

DO AS I CAN, NOT AS I SAY: GROUNDING LANGUAGE IN ROBOTIC AFFORDANCES

¹ Michael Ahn*, Anthony Brohan*, Noah Brown*, Yevgen Chebotar*, Omar Cortes*, Byron David*, Chelsea Finn*, Keerthana Gopalakrishnan*, Karol Hausman*, Alex Herzog[†], Daniel Ho[†], Jasmine Hsu*, Julian Ibarz*, Brian Ichter*, Alex Irpan*, Eric Jang*, Rosario Jauregui Ruano*, Kyle Jeffrey*, Sally Jesmonth*, Nikhil J Joshi*, Ryan Julian*, Dmitry Kalashnikov*, Yuheng Kuang*, Kuang-Huei Lee*, Sergey Levine*, Yao Lu*, Linda Luu*, Carolina Parada*, Peter Pastor[†], Jornell Quiambao*, Kanishka Rao*, Jarek Rettinghouse*, Diego Reyes*, Pierre Sermanet*, Nicolas Sievers*, Clayton Tan*, Alexander Toshev*, Vincent Vanhoucke*, Fei Xia*, Ted Xiao*, Peng Xu*, Sichun Xu[†], Mengyuan Yan[†]

*Robotics at Google, [†]Everyday Robots

ABSTRACT

Large language models can encode a wealth of semantic knowledge about the world. Such knowledge could be extremely useful to robots aiming to act upon high-level, temporally extended instructions expressed in natural language. However, a significant weakness of language models is that they lack real-world experience, which makes it difficult to leverage them for decision making within a given embodiment. For example, asking a language model to describe how to clean a spill might result in a reasonable narrative, but it may not be applicable to a particular agent, such as a robot, that needs to perform this task in a particular environment. We propose to provide real-world grounding by means of pretrained skills, which are used to constrain the model to propose natural language actions that are both feasible and contextually appropriate. The robot can act as the language model’s “hands and eyes,” while the language model supplies high-level semantic knowledge about the task. We show how low-level skills can be combined with large language models so that the language model provides high-level knowledge about the procedures for performing complex and temporally extended instructions, while value functions associated with these skills provide the grounding necessary to connect this knowledge to a particular physical environment. We evaluate our method on a number of real-world robotic tasks, where we show the need for real-world grounding and that this approach is capable of completing long-horizon, abstract, natural language instructions on a mobile manipulator. The project’s website and the video can be found at say-can.github.io.

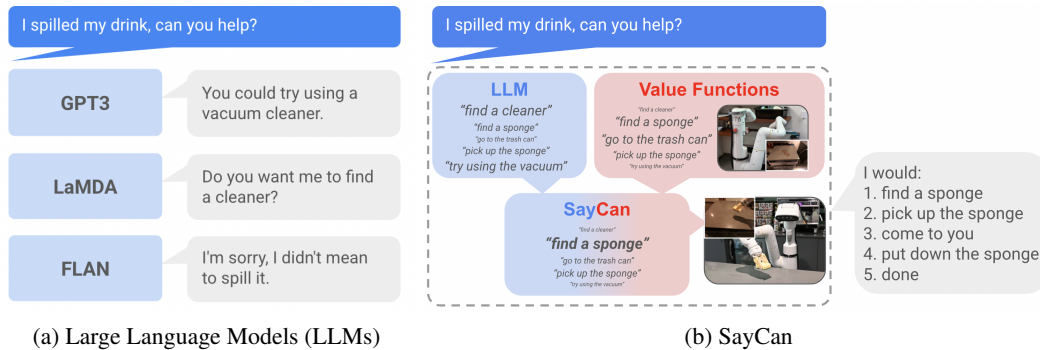


Figure 1: The nature of LLM training is such that they have never interacted with its environment and observed the outcome of its responses, and thus are not grounded in the world. SayCan grounds LLMs through value functions of pretrained skills, which act as an LLM’s “hands and eyes”, allowing them to be used to execute real-world, abstract, long-horizon commands on robots.

¹Authors listed in alphabetical order. Contributions in Appendix A.

1 INTRODUCTION

Recent progress in training large language models (LLMs) has led to systems that can generate long text based on prompts, answer questions, or even engage in a dialog on a wide range of topics. To do so, these models absorb vast quantities of knowledge from text corpora mined from the web. It is natural to consider whether the knowledge of everyday tasks that is encoded in such models can be used to broaden the set of tasks robots are able to plan and execute, e.g., enabling a robot to perform abstract, temporally-extended textual instructions. This bears the question: how can embodied agents extract and harness the knowledge of LLMs for physically-grounded tasks?

This question poses a major challenge – LLMs are inherently not grounded in the physical world. The nature of the training process means that an LLM has never interacted with its environment, has never observed the outcomes of its responses, and lacks any ability to observe its environment beyond what is presented to it in text (Bender & Koller, 2020). This can lead LLMs to not only make mistakes that seem unreasonable or humorous to people, but also to interpret instructions in ways that are nonsensical or unsafe for a particular embodiment or situation. Figure 1 shows an example – a kitchen robot capable of executing skills such as “pick up the sponge”, “pick up the coke can”, or “go to the table” may be asked for help cleaning up a spill (“I spilled my drink, can you help?”). A language model responds with a reasonable narrative, but one that is not feasible or useful for the robot (“You could try using a vacuum cleaner” may be impossible if there is no vacuum in the scene or if the robot is incapable of operating a vacuum). The LLM has no knowledge of the current environment it is operating in, nor the capabilities of the agent that is following the instructions. With prompt engineering, a language model may be capable of splitting the high-level instruction into sub-tasks, but it cannot effectively do so without the context of what the robot is capable of given its abilities *and* the current state of the robot *and* the environment. In short, an LLM without grounding is divorced from reality. Unlike language models, robotic systems deal directly with perception and interaction, which leads to the question of whether such systems can effectively ground the semantic knowledge in language models.

Motivated by this question, we study the problem of how to extract the knowledge in LLMs for enabling an embodied agent, such as a robot, to follow high-level textual instructions. The robot is equipped with a repertoire of learned skills for basic “atomic” behaviors, and these skills are capable of low-level perception and control. Our main insight is that, if we can simply make the LLM aware of the available and feasible repertoire of skills, this can provide it with an awareness of *both* the agent’s capabilities and the current state of the environment. In this sense, the skills will act as the LLMs “hands and eyes.”

The main principle that we use to connect LLMs to physical tasks is to observe that, in addition to asking the LLM to simply interpret an instruction, we can use it to score the likelihood that an individual skill makes progress towards completing the high-level instruction. Furthermore, if each skill has an accompanying affordance function that quantifies how likely it is to succeed from the current state (such as a learned value function), its value can be used to weight the skill’s likelihood. This results in an elegant probabilistic interpretation of our method, where the LLM describes the probability that each skill aids in completing the instruction, and the affordance function describes the probability that each skill will succeed – combining the two provides the probability that each skill will perform the instruction *successfully*. The affordance functions make the LLM aware of the current scene, and constraining the completions to the skill descriptions makes the LLM aware of the robot’s capabilities. Furthermore, this combination results in a fully explainable sequence of steps that the robot will execute to accomplish an instruction – an interpretable plan that is expressed through natural language.

In this work we present SayCan, an algorithm that extracts and leverages the knowledge within LLMs in physically-grounded tasks. The large language model (Say) provides a task-grounding to determine useful actions to accomplish a high-level goal and the robot-learned affordance functions (Can) provide a world-grounding to determine what is possible and to execute upon the plan. In particular, we use reinforcement learning (RL) as a way to learn language-conditioned value functions that provide affordances of what is possible in the world. We evaluate the proposed approach on about 100 real-world robotic tasks that involve a mobile robot accomplishing a large set of language instructions in a real kitchen in a zero-shot fashion. Our experiments validate that SayCan can execute temporally-extended, complex, and abstract instructions. Grounding the LLM in the real-world via affordances nearly doubles the performance over the non-grounded baselines.

2 PRELIMINARIES

Large Language Models. Language models seek to model the probability $p(W)$ of a text $W = \{w_0, w_1, w_2, \dots, w_n\}$, a sequence of strings w . This is generally done through factorizing the probability via the chain rule to be $p(W) = \prod_{j=0}^n p(w_j | w_{<j})$, such that each successive string is predicted from the previous. Recent breakthroughs initiated by neural network-based Attention architectures (Vaswani et al., 2017) have enabled efficient scaling of so-called Large Language Models (LLMs). Such models include Transformers (Vaswani et al., 2017), BERT (Devlin et al., 2018), T5 (Raffel et al., 2019), GPT-3 (Brown et al., 2020), Gopher (Rae et al., 2021), LAMDA (Thopilan et al., 2022), FLAN (Wei et al., 2021), and PaLM Chowdhery et al. (2022), each showing increasingly large capacity (billions of parameters and terabytes of text) and subsequent ability to generalize across tasks.

In this work, we utilize the vast semantic knowledge contained in LLMs to determine useful tasks for solving high-level instructions.

Value functions and RL. Our goal is to be able to accurately predict whether a skill (given by a language command) is feasible at a current state. We use temporal-difference-based (TD) reinforcement learning to accomplish this goal. In particular, we define a Markov decision process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where \mathcal{S} and \mathcal{A} are state and action spaces, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}_+$ is a state-transition probability function, $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function and γ is a discount factor. The goal of TD methods is to learn state or state-action value functions (Q-function) $Q^\pi(s, a)$, which represents the discounted sum of rewards when starting from state s and action a , followed by the actions produced by the policy π , i.e. $Q^\pi(s, a) = \mathbb{E}_{a \sim \pi(a|s)} \sum_t R(s_t, a_t)$. The Q-function, $Q^\pi(s, a)$ can be learned via approximate dynamic programming approaches that optimize the following loss: $L_{TD}(\theta) = \mathbb{E}_{(s, a, s') \sim \mathcal{D}} [R(s, a) + \gamma \mathbb{E}_{a' \sim \pi} Q_\theta^\pi(s', a') - Q_\theta^\pi(s, a)]$, where \mathcal{D} is the dataset of states and actions and θ are the parameters of the Q-function.

In this work, we utilize TD-based methods to learn said value function that is additionally conditioned on the language command and utilize those to determine whether a given command is feasible from the given state. It is worth noting that in the sparse reward case, where the agent receives the reward of 1.0 at the end of the episode if it was successful and 0.0 otherwise, the value function trained via RL corresponds to an affordance function (Gibson, 1977) that specifies whether a skill is possible in a given state. We leverage that intuition in our setup and express affordances via value functions of sparse reward tasks.

3 SAYCAN: DO AS I CAN, NOT AS I SAY

In this section we present our algorithm, SayCan, which adapts the knowledge contained in large language models (Say) to embodied tasks by leveraging real-world-grounded value functions (Can). The core idea behind our method is to use LLMs to provide task-grounding to an agent that is capable of executing short-horizon commands and, at the same time, use the agent’s value functions to provide the real-world grounding necessary for making the output of a LLM useful to the agent.

At a high level, we start by querying the LLM about how it would split an instruction commanded by a user into smaller tasks that the agent is capable of executing. The initial user instruction can be ambiguous, abstract, or temporally extended, and it is the role of the LLM to disambiguate it to the level that it can be turned into a sequence of atomic steps that the robot can execute. We consider a set of available skills that the agent is able to perform and use their value functions that are able to specify which proposal skills are feasible given the current state of the agent and its environment. This effective filtering mechanism allows us to ground the knowledge extracted from the LLM to the real world of the agent’s embodiment.

3.1 PROBLEM STATEMENT

Our system receives a user-provided natural language instruction i that describes a task that the robot should execute. The instruction can be long, abstract, or ambiguous. We also assume that we are given a set of skills Π , where each skill $\pi \in \Pi$ is a policy that performs some short skill, such as picking up a particular object. Each skill $\pi \in \Pi$ has a short language description ℓ_π (e.g., “find

a sponge”), and an accompanying affordance function $p(c_\pi|s, \ell_\pi)$ that indicates the probability of c -ompleting the skill with description ℓ_π successfully from state s . Intuitively, $p(c_\pi|s, \ell_\pi)$ means “if I ask the robot to do ℓ_π , will it do it?”. In RL terminology, $p(c_\pi|s, \ell_\pi)$ corresponds to the value function for the skill if we take the reward to be 1 for successful completion and 0 otherwise, though more generally it can be trained as an affordance classifier. In our experiments, we use affordance functions obtained from RL value functions.

As mentioned above, ℓ_π denotes the textual label of skill π and $p(c_\pi|s, \ell_\pi)$ denotes the probability that skill π with textual label ℓ_π successfully completes if executed from state s , where c_π is a Bernoulli random variable. The LLM provides us with $p(\ell_\pi|i)$, the probability that a skill’s textual label is a valid next step for the user’s instruction. However, what we are interested in is the probability that a given skill successfully makes progress toward actually completing the instruction, which we denote as $p(c_i|i, s, \ell_\pi)$. Assuming that a skill that succeeds makes progress on i with probability $p(\ell_\pi|i)$ (i.e., its probability of being the right skill), and a skill that fails makes progress with probability zero, we can factorize this as $p(c_i|i, s, \ell_\pi) \propto p(c_\pi|s, \ell_\pi)p(\ell_\pi|i)$. This corresponds to multiplying the probability of the language description of the skill given the instruction $p(\ell_\pi|i)$, which we refer to as task-grounding, and the probability of the skill being possible in the current state of the world $p(c_\pi|s, \ell_\pi)$, which we refer to as world-grounding.

3.2 CONNECTING LARGE LANGUAGE MODELS TO ROBOTS

While large language models can draw on a wealth of knowledge learned from copious amounts of text, they will not necessarily break down high-level commands into low-level instructions that are suitable for robotic execution. If a language model were asked “how would a robot bring me an apple”, it may respond “a robot could go to a nearby store and purchase an apple for you”. Though this response is a reasonable completion for the prompt, it is not necessarily actionable to an embodied agent, which may have a fixed set of abilities and a narrow operational regime. Therefore, to adapt language models to our problem statement, we must somehow inform them that we specifically want the high-level instruction to be broken down into sequences of available low-level skills, and these skills should be appropriate based on a current observation of the world. One approach to alleviate such issues is careful prompt engineering (Brown et al., 2020; Ouyang et al., 2022), a technique to coax a language model to a specific response structure. Prompt engineering provides examples in the context text (“prompt”) for the model that specify the task and the response structure; the prompt used in this work is shown in Appendix C.2 along with experiments ablating it. This makes solutions to the task more likely than superfluous completions. However, this is not enough to fully constrain the output to admissible primitive skills for an embodied agent, and indeed at times it can produce inadmissible actions or language that is not formatted in a way that is easy to parse into individual steps.

Scoring language models open an avenue to constrained responses by outputting the probabilities assigned by a language model to fixed outputs. Recall that a language model represents a *distribution* over potential completions $p(w_k|w_{<k})$, where w_k is a word that appears at a k^{th} position in a text. While typical generation applications (e.g., conversational agents) might sample from this distribution or decode the maximum likelihood completion, we can also use the model to *score* a candidate completion selected from a set of available options. Formally in SayCan, given a set of low-level skills Π , their language descriptions ℓ_Π and an instruction i , we compute the probability of a language description of a skill $\ell_\pi \in \ell_\Pi$ making progress towards executing the instruction i : $p(\ell_\pi|i)$, which corresponds to querying the model over potential completions. The optimal skill according to the language model is computed via $\ell_\pi = \arg \max_{\ell_\pi \in \ell_\Pi} p(\ell_\pi|i)$. Once selected, the process proceeds by iteratively selecting a skill and appending it to the instruction. Practically, in this work we structure the planning as a dialog between a user and a robot, in which a user provides the high level-instruction (e.g., “How would you bring me a coke can?”) and the language model responds with an explicit sequence (“I would: 1. ℓ_π ”, e.g., “I would: 1. find a coke can, 2. pick up the coke can, 3. bring it to you”).

This has the added benefit of interpretability, as the model not only outputs likely generative responses, but also gives a notion of likelihood across many possible responses. Figure 2 shows this full process of forcing the language model into a language pattern, where the set of tasks are the skills the low-level policy is capable of performing and prompt engineering shows several plan examples and dialog between the user and the robot. With this approach, we find that we are able

to effectively extract knowledge from the language model, but it leaves a major issue: while the decoding of the instruction obtained in this way always consists of skills that are available to the robot, these skills may not necessarily be appropriate for executing the desired high-level task in the *specific* situation that the robot is currently in. For example, if I ask a robot to “bring me an apple”, the optimal set of skills changes if there is no apple in view or if it already has one in its hand.

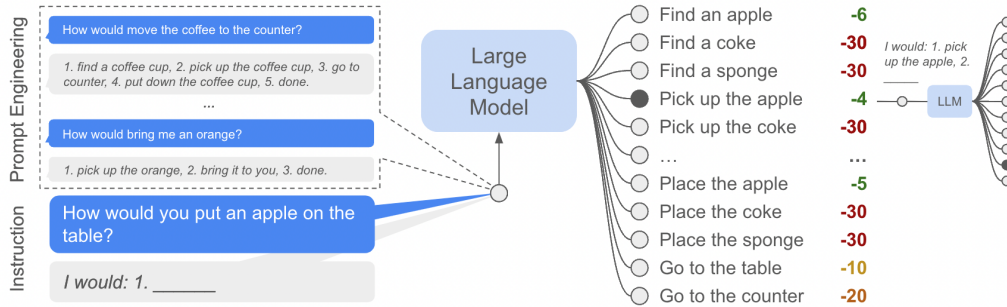


Figure 2: A scoring language model is queried with a prompt-engineered context of examples and the high-level instruction to execute and outputs the probability of each skill being selected. To iteratively plan the next steps, the selected skill is added to the natural language query and the language model is queried again.

SayCan. The key idea of SayCan is to ground large language models through value functions – affordance functions that capture the log likelihood that a particular skill will be able to succeed in the current state. Given a skill $\pi \in \Pi$, its language description ℓ_π and its corresponding value function, which provides $p(c_\pi|s, \ell_\pi)$, the probability of c -ompletion for the skill described by ℓ_π in state s , we form an affordance space $\{p(c_\pi|s, \ell_\pi)\}_{\pi \in \Pi}$. This value function space captures affordances across all skills (Shah et al., 2022). For each skill, the affordance function and the LLM probability are then multiplied together and ultimately the most probable skill is selected, i.e. $\pi = \arg \max_{\pi \in \Pi} p(c_\pi|s, \ell_\pi)p(\ell_\pi|i)$. Once the skill is selected, the corresponding policy is executed by the agent and the LLM query is amended to include ℓ_π and the process is run again until a termination token (e.g., “done”) is chosen. A full description of this process is provided in Algorithm 1 and Figure 4. These two mirrored processes together lead to a probabilistic interpretation of SayCan, where the LLM provides probabilities of a skill being useful for the high-level instruction and the affordances provide probabilities of successfully executing each skill. Combining these two probabilities together provides a probability that this skill furthers the execution of the high-level instruction commanded by the user.

4 IMPLEMENTING SAYCAN IN A ROBOTIC SYSTEM

In this section, we describe our implementation of SayCan in a robotic system that can perform a variety of user instructions in a kitchen environment. We first discuss the language conditioning for the control policies, then present details of how the policies and their corresponding value functions

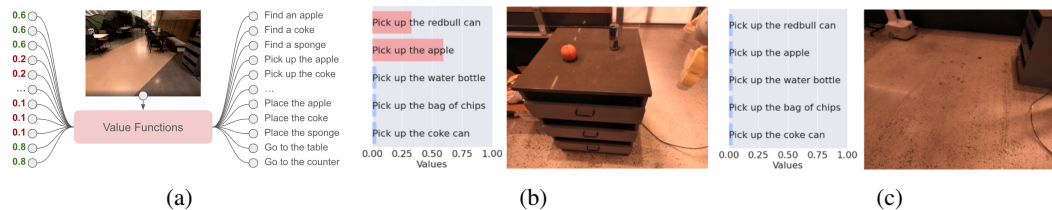


Figure 3: A value function module (a) is queried to form a value function space of action primitives based on the current observation. Visualizing “pick” value functions, in (b) “Pick up the red bull can” and “Pick up the apple” have high values because both objects are in the scene, while in (c) the robot is navigating an empty space, and thus none of the pick up actions receive high values.

Algorithm 1 SayCan

Given: A high level instruction i , state s_0 , and a set of skills Π and their language descriptions ℓ_Π

```
1:  $n = 0, \pi = \emptyset$ 
2: while  $\ell_{\pi_{n-1}} \neq \text{“done”}$  do
3:    $\mathcal{C} = \emptyset$ 
4:   for  $\pi \in \Pi$  and  $\ell_\pi \in \ell_\Pi$  do
5:      $p_\pi^{\text{LLM}} = p(\ell_\pi | i, \ell_{\pi_{n-1}}, \dots, \ell_{\pi_0})$  ▷ Evaluate scoring of LLM
6:      $p_\pi^{\text{affordance}} = p(c_\pi | s_n, \ell_\pi)$  ▷ Evaluate affordance function
7:      $p_\pi^{\text{combined}} = p_\pi^{\text{affordance}} p_\pi^{\text{LLM}}$ 
8:      $\mathcal{C} = \mathcal{C} \cup p_\pi^{\text{combined}}$ 
9:   end for
10:   $\pi_n = \arg \max_{\pi \in \Pi} \mathcal{C}$ 
11:  Execute  $\pi_n(s_n)$  in the environment, updating state  $s_{n+1}$ 
12:   $n = n + 1$ 
13: end while
```

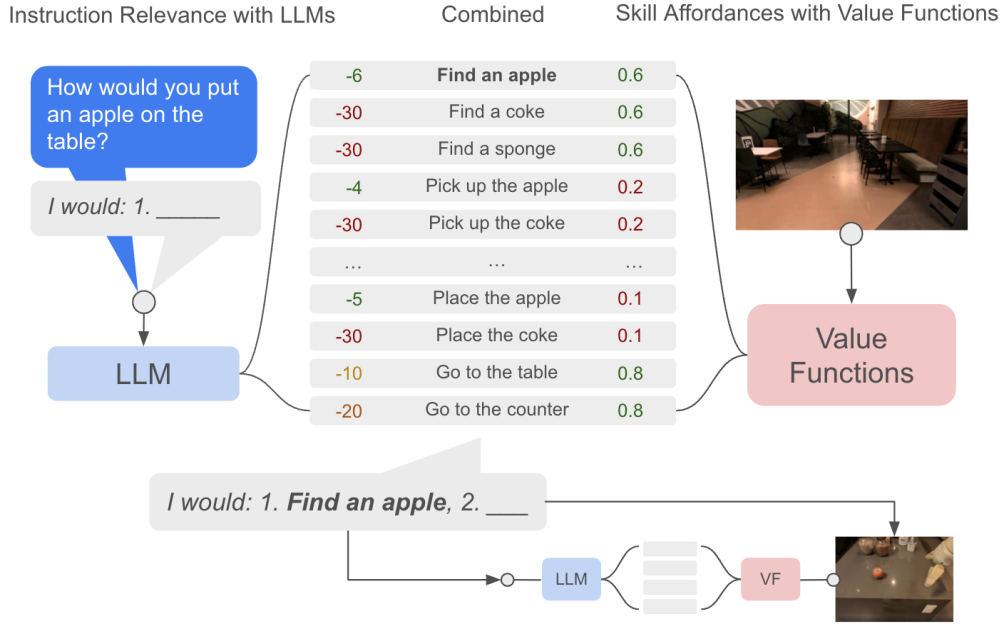


Figure 4: Given a high-level instruction, SayCan combines probabilities from a language model (representing the probability that a skill is useful for the instruction) with the probabilities from a value function (representing the probability of successfully executing said skill) to select the skill to perform. This emits a skill that is both possible and useful. The process is repeated by appending the selected skill to the robot response and querying the models again, until the output step is to terminate.

are trained, and finally describe the specifics of the real-world robotic system and the skills in our experimental evaluation.

4.1 LANGUAGE-CONDITIONED ROBOTIC CONTROL POLICIES

To instantiate SayCan, we must provide it with a set of skills, each of which has a policy, a value function, and a short language description (e.g., “pick up the can”). These skills, value functions, and descriptions can be obtained in a variety of different ways. In our implementation, we train the individual skills either with image-based behavioral cloning, following the BC-Z method (Jang et al., 2021), or reinforcement learning, following MT-Opt (Kalashnikov et al., 2021). Regardless of how the skill’s policy is obtained, we utilize value functions trained via TD backups as the affordance model for that skill. While we find that the BC policies achieve higher success rates at the current

stage of our data collection process, the value functions provided by the RL policies are crucial as an abstraction to translate control capabilities to a semantic understanding of the scene. In order to amortize the cost of training many skills, we utilize multi-task BC and multi-task RL, respectively, where instead of training a separate policy and value function per skill, we train multi-task policies and models that are *conditioned* on the language description. Note, however, that this description only corresponds to *low level* skills – it is still the role of the LLM in SayCan to interpret the high-level instruction and break it up into individual low level skill descriptions.

To condition the policies on language, we utilize a pre-trained large sentence encoder language model (Cer et al., 2018). We freeze the language model parameters during training and use the embeddings generated by passing in text descriptions of each skill. These text embeddings are used as the input to the policy and value function that specify which skill should be performed (see the details of the architectures used in the Appendix B.1). Since the language model used to generate the text embeddings is not necessarily the same as the language model used for planning, SayCan is able to utilize different language models well suited for different abstraction levels – understanding planning with respect to many skills as opposed to expressing specific skills more granularly.

4.2 TRAINING THE LOW-LEVEL SKILLS

We utilize both BC and RL policy training procedures to obtain the language-conditioned policies and value functions, respectively. In the following, we describe the corresponding training procedures for each of the methods.

Skill specification, reward functions, and action space. To complete the description of the underlying MDP that we consider, we provide the reward function as well as the skill specification that is used by the policies and value functions. As mentioned previously, for skill specification we use a set of short, natural language descriptions that are represented as language model embeddings. We utilize sparse reward functions with reward values of 1.0 at the end of an episode if the language command was executed successfully, and 0.0 otherwise. The success of language command execution is rated by humans where the raters are given a video of the robot performing the skill, together with the given instruction. If two out of the three raters agree that the skill was accomplished successfully, the episode is labelled with a positive reward.

To additionally process the data, we also ask the raters to mark the episodes as unsafe (i.e., if the robot collided with the environment), undesirable (i.e., if the robot perturbed objects that were not relevant to the skill) or infeasible (i.e., if the skill cannot be done or is already accomplished). If any of these conditions are met, the episode is excluded from training.

The action space of our policies includes the six degrees of freedom of the end-effector pose as well as gripper open and close commands, x-y position and yaw orientation delta of the mobile base of the robot, and the *terminate* action.

BC training. We use 68000 teleoperated demonstrations that were collected over the course of 11 months using a fleet of 10 robots. The operators use VR headset controllers to track the motion of their hand, which is then mapped onto the robot’s end-effector pose. The operators can also use a joystick to move the robot’s base. We additionally expand the demonstration dataset with 276000 autonomous episodes of learned policies which are later success-filtered and included in BC training, resulting in an additional 12000 successful episodes. To learn language-conditioned BC policies at scale in the real world, we build on top of BC-Z (Jang et al., 2021) and use a similar policy-network architecture (shown in Fig. 9). The details of the BC training procedure are in Appendix B.2.

RL training. In addition to using demonstrations in the BC setup, we also learn language-conditioned value functions with RL. For this purpose, we complement our real robot fleet with a simulated version of the skills and environment. To reduce the simulation-to-real gap we transform robot images via RetinaGAN (Ho et al., 2021) to look more realistic while preserving general object structure. In order to learn a language-conditioned RL policy, we utilize MT-Opt (Kalashnikov et al., 2021) in the Everyday Robots simulator using said simulation-to-real transfer. We bootstrap the performance of simulation policies by utilizing simulation demonstrations to provide initial successes, and then continuously improve the RL performance with online data collection in simulation. We

use a network architecture similar to MT-Opt (shown in Fig. 8) and provide the details of the RL training procedure in Appendix B.2.

4.3 ROBOTIC SYSTEM AND SKILLS

For the underlying control policies, we study a diverse set of manipulation and navigation skills using a mobile manipulator robot. Inspired by common skills one might pose to a robot in a kitchen environment, we propose 551 skills that span a set of seven skill families and 17 objects, which include picking and placing objects, rearranging objects, opening and closing drawers, navigating to various locations, and placing objects in a specific configurations. In this study we utilize the skills that are most amenable to more complex behaviors via composition and planning as well as those that have high performance at the current stage of data collection; for more details, see Appendix C.

5 EXPERIMENTAL EVALUATION

In our experiments, we evaluate SayCan with a mobile manipulator and a set of object manipulation and navigation skills in two office kitchen environments. We evaluate the full system and compare to baseline methods that lack the embodied grounding of value functions or that lack the semantic reasoning of a large language model. We find that SayCan can perform even temporally extended tasks from high-level instructions, and correctly adjusts the selected behaviors to the particular situation and embodiment. Additional experiments to ablate parameters and to evaluate the low-level policies are presented in Appendix D.

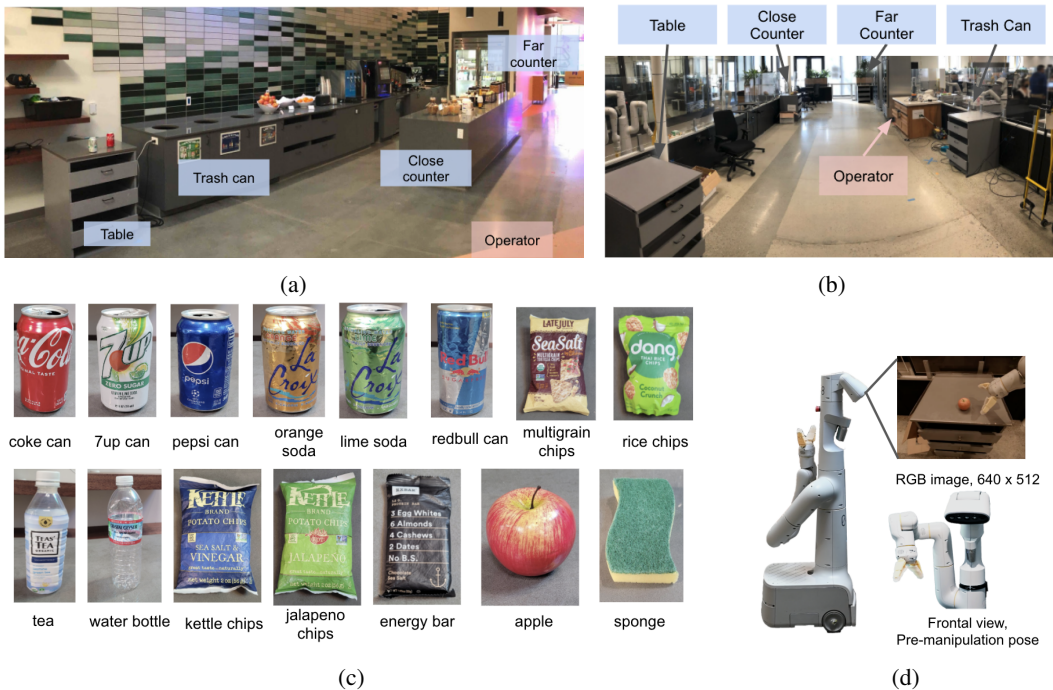


Figure 5: The experiments were performed in an office kitchen (a) and a mock office kitchen mirroring this setup (b), with 5 predicted locations and 15 objects (c). The robot is a mobile manipulator with a 7 degree-of-freedom arm and a gripper with policies trained from an RGB observation (d).

5.1 EXPERIMENTAL SETUP

Figure 5 shows the environment setup and the robot. We use 15 objects (Figure 5c) commonly found in an office kitchen and 5 known locations with semantic meaning (two counters, a table, a trash can, and the user location). We test our method in two environments: a real office kitchen (Figure 5a) and a mock environment mirroring the kitchen (Figure 5b), which is also the environment in which

the robot’s skills were trained. The robot used is a mobile manipulator from [Everyday Robots](https://everydayrobots.com/)² with a 7 degree-of-freedom arm and a two-fingered gripper, illustrated in Figure 5d. The LLM used is FLAN (Wei et al., 2021).

We also note a few practical considerations for setting up our affordance functions and policies. The flexibility of our approach allows us to mix and match policies and affordances from different methods. For the pick manipulation skills we use a single multi-task, language-conditioned policy, for the place manipulation skills we use a scripted policy with an affordance based on the gripper state, and for navigation policies we use a planning-based approach which is aware of the locations where specific objects can be found and a distance measure. In order to avoid a situation where a skill is chosen but has already been performed or will have no effect, we set a cap for the affordances indicating that the skill has been completed and the reward received. More details on parameters (such as the prompt) and policies can be found in Appendix C.

5.2 INSTRUCTIONS

To evaluate SayCan, we test across 101 instructions from 7 instruction families, summarized in Table 1 and fully enumerated in Appendix D.1. These were developed to test various aspects of SayCan’s performance and were inspired by crowd sourcing via Amazon Mechanical Turk and in-person kitchen users, as well as benchmarks such as ALFRED (Shridhar et al., 2020) and BEHAVIOR (Srivastava et al., 2022). Most instructions are designed to illicit a dialog language response (“How would you”), but we also test more natural imperative queries inspired by crowd-sourcing – we find no loss in performance. The evaluation instructions span multiple axes of variation: time-horizon (from single primitives to 10+ in a row), language complexity (from structured language to fully crowd-sourced natural requests), and embodiment (variations over the robot and environment state). Table 1 details examples for each family and Appendix D.1 contains all tasks, which fall into the following families:

- **Natural Language (NL) Single Primitive.** Given a natural language command corresponding to performing a single primitive, can SayCan recover that primitive skill and terminate?
- **NL Noun.** Given a natural language query that replaces a noun (typically an object or location) with a synonym, can SayCan execute the appropriate sequence?
- **NL Verbs.** Given a natural language query that replaces a verb (typically an action) with a synonym, can SayCan execute an appropriate sequence?
- **Structured Language.** Given a structure language query that mirrors the NL Verbs and spells out the sequence of commands, how well can SayCan plan compared to NL Verbs? This acts as an ablation to see the performance loss of understanding a natural language query over an explicit solution.
- **Embodiment.** Given a query with different environment and robot states, can SayCan still execute at a high rate? This tests the performance of SayCan’s affordance model and the LLM’s ability to reason within it.
- **Crowd-Sourced.** These queries were crowd sourced from Mechanical Turk by giving humans a description of what occurred (e.g., an apple was moved in front of you) and asking them what they would ask the robot to do. They were also crowd sourced by asking humans in a real office kitchen to command the robot to perform tasks (given knowledge of the robot’s abilities). This tests SayCan’s performance with natural requests.
- **Long-Horizon.** These challenging queries require SayCan to reason over temporally extended instructions to investigate how well it scales to such regimes.

5.3 METRICS

To understand the performance of the proposed method we measure two main metrics. The first is **plan success rate**, which measures whether the skills selected by the model are correct for the instruction, regardless of whether or not they actually successfully executed. We ask 3 human raters

²<https://everydayrobots.com/>

Instruction Family	Count	Explanation	Example Instruction
NL Single Primitive	15	Natural language queries for a single primitive	How would you let go of the coke can?
NL Nouns	15	Natural language queries focused on abstract nouns	How would you bring me a fruit?
NL Verbs	15	Natural language queries focused on abstract verbs	How would you restock the rice chips on the far counter?
Structured Language	15	Explicit, structured language queries that mirror NL Verbs	How would you move the rice chips to the far counter?
Embodiment	11	Queries to test SayCan’s understanding of the current state of the environment and robot	How would you put the coke can down on the far counter? (starting from different completion stages of the instruction)
Crowd-Sourced	15	Queries in unstructured formats	1. My favorite drink is redbull, bring one to me. 2. Please, move the pepsi to the close counter.
Long-Horizon	15	Long-horizon queries that require many steps of reasoning	I spilled my coke on the table, how would you throw it away and then bring me something to help clean?

Table 1: **List of instruction family definitions:** We evaluate the algorithm on 101 instructions. We group the instructions into different families, with each family focusing on testing one aspect of the proposed method.

to indicate whether the plan generated by the model can achieve the instruction, and if 2 out of 3 raters agree that the plan is valid, it is marked a success. Note that for many instructions there may be multiple valid solutions. For example if the instruction is to “bring a sponge and throw away the soda can”, the plan can choose to bring sponge first or throw away the soda can first.

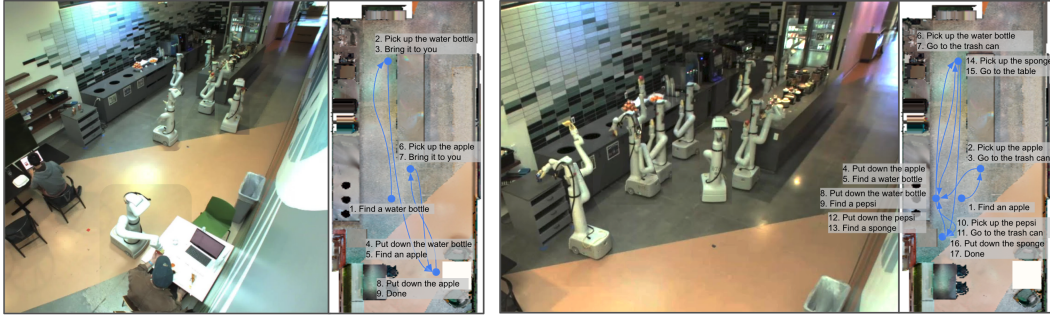
The second metric is **execution success rate**, which measures whether the full SayCan system actually performs the desired instruction successfully. We ask 3 human raters to watch the robot execution. The raters are asked to answer the question “whether the robot achieves the task specified by the task string?” We mark an execution successful if 2 out of 3 raters agree that the execution is successful.

5.4 RESULTS

Table 2 shows the performance of SayCan across 101 tasks. Across all instruction families in the mock kitchen, SayCan achieved a planning success rate of 70% and an execution rate of 61%. In Table 3 we further verify the performance of SayCan out of the lab setting and in the real kitchen on a subset of the instructions, particularly to verify the performance of the policies and value functions in this setting. We find no significant loss of performance between the two settings, indicating SayCan and the underlying policies generalize well to the full kitchen. The full task list and results can be found in the Appendix Table 5 and videos of experiment rollouts and the decision making process can be found on the project website: say-can.github.io.

Figure 6 shows two long-horizon queries and the resulting rollouts. These tasks require SayCan to plan many steps without error and for the robot to navigate and interact with a significant portion of the kitchen. Each query requires SayCan to understand context implicit within the instruction. In Figure 6a, the algorithm must understand the operator has asked for something to “recover from a workout”, i.e. something healthy, and thus it brings water and an apple rather than, e.g., a soda and chips. Furthermore, the algorithm must understand ordering and history, that it has already brought a drink and now must bring a snack before terminating. In Figure 6b, SayCan must track which objects are the “them” that need to be disposed of and where the sponge should be brought.

Figure 7 highlights SayCan’s decision making, along with another feature – its interpretability. The decision making process can be naturally understood as it solves high-level instructions by visualizing what the two sides of the algorithm output. This allows a user to understand what options



(a) “I just worked out, can you bring me a drink and a snack to recover?” (b) “I left out a coke, apple, and water, can you throw them away and then bring me a sponge to wipe the table?”

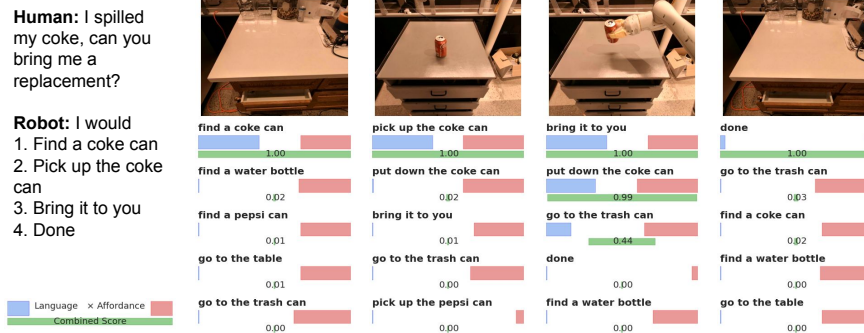
Figure 6: Timelapse of rollouts to two long-horizon queries. The robot interacts with a large portion of the kitchen environment and successfully performs sequences of manipulation and navigation skills.

SayCan is considering as language completions as well as what it believes is possible. Through this we find that sequence order is understood (approaching objects before picking them up and picking them up before bringing them) as well as context. For example, in Figure 7a, to bring a replacement drink the LLM considers multiple drink options, but understands the coke is the best replacement. Between Figures 7a and 7b, though the query is semantically similar (there has been a spill), SayCan understands the important object is the coke if the user wants a replacement and the sponge if the user wants to clean up. Figure 7c illustrates that SayCan understands generic object families and what belongs to them by choosing the apple as a fruit among the objects in the scene; this is similarly the case for “soda” or “bag of chips” which will apply to many objects in the scene. Figure 7d illustrates the ability of the affordances to *override* the language model; though the language model believes picking up the sponge is the right skill, the affordances are aware this isn’t possible and instead “find a sponge” is chosen. We believe such real-time and clear interpretability opens avenues to more interactive operation and safety.

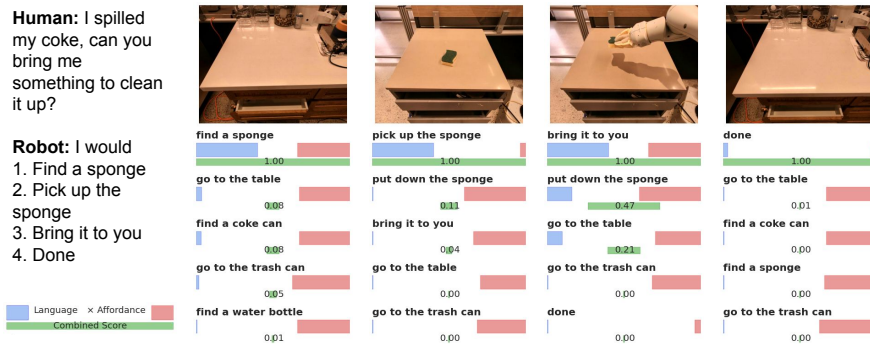
Family	Num	Plan	Execution
NL Single Primitive	15	67%	67%
NL Nouns	15	60%	53%
NL Verbs	15	80%	67%
Structured Language	15	100%	87%
Embodiment	11	64%	55%
Crowd Sourced	15	73%	67%
Long-Horizon	15	47%	33%
Total	101	70%	61%

Table 2: Plan and execution success rate of instructions by family. The proposed approach achieves an overall plan success rate of 70% and execution success rate of 61%.

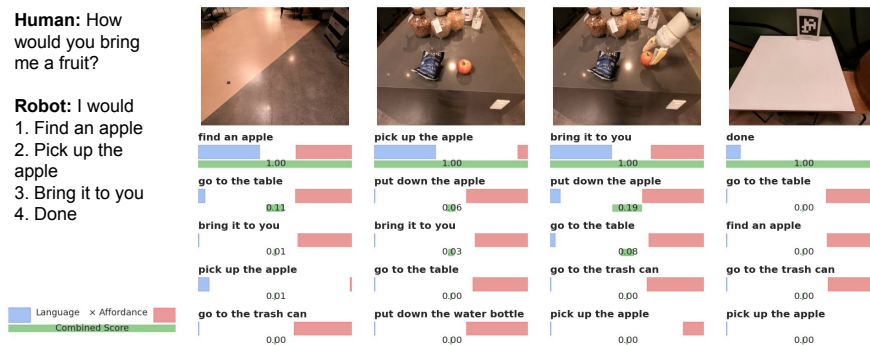
When comparing the performance of different instruction families in Table 2 and 3 (see Table 1 for an explanation), we see that the natural language nouns performed worse than natural language verbs, due to the number of nouns possible (15 objects and 5 locations) versus number of verbs (6). The structured language tasks (created to ablate the performance loss of spelling out the solution versus understanding the query) were planned correctly 100% of the time, while their natural language verb counterparts were planned correctly 80%. This indicates that the sequence is reasonable for the planner, but that the language was challenging to parse. The embodiment tasks were planned correctly 64% of the time, generally with failures as a result of affordance function misclassification. SayCan planned and executed crowd-sourced natural queries with performance on par with other instruction families. SayCan performed worst on the most challenging long-horizon tasks, where



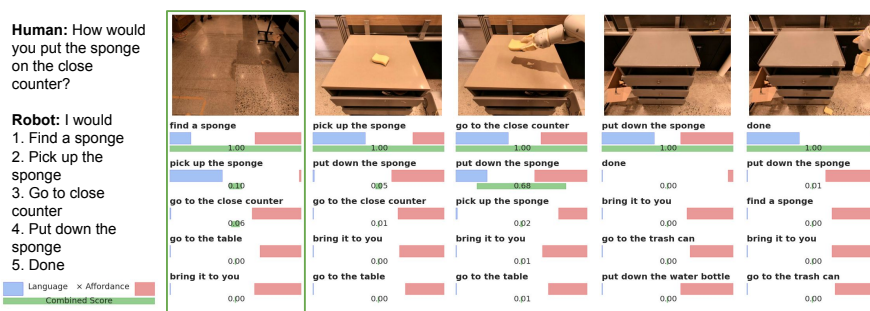
(a)



(b)



(c)



(d)

Figure 7: Visualization of the decision making process of SayCan shows its interpretability and successful temporally extended execution, where the top combined score chooses the correct skill.

most failures were a result of early termination by the LLM (e.g., bringing one object but not the second). We also find that SayCan struggles with negation (e.g., “bring me something that isn’t an apple from the table”) and ambiguous references (e.g. asking for drinks with caffeine, sugary drinks), which is a known issue inherited from underlying language models (Hosseini et al., 2021). Overall, 65% of the errors were a result of LLM failures and 35% were affordance errors.

Returning to our initial example, “I spilled something, can you help?”, an ungrounded language model would respond with statements like “Sure I can help”, “I can call you a cleaner”, and “I can vacuum that up for you”, which given our robot embodiment are unreasonable and impossible. We have shown that SayCan on the other hand responds “I would: 1. find a sponge, 2. pick up the sponge, 3. bring it to you, 4. done” and is able execute this sequence on the robot in a real kitchen. This requires long-horizon reasoning over a required order, an abstract understanding of the instruction, and knowledge of both the environment and robot’s capabilities.

Family	Num	Kitchen		Mock Kitchen	
		Plan	Execution	Plan	Execution
NL Single Primitive	15	87%	87%	67%	67%
NL Nouns	8	63%	50%	63%	50%
NL Verbs	6	83%	67%	83%	83%
Structured Language	6	83%	50%	100%	83%
Crowd Sourced	5	80%	80%	80%	80%
Long-Horizon	6	33%	33%	33%	33%
Total	46	74%	65%	70%	65%

Table 3: Evaluations of SayCan on a subset of tasks in the Kitchen and the Mock Kitchen. The performance is comparable in the two scenes.

5.5 ABLATIONS: REMOVING LANGUAGE AND AFFORDANCE GROUNDINGS

Ablating Language. To study the importance of the language model for interpreting complex instructions, we conduct an ablation experiment where we directly use the language-conditioned primitives (see Sections 4.1-4.2) and feed the full instruction i directly into the text-conditioned primitive policy as though it were a skill label. Note that this approach is representative of standard RL or BC-based instruction following methods (Jang et al., 2021; Stepputtis et al., 2020; Nair et al., 2021; Lynch & Sermanet, 2020), which train a policy conditioned on a natural language command. The results in Table 4 illustrate the necessity of the language grounding, in which we look at our learned BC-based policies alone. These policies were trained with a fixed set of language commands, and indeed we find anything outside of these commands fails completely, resulting in a 0% success rate even in the NL Single Primitive family. To alleviate this issue, we also try projecting the high-level instruction into the set of known language commands via the Universal Sentence Encoder (USE) embeddings (Cer et al., 2018). To do this, we embed the instruction and all the tasks, along with the combinatorial set of sequences tasks (i.e., we consider “pick coke can” as well as “1. find coke can, 2. pick coke can” and so on), and select the highest cosine similarity trained instruction. This achieves 60% on the NL Single Primitives family, but 0% across all the other families.

Ablating Value Functions. The results in Table 4 illustrate the necessity of the affordance grounding through value functions. We compare SayCan to (1) *No VF*, which removes the value function grounding (i.e., choosing the maximum language score skill) and to (2) *Generative*, which uses the generative output of the LLM and then projects each planned skill to its maximal cosine similarity trained skill via pretrained Universal Sentence Encoder embeddings (Cer et al., 2018). The latter in effect compares to the algorithm proposed in Huang et al. (2022), which we note loses the explicit probabilities across options, and thus is less interpretable and cannot be combined with affordance probabilities. For *Generative* we also tried BERT embeddings (Devlin et al., 2018), but found the performance of USE embeddings to be superior. The *No VF* and *Generative* approaches performed similarly in aggregate, achieving 38% and 43% planning success rate respectively, and worse than SayCan’s 70%. Both baselines performed similarly for the Long-Horizon, Embodiment, Crowd

Sourced, and NL Single Primitive instruction families. These methods perform particularly poorly in the Embodiment experiments and in the Long-Horizon tasks, where the affordances are important in eliminating impossible skills that will lead to failure if triggered. Interestingly, the two ablations were not uniformly worse across all instruction families. In the instruction families NL Nouns, NL Verbs, and Structured Language, we find that the sequencing is more consistent with *Generative*, and that the limited number of verbs compared to objects means the loss in performance due to the projection is smaller. The nouns however are worse with the *Generative* approach, because the projection loses information about the overall query. For example, in the query “I just worked out, can you bring me a drink and a snack to recover?”, the generative model produces the sequence {find a drink, pick up the drink, go to the table, put down the drink, find a snack, pick up the snack, go to the table, put down the snack, done}. This sequence is reasonable, but when projecting “find a drink” it loses the earlier context that this should be a healthy drink and instead chooses “find a coke can”. This loss of context is not an issue for the scoring model however, which evaluates the probability of the skill given the full sequence so far, including the full original instruction.

Family	Num	SayCan	No VF	Generative	BC (NL)	BC (USE)
NL Single Primitive	15	67%	33%	40%	0%	60%
NL Nouns	15	60%	53%	40%	0%	0%
NL Verbs	15	80%	53%	87%	0%	0%
Structured Language	15	100%	20%	73%	0%	0%
Embodiment	11	64%	18%	0%	0%	0%
Crowd Sourced	15	73%	60%	47%	0%	0%
Long-Horizon	15	47%	20%	0%	0%	0%
Total	101	70%	38%	43%	0%	9%

Table 4: Planning rate across SayCan and various ablations. *No VF* uses the maximum language score skill from the LLM, *Generative* uses a generative LLM and then projects to the nearest skill via USE embeddings, *BC (NL)* uses the low-level policy with the natural language instruction, and *BC (USE)* uses the low-level policy with the natural language instruction projected to the nearest skill via USE embeddings.

6 RELATED WORK

Grounding Language Models. A significant body of work has focused on grounding language models to apply them to novel settings (Siskind, 1994; Winograd, 1972). A common approach is to train them to accept additional environment inputs (Sun et al., 2019; Li et al., 2019; Lu et al., 2019; Zellers et al., 2021b; Radford et al., 2021) or to directly output actions (Suglia et al., 2021; Pashevich et al., 2021; Sharma et al., 2021). Others are grounded in an environment through prompt engineering (Wei et al., 2022). Concurrently with SayCan, Huang et al. (2022) uses such a prompt engineering approach to extract temporally extended plans, but without any additional mechanism to ensure grounding, roughly corresponding to the “Generative” baseline in our experiments. The above methods are all trained without interaction with a physical environment, thus limiting their ability to reason over embodied interactions. One approach to grounding language models in interaction is through learning downstream networks with pre-trained LLM representations (Humphreys et al., 2022; Hill et al., 2020; Blukis et al., 2022; Nair et al., 2022; Akakzia et al., 2020; Zellers et al., 2021a). Another approach finetunes language models with interactive data, such as rewards or ranking feedback of the interaction (Ouyang et al., 2022; Reid et al., 2022; Li et al., 2022). In our work, SayCan is able to ground language models in the given environment through previously-trained value functions, enabling general, long-horizon behaviors in a zero-shot manner, i.e., without additional training.

Learning Language-Conditioned Behavior. There is a long history of research studying how to connect language and behavior (MacMahon et al., 2006; Kollar et al., 2010; Tellex et al., 2011; Luketina et al., 2019; Tellex et al., 2020). A large number of prior works have learned language-conditioned behavior via imitation learning (Mei et al., 2016; Lynch & Sermanet, 2020; Stepputtis

et al., 2020; Jang et al., 2021; Shridhar et al., 2022; Sharma et al., 2021) or reinforcement learning (Misra et al., 2017; Hermann et al., 2017; Luketina et al., 2019; Jiang et al., 2019; Cideron et al., 2019; Goyal et al., 2020; Nair et al., 2021; Akakzia et al., 2020). Most of these prior works focus on following low-level instructions, such as for pick-and-place tasks and other robotic manipulation primitives (Stepputtis et al., 2020; Lynch & Sermanet, 2020; Goyal et al., 2020; Nair et al., 2021; Jang et al., 2021; Shridhar et al., 2022), though some methods address long-horizon, compound tasks in simulated domains (Oh et al., 2017; Andreas et al., 2017; Jiang et al., 2019). Like these latter works, we focus on completing temporally extended tasks. However, a central aspect of our work is to solve such tasks by extracting and leveraging the knowledge in large language models. While prior works have studied how pre-trained language embeddings can improve generalization to new instructions (Hill et al., 2020; Lynch & Sermanet, 2020; Nair et al., 2021) and to new low-level tasks (Jang et al., 2021), we extract much more substantial knowledge from LLMs by grounding them within the robot’s affordances. This allows robots to use language models for planning.

Task Planning and Motion Planning. Task and Motion Planning (Kaelbling & Lozano-Pérez, 2010; Srivastava et al., 2014) is a fundamental problem in which one must sequence tasks to solve a high-level problem, while ensuring the feasibility given an embodiment (Task Planning (Fikes & Nilsson, 1971; Sacerdoti, 1975; Nau et al., 1999); Motion Planning (LaValle, 2006)). Classically, this problem has been solved through symbolic planning approaches (Fikes & Nilsson, 1971; Nau et al., 1999) or optimization (Toussaint, 2015; Toussaint et al., 2018), but these typically require explicit primitives and constraints. Machine learning has recently been applied to enable abstract task specification, allow general primitives, or relax explicit constraints (Xu et al., 2018; 2019; Huang et al., 2019; Eysenbach et al., 2019; Savinov et al., 2018; Ichter et al., 2021; Matuszek et al., 2012; Silver et al., 2022; Garrett et al., 2020; Zhu et al., 2017; Misra et al., 2016; Wu et al., 2021). Others learn to hierarchically solve such long-horizon problems (Nair & Finn, 2020; Xia et al., 2021; Shah et al., 2022; Li et al., 2020; Jiang et al., 2019). Going beyond prior methods, SayCan leverages the large capacity of LLMs, where their semantic knowledge about the world can be useful both for interpreting instructions *and* understanding how to execute them (e.g. understanding of sequences and order). The combination of LLMs and generality of learned low-level policies enables long-horizon, abstract tasks that scale effectively to the real world, as demonstrated in our extensive on robot experimental evaluation.

7 CONCLUSIONS, LIMITATIONS AND FUTURE WORK

We presented SayCan, a method that enables leveraging and grounding the rich knowledge in large language models to complete embodied tasks. For real-world grounding, we leverage pre-trained skills, which are then used to condition the model to choose natural language actions that are both feasible and contextually appropriate. More specifically, we use reinforcement learning as a way to learn value functions for the individual skills that provide affordances of what is possible in the world, and then use textual labels for these skills as potential responses that are scored by a language model. This combination results in a symbiotic relationship where the skills and their value functions can act as the language model’s “hands and eyes,” while the language model supplies high-level semantic knowledge about how to complete a task. We evaluated the proposed approach on a number of real-world robotic tasks that involve a mobile manipulator robot accomplishing a large set of long-horizon natural language instructions in a real kitchen.

While SayCan presents a viable way to ground language models in agents’ affordances, it has a number of limitations. First, we expect this method to inherit the limitations and biases of LLMs (Bender et al., 2021; Bommasani et al., 2021), including the dependence on the training data. Secondly, we observe that even though SayCan allows the users to interact with the agents using natural language commands, the primary bottleneck of the system is in the range and capabilities of the underlying skills. To illustrate this, we present representative failure cases in Appendix D. Future work that extends the repertoire of skills and improves their robustness would mitigate this limitation. In addition, at the current stage, the system is not easily able to react to situations where individual skills fail despite reporting a high value, though this could potentially be addressed by appropriate prompting of the language model for a correction.

There are many other potential avenues for future work. A natural question that this work raises is how the information gained through grounding the LLM via real-world robotic experience can

be leveraged to improve the language model itself, both in terms of its factuality and its ability to perform common-sense reasoning about real-world environments and physics. Furthermore, since our method uses generic value functions to score affordances, it is intriguing to consider what other sources of grounding could be incorporated in the same manner, such as non-robotic contexts.

In the future, it is also interesting to examine whether *natural language is the right ontology to use to program robots*: natural language naturally incorporates contextual and semantic cues from the environment, and provides a level of abstraction which enables robots to decide on how to execute a strategy based on its own perception and affordances. At the same time, as opposed to, e.g., hindsight goal images (Chebotar et al., 2021), it requires supervision and it might not be the most descriptive medium for certain tasks.

Lastly, SayCan presents a particular way of connecting and factorizing the challenges of language understanding and robotics, and many further extensions can be proposed. Ideas such as combining robot planning and language (Tellex et al., 2014), using language models as a pre-training mechanism for policies (Reid et al., 2022) and many other ways of combining language and interaction (Nair et al., 2021; Lynch & Sermanet, 2020; Hill et al., 2020; Wang et al., 2016) are exciting avenues for future research.

ACKNOWLEDGMENTS

The authors would like to thank Yunfei Bai, Matt Bennice, Maarten Bosma, Bill Byrne, Kendra Byrne, Noah Constant, Pete Florence, Laura Graesser, Rico Jonschkowski, Daniel Kappler, Hugo Larochelle, Benjamin Lee, Adrian Li, Suraj Nair, Krista Reymann, Jeff Seto, Dhruv Shah, Ian Storz, Razvan Surdulescu, Andy Zeng, and Vincent Zhao for their help and support in various aspects of the project.

REFERENCES

- Ahmed Akakzia, Cédric Colas, Pierre-Yves Oudeyer, Mohamed Chetouani, and Olivier Sigaud. Grounding language to autonomously-acquired skills via goal generation. *arXiv preprint arXiv:2006.07185*, 2020.
- Jacob Andreas, D. Klein, and Sergey Levine. Modular multitask reinforcement learning with policy sketches. *ArXiv*, abs/1611.01796, 2017.
- Emily M Bender and Alexander Koller. Climbing towards nlu: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 610–623, 2021.
- Valts Blukis, Chris Paxton, Dieter Fox, Animesh Garg, and Yoav Artzi. A persistent spatial semantic representation for high-level natural language instruction execution. In *Conference on Robot Learning*, 2022.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*, 2018.
- Yevgen Chebotar, Karol Hausman, Yao Lu, Ted Xiao, Dmitry Kalashnikov, Jacob Varley, A. Irpan, Benjamin Eysenbach, Ryan C. Julian, Chelsea Finn, and Sergey Levine. Actionable models: Unsupervised offline reinforcement learning of robotic skills. *ArXiv*, abs/2104.07749, 2021.

- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, et al. Palm: Scaling language modeling with pathways. 2022. URL <https://storage.googleapis.com/pathways-language-model/PaLM-paper.pdf>.
- Geoffrey Cideron, Mathieu Seurin, Florian Strub, and O. Pietquin. Self-educated language agent with hindsight experience replay for instruction following. *ArXiv*, abs/1910.09451, 2019.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Ben Eysenbach, Russ R Salakhutdinov, and Sergey Levine. Search on the replay buffer: Bridging planning and reinforcement learning. *Advances in Neural Information Processing Systems*, 2019.
- Richard E Fikes and Nils J Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 1971.
- Caelan Reed Garrett, Chris Paxton, Tomás Lozano-Pérez, Leslie Pack Kaelbling, and Dieter Fox. Online replanning in belief space for partially observable task and motion problems. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- James J Gibson. The theory of affordances. *The Ecological Approach to Visual Perception*, 1977.
- Prasoon Goyal, S. Niekum, and R. Mooney. Pixl2r: Guiding reinforcement learning using natural language by mapping pixels to rewards. *ArXiv*, abs/2007.15543, 2020.
- K. Hermann, Felix Hill, Simon Green, F. Wang, R. Faulkner, Hubert Soyer, David Szepesvari, Wojciech Czarnecki, Max Jaderberg, Denis Teplyashin, M. Wainwright, C. Apps, D. Hassabis, and P. Blunsom. Grounded language learning in a simulated 3d world. *ArXiv*, abs/1706.06551, 2017.
- Felix Hill, Sona Mokra, Nathaniel Wong, and Tim Harley. Human instruction-following with deep reinforcement learning via transfer-learning from text. *arXiv preprint arXiv:2005.09382*, 2020.
- Daniel Ho, Kanishka Rao, Zhuo Xu, Eric Jang, Mohi Khansari, and Yunfei Bai. Retinagan: An object-aware approach to sim-to-real transfer. *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 10920–10926, 2021.
- Arian Hosseini, Siva Reddy, Dzmitry Bahdanau, R Devon Hjelm, Alessandro Sordani, and Aaron Courville. Understanding by understanding not: Modeling negation in language models. *arXiv preprint arXiv:2105.03519*, 2021.
- De-An Huang, Suraj Nair, Danfei Xu, Yuke Zhu, Animesh Garg, Li Fei-Fei, Silvio Savarese, and Juan Carlos Niebles. Neural task graphs: Generalizing to unseen tasks from a single video demonstration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *arXiv preprint arXiv:2201.07207*, 2022.
- Peter C Humphreys, David Raposo, Toby Pohlen, Gregory Thornton, Rachita Chhaparia, Alistair Muldal, Josh Abramson, Petko Georgiev, Alex Goldin, Adam Santoro, et al. A data-driven approach for learning to control computers. *arXiv preprint arXiv:2202.08137*, 2022.
- Brian Ichter, Pierre Sermanet, and Corey Lynch. Broadly-exploring, local-policy trees for long-horizon task planning. *Conference on Robot Learning (CoRL)*, 2021.
- Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, pp. 991–1002. PMLR, 2021.
- Yiding Jiang, Shixiang Gu, K. Murphy, and Chelsea Finn. Language as an abstraction for hierarchical deep reinforcement learning. In *NeurIPS*, 2019.

- Leslie Pack Kaelbling and Tomás Lozano-Pérez. Hierarchical planning in the now. In *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- Dmitry Kalashnikov, Jake Varley, Yevgen Chebotar, Ben Swanson, Rico Jonschkowski, Chelsea Finn, Sergey Levine, and Karol Hausman. Mt-opt: Continuous multi-task robotic reinforcement learning at scale. *arXiv*, 2021.
- T. Kollar, Stefanie Tellex, D. Roy, and N. Roy. Toward understanding natural language directions. In *HRI 2010*, 2010.
- Steven M LaValle. *Planning algorithms*. 2006.
- Chengshu Li, Fei Xia, Roberto Martin-Martin, and Silvio Savarese. Hrl4in: Hierarchical reinforcement learning for interactive navigation with mobile manipulators. In *Conference on Robot Learning*, 2020.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- Shuang Li, Xavier Puig, Yilun Du, Clinton Wang, Ekin Akyurek, Antonio Torralba, Jacob Andreas, and Igor Mordatch. Pre-trained language models for interactive decision-making. *arXiv preprint arXiv:2202.01771*, 2022.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 2019.
- Jelena Luketina, Nantas Nardelli, Gregory Farquhar, Jakob N. Foerster, Jacob Andreas, Edward Grefenstette, S. Whiteson, and Tim Rocktäschel. A survey of reinforcement learning informed by natural language. In *IJCAI*, 2019.
- Corey Lynch and Pierre Sermanet. Grounding language in play. 2020.
- Matt MacMahon, Brian Stankiewicz, and Benjamin Kuipers. Walk the talk: Connecting language, knowledge, and action in route instructions. 01 2006.
- Cynthia Matuszek, Nicholas FitzGerald, Luke Zettlemoyer, Liefeng Bo, and Dieter Fox. A joint model of language and perception for grounded attribute learning. *arXiv preprint arXiv:1206.6423*, 2012.
- Hongyuan Mei, Mohit Bansal, and Matthew R. Walter. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *AAAI*, 2016.
- Dipendra K Misra, Jaeyong Sung, Kevin Lee, and Ashutosh Saxena. Tell me dave: Context-sensitive grounding of natural language to manipulation instructions. *The International Journal of Robotics Research*, 2016.
- Dipendra Kumar Misra, J. Langford, and Yoav Artzi. Mapping instructions and visual observations to actions with reinforcement learning. In *EMNLP*, 2017.
- Suraj Nair and Chelsea Finn. Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. *ArXiv*, abs/1909.05829, 2020.
- Suraj Nair, Eric Mitchell, Kevin Chen, Brian Ichter, Silvio Savarese, and Chelsea Finn. Learning language-conditioned robot behavior from offline data and crowd-sourced annotation. In *Conference on Robot Learning*, pp. 1303–1315. PMLR, 2021.
- Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022.
- Dana Nau, Yue Cao, Amnon Lotem, and Hector Munoz-Avila. Shop: Simple hierarchical ordered planner. 1999.
- Junhyuk Oh, Satinder Singh, Honglak Lee, and P. Kohli. Zero-shot task generalization with multi-task deep reinforcement learning. *ArXiv*, abs/1706.05064, 2017.

- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Preprint*, 2022.
- Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic transformer for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, pp. 1–67, 2019.
- Machel Reid, Yutaro Yamada, and Shixiang Shane Gu. Can wikipedia help offline reinforcement learning. *arXiv preprint arXiv:2201.12122*, 2022.
- Earl D Sacerdoti. A structure for plans and behavior. Technical report, SRI International, Menlo Park California Artificial Intelligence Center, 1975.
- Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*, 2018.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *ICLR*, 2016.
- Dhruv Shah, Peng Xu, Yao Lu, Ted Xiao, Alexander Toshev, Sergey Levine, and Brian Ichter. Value function spaces: Skill-centric state abstractions for long-horizon reasoning. *ICLR*, 2022.
- Pratyusha Sharma, Antonio Torralba, and Jacob Andreas. Skill induction and planning with latent language. *arXiv preprint arXiv:2110.01517*, 2021.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10740–10749, 2020.
- Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Cliport: What and where pathways for robotic manipulation. In *Conference on Robot Learning*, 2022.
- Tom Silver, Rohan Chitnis, Nishanth Kumar, Willie McClinton, Tomas Lozano-Perez, Leslie Pack Kaelbling, and Joshua Tenenbaum. Inventing relational state and action abstractions for effective and efficient bilevel planning. *arXiv preprint arXiv:2203.09634*, 2022.
- Jeffrey Mark Siskind. Grounding language in perception. *Artificial Intelligence Review*, 1994.
- Sanjana Srivastava, Chengshu Li, Michael Lingelbach, Roberto Martín-Martín, Fei Xia, Kent Elliott Vainio, Zheng Lian, Cem Gokmen, Shyamal Buch, Karen Liu, et al. Behavior: Benchmark for everyday household activities in virtual, interactive, and ecological environments. In *Conference on Robot Learning*, pp. 477–490. PMLR, 2022.
- Siddharth Srivastava, Eugene Fang, Lorenzo Riano, Rohan Chitnis, Stuart Russell, and Pieter Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *2014 IEEE international conference on robotics and automation (ICRA)*, 2014.
- Simon Stepputtis, J. Campbell, Mariano Phielipp, Stefan Lee, Chitta Baral, and H. B. Amor. Language-conditioned imitation learning for robot manipulation tasks. *ArXiv*, abs/2010.12083, 2020.

- Alessandro Suglia, Qiaozhi Gao, Jesse Thomason, Govind Thattai, and Gaurav Sukhatme. Embodied bert: A transformer model for embodied, language-guided visual task completion. *arXiv preprint arXiv:2108.04927*, 2021.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew Walter, Ashis Banerjee, Seth Teller, and Nicholas Roy. Understanding natural language commands for robotic navigation and mobile manipulation. volume 2, 01 2011.
- Stefanie Tellex, Ross Knepper, Adrian Li, Daniela Rus, and Nicholas Roy. Asking for help using inverse semantics. 2014.
- Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. Robots that use language. *Annual Review of Control, Robotics, and Autonomous Systems*, 2020.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.
- Marc Toussaint. Logic-geometric programming: An optimization-based approach to combined task and motion planning. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- Marc A Toussaint, Kelsey Rebecca Allen, Kevin A Smith, and Joshua B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. 2018.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Sida I Wang, Percy Liang, and Christopher D Manning. Learning language games through interaction. *arXiv preprint arXiv:1606.02447*, 2016.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.
- Terry Winograd. Understanding natural language. *Cognitive psychology*, 1972.
- Bohan Wu, Suraj Nair, Li Fei-Fei, and Chelsea Finn. Example-driven model-based reinforcement learning for solving long-horizon visuomotor tasks. In *5th Annual Conference on Robot Learning*, 2021.
- Fei Xia, Chengshu Li, Roberto Martín-Martín, Or Litany, Alexander Toshev, and Silvio Savarese. Relmogen: Integrating motion generation in reinforcement learning for mobile manipulation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- Ted Xiao, Eric Jang, Dmitry Kalashnikov, Sergey Levine, Julian Ibarz, Karol Hausman, and Alexander Herzog. Thinking while moving: Deep reinforcement learning with concurrent control. In *International Conference on Learning Representations*, 2019.
- Danfei Xu, Suraj Nair, Yuke Zhu, Julian Gao, Animesh Garg, Li Fei-Fei, and Silvio Savarese. Neural task programming: Learning to generalize across hierarchical tasks. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- Danfei Xu, Roberto Martín-Martín, De-An Huang, Yuke Zhu, Silvio Savarese, and Li F Fei-Fei. Regression planning networks. *Advances in Neural Information Processing Systems*, 32, 2019.

Rowan Zellers, Ari Holtzman, Matthew Peters, Roozbeh Mottaghi, Aniruddha Kembhavi, Ali Farhadi, and Yejin Choi. Piglet: Language grounding through neuro-symbolic interaction in a 3d world. *arXiv preprint arXiv:2106.00188*, 2021a.

Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. Merlot: Multimodal neural script knowledge models. *Advances in Neural Information Processing Systems*, 2021b.

Yuke Zhu, Daniel Gordon, Eric Kolve, Dieter Fox, Li Fei-Fei, Abhinav Gupta, Roozbeh Mottaghi, and Ali Farhadi. Visual semantic planning using deep successor representations. In *Proceedings of the IEEE international conference on computer vision*, 2017.

A CONTRIBUTIONS

A.1 BY TYPE

- **Designed and built distributed robot learning infrastructure:** Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Byron David, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Alex Irpan, Eric Jang, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Yao Lu, Peter Pastor, Kanishka Rao, Nicolas Sievers, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, and Mengyuan Yan.
- **Designed, implemented, or trained the underlying manipulation policies:** Yevgen Chebotar, Keerthana Gopalakrishnan, Karol Hausman, Julian Ibarz, Alex Irpan, Eric Jang, Nikhil Joshi, Ryan Julian, Kuang-Huei Lee, Yao Lu, Kanishka Rao, and Ted Xiao.
- **Designed or implemented the data generation and curation or collected data:** Noah Brown, Omar Cortes, Jasmine Hsu, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Linda Luu, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Clayton Tan, and Sichun Xu.
- **Designed or implemented SayCan:** Karol Hausman, Brian Ichter, Sergey Levine, Alexander Toshev, and Fei Xia.
- **Managed or advised on the project:** Chelsea Finn, Karol Hausman, Eric Jang, Sally Jesmonth, Sergey Levine, Yao Lu, Carolina Parada, Kanishka Rao, Alexander Toshev, and Vincent Vanhoucke.
- **Ran evaluations or experiments:** Noah Brown, Omar Cortes, Brian Ichter, Rosario Jauregui Ruano, Kyle Jeffrey, Linda Luu, Jornell Quiambao, Jarek Rettinghouse, Diego Reyes, Clayton Tan, and Fei Xia.
- **Scaled simulation infrastructure:** Nikhil J Joshi, Yao Lu, Kanishka Rao, and Ted Xiao.
- **Wrote the paper:** Chelsea Finn, Karol Hausman, Brian Ichter, Alex Irpan, Sergey Levine, Fei Xia, and Ted Xiao.

A.2 BY PERSON

Michael Ahn developed the deployment system that enabled the ability to scale up data collection on real robots.

Anthony Brohan implemented the logging system for the project and designed and implemented the data labeling pipelines.

Noah Brown led and coordinated the real-robot operations including data collection with teleoperators, evaluations and the real-world setup.

Yevgen Chebotar designed and implemented multiple offline RL methods allowing the manipulation policies to process data coming from different sources.

Omar Cortes collected data on the robots and ran and supervised real-world evaluations.

Byron David developed simulation assets and performed system identification.

Chelsea Finn advised on the project, helped set the research direction and wrote parts of the paper.

Keerthana Gopalakrishnan provided multiple infrastructure contributions that allowed for scalable learning of manipulation policies.

Karol Hausman co-lead the project as well as developed SayCan, helped set the research direction, trained the underlying manipulation policies, and wrote the paper.

Alex Herzog developed the teleoperation tools and implemented multiple infrastructure tools that allowed for continuous robot operation.

Daniel Ho helped develop sim-to-real pipelines for manipulation policies.

Jasmine Hsu provided logging and monitoring infrastructure tools as well as data labeling pipelines.

Julian Ibarz provided multiple contributions that enabled scaling learning algorithms for manipulation policies, and helped set the research direction.

Brian Ichter initiated and led the SayCan algorithm, combined the manipulation and navigation skills, ran experiments for the paper, and wrote the paper.

Alex Irpan set up and led the autonomous data collection effort as well as verified the data collected by the robots, and wrote parts of the paper.

Eric Jang helped set the research and team direction, managed the data for learning, developed the behavioral cloning manipulation policies, and wrote parts of the paper.

Rosario Jauregui Ruano collected data on the robots and ran and supervised real-world evaluations.

Kyle Jeffrey collected data on the robots and ran and supervised real-world evaluations.

Sally Jesmonth was the program manager for the project.

Nikhil J Joshi developed a number of simulation and infrastructure tools that allowed to scale up simulation training.

Ryan Julian developed multi-modal network architectures and trained manipulation policies.

Dmitry Kalashnikov contributed infrastructure pieces that enabled training from logged data.

Yuheng Kuang implemented the logging system for the project and designed and implemented the data labeling pipelines

Kuang-Huei Lee made improvements to training algorithms for manipulation policies.

Sergey Levine advised on the project, helped set the research direction, developed SayCan, and wrote parts of the paper.

Yao Lu led and designed the robot learning infrastructure for the project providing most of the tools and improving manipulation policies.

Linda Luu ran multiple evaluations, collected data and helped establish real-robot operations.

Carolina Parada advised on the project, managed the team, helped write the paper, and helped set the research direction.

Peter Pastor provided infrastructure tools that allowed for continuous robot operations.

Jornell Quiambao collected data on the robots and ran and supervised real-world evaluations.

Kanishka Rao co-lead the project, managed the team, helped set the research direction and contributed to training manipulation policies.

Jarek Rettinghouse collected data on the robots and ran and supervised real-world evaluations.

Diego Reyes collected data on the robots and ran and supervised real-world evaluations.

Pierre Sermanet set up the crowd compute rating pipeline.

Nicolas Sievers provided simulation assets and environments used for simulation training.

Clayton Tan collected data on the robots and ran and supervised real-world evaluations and helped establish real-robot operations.

Alexander Toshev advised on the project, developed SayCan, helped write the paper, and helped set research direction.

Vincent Vanhoucke advised on the project, managed the team, and helped write the paper.

Fei Xia developed, implemented, and led on-robot SayCan, ran the experiments for the paper, created the demos, and wrote the paper.

Ted Xiao led the scaling of manipulation skills, designed and developed learning from simulation for manipulation skills, and developed multi-modal network architectures.

Peng Xu was the engineering lead for integrating manipulation and navigation and developed the underlying infrastructure for SayCan.

Sichun Xu developed remote teleoperation tools that allowed scaling up data collection in simulation.

Mengyuan Yan implemented infrastructure and learning tools that allowed for learning manipulation policies from different data sources.

A.3 CORRESPONDING EMAILS:

{*ichter, xiafei, karolhausman*}@google.com

B RL AND BC POLICIES

B.1 RL AND BC POLICY ARCHITECTURE

The RL models use an architecture similar to MT-Opt (Kalashnikov et al., 2021), with slight changes to support natural language inputs (see Fig. 8 for the network diagram). The camera image is first processed by 7 convolutional layers. The language instruction is embedded by the LLM, then concatenated with the robot action and non-image parts of the state, such as the gripper height. To support asynchronous control, inference occurs while the robot is still moving from the previous action. The model is given how much of the previous action is left to execute (Xiao et al., 2019). The conditioning input goes through FC layers, then tiled spatially and added to the conv. volume, before going through 11 more convolutional layers. The output is gated through a sigmoid, so the Q-value is always in $[0, 1]$.

The BC models use an architecture similar to BC-Z (Jang et al., 2021) (see Fig. 9 for the network diagram). The language instruction is embedded by a universal sentence encoder (Cer et al., 2018), then used to FiLM condition a Resnet-18 based architecture. Unlike the RL model, we do not provide the previous action or gripper height, since this was not necessary to learn the policy. Multiple FC layers are applied to the final visual features, to output each action component (arm position, arm orientation, gripper, and the termination action).

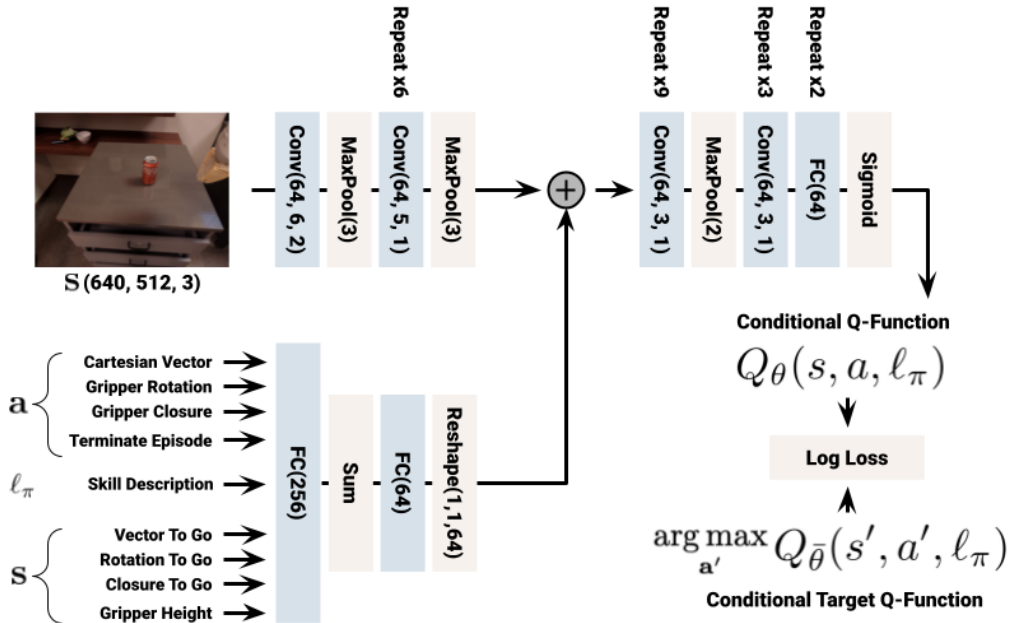


Figure 8: Network architecture in RL policy

B.2 RL AND BC POLICY TRAINING

Both RL and BC models were trained using 16 TPUv3 chips. RL models were trained for about 100 hours, and BC models were trained for about 27 hours. Both methods applied standard image augmentations (random brightness and contrast) as well as random cropping. The 640×512 input image was padded by 100 pixels left-right and 40 pixels top-down, then cropped back down to a 640×512 image, so as to allow for random spatial shifts without limiting the field of view.

The RL model is trained via MT-Opt (Kalashnikov et al., 2021), using a pool of 3000 CPU workers to collect episodes in the Everyday Robots simulator. Another 3000 CPU workers are used to compute target Q-values. Computing target Q-values outside the TPU allows the TPU to be used solely for computing gradient updates. Episode rewards are sparse and always 0 or 1, so the Q-function is updated using a log loss. Models were trained using prioritized experience replay (Schaul et al.,

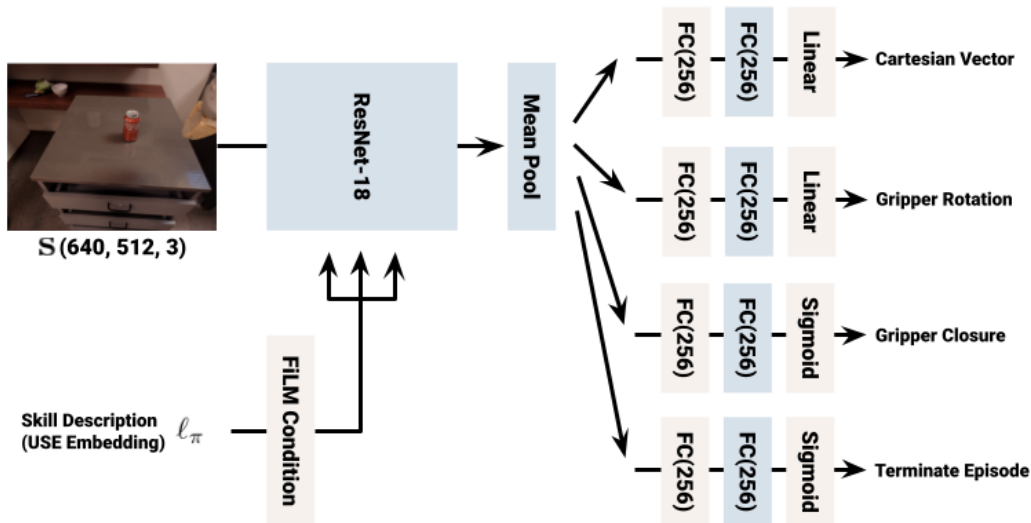


Figure 9: Network architecture in BC policy

2016), where episode priority was tuned to encourage replay buffer training data for each skill to be close to 50% success. Episodes were sampled proportionally to their priority, defined as $1 + 10 \cdot |p - 0.5|$, where p is the average success rate of episodes in the replay buffer.

The BC model is trained with an MSE loss for the continuous action components, and a cross-entropy loss for the discrete action components. Each action component was weighted evenly. For faster iteration speeds with negligible training performance reduction, image inputs were down sampled to half-size (256 x 320 images). Note that the affordance value functions were trained with full-size images, since half-size images did not work as well when learning $Q(s, a, \ell_\pi)$.

B.3 RL AND BC POLICY EVALUATIONS

In order to obtain the best possible manipulation capabilities for use in SayCan, we use a separate evaluation protocol for iterating on the RL and BC policies in the Mock Office Kitchen stations. Evaluations are divided by skill (pick up, knock over, place upright, open/close drawers, move object close to another one), and within each skill, 18-48 skills are sampled from a predetermined set of three objects. Object positions are randomized on each episode, with one or two objects serving as a distractor.

The episode ends when 50 actions have been taken or the policy samples a terminate action. A human operator supervises multiple robots performing evaluation and performs scene resets as needed, and records each episode as a success or failure. Models whose per-skill performance outperforms prior models are "graduated" to the same evaluation protocol in the real kitchen, and then integrated into SayCan. We found that despite the domain shift from Mock Office Kitchen stations to the actual kitchen counter and drawers, higher success rates on mock stations usually corresponded to higher success rates in the real kitchen setting.

Figure 10 shows the development of the manipulation skills over time. It reports the per-skill success rate, the average success rate across all skills, and the number of instructions the policy was trained on. Over the course of the project, we increased the number of skills evaluated, from 1 instruction in April 2021 to hundreds of instructions at time of publication over the course of 366 real-world model evaluations.

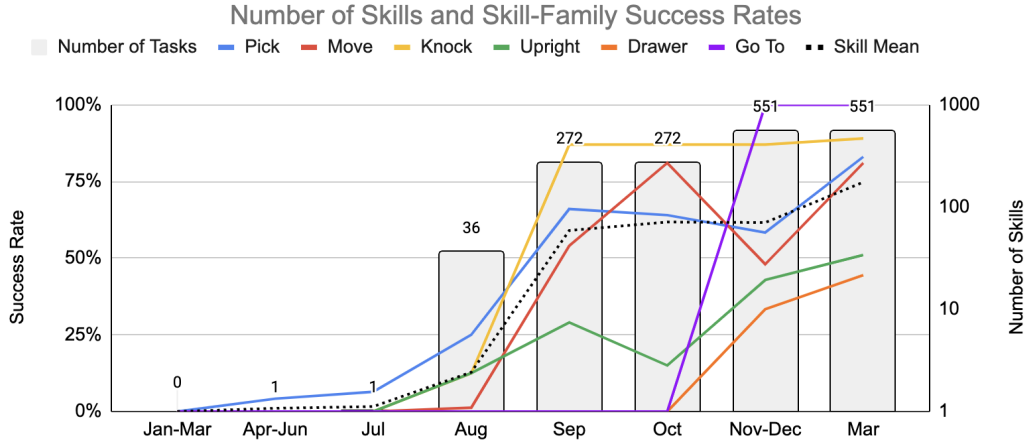


Figure 10: Per-skill evaluation performance of the best policies and number of skills over the duration of the project. The performance as well as the number of skills that the robots are able to handle grow over time due to the continuous data collection efforts as well as improving the policy training algorithms.

C SAYCAN PARAMETERS

C.1 POLICIES AND AFFORDANCE FUNCTIONS

SayCan is capable of incorporating many different policies and affordance functions through its probability interface. Though in principle each type of skill has been trained with the pipeline described in Appendix B, to the success rates seen in Figure 10, we wish to show the generality of SayCan to different policies and affordance functions as well as the robustness of other functions (e.g. distance for navigation). Furthermore, some skills (such as the manipulation skill “move object near object” and “knock object over”) are not naturally part of long-horizon tasks and thus we do not utilize them. Other skills, such as drawer opening, were not consistent enough for long-horizon planning and thus unused. However, we note that as skills become performant or as new skills are learned, it is straightforward to incorporate these skills by adding them as options for LLM scoring and as examples in the prompt. We use the following for each skill family:

- **Pick.** For pick we use the learned policies in Appendix B and Section 4.2 with actions from BC and value functions from RL trained on the same skill. In natural language these are specified as “pick up the object”.
- **Go to.** Since the focus of this work is mainly on planning, we assume the location of objects are known. Thus any navigation skill maps to the coordinate of the object with a classical planning-based navigation stack. In natural language these are specified as “go to location” and “find object”.
- **Place.** Though our manipulation policies have a “place upright” skill, this skill only applies to objects that have a canonical upright direction, e.g., a water bottle but not a bag of chips. One could also train a universal “place” command, but our current policies are trained in a setup-free environment and thus are not amenable to an initial pick. Thus to have a consistent place policy across all objects we use a classical motion planning policy. We use Cartesian space motion planning to plan a path from pre-grasp pose shown in Figure 5d to a gripper release pose. The robot executes that path until the gripper is in contact with a supporting surface, and then the gripper opens and releases the object. In natural language these are specified as “put down the object”.

Recall that we wish to find the affordance function $p(c_\pi | s, \ell_\pi)$, which indicates the probability of c -ompleting the skill with description ℓ_π successfully from state s . Our learned policies produce a Q-function, $Q^\pi(s, a)$. Given $Q^\pi(s, a)$ with action a and state s , value $v(s) = \max_a Q^\pi(s, a)$ is found through optimization via the cross entropy method, similar to MT-Opt. For brevity below we refer to

the value functions by their skill-text description ℓ_π as v^{ℓ_π} and the affordance function as $p_{\ell_\pi}^{\text{affordance}}$. Due to artifacts of training and each implementation, the value functions require calibration to be directly applied as a probability. The parameters used for calibration are determined empirically. Furthermore, SayCan enforces logic that if a skill that has already been completed and the reward received (e.g., navigating to the table the robot is already in front of) then it should not be performed.

- **Pick.** We find the trained value functions generally have a minimum value for when a skill is not possible and a maximum when the skill is successful and thus we normalize the value function to get a affordance function with

$$p_{\text{pick}}^{\text{affordance}} = \text{clamp}\left(\frac{v_{\text{pick}}^{\text{pick}} - v_{\text{pick}}^{\text{min}}}{v_{\text{pick}}^{\text{max}} - v_{\text{pick}}^{\text{min}}}, 0, 1\right), \text{ where } v_{\text{pick}}^{\text{max}} = 0.5, v_{\text{pick}}^{\text{min}} = 0.2.$$

- **Go to.** The affordance function of go to skills are based on the distance d (in meters) to the location. We use

$$p_{\text{goto}}^{\text{affordance}} = \text{clamp}\left(\frac{d_{\text{goto}}^{\text{max}} - d_{\text{goto}}}{d_{\text{goto}}^{\text{max}} - d_{\text{goto}}^{\text{min}}}, 0, 1\right), \text{ where } d_{\text{goto}}^{\text{max}} = 100, d_{\text{goto}}^{\text{min}} = 0.$$

- **Place.** We assume place is always possible, $p_{\text{place}}^{\text{affordance}} = 1.0$, since we find language is sufficient to understand place is only possible after a pick. In the future work having an affordance function module for place could further improve the performance of SayCan.
- **Terminate.** We give terminate a small affordance value, to make sure the planning process terminates when there is no feasible skills to choose from. $p_{\text{terminate}}^{\text{affordance}} = 0.1$.

C.2 LLM

The LLM uses prompt engineering and a strict response structure to score skills. But, as SayCan as a whole requires affordances from a world embodiment, it is not straightforward to optimize this structure and tune parameters quickly. Thus we built a language-based simulator which, given a query and a solution sequence of skills, outputs affordances consistent with the query and solution. It also generates consistent distractor affordances to ensure robustness. The simulator then verifies that SayCan recovers the correct solution and tests how confident SayCan is in the correct solutions. In Table 5 we test the effect of the number of examples in the prompt on the planning success rate in the language-based simulator (over 50 demonstrative instructions). We show a success rate with and without requiring the plan to terminate; without examples we found the LLM was unlikely to issue a “done” phase. With no examples SayCan is able to successfully plan 54% without the done condition, but only 10% with the done condition. Though it makes mistakes, clearly some information is already imbued within the language model. It is able to correctly solve “Can I have a redbull please?” and “Move the chips bag from the table to the counter?”. With only one example the LLM quickly improves in both planning rates, though still fails to terminate the plan occasionally. After only four examples the LLM is performant, planning 82% of the queries correctly, though the remaining errors are largely within a single instruction family: Long-Horizon. Finally, the prompt used in this work, Listing 1, involved 17 examples and recovered 88% of the solutions correctly.

We note here briefly a few lessons learned in prompt engineering and structuring the final prompt. Providing explicit numbers between steps (e.g., 1., 2., instead of combining skills with “and then” or other phrases) improved performance, as did breaking each step into a separate line (e.g. adding a “\n” between steps). Examples which overly include objects used in the actual planning tend to bias results to those objects (e.g., if every example is about apples then the apple scoring will be off in planning). Phrasing of the natural language names of skills and objects is important due to the auto-regressive nature of the LLM scoring – skills and objects should be naturally named and errors such as misspellings or mismatches in “a” vs “an” can be problematic. Notably, since user generated instructions are taken as given such fragility is not issues for the input, allowing a robustness to user queries. For our language model, FLAN (Wei et al., 2021), structuring the interaction as dialog (How would you - I would) was both more natural and performant. Although dialog is used as prompt, the model generalized to imperative sentences at deployment time.

Listing 1: Prompt in SayCan.

```
Robot: Hi there, I'm a robot operating in an office kitchen.
```

Num Examples	Require Termination	No Termination Required
0	10%	52%
1	64%	74%
2	68%	76%
4	82%	84%
8	80%	80%
Full Prompt (17)	88%	88%

Table 5: Number of prompt engineering examples versus planning rate (when requiring the plan to know when to terminate, e.g., issue a “done” command, or not).

Robot: You can ask me to do various tasks and I’ll tell you the sequence of actions I would do to accomplish your task.

Human: How would you hold the snickers?

Robot: 1. pick up the snickers, 2. done.

Human: How would you put the trailmix here?

Robot: 1. put down the trailmix, 2. done.

Human: How would you put a water bottle and an oatmeal next to the microwave?

Robot: 1. find a water bottle, 2. pick up the water bottle, 3. go to microwave, 4. put down the water bottle, 5. find an oatmeal, 6. pick up the oatmeal, 7. go to the microwave, 8. put down the oatmeal, 9. done.

Human: How would you put a grapefruit from the table into the bowl?

Robot: 1. find a grapefruit, 2. pick up the grapefruit, 3. go to the bowl, 4. put down the grapefruit, 5. done.

Human: How would you get a sponge from the counter and put it in the sink?

Robot: 1. find a sponge, 2. pick up the sponge, 3. go to the sink, 4. put down the sponge, 5. done.

Human: How would you move the water bottle from the table to the counter?

Robot: 1. find a water bottle, 2. pick up the water bottle, 3. go to the counter, 4. put down the water bottle, 5. done.

Human: How would you bring me some snacks?

Robot: 1. find a twix, 2. pick up the twix, 3. bring it to you, 4. put down the twix, 5. find an dried fruit, 6. pick up the dried fruit, 7. bring it to you, 8. put down the dried fruit, 9. done.

Human: How would you bring me an orange?

Robot: 1. find an orange, 2. pick up the orange, 3. bring it to you, 4. put down the orange, 5. done.

Human: How would you deliver a can of coke to me?

Robot: 1. find a coke can, 2. pick up the coke can, 3. bring it to you, 4. put down the coke can, 5. done.

Human: I finished my drink, can you throw it away and bring me a mountain dew?

Robot: 1. find a sprite can, 2. pick up the sprite can, 3. go to the trash can, 4. put down the sprite can, 5. find a mountain dew can, 6. pick up the mountain dew can, 7. bring it to you, 8. put down the mountain dew can, 9. done.

Human: How would you let go of the banana?

Robot: 1. put down the banana, 2. done.

Human: How would you put the grapes in the bowl and then move the cheese to the table?

Robot: 1. find the grapes, 2. pick up the grapes, 3. go to the bowl, 4. put down the grapes, 5. pick up the cheese, 6. go to the table, 7. put down the cheese, 8. tell you I’m done.

Human: How would you pick up the peanuts?

Robot: 1. pick up the peanuts, 2. done.

Human: How would you bring me the peanuts?

Robot: 1. pick up the peanuts, 2. bring it to you, 3. put down the peanuts, 4. done.

Human: How would you throw away a coffee cup?

Robot: 1. find a coffee cup, 2. pick up the coffee cup, 3. go to trash can, 4. put down the coffee cup, 5. done.

Human: How would you place a knife and a banana to the table?

Robot: 1. find a knife, 2. pick up the knife, 3. go to the table, 4. put down the knife, 5. find a banana, 6. pick up the banana, 7. go to the table, 8. put down the banana, 9. done .

D EXPERIMENTS

D.1 TASKS

Below we include every instruction run, which environment it was run in, and its planning and execution success rate. Table 5 shows all instructions as broken down by instruction family, which are defined in Section 5.2 and Table 1.

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
How would you pick up the coke can	1	1	1	1
How would you put the coke can in the your gripper	0	0	0	0
How would you grasp the coke can	1	1	1	1
How would you hold onto the coke can	1	1	1	1
How would you lift and hold the coke can up	0	0	0	0
How would you put the coke can down	1	1	1	1
How would you place the coke can on the table	1	1	1	1
How would you let go of the coke can	0	0	1	1
How would you release the coke can	0	0	1	1
How would you place the coke can	0	0	1	1
How would you move to the table	1	1	1	1
How would you go to the table	1	1	1	1
How would you park at the table	1	1	1	1
How would you come to the table	1	1	1	1
How would you navigate to the table	1	1	1	1

(a) NL Single Primitive

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
How would you throw away the apple	1	1	1	0
How would you bring me a sponge?	1	1	1	1
How would you bring me a coke can	1	1	1	1
How would you grab me an apple	1	1	-	-
How would you grab me a 7up from the table	1	1	-	-
How would you deliver the red bull to the close counter	1	0	-	-
How would you throw away the jalapeno chips	1	1	-	-
How would you restock the rice chips on the far counter	1	1	1	1
How would you recycle the coke can	0	0	-	-
How would you throw away the water bottle	1	1	1	1
How would you bring me something hydrating	1	1	-	-
How would you put the apple back on the far counter	1	1	-	-
How would you recycle the 7up	0	0	0	0
How would you throw away jalapeno chips	0	0	-	-
How would you compost the apple	1	0	-	-

(b) NL Verb

D.2 ADDITIONAL RESULTS

Additional results are shown in Figure 11 and some failure cases in Figure 12.

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
How would you bring me lime drink	1	1	1	1
How would you bring me something to clean the kitchen with	1	1	0	0
How would you bring me something to eat	1	1	-	-
How would you put the grapefruit drink on the close counter	0	0	-	-
How would you move the sugary drink to the far counter	0	0	-	-
How would you move something with caffeine from the table to the close counter	0	0	-	-
How would you bring me an energy bar	1	1	-	-
How would you bring me something to quench my thirst	1	1	1	1
How would you bring me a fruit	1	0	1	1
How would you bring me a fruit from the close counter	1	1	-	-
How would you bring me something that is not a fruit from the close counter	0	0	-	-
How would you bring me a soda from the table	0	0	0	0
How would you bring me a soda	1	1	1	1
How would you bring me a bag of chips from close counter	0	0	0	0
How would you bring me a snack	1	1	1	0

(c) NL Nouns

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
How would you pick up the apple and move it to the trash	1	0	1	0
How would you pick up the sponge and bring it to me	1	1	1	1
How would you pick up the coke can and bring it to me	1	1	1	1
How would you pick up the apple and bring it to me	1	1	-	-
How would you pick up the 7up and bring it to me	1	1	-	-
How would you pick up the redbull and move it to the close counter	1	0	-	-
How would you pick up the jalapeno chips and move it to the trash	1	1	-	-
How would you pick up the rice chips and move it to the far counter	1	1	1	0
How would you pick up the coke can and move it to the trash	1	1	-	-
How would you pick up the water bottle and move it to the trash	1	1	1	1
How would you pick up the grapefruit soda and bring it to me	1	1	-	-
How would you pick up the apple and move it to the far counter	1	1	-	-
How would you pick up the 7up and move it to the trash	1	1	0	0
How would you pick up the jalepeno chips and move it to the trash	1	1	-	-
How would you pick up the apple and move it to the trash	1	1	-	-

(d) Structured Language

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
How would you put the coke can down on the far counter(with operator)	1	1	-	-
How would you put the coke can down on the far counter(at table)	1	0	-	-
How would you put the coke can down on the far counter(at table with coke can in hand)	1	1	-	-
How would you put the coke can down on the far counter(at far counter with coke can in hand)	1	1	-	-
How would you put the sponge on the close counter(with operator)	1	1	-	-
How would you put the sponge on the close counter(at far counter)	1	1	-	-
How would you put the sponge on the close counter(at far counter with sponge in hand)	0	0	-	-
How would you put the sponge on the close counter(at close counter with coke can in hand)	0	0	-	-
How would you pick up the drink from the far counter	1	1	-	-
I left something on the table, can you throw it away?	0	0	-	-
I left something on the table or the counter, can you bring it to me?	0	0	-	-

(e) Embodiment

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
I opened a pepsi earlier. How would you bring me an open can?	1	1	-	-
I spilled my coke, can you bring me a replacement?	1	1	1	1
I spilled my coke, can you bring me something to clean it up?	1	1	1	1
I accidentally dropped that jalapeno chip bag after eating it. Would you mind throwing it away?	0	0	-	-
I like fruits, can you bring me something I'd like?	1	1	-	-
There is a close counter, far counter, and table. How would you visit all the locations?	1	1	-	-
There is a close counter, trash can, and table. How would you visit all the locations?	1	1	-	-
Redbull is my favorite drink, can I have one please?	1	1	1	1
Would you bring me a coke can?	1	1	1	1
Please, move the pepsi to the close counter	1	1	-	-
Please, move the ppsi(intentional typo) to the close counter	0	0	-	-
Can you move the coke can to the far counter?	1	1	-	-
Can you move coke can to far counter?	1	0	-	-
Would you throw away the bag of chips for me?	0	0	0	0
Would you throw away the bag of chips(intentional typo) for me?	0	0	-	-

(f) Crowd-Sourced

Instruction	Mock Kitchen		Kitchen	
	Plan	Execution	Plan	Execution
How would you put an energy bar and water bottle on the table	0	0	0	0
How would you bring me a lime soda and a bag of chips	0	0	0	0
Can you throw away the apple and bring me a coke	1	1	-	-
How would you bring me a 7up can and a tea?	0	0	-	-
How would throw away all the items on the table?	0	0	0	0
How would you move an multigrain chips to the table and an apple to the far counter?	1	1	-	-
How would you move the lime soda, the sponge, and the water bottle to the table?	1	0	-	-
How would you bring me two sodas?	0	0	-	-
How would you move three cokes to the trash can?	0	0	-	-
How would you throw away two cokes?	1	1	1	1
How would you bring me two different sodas?	1	1	-	-
How would you bring me an apple, a coke, and water bottle?	1	0	-	-
I spilled my coke on the table, how would you throw it away and then bring me something to help clean?	1	1	1	1
I just worked out, can you bring me a drink and a snack to recover?	0	0	-	-
How would you bring me a fruit, a soda, and a bag of chips for lunch	0	0	0	0

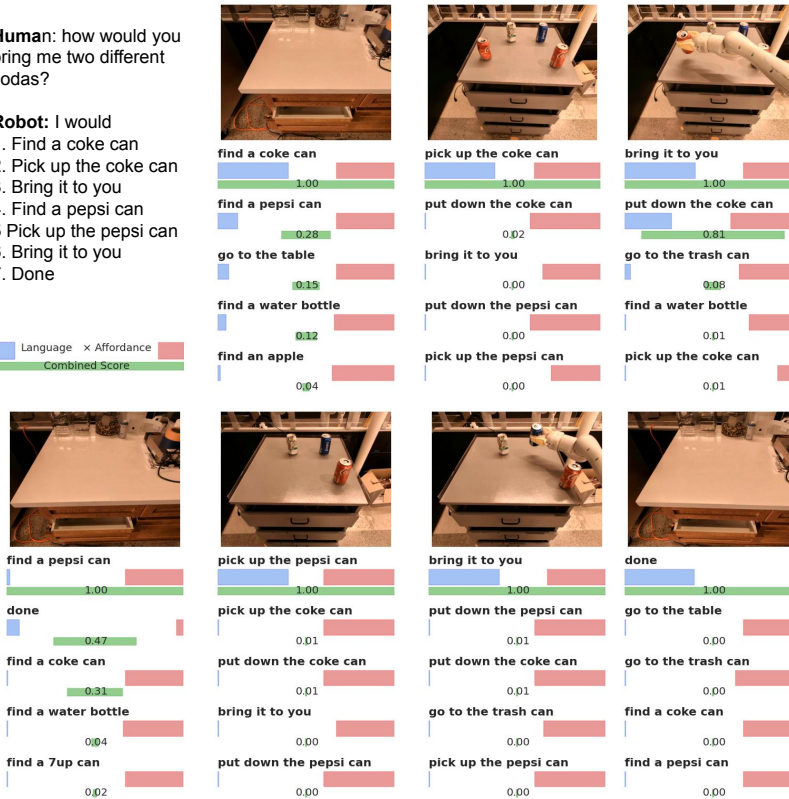
(g) Long-Horizon

Table 5: **List of all instruction and experiment results** We evaluate the algorithm on 101 instructions on 2 scenes, 147 experiments in total. The metrics and success definitions can be found in Sec. 5.3. A 0 indicates failure, 1 indicates success, and - indicates the experiment was not ran in the particular setup.

Human: how would you bring me two different sodas?

Robot: I would
 1. Find a coke can
 2. Pick up the coke can
 3. Bring it to you
 4. Find a pepsi can
 5. Pick up the pepsi can
 6. Bring it to you
 7. Done

Language × Affordance
 Combined Score

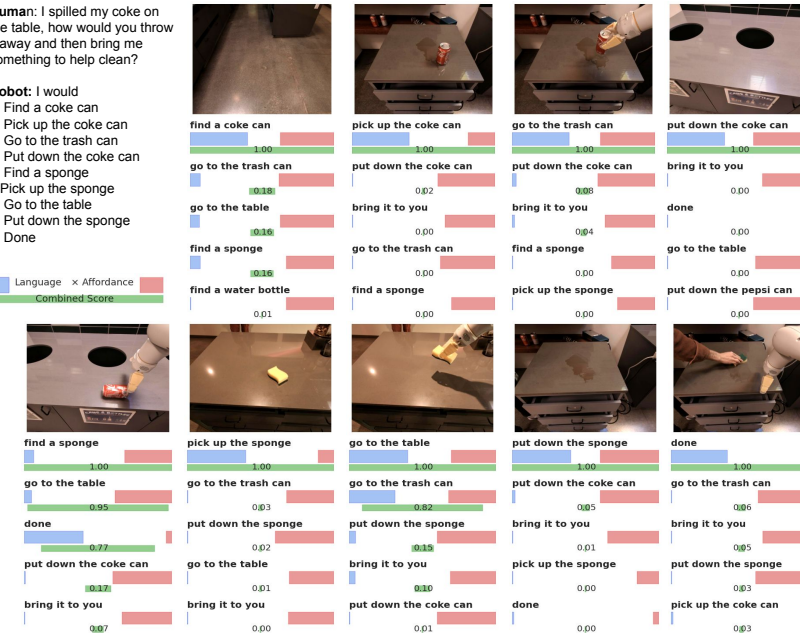


(a) In this long-horizon task, the language model gives high score to the two sodas. After the coke is delivered, the language model scores pepsi higher. The affordance rating overcomes potential early termination after the first can has been delivered.

Human: I spilled my coke on the table, how would you throw it away and then bring me something to help clean?

Robot: I would
 1. Find a coke can
 2. Pick up the coke can
 3. Go to the trash can
 4. Put down the coke can
 5. Find a sponge
 6. Pick up the sponge
 7. Go to the table
 8. Put down the sponge
 9. Done

Language × Affordance
 Combined Score



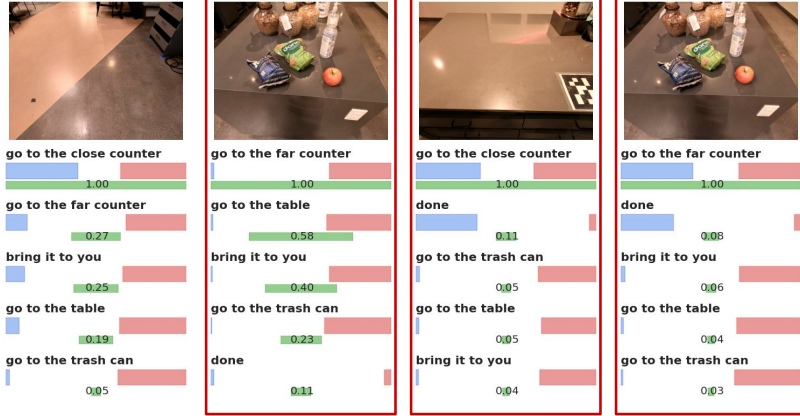
(b) In this task, the model completes a 9-step plan. It narrowly avoids an early termination at step 5.

Figure 11: Long horizon sequences, see the video on our website say-can.github.io for more.

Human: How would you bring me a bag of chips from close counter?

Robot: I would
 1. Go to the close counter
 2. Go to the far counter
 3. Go to the close counter
 4. Go to the far counter ...

Language × Affordance
 Combined Score

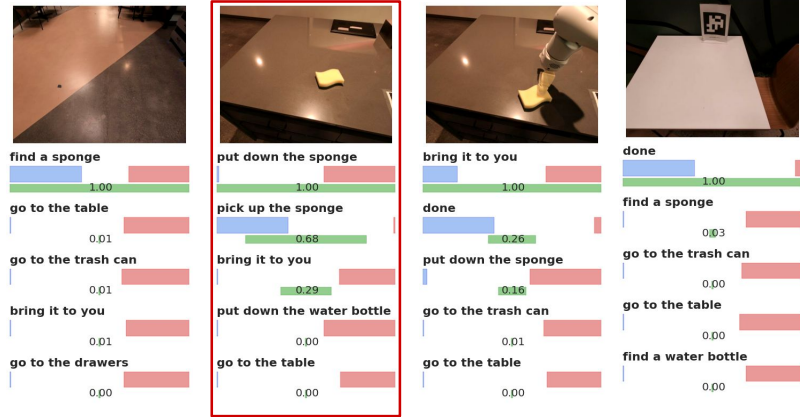


(a) The affordance model fails to identify either bag of chips as pickable, though the language model approaches the counter twice.

Human: How would you bring me something to clean the kitchen with?

Robot: I would
 1. Find a sponge
 2. Put down the sponge
 3. Bring it to you
 4. Done

Language × Affordance
 Combined Score

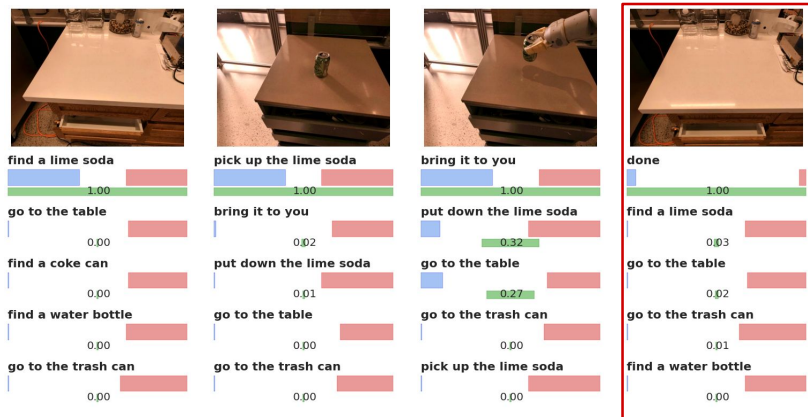


(b) The affordance model fails to identify the sponge as pickable.

Human: How would you bring me a lime soda and a bag of chips?

Robot: I would
 1. Find a lime soda
 2. Pick up the lime soda
 3. Bring it to you
 4. Done

Language × Affordance
 Combined Score



(c) Language model terminates a long-horizon task prematurely.

Figure 12: Failure cases. The planning success rate was 70%. Of the errors, 65% were a result of an LLM error and 35% were affordance errors.