policies that are executed following a certain structure, such as some appropriate sequential order, or corresponding to a data structure like a behavior tree. In situations when there is access to a team of experts, it is reasonable to assume that they are capable to define some reward function that measures appropriate compositions of atomic actions, i.e. behaviors. With access to experts, we can reasonably assume some reward or fitness function that specifies an appropriate composition of these policies, in which case we could explore methods based on hierarchical reinforcement learning [1]. However, in our case, absent expert guidance, the robotic system must rely on environmental observations o and non-expert human input h to guide policy selection. Also, we assume the input from the human is given by text.

Assuming that an appropriate data representation (e.g. a sequence or a tree) has been selected, it denotes a specific configuration of atomic action composition by b and terms this as a *behavior*. For example, in the case of a sequence, given an atomic action library with $N = 3$ policies, we could choose $b = \{3, 1, 2, 1\}$ meaning we should start executing policy $\hat{\mu}_3$ first, then $\hat{\mu}_1$, $\hat{\mu}_2$, and finally $\hat{\mu}_1$. As mentioned above, this choice of composition b is based on the observation o and human input h . Thus, we define the mapping $\pi(b|o, h)$ as a *behavior policy* - this is our AI agent.

When a behavior b is chosen by the behavior policy π , it may be the case that one of these policies fails or all may succeed. In the case of failure, it is typically not worthwhile to proceed. For example, if a robot reaches to grab a mug but knocks it over when attempting to grasp due to misalignment, then it would not make sense to continue to try to place the mug. Thus, in this case, we consider that the result of the behavior b is given by $\bar{f} = 1$. When all the atomic actions complete successfully, then the result of the behavior b is $\bar{f} = 0$. Note, the bar \bar{f} indicates the failure result of the behavior time step τ . Another way to conceptualize this idea is that the final atomic action failure flag is given as the failure flag of the behavior b .

To enhance the selection of optimal behaviors b and improve the behavior policy π , the agent requires access to a measure of goodness. As previously noted, we refrain from relying on domain experts to formulate new reward functions tailored to specific tasks. Under our premise, where access to a failure indicator f is assumed, a clear definition is needed. In our current work, we adopt the simplifying assumption that the failure indicator can be hard-coded. For instance, during the execution of the “grasp object” atomic action, the final gripper width serves as an indicator: zero width implies failure to grasp the object, while a width greater than zero indicates success. This approach is constrained by the reliance on straightforward rules or expert input to establish the failure indicator. Recent work, however, suggest promise in training neural networks to classify multimodal failures in manipulation tasks [20]. Of interest in our future work is to explore the use of multi-modal transformer networks (e.g. vision-language models) to model the failure indicator f . Ideally, therefore, the agent should choose b to maximize success and also minimize time. Since, by assumption, we have access to the failure flags \hat{f} we define for the behavior policy π a return function

$$R_\pi := \mathbb{E} \left[\sum_{\tau} -\beta^\tau (1 + \hat{f}_\tau) \right] \quad (1)$$

where τ is a time step for the behavior policy π , and $0 < \beta \leq 1$ is a scalar discount factor. Notice that we add 1 to the the failure flag \hat{f}_τ . This is required to indicate to the agent that as few atomic actions should be used to construct a behavior as possible. Our goal in this work, is to develop a framework that is capable of integrating with real systems that addresses the above novel formulation.

3 The ROS-LLM Framework

In this section, we provide an overview of our framework, highlighting specific design considerations tailored to the integration of embodied AI in robotics. We show on Figure 2 an illustrative summary of the proposed framework.

3.1 Atomic action library

In both industrial and research settings, the development of robotic tasks often involves breaking down complex behaviors into simpler atomic actions. For instance, it is possible to complete a longer-horizon task such as pick-and-place by breaking it down into (i) acquire target, (ii) reach, (iii) grasp, (iv) acquire new target, (v) reach, and (vi) place. While atomic actions commonly include physical movements, such as reaching or grasping, they can also include perceptual tasks like object detection or localization. Over time, a repertoire of tasks expands and the resulting library of atomic actions can facilitate the reusability of robotic behaviors. While an initial set of these actions can be initially provided, it is unlikely to cover all potential tasks. Therefore, our system includes a facility allowing non-expert users to add extra atomic actions to the library via imitation learning (see Section 3.6 for further details).

In our framework, we integrate the concept of atomic actions with ROS and implement each atomic action as either a ROS action or a ROS service. Additionally, for each action or service within the library, we provide a textual description to convey its intended functionality and usage. This information is stored in a JSON file with the following fields: ‘name’ of the ROS action/service, ‘type’ that specifies if the atomic action is a ROS action or ROS service, ‘description’ of the atomic action and its input/output. We also designed tools, exposed in the ROS environment, to readily retrieve the action library description in the form of a readable string that can be exposed to the LLM through prompting. These atomic actions can then be combined and orchestrated by the LLM agent to execute complex behaviors.

In addition, we provide functionality for the atomic action library, which is described by a code API that interfaces with the system. In this case, the JSON field ‘type’ is given the value ‘code’, and the behavior output from the language model can be an executable Python script. Note that in this case, it is possible to have an action library combining ROS actions, ROS services, and a code API. However, the output of the language model is restricted to executable Python scripts.

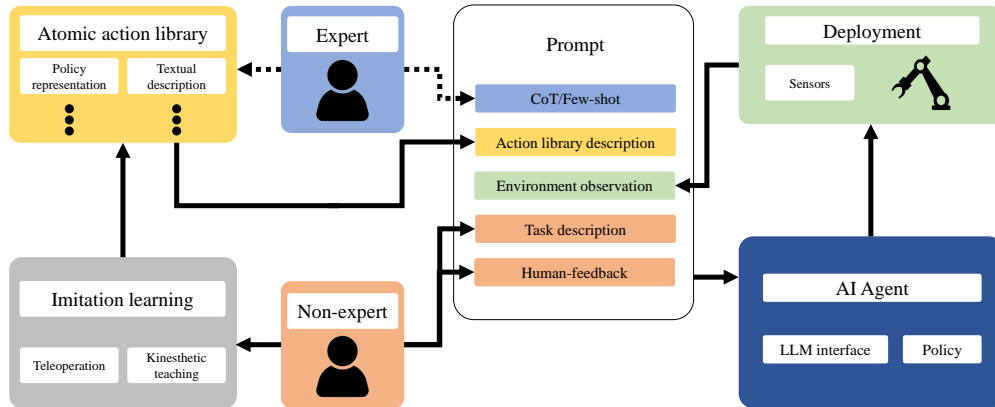


Figure 2: Our proposed ROS-LLM framework overview illustrates the integration of several components. In this figure, the dashed lines denote elements that are only introduced once, such as the initial version of the atomic action library and CoT/Few-shot prompts by an expert.

3.2 Environment observation

In our framework, we operate under the assumption that the robot has access to various sensors capable of observing changes in the environment. These sensors provide valuable information about the state of the surroundings, which is essential for the robot to make informed decisions and adapt its behavior accordingly. Given that our agent is a language model, we require environmental observations to be represented in textual form to serve as input to the model.

To facilitate this process, we implement a ROS package called the ‘`observation_manager`’. This package implements tools that are responsible for querying multiple sensors and gathering textual observations of the environment. Upon initialization, a ROS node is configured with a list of services to call whenever an observation is needed. This modular design allows a user of our framework to customize the set of sensors they want to use based on their specific application requirements.

To set up our framework, we require users to implement these services that map observations from sensors to textual representations. For instance, a service may output descriptions such as “the gripper is open” or “the blue box is detected in the camera view”. By standardizing the format of environmental observations as text, we promote modularity and interoperability within the framework.

3.3 Human non-expert interface

We provide a chat-based interface to our framework to allow easy adoption from non-expert human users. Each environment step is executed after the human feedback is received from the interface, and then once the execution is over on the system (ending either with a success or a failure), we ask the human to input a new textual entry. At the beginning, we let the system interpret the first human input as the task description. The task description should outline the goal or objective to assign to the robot, providing context for the subsequent actions to generate. Thereafter, the system treats human input prompts as a feedback, which may contain suggestions for corrective behavior or suggestions for alternative approaches for the robot to complete the task.

Another potential interaction mode could be via speech, which would have the potential to be even more intuitive for non-experts. We actually plan to implement a microphone into our setup and use an off-the-shelf audio-to-text package for parsing the input. This functionality will be incorporated into our main code-base in the future.

3.4 Prompt generation

The prompt provided to the language model serves as input to generate behavior representations that can be executed on the system. At each environment step τ , the prompt is updated, ensuring that the language model receives the latest information necessary for decision-making. We show in the central part of Figure 2 the different elements that we expect in a prompt to shape the behavior of the system.

The prompt includes a task description that is provided by the user, as described in the previous sub-section. After the first environment step τ , the non-expert provides feedback that the system uses to correct its behavior. A description of the atomic action library is also included to provide context on the admissible behaviors of the system, as described in Section 3.1. Moreover, the prompt contains an observation of the environment that is collected by mapping several sensor readings to text. Several well-known prompt engineering strategies are utilized to aid the language model construct a behavior for the system, namely chain-of-thought and few-shot prompting. These additional portions of the prompt are assumed to be given as part of all observations o . Finally, some additional notes are written in the prompt, such as how the language model should format the behavior output (e.g. Python or XML).

Overall, the prompt generation process gathers information from both the ROS environment and the human interface, ensuring that the language model receives comprehensive input to guide its decision-making process. Once the prompt is constructed, it is then passed to the language model. We consider the output of the language model to represent the desired behavior for the system. We describe next the different formatting options for the output of the language model.

3.5 Behavior representation

We call a *behavior* the combination of atomic actions that is extracted from the textual output of the language model. To represent a behavior, the LLM generates either a Python or XML code. When Python format is used, a Python terminal exposed to the ROS environment executes the LLM output. In the case of XML, the LLM response is interpreted as a behavior tree¹. We use regular expressions to easily identify parse LLM output that should encapsulate the code into `''python...''`, `''json...''`, or `''xml...''`.

Python output We expect the use of Python code when the action library is a set of Python function that exposes the various functionality of the system. The library can also contain ROS actions and ROS services that can be interfaced with the script.

JSON output When the JSON format is used, a behavior representation called an action sequence is used. In this case, the specified actions in the sequence are executed one after the other, and we expect each action to be a ROS service of the type `rosl1m_srvs/AtomicAction`. This service returns a string called `output` and takes as input one string argument called `input`, which takes input per action, and another string argument `prev_output`, which is the output from the previous action.

XML output A behavior tree, represented by XML code, is a hierarchical model that describes the behavior of autonomous agents or robots. It consists of nodes that define specific actions or conditions, and it is organized in a tree-like structure. At the root of the tree, the behavior selector node determines the order in which to evaluate and execute child nodes. These child nodes can include sequences of nodes, which execute their child nodes sequentially until one fails, or parallel nodes, which execute their child nodes simultaneously. Other types of nodes include conditional nodes, action nodes, and decorator nodes, each serving distinct roles in controlling the agent’s behavior. In this case, the language model is tasked with producing the XML code defining a behavior tree.

3.6 Updating the atomic action library via imitation learning

In our framework, the atomic action library may not encompass all the tasks that non-expert users want to realize. To address this limitation, we provide a mechanism for non-experts to update the

¹The behavior tree interface has recently been added to the code base. We plan in the future to run experiments and report the results in an updated version of this article.

action library using imitation learning techniques. This section details our approach to integrating imitation learning into our framework and describes the process of updating the action library through kinesthetic teaching or teleoperation.

As a first step, we integrate to our framework a connection to dynamic movement primitives (DMP), which is a widely-used technique requiring a minimal number of demonstrations to learn smooth representations of tasks. However, our framework remains flexible and can accommodate other imitation learning pipelines, including behavior cloning with deep learning, depending on the specific requirements of the task.

When a non-expert provides a demonstration of a task along with a textual description, our framework automatically translates this information into a ROS service and adds it to the atomic action library. The demonstration is captured using kinesthetic teaching or teleoperation, allowing the non-expert to guide the robot through the desired behavior. Subsequently, the learned task is represented as an atomic action within the library, while the accompanying textual description is included in the action library descriptions. This process enables the framework to expand its repertoire of available actions based on user demonstrations and textual annotations, thereby enhancing its adaptability and versatility in addressing a wide range of tasks.

By empowering non-experts to contribute to the expansion of the action library through imitation learning, our framework promotes collaborative human-robot interaction, which could accelerate the development of robot capabilities tailored to specific users' needs.

4 Experiments

In this section, we present a series of experiments to evaluate the performance and capabilities of our framework for intuitive robot programming. These experiments encompass diverse scenarios, ranging from long-horizon tasks to dynamic environment adaptations and intercontinental supervisory control. Through these experiments, we assess the system's ability to comprehend natural language prompts, to generate accurate action sequences, to adapt to changing environments, and to collaborate effectively with human operators. The experiments showcase our framework's robustness, scalability, and versatility, enabling non-experts to program robots effortlessly. Thus, we are advancing the state-of-the-art in embodied AI and human-robot collaboration.

4.1 Experiments setup

Our real-world robot setup, shown in Figure 3, comprises a Universal Robots UR5 arm equipped with a Robotiq 2-finger gripper (2F-85) attached at the end-effector. Situated within a kitchen-like environment, our setup includes various objects such as boxes, a bowl, utensils (e.g., a spoon), a sink, and a cabinet with a door. We strategically place these objects around the robot within its workspace to simulate real-world scenarios. The robot is controlled using the ROS Noetic, utilizing standard packages tailored for UR/Robotiq robots. Our control system runs on a laptop powered by Ubuntu 20.04, featuring a 16-core i9 Intel CPU, which ensures robust performance and efficient task execution.

For language processing and task execution, we utilize the Deepseek 7B Coder, an open-source LLM deployed on a virtual LLM (vLLM) server located in Europe. Communication between the robot setup in Asia and the vLLM server is established via the internet, with requests transmitted over HTTP protocol.

To facilitate seamless interaction and task execution, we have developed a chat interface integrated with an AI agent, a novel framework for integrating and learning structured reasoning into AI agents' policies. This chat interface enables intuitive communication with the robot system, allowing users to articulate task requirements and provide feedback effortlessly. Furthermore, we incorporate Whisper [42], a tool that translates audio into text, offering an alternative interface through spoken language.

For perception and state estimation, we use an Intel Realsense RGBD camera positioned at a fixed location facing the robot and scene. Despite the camera's depth-sensing capabilities, we use only RGB data channels for computational efficiency to simplify our implementation. In the environment, we add fiducial marker tags and utilize capabilities from the AprilTag library [38] for sensing their




Figure 3: Real-world laboratory setup used in our experiments.

poses with respect to the camera frame. To estimate the transformation between the camera the and robot base frame, we use standard eye-to-base calibration packages from the Moveit library.

In summary, our real-world robot setup leverages state-of-the-art hardware and software components, coupled with advanced language processing capabilities, to enable intuitive interactions and successful task execution within a dynamic and realistic environment.

4.2 Long-horizon tasks

The experiment on long-horizon tasks serves to evaluate the system’s ability to comprehend and execute complex multistep tasks, exemplified by the request “can you make me a coffee”. This experiment is essential as it represents a real-world scenario where robots need to perform sequential actions to accomplish a goal, requiring a combination of perception, manipulation, and decision-making skills.

In this experiment, we employ an open-source language model(Deepseek 7B Coder) to generate action sequences based on natural language prompts. The need for a detailed, natural language description of the task steps arises from the intricacies of the language model’s comprehension capabilities.

We provide the system with a comprehensive atomic action library encompassing a range of fundamental actions necessary for coffee preparation. These include reaching, opening/closing doors, picking up objects, placing objects, switching machines on/off, opening/closing the coffee machine cover, taking objects out of a cabinet, putting objects back in the cabinet, spooning coffee grounds from the bowl with a spoon, inserting the mug into the coffee machine, drop objects in the sink.

To execute the task of making coffee, the system first identifies and locates the necessary objects and equipment in the environment, such as the coffee machine, mug, coffee grounds, and spoon. Subsequently, it plans a sequence of actions to perform each step of the coffee-making process. This entails grasping and manipulating objects, navigating the environment, and interacting with the coffee machine and other appliances.

The system takes as input a task prompt that consists of a natural language description of the task steps, resembling a recipe for making coffee. This contextual information is crucial for guiding the system’s understanding and ensuring the generation of accurate action sequences.

In Figure 4, we show the key steps of the experiment, with corresponding photos in Figure 5, illustrating the systematic execution of actions by the robot in the real-world. The successful outcome

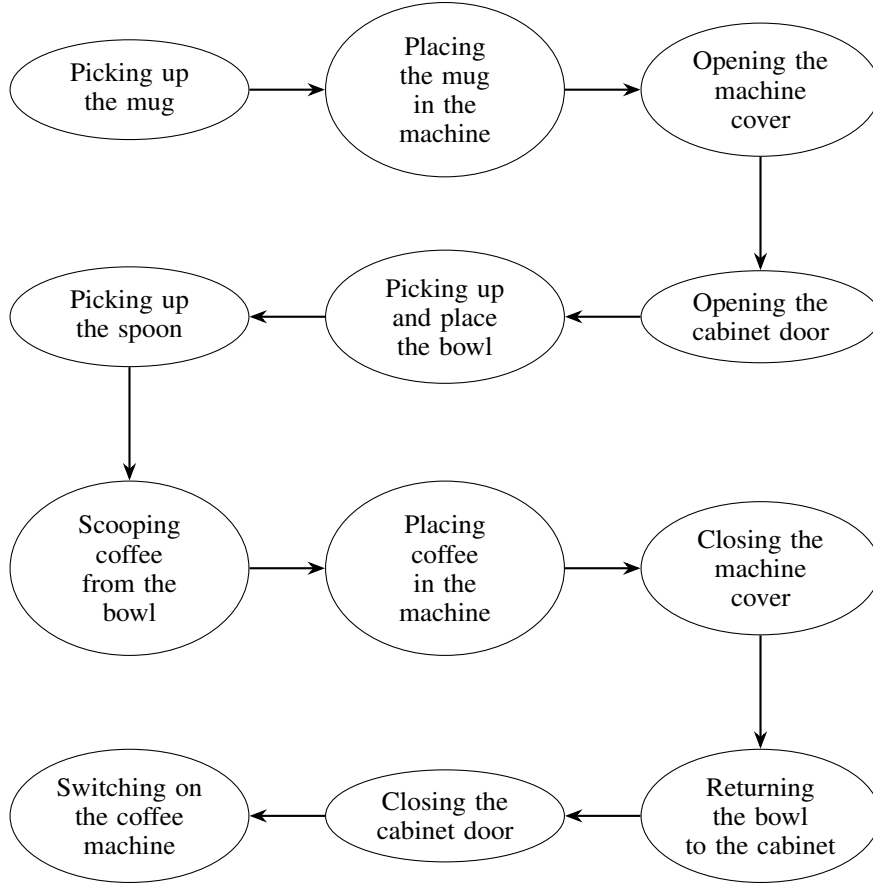


Figure 4: Detailed steps in the coffee-making process arranged in a modified Z-shaped flow across four rows.

of the experiment demonstrates the system’s capability to comprehend and execute long-horizon tasks accurately and efficiently.

Overall, this experiment underscores the effectiveness of our framework in enabling robots to tackle complex, multistep tasks through intuitive natural language prompts and comprehensive action libraries.

4.3 Policy correction via human feedback

The ability to correct policy decisions based on human feedback is crucial to enhance the adaptability and robustness of autonomous systems. In this section, we present experiments assessing the effectiveness of human feedback in correcting policy errors generated by the LLM. Our objective is to demonstrate that, as task complexity increases, the performance of the LLM decreases, but that with the incorporation of human feedback, the success rate can be maintained.

We designed a tabletop rearrangement task involving several colored boxes placed on a table in front of the robot, with a bowl and a sink nearby. Each task prompt specifies instructions for rearranging the boxes, such as placing specific colors into the sink or bowl and stacking others. The difficulty of the task is determined by the number of boxes present in the scene, ranging from 2 to 8 boxes, with 5 unique tasks generated for each number of boxes. In total, 35 tasks were manually created.

For each trial, the LLM generates an action sequence and executes it on the real robot. If the action sequence is correct, the trial is considered a success; otherwise, human feedback is provided, the scene is reset, and the policy is executed again. We determine success by comparing the robot’s actions with the known solution for each task. In cases where human feedback is provided, a second execution of the policy is performed, and we measure success in the same way.

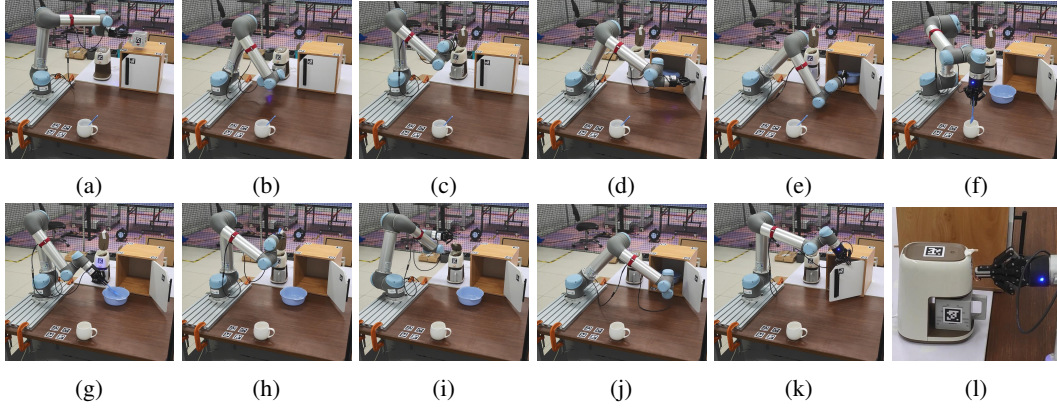


Figure 5: Detailed steps in the coffee-making process are depicted across twelve images: (a) picking up the mug, (b) placing the mug in the coffee machine, (c) opening the coffee machine cover, (d) opening the cabinet door, (e) picking up the bowl, (f) picking up the spoon, (g) scooping coffee from the bowl, (h) placing the coffee in the machine, (i) closing the coffee machine cover, (j) returning the bowl to the cabinet, (k) closing the cabinet door, (l) switching on the coffee machine.




Figure 6: Sequence for the 6-cube task, depicted over ten stages: (a) starting, (b) picking up a green cube, (c) placing the green cube on top of a red cube, (d) picking up an orange cube, (e) dropping the orange cube into the sink, (f) picking up a blue cube, (g) placing the blue cube on top of another green cube, (h) picking up another orange cube, (i) placing the orange cube on top of the blue cube, (j) finishing.

Figure 6 provides a sequence of images demonstrating our system executing a task, while Figure 7 shows the results of our experiments. As expected, we observe a decrease in performance without human feedback as task difficulty increases. One example of human feedback is when the robot picks up the wrong cube, and the human indicates that it is not the correct target cube and reiterates the correct task order. Another example is when the robot fails to locate the target cube and raises an error due to the arm blocking the camera; the human provides the feedback as “Please home the arm before looking for the cube.” Conversely, the inclusion of human feedback generally leads to improved task success rates across varying levels of difficulty. However, we note that sometimes human feedback does not result in a successful task completion, highlighting potential limitations or challenges in the correction process.

Overall, our experiments demonstrate the effectiveness of human feedback in correcting policy errors generated by the LLM, thereby enhancing the adaptability and robustness of autonomous systems.




Figure 7: Results depicting policy correction through human feedback, where orange indicates task success with LLM planning and blue indicates task success after correction with human feedback. The graph illustrates a decrease in the success rate of the LLM as task complexity rises, countered by sustained success rates with the incorporation of human feedback for action sequence correction.

4.4 Updating the action library with imitation learning

Enhancing robotic skill sets to include a variety of tasks is crucial, especially in household environments where the required actions may not be sufficiently covered by existing action libraries. To meet this need, we present an experiment that demonstrates the system’s adaptability and capacity for continual learning. By showcasing its ability to learn new atomic actions from human demonstrations and seamlessly integrate them into the action library, we aim to highlight the system’s versatility and its potential to evolve in response to user needs.

The experiment entails the demonstration of various tasks through kinaesthetic teaching, wherein a human guides the robot through the execution of specific actions. These actions include stirring, pouring, tossing the pan, wiping the table, seasoning food, and grating cheese. Each demonstration is captured and represented using a DMP representation, facilitating the learning process. Subsequently, the human provides a descriptive label for each atomic action, enriching the action library with contextual information.

Following the addition of new atomic actions to the library, we challenge the system with a long-horizon task: “make me pasta.” Leveraging the updated action library, the system orchestrates a sequence of atomic actions to fully fulfill the task requirements.

This experiment highlights the system’s proficiency in continual learning and adaptation, demonstrated through its ability to learn diverse atomic actions from human demonstrations and seamlessly incorporate them into the action library. The successful execution of the “make me pasta” task exemplifies the system’s capability to effectively utilize learned atomic actions for complex, long-horizon tasks. Additionally, the inclusion of descriptive labels improves the action library’s interpretability and usability, facilitating intuitive task specification by non-experts. This capability enhances the system’s versatility and applicability in various real-world scenarios, emphasizing its potential as a valuable tool for intuitive robot programming and enhancing human-robot collaboration in household environments.

4.5 Adapting to a changing environment and continual learning

The adaptability of robotic systems to dynamic and unpredictable changes in the environment is crucial for their effective operation in real-world scenarios. In this experiment, we aim to demonstrate




Figure 8: Steps in the pasta-making process, depicted in five stages: (a) grating cheese, (b) pouring sauce, (c) stirring ingredients, (d) tossing the contents in the wok, (e) adding seasoning.

the capability of our framework to adapt and recover from such changes through continual learning facilitated by human feedback. The primary goal is to showcase the system’s ability to utilize human feedback to recover from unforeseen disruptions in task execution and subsequently learn from these experiences to handle similar changes in the future autonomously.

The experimental protocol is as follows.

1. **Task Specification:** The system is tasked with “pick and place the box,” a common robotic manipulation task.
2. **Action Sequence Generation:** Using the provided task description, the language model generates an action sequence for execution.
3. **Execution and Disruption:** As the robot executes the task, the environment is intentionally perturbed by moving the target box, leading to a failure to grasp the object.
4. **Human Feedback:** Upon observing the failure, the human provides corrective feedback, advising the system to ensure the box’s proximity before grasping and then to retry the task.
5. **Recovery and Adaptation:** Leveraging the feedback, the system adapts its approach and successfully completes the task, demonstrating its ability to recover from environmental changes through human-guided learning.

To evaluate the system’s capability for continual learning, a second trial is conducted under identical conditions, with the box moved simultaneously as in the previous trial. However, this time, the human feedback from the initial trial is incorporated into the task prompt. By doing so, we aim to assess whether the system autonomously applies the learned corrective action to handle similar environmental changes without human intervention.

The results of the experiment demonstrate the efficacy of our framework in adapting to changing environments and leveraging human feedback for continual learning. By successfully recovering from unforeseen disruptions and autonomously applying learned strategies in subsequent trials, the system showcases its resilience and ability to evolve through experience. These findings highlight the potential of our framework to enhance robustness and adaptability in real-world robotic applications, paving the way for more reliable and versatile autonomous systems.

4.6 Remote supervisory control




Figure 9: Remote supervisory control using (a) language interfaces, depicted through continuous actions with two interfaces: (b) picking up a cube, (c) avoiding obstacles, (d) placing the cube in a bowl.

The capability for remote supervisory control of robotic systems holds significant promise for a variety of applications, particularly in environments where direct human presence is impractical or unsafe, e.g. [35]. Tasks such as offshore inspection and maintenance, search and rescue operations, mining activities, and space robotics often necessitate remote operations due to safety concerns and logistical challenges. Moreover, the integration of 5G technology further enhances the feasibility and effectiveness of remote control systems, offering high-speed, low-latency communication capabilities that are essential for near real-time interaction between operators and robots over vast distances. For instance, handling obstacles in the environment, such as debris in search and rescue, may pose challenges for the robot’s perception algorithms, requiring human intervention to guide the robot effectively.

In our experiment, conducted with the operator located in Europe and the robot system situated in Asia, participants were tasked with controlling the robot to perform pick-and-place manoeuvres with obstacles present in the scene. The action library provided to the operator included basic movement commands such as `move_left`, `move_right`, `move_up`, `move_down`, `move_forward`, `move_backward`, `open_gripper`, and `close_gripper`. The operator communicated commands to the robot using natural language, describing where they wanted the robot to move and manipulate objects in the environment. A camera feed provided visual feedback to the operator, enabling them to perceive the scene remotely. Remote supervisory control experiments using the language interface are shown in Figure 9a.

A significant lag in the system, measured at approximately 2–3 seconds, was observed due to the long distance between the operator and the robot system. Despite this latency, participants successfully completed the assigned tasks using the language-based interface, demonstrating the effectiveness of our framework for remote supervisory control.

Seven volunteers participated in the experiment. Each participant was required to complete two tasks. The first task involved picking up the blue box and placing it in the blue bowl, and the second task involved picking up a red box and placing it on the white area of the table while avoiding an obstacle in the middle, as shown in Figure 9 and 10. To evaluate the performance and usability of the human-robot interaction [54], we recorded the time taken to complete each task. Furthermore, usability was assessed using the NASA Task Load Index (NASA-TLX) [15] and pairwise subscale comparisons, which provided a comprehensive view of user experience across various dimensions. Seven participants’ results yielded valuable insights into the method’s performance and usability, underscoring our framework as a practical tool for developing supervisory control using natural language in real-world scenarios. Detailed results are shown in Table 1 and Figure 11 for further analysis and discussion.

5 Discussion

The results of our experiments highlight the potential of this approach to simplify the interaction between non-expert users and robotic systems through natural language. This section discusses the implications of our findings, the limitations and strengths of the current framework, and future directions for research.

Table 1: Task time for the remote supervisory control experiment.

No.	LLM	
	Task 1	Task 2
1	5:09	9:40
2	6:10	9:55
3	5:06	7:31
4	11:40	15:42
5	5:05	4:51
6	5:32	7:45
7	4:27	7:37
Avg	6:09	9:00




Figure 10: Experiment setup for the human study.

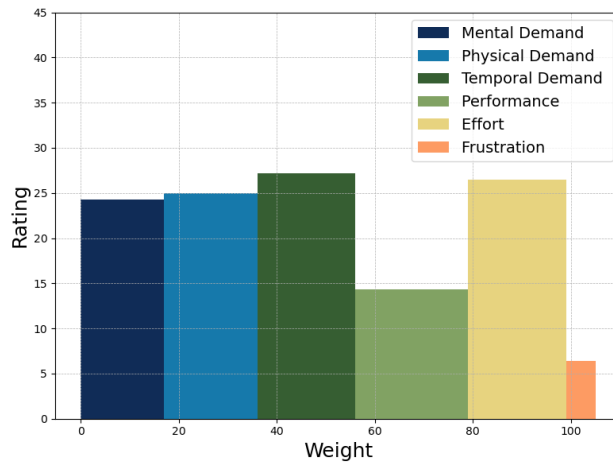


Figure 11: Weighted NASA TLX results for remote supervisory control.

5.1 Performance observations in long-horizon tasks

Alongside the promising outcomes, our extensive experiments have uncovered a range of interesting phenomena, showcasing both notable strengths and areas in need of improvement within the current framework. These findings highlight the robust capabilities of our approach while also pinpointing challenges that guide further development. This balanced perspective not only underscores the potential of our work but also clearly delineates paths for enhancing performance and reliability.

5.1.1 Sensitive prompt

During our experiments, we observed that minor variations in prompt wording significantly impact the model’s output, revealing both vulnerabilities and strengths in its language processing capabilities. For example, the model generated the correct action sequence when the prompt used “another cube” to indicate the second cube but struggled with “the other cube.” Similarly, discrepancies arose with the prompts “put box A on box B” versus “put box A on top of box B,” where the addition of “on top of” altered the generated actions.

Conversely, the model demonstrated robustness to other variations, such as replacing “put box” with “move box” and changing phrases from “in the bowl” to “to the bowl.” These instances did not affect the model’s performance, indicating a certain level of linguistic flexibility.

Possible explanations for these observed behaviors might include:

- **Training Data Variance:** The model’s differential response could be attributed to the frequency and context of phrase occurrences in its training data. Phrases like “move box” and “put box” might appear interchangeably across diverse contexts, helping the model learn their equivalence.
- **Semantic Parsing Differences:** For the phrases that cause discrepancies, it may be that the model perceives “on top of” as indicating a more precise spatial relationship, thus requiring more specific handling in action generation. This points to a nuanced understanding of prepositions and their implications in task execution.
- **Context Sensitivity:** The model might be particularly sensitive to certain keywords or phrases that imply a change in action complexity or specificity, such as spatial relations that require precise positioning.

These findings underscore the need for enhanced natural language understanding within robotic AI systems. Improving the model’s ability to generalize across varied linguistic inputs is crucial without losing contextual accuracy. Enhancements should focus on better handling of semantic nuances and increasing the robustness against minor linguistic variations to develop more reliable and capable robotic assistants.

5.1.2 Confusing example

During our experiments, we encountered specific issues where including example action sequences in the prompts led to confusion for the model. Notably, the model occasionally attempted to execute actions on objects mentioned in the example sequences, but that were not actually present in the real-world environment state. This confusion stems from the model’s inability to distinguish between illustrative examples provided for context with current environment observation and actual commands.

This issue highlights a significant challenge in the design of natural language interfaces for robotic systems: the need for improved parsing and contextual understanding. The model’s current parsing mechanisms may not effectively differentiate between descriptive content (intended to enhance understanding or provide background information) and imperative content (direct commands).

Possible strategies to mitigate this issue could include:

- **Contextual Tagging:** Implementing a tagging system in the training phase where example actions and real commands are tagged differently could help the model learn to distinguish between these types of content more effectively.
- **Enhanced Semantic Analysis:** Developing more advanced semantic analysis capabilities that can interpret the context and intent behind each phrase more accurately. This could involve deeper training on linguistic cues that indicate hypothetical or illustrative scenarios versus actionable instructions.
- **User Feedback Integration:** Incorporating a feedback loop where the system asks for confirmation or clarification when it detects potential ambiguities in the prompt could prevent incorrect actions based on misinterpretations.

Addressing these parsing and contextual understanding challenges is crucial for advancing the usability and reliability of natural language interfaces in robotic systems. By enhancing the model’s ability to discern and segregate different types of linguistic inputs accurately, we can significantly reduce errors and improve the system’s overall performance in real-world tasks.

5.1.3 Robust long-term planning

An interesting aspect of our experiments was evaluating the model’s long-horizon planning capabilities. Typically, the success rate of tasks is expected to decrease with an increase in the number of steps involved due to the accumulation of potential errors at each stage. However, our findings challenge

this expectation. The model demonstrated a notable consistency in success rates across experiments involving 4 to 8 cubes. This observation suggests that the model possesses robust error-handling and planning capabilities that maintain performance even as task complexity increases, as shown in Figure 7.

These results are encouraging for the deployment of this framework in scenarios requiring complex sequential task execution, highlighting its potential reliability and effectiveness in practical applications.

5.2 Enhancing policy correction via human feedback

Targeted human feedback has shown the potential to mitigate this degradation by correcting erroneous policy decisions dynamically. We noticed that due to the variety of mistakes made by the LLM, there is no one-size-fits-all feedback correction.

5.2.1 Feedback implementation

Our findings underscore that not all feedback is equally effective. Specific, actionable feedback that directly corrects the decision-making process or clarifies the task’s objectives tends to result in more favourable outcomes. For instance, if the LLM incorrectly sequences the actions (e.g. attempting to place a cube before picking it up), repeating the command with the correct sequence often resolves the error, effectively reprogramming the LLM’s task strategy.

Also, the human feedback needs to indicate the correction action or action sequence instead of only pointing out the mistake itself. For example, if the LLM erroneously attempts to place a blue cube on a green cube before securing the blue cube, the provided feedback would explicitly instruct to “Pick up the blue cube first, then place it on the green cube.” instead of “You picked up the cube in a wrong order.”

5.3 Usability for remote supervisory control

From the remote supervisory control experiment, we observed notable phenomena related to the human-robot interface. Overall, we noticed that participants experienced similar mental, physical, and temporal demands during the operation, as shown in Figure 11. This suggests that the system maintains a balanced workload distribution.

5.3.1 Stability of LLM

During the experiment, we found that the LLM interface exhibited stability issues. Operators’ freedom to input any command increased the likelihood of ambiguous instructions and typographical errors. For instance, operators have placed commands outside the giving action, including “go but avoid the water” and “move the hand as close to the table as possible”. This highlights the need for further research to improve the robustness of LLMs in managing ambiguous inputs during robotic operations.

Also, even when commands were entered correctly, there was still a small chance of execution errors. There were instances where the robot executed a command twice, repeated a command continuously without stopping, or misinterpreted the number of steps required by a command.

One participant left the following feedback: “Language is easier to use due to the ability to execute sequences, although occasionally the sequence was not executed entirely correctly.”

5.3.2 Limitations of the experiment

In addition to the LLM being tested, visual feedback significantly impacts the system’s performance and workload. Participant feedback indicates that video resolution plays a more crucial role than time delay. Furthermore, only providing a single camera view for the operator in the experiment made it challenging to accurately determine the distance to an object. However, with higher resolution, participants were able to infer the relative position of other objects through shadows and reflections.

Additionally, mistakes can significantly impact task duration. The most common error is dropping the cube outside the bowl, requiring the robot to pick up the cube again, which is time-consuming. In

some cases, when the cube is dropped very close to the bowl, it requires fine movement to get the end-effector to the right position.

Another common mistake is perspective reversal. However, the participant usually can recover quickly from the mistake, and it only has a small impact on task time.

6 Related work

This section reviews various approaches that enhance the interaction between robots and human language, focusing on task planning through action composition with the integration of language models and robotic systems. The comparative analysis presented in Table 2 highlights distinct contributions in the existing approaches within the domain of robotic task execution using language models, including utilising open-source model, correcting mistakes on the fly with human and environment feedback, validating with real robot experiment, capability for ROS interface and fine-tuning, and integrating multimodal.

Kim et al. [25] conducted a user study on LLM-powered human-robot interaction. Their findings indicate that LLM-powered robots excel in tasks involving connection-building and deliberation but face challenges in logical communication and inducing anxiety.

6.1 Action composition from language

The integration of language models with robotic systems for task planning through action composition has been explored in various studies. Zeng et al. [59] developed Socratic Models to enhance multimodal capabilities without the need for fine-tuning. Li et al. [28] introduced Chain of Code, which leverages code execution to bolster reasoning processes. Kwon et al. [27] demonstrated the use of language models to predict robot end-effector poses based on visual inputs and task descriptions, showcasing direct application in robotic control. Silver et al. [44] explored how closed-sourced models like GPT-4 can generate Python programs for task planning within Planning Domain Definition Language (PDDL) domains, a framework with a long history in the field of automated planning and scheduling.

Song et al. [48] introduced the LLM-Planner for embodied agents, focusing on few-shot grounded planning to leverage LLMs, and emphasized the need for dynamic planning adaptability in complex environments. This work aligns closely with traditional symbolic planning [4] but introduces a novel integration of natural language processing, addressing both the symbolic and dynamic aspects of task planning. Wang et al. [55] developed LaMI, an LLM-based system designed to enhance multi-modal human-robot interaction. This system integrates high-level linguistic guidance, atomic actions, and multi-modal expressions to regulate robot behavior.

Izzo et al. [22] and Yang et al. [57] take a structured approach by translating natural language into behavior trees and state machines, respectively, which guides robotic behavior in a fixed and somewhat rigid framework. In contrast, our work emphasizes a more dynamic and flexible generation of executable action sequences directly from natural language prompts. This not only simplifies the user interface but also enhances the adaptability of the system to a wider range of tasks and environments without the constraints of pre-defined structures of behavior trees or state machines. Our method also corrects its mistakes based on human and environmental feedback, crucial for practical applications in dynamic environments, setting it apart from these more static and predefined approaches.

6.2 Incorporating human feedback

Most current systems operate in an open-loop configuration without the capability for error recovery. Shi et al. [43] introduce a framework where robots anticipate failures and proactively request human help to refine their task execution. Han et al. [14] develop a system that allows robots to explain their actions and receive corrective feedback from users, enhancing transparency and trust. Singh et al. [45] create an interactive prompting system that uses structured feedback to correct robot tasks, catering especially to technical and educational settings. Liu et al. [30] introduced OLAF (Operation-relabeled Learning with Language Feedback), a system that learns robot policies interactively using verbal corrections. It updates the robot’s visuomotor neural policy based on verbal feedback to avoid

Table 2: Comparison of related works on the Use of Open-Sourced Models, Human Feedback, Real Robot Experiments, and ROS Integration.

	Open-Source Model	Human Feedback	Environment Feedback	Real Robot Experiment	Remote Control	ROS Capability	Fine-Tuning Capability	Multimodal Integration
Ours	✓	✓	✓	✓	✓	✓	✓	WIP
Btgenbot [22]				✓		✓	✓	
GLM2FSA [57]			✓					
Socratic models [59]								✓
Chain of code [28]								
To help or not to help [49]			✓	✓				
Kwon et al. [27]				✓				
Statler [58]				✓				
Embodiedgpt [37]	✓						✓	✓
Huang, Abbeel, et al. [17]								
Yell at your robot [43]		✓	✓	✓			✓	✓
Rt-h [3]				✓			✓	✓
Vima [23]	✓						✓	✓
Natural language as policies [33]								✓
Clipswarm [40]	✓							✓
Progprompt [45]		✓		✓				
Llm-brain [31]	✓						✓	✓
Cao and Lee[6]								
Grounded decoding [18]				✓				✓
Silver, Tom, et al. [44]								
CLMASP [29]								
InterPreT [14]		✓		✓				
LLM-Planner [48]			✓					
OLAF [30]		✓		✓			✓	
LaMI [55]				✓		✓		✓

repeating mistakes. Our framework enables real-time interaction and task execution, facilitating natural user interaction between steps without requiring prior coding knowledge.

6.3 Utilizing open-source models

The use of open-source models in robotics research greatly facilitates accessibility and reproducibility, essential for advancing the field. Open-source frameworks allow researchers and developers to replicate studies and verify results, enhancing the credibility and scalability of the technologies developed. Mu et al. [37] and Huang [18] harness various open-source language models to drive their research, while Pueyo et al. [40] utilize the multimodal capabilities of the CLIP model to explore new robotic functionalities.

Our approach builds on these foundations by exclusively employing open-source models, which ensures that our methodologies are transparent and easily accessible to the robotics community. This commitment supports more stable development environments and reproducibility of results. Moreover, by integrating these models with real-world feedback mechanisms, our framework enhances the dynamic control capabilities necessary for practical applications, addressing critical gaps observed in previous studies like those by Cao and Lee [6], which lack empirical validation on actual robotic platforms. Thus, our work not only leverages the foundational benefits of open-source software to

ensure reproducibility and enhance community engagement but also pushes the boundaries of what these models can achieve in real-world settings.

6.4 Experiment on real robots

Experimenting with real robots is a crucial step in robotics research, as it allows for the validation and refinement of theoretical models. Several studies have implemented their methodologies on real robots, but the integration with widely used platforms like the Robot Operating System (ROS) remains limited. Tanneberg et al. [49], Kwon et al. [27], and others have demonstrated practical applications, yet the potential for easier study and development through ROS integration is largely untapped. Our framework aims to fill this gap, providing a robust solution that combines the benefits of open-source models, human feedback, and easy deployment. Furthermore, contrasting with simulation-only studies such as Pueyo et al. [40], our framework emphasizes practical, real-world testing to ensure operational reliability and scalability.

7 Use cases

In this section, we report use-cases of our framework in the wild.

7.1 LEJU Robotics Humanoid Kitchen Demonstration

The proposed framework has been recently adapted by LEJU Robotics to demonstrate their humanoid robot working in a kitchen setting, shown in Figure 12. The robot was tasked with cooking a meal based on user input, i.e. the human asked the robot to cook a particular dish. The system hardware includes the humanoid robot from LEJU, a Shiwan 3.0 Pro intelligent cooking machine from TINECO, and several pre-prepared ingredients placed in bowls. The humanoid robot is equipped with an Intel Realsense RGB-D camera on its head.

Atomic actions were developed for both the humanoid robot and the cooking machine system. The cooking machine actions are pre-set by the manufacturer, e.g. “heat pan to {TEMPERATURE}” where {TEMPERATURE} is the desired temperature. In the case of the humanoid robot, atomic actions are developed solely with imitation learning. For example, a control policy was learned for the atomic action with textual description “put the {INGREDIENT} in the bowl” where {INGREDIENT} is replaced with an ingredient such as “broccoli”. Several demonstrations of each atomic action are provided via a teleoperation interface. Once a database of demonstrations is collected, a policy is learned that maps the textual description and RGB image to the robot’s end-effector displacement.



Figure 12: LEJU Robotics Humanoid setup for a kitchen demonstration.

Subsequently, an inverse-kinematics controller translates these end-effector displacements into specific joint commands. Once refined, this learned control policy is systematically integrated into the library as a predefined atomic action.

The voice input from the user is recorded using a microphone and converted into text which is used as the task description. Given the task description, the language model generates an appropriate sequence for the atomic actions that is then executed on the system. Future development plans are to increase the atomic library, improve capabilities of the humanoid system, and replace the language model with a vision-language model.

7.2 Robotics Air Hockey Challenge

Bridging the simulation-reality gap is crucial for achieving high-performing embodied intelligence. Deploying learning approaches on real-world robots requires addressing practical challenges such as disturbances, observation noise, safety, model mismatches, delays, actuator limitations, physical feasibility, and limited real-world interactions. Additionally, future robots must dynamically react to their environments, execute agile movements, and engage in long-horizon planning. The Robot Air Hockey Challenge, shown in Figure 13, was developed as a collaborative platform for researchers to tackle a realistic robotic task. Teams designed and built air hockey agents, competed in various sub-tasks (in simulation) and finished in full games (both in simulation and the real world).

When humans learn a new skill (such as air hockey), it is common that a tutor/mentor will provide feedback and suggestions for their improvement. The learner internalizes this feedback and augments their behavior in such a way to improve their future performance. In contrast, current robot systems require experts to update a control policy directly to improve the model performance. We are utilizing our framework as an intuitive interface to enable a human non-expert to modify the control policy of the robot in order to improve the systems' performance within the context of the Robot Air Hockey Challenge. The goal being to evaluate the low-level atomic action execution performance across diverse scenarios. Experiments will be reported in future updates of this article.

8 Conclusions

In this section, we provide an overview of our main conclusions, acknowledge the current limitations of our framework and highlight potential future directions.

8.1 Overview

This work presents a ROS framework for intuitive robot programming, leveraging natural language prompts and contextual information from the robot environment and human feedback. Through a series of experiments conducted on a real-world robot setup, we have demonstrated the efficacy and versatility of our framework in enabling non-experts to program robots. The integration of open-source language models and common tools such as ROS with an AI agent represents a step towards realizing automated robotic solutions that can address real-world challenges in research and industry. We have reported two case studies that indicate how our framework is demonstrably accelerating companies and research groups into useful deployment in different types of environments and scales.

8.2 Limitations

It is important to acknowledge the current limitations of our system. Presently, our experiments rely solely on action sequences, which are not adaptable without human feedback. For example, if the robot must open a door and navigate through, then it may be worthwhile to attempt opening the door several times in case of failure of the first or subsequent attempts. This repetition could be achieved with a behavior tree, for example. We have recently added an interface that is able to execute behavior trees generated by the language model in the form of an XML file - this feature is already usable in our code base. Future work will investigate the benefits of using behavior trees in real-world laboratory setups.

Currently, we provide feedback, i.e. an observation, only in text form. This neglects potentially informative data such as images and interaction forces with humans.




Figure 13: The setup for the Robot Air Hockey Challenge with two KUKA IIWA robot arms.

Our framework provides a reward function in the problem formulation. One limitation is that the reward function is heavily dependent on the failure flag. In some cases of command, such as opening a door without specifying an angle, it can be difficult to determine whether the action is successful. Also, the design of the reward function focused on achieving task-specific outcomes. As tasks become more complex, involving a sequence of atomic actions, the dependency on a single reward signal at the end of the sequence may not provide enough guidance to the agent.

8.3 Future work

There are several avenues for future research and development we are interested in pursuing. We plan to extend our framework to more versatile robotic platforms, including quadrupeds and quadrupeds with an arm. Additionally, we aim to integrate our system with vision-language models to enhance its perceptual capabilities.

Moreover, while our current system supports ROS 1 (Noetic), we intend to provide support for ROS 2 (Humble) to broaden its compatibility and applicability. Additionally, we plan to explore the use of other behavior representations, such as state machines and behavior trees, to further enhance the flexibility and adaptability of our framework.

Whilst we have shown promise that pre-trained language models are capable of orchestrating action sequences for several tasks, there is an opportunity to fine-tune the language model. However, there are significant challenges associated with fine-tuning language models for robotics, notably the issue of the gap between simulation and reality. We aim to investigate the potential of fine-tuning language models within our framework.

Finally, as mentioned in the previous sub-section, our reward function is limited. Future work will additionally investigate reward shaping for more diversified tasks, including nonbinary rewards.

The standardization and versatility of our framework in a variety of industries and research settings are creating opportunities for new collaborations, faster development, and propelling newly developed technologies forward. This trend will likely continue to manifest in the coming years as the ROS-LLM framework continues to mature.

In our approach, atomic actions are pre-trained or pre-defined policies. In future work, we may be able to fine-tune these individual policies to have faster and more effective atomic actions.

References

- [1] Andrew G. Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4):341–379, Oct 2003.
- [2] Patrick Beeson and Barrett Ames. Trac-ik: An open-source library for improved solving of generic inverse kinematics. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 928–935, 2015.
- [3] Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. Rt-h: Action hierarchies using language, 2024.
- [4] Calin Belta, Antonio Bicchi, Magnus Egerstedt, Emilio Frazzoli, Eric Klavins, and George J. Pappas. Symbolic planning and control of robot motion [grand challenges of robotics]. *IEEE Robotics & Automation Magazine*, 14(1):61–70, 2007.
- [5] H. Bruyninckx. Open robot control software: the orocos project. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, volume 3, pages 2523–2528 vol.3, 2001.
- [6] Yue Cao and C. S. George Lee. Robot behavior-tree-based task generation with large language models, 2023.
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. 2021.
- [8] Sachin Chitta, Eitan Marder-Eppstein, Wim Meeussen, Vijay Pradeep, Adolfo Rodríguez Tsouroukdissian, Jonathan Bohren, David Coleman, Bence Magyar, Gennaro Raiola, Mathias Lüdtke, and Enrique Fernandez Perdomo. ros_control: A generic and simple control framework for ros. *Journal of Open Source Software*, 2(20):456, 2017.
- [9] David Coleman, Ioan Sucan, Sachin Chitta, and Nikolaus Correll. Reducing the barrier to entry of complex robotic software: a moveit! case study. *arXiv preprint arXiv:1404.3785*, 2014.
- [10] John G. Everett and Alexander H. Slocum. Automation and robotics opportunities: Construction versus manufacturing. *Journal of Construction Engineering and Management*, 120(2):443–452, 1994.
- [11] Tully Foote. tf: The transform library. In *2013 IEEE Conference on Technologies for Practical Robot Applications (TePRA)*, pages 1–6, 2013.
- [12] Simon Frieder, Luca Pinchetti, Alexis Chevalier, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiwicz, Philipp Christian Petersen, and Julius Berner. Mathematical capabilities of chatgpt, 2023.
- [13] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming – the rise of code intelligence, 2024.
- [14] Muzhi Han, Yifeng Zhu, Song-Chun Zhu, Ying Nian Wu, and Yuke Zhu. Interpret: Interactive predicate learning from language feedback for generalizable task planning, 2024.
- [15] Sandra G Hart. NASA task load index (TLX). 1986.
- [16] Albert S. Huang, Edwin Olson, and David C. Moore. Lcm: Lightweight communications and marshalling. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4057–4062, 2010.
- [17] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents, 2022.

- [18] Wenlong Huang, Fei Xia, Dhruv Shah, Danny Driess, Andy Zeng, Yao Lu, Pete Florence, Igor Mordatch, Sergey Levine, Karol Hausman, and Brian Ichter. Grounded decoding: Guiding text generation with grounded models for embodied agents, 2023.
- [19] Martin Huber, Christopher E. Mower, Sebastien Ourselin, Tom Vercauteren, and Christos Bergeles. Lbr-stack: Ros 2 and python integration of kuka fri for med and iiwa robots, 2024.
- [20] Arda Inceoglu, Eren Erdal Aksoy, and Sanem Sariel. Multimodal detection and classification of robot manipulation failures. *IEEE Robotics and Automation Letters*, 9(2):1396–1403, 2024.
- [21] Matteo Iovino, Edvards Scukins, Jonathan Styrod, Petter Ögren, and Christian Smith. A survey of behavior trees in robotics and ai. *Robotics and Autonomous Systems*, 154:104096, 2022.
- [22] Riccardo Andrea Izzo, Gianluca Bardaro, and Matteo Matteucci. Btgenbot: Behavior tree generation for robotic tasks with lightweight llms, 2024.
- [23] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. Vima: General robot manipulation with multimodal prompts, 2023.
- [24] Ehsan Kamaloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. HA-GRID: A human-llm collaborative dataset for generative information-seeking with attribution. *arXiv:2307.16883*, 2023.
- [25] Callie Y Kim, Christine P Lee, and Bilge Mutlu. Understanding large-language model (llm)-powered human-robot interaction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, pages 371–380, 2024.
- [26] Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2149–2154, Sendai, Japan, Sep 2004.
- [27] Teyun Kwon, Norman Di Palo, and Edward Johns. Language models as zero-shot trajectory generators, 2023.
- [28] Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, and Brian Ichter. Chain of code: Reasoning with a language model-augmented code emulator, 2023.
- [29] Xinrui Lin, Yangfan Wu, Huanyu Yang, Yu Zhang, Yanyong Zhang, and Jianmin Ji. Cmasp: Coupling large language models with answer set programming for robotic task planning, 2024.
- [30] Huihan Liu, Alice Chen, Yuke Zhu, Adith Swaminathan, Andrey Kolobov, and Ching-An Cheng. Interactive robot learning from verbal correction, 2023.
- [31] Artem Lykov and Dzmitry Tsetserukou. Llm-brain: Ai-driven fast generation of robot behaviour tree based on large language model, 2023.
- [32] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot operating system 2: Design, architecture, and uses in the wild. *Science Robotics*, 7(66):eabm6074, 2022.
- [33] Yusuke Mikami, Andrew Melnik, Jun Miura, and Ville Hautamäki. Natural language as policies: Reasoning for coordinate-level embodied control with llms, 2024.
- [34] Christopher Mower, Theodoros Stouraitis, Joao Moura, Christian Rauch, Lei Yan, Nazanin Zamani Behabadi, Michael Gienger, Tom Vercauteren, Christos Bergeles, and Sethu Vijayakumar. Ros-pybullet interface: A framework for reliable contact simulation and human-robot interaction. In *Conference on Robot Learning*, pages 1411–1423. PMLR, 2023.
- [35] Christopher E Mower, Joao Moura, and Sethu Vijayakumar. Skill-based Shared Control. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.
- [36] Christopher E. Mower, João Moura, Nazanin Zamani Behabadi, Sethu Vijayakumar, Tom Vercauteren, and Christos Bergeles. Optas: An optimization-based task specification library for trajectory optimization and model predictive control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9118–9124, 2023.
- [37] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought, 2023.

- [38] Edwin Olson. Apriltag: A robust and flexible visual fiducial system. In *2011 IEEE International Conference on Robotics and Automation*, pages 3400–3407, 2011.
- [39] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Curry, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lillian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024.
- [40] Pablo Pueyo, Eduardo Montijano, Ana C. Murillo, and Mac Schwager. Clipswarm: Generating drone shows from text prompts with vision-language models, 2024.
- [41] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009.
- [42] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pages 28492–28518. PMLR, 2023.

- [43] Lucy Xiaoyang Shi, Zheyuan Hu, Tony Z. Zhao, Archit Sharma, Karl Pertsch, Jianlan Luo, Sergey Levine, and Chelsea Finn. Yell at your robot: Improving on-the-fly from language corrections, 2024.
- [44] Tom Silver, Soham Dan, Kavitha Srinivas, Joshua B. Tenenbaum, Leslie Pack Kaelbling, and Michael Katz. Generalized planning in pddl domains with pretrained large language models, 2023.
- [45] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. Progprompt: Generating situated robot task plans using large language models, 2022.
- [46] Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkumar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, Aug 2023.
- [47] R. Smits. KDL: Kinematics and Dynamics Library. <http://www.orocos.org/kdl>.
- [48] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2998–3009, 2023.
- [49] Daniel Tanneberg, Felix Ocker, Stephan Hasler, Joerg Deigmoeller, Anna Belardinelli, Chao Wang, Heiko Wersing, Bernhard Sendhoff, and Michael Gienger. To help or not to help: Llm-based attentive support for human-robot group interactions, 2024.
- [50] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023.
- [51] Trieu H. Trinh, Yuhuai Wu, Quoc V. Le, He He, and Thang Luong. Solving olympiad geometry without human demonstrations. *Nature*, 625(7995):476–482, Jan 2024.
- [52] Rasul Tutunov, Antoine Grosnit, Juliusz Ziomek, Jun Wang, and Haitham Bou-Ammar. Why can large language models generate correct chain-of-thoughts? 2023.
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [54] Yuhui Wan, Jingcheng Sun, Christopher Peers, Joseph Humphreys, Dimitrios Kanoulas, and Chengxu Zhou. Performance and usability evaluation scheme for mobile manipulator teleoperation. *IEEE Transactions on Human-Machine Systems*, 2023.
- [55] Peng Wang, Mattia Robbiani, and Zhihao Guo. Llm granularity for on-the-fly robot control, 2024.
- [56] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. 2022.
- [57] Yunhao Yang, Jean-Raphael Gaglione, Cyrus Neary, et al. Large language models for verifiable sequential decision-making in autonomous systems. In *2nd Workshop on Language and Robot Learning: Language as Grounding*, 2023.
- [58] Takuma Yoneda, Jiading Fang, Peng Li, Huanyu Zhang, Tianchong Jiang, Shengjie Lin, Ben Picker, David Yunis, Hongyuan Mei, and Matthew R. Walter. Statler: State-maintaining language models for embodied reasoning, 2023.
- [59] Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, and Pete Florence. Socratic models: Composing zero-shot multimodal reasoning with language, 2022.