

A Survey of Robotic Language Grounding: Tradeoffs between Symbols and Embeddings

Vanya Cohen¹, Jason Xinyu Liu², Raymond Mooney¹, Stefanie Tellex^{2,3} and David Watkins³

¹UT Austin

²Brown University

³The AI Institute

{vanya, mooney}@cs.utexas.edu, {xliu53, stefie10}@cs.brown.edu,
{stefie10, dwatkins}@theaiinstitute.com

Abstract

With large language models, robots can understand language more flexibly and more capable than ever before. This survey reviews and situates recent literature into a spectrum with two poles: 1) mapping between language and some manually defined formal representation of meaning, and 2) mapping between language and high-dimensional vector spaces that translate directly to low-level robot policy. Using a formal representation allows the meaning of the language to be precisely represented, limits the size of the learning problem, and leads to a framework for interpretability and formal safety guarantees. Methods that embed language and perceptual data into high-dimensional spaces avoid this manually specified symbolic structure and thus have the potential to be more general when fed enough data but require more data and computing to train. We discuss the benefits and tradeoffs of each approach and finish by providing directions for future work that achieves the best of both worlds.

1 Introduction

Large language models (LLMs) have fueled a surge of interest in the problem of making robots understand natural language commands. Solving this problem requires mapping between words in language and actions or behaviors taken by the robot. [Harnad, 1990; Harnad, 2007] defined the symbol grounding problem as constructing a mapping from symbols of a symbolic system or words in an utterance to sensorimotor substrates in the physical world. Large language models (LLMs) semantically represent concepts without explicit higher-order symbols beyond the words in the text, leading many to try end-to-end approaches for robotic language understanding. Yet many recent works in robotic language understanding leverage large language models hand-in-hand with formal symbolic representations. For example, Code as Policies [Liang *et al.*, 2022] generates Python code with predefined Python APIs. SayCan [Ahn *et al.*, 2022] grounds natural language commands to predefined discrete skills implemented using a deep neural network. Other approaches

take a more end-to-end approach, such as VIMA [Jiang *et al.*, 2023], which learns a mapping from vision and language instructions to low-level robot actions such as joint states.

This survey paper evaluates work that grounds natural language to robot behavior. We observe that these approaches can be situated on a spectrum ranging between two high-level approaches: mapping between language and a manually defined formal representation and mapping between language and high-dimensional vector spaces that translate directly to low-level robot policy. We define the advantages and limitations of each approach.

Using a formal representation constrains the search space during training and inference and may require less training data. It also provides a natural framework for strong interpretability and formal safety guarantees. Well-defined model-checking tools exist to check whether the model of a system meets a given logical specification [Baier and Katoen, 2008]. Formal methods can also synthesize correct-by-construction robot controllers given a logical specification and the system model and provide counterexamples to explain failure cases [Kress-Gazit *et al.*, 2018]. However, a formal representation constrains the space of possible models that can be learned, limiting the system’s ability to represent the meanings of what a person may say. With the advent of large language models, mapping between human language and a formal language is much easier; the research questions then need to focus on what formal language to use, where it comes from, and how it connects to the physical world. There are opportunities to more easily use existing representations such as the planning domain definition language (PDDL), linear temporal logic (LTL), or motion planners without collecting large training sets. Many approaches at the border, such as SayCan [Ahn *et al.*, 2022], map language to a manually specified vocabulary of robot skills but give their formal language relatively little attention despite it playing a critical role in the system.

End-to-end approaches may require more data to train but are more flexible in representing a user’s intended meaning and translating it to robot behavior. [Ng and Jordan, 2001] observed that structured generative models perform better with less training data. In contrast, discriminative models with more parameters can perform better when given lots of data because they make fewer assumptions about the struc-

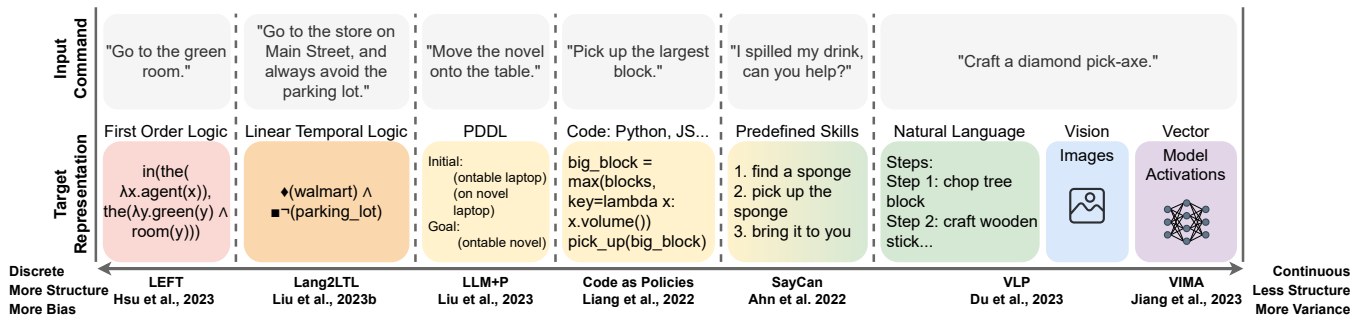


Figure 1: Approaches to representing natural language for robotics fall along a spectrum from more symbol-like representations to more continuous embedding-like representations. However, most approaches use a mixture of both. SayCan uses a fixed ontology of predefined skills but implements these as neural value functions conditioned on language.

ture of the learned model. Similarly, end-to-end neural embedding approaches require more data but can generalize better than formal methods because they place fewer constraints on the learned model. True “pixels to torque” approaches [Wahlström *et al.*, 2015] learn to produce motor torques directly from sensor input; however, many end-to-end approaches use intermediate outputs such as end-effector poses or joint states. The challenge with end-to-end approaches is acquiring enough training data, processing this data, generalizing outside the training set, chaining policies together to produce long-term behaviors, explaining robot behaviors, and providing safety guarantees.

We situate the robotic language grounding works on a spectrum in Figure 1 and describe more formal work in Section 2 and more end-to-end work in Section 3. More formal work uses discrete and structured representations that introduce bias into the learning. We begin with the most structured, abstract representations and successively review less structured, lower-level, continuous representations. Many of the implementations surveyed combine aspects of formal and end-to-end representations. In Section 4, we survey output representations for both approaches, review available datasets, and discuss the scope of natural language commands understood in theory and practice by different approaches. We conclude by identifying key open problems and recommending future research directions.

This new review paper focuses on the transformative role LLMs have played in this space. Other relevant reviews are:

- [Tellex *et al.*, 2020] is a review that surveys robot and language grounding work predating LLMs’ emergence.
- [Zhang *et al.*, 2023] reviews large language models more broadly used in human-robot interaction, including question answering, social robotics, and instruction following.
- [Zeng *et al.*, 2023b] reviews LLMs applied to robotics broadly, including related technologies, but does not focus on the spectrum from formal methods to high-dimensional vectors as this paper does.
- [Wang *et al.*, 2024] reviews applications of LLMs to robotics but does not situate the work on a spectrum and focuses on a broader set of tasks than command understanding.

Our work, in contrast, focuses specifically on the problem of command understanding, situating work along a spectrum based on formal methods.

2 Mapping from Natural Language to a Formal Representation

Works closer to the formal end of the spectrum map natural language commands from humans to a manually defined formal representation, e.g., temporal logic, planning domain definition language (PDDL), computer code, or some predefined skills. The symbols in the formal representation are then grounded to robot percepts and control by predefined detectors and controllers, respectively. Unlike machine translation of natural languages, where vast training data is available online, translating natural language to logic often lacks labeled pairs of natural language commands and logic formulas. Most recent works leverage few-shot learning or fine-tuning of large language models (LLMs) for various parts of the language grounding system to address the lack of training data.

Many formal representations used to ground natural language are Turing-complete and thus can be translated from one to another. However, depending on the language used, this translation may be direct or indirect and require significantly longer or more complex expressions. For example, linear temporal logic (LTL) can naturally represent English sentences such as “Avoid the red room” with a short, direct expression. Python can represent the same command by defining an “avoid” function but may need a significantly longer program if the function is not provided. Thus, in our review, we order these representations from those with more structure and bias to those with less structure and less bias, as typically used, shown in Figure 1.

More structured methods often map directly to certain natural language commands and express the goal or constraint directly rather than imperatively how to achieve it. Goal-based representations specify what state the world should be in but not what actions the robot should take to attain that state. In contrast, action-based representations specify a sequence of actions but not necessarily the goal or results of those actions. A goal-based representation for “avoid the red room” might be a logical formula such

	Formal vs. Intermediate vs. End-to-End	Intermediate Grounding	Final System Output	Executing Output on Robots	Domains
LEFT [Hsu <i>et al.</i> , 2023]	Formal	N/A	FOL formula	Low-level controllers	Pick-and-Place (Real)
Lang2LTL [Liu <i>et al.</i> , 2023b]	Formal	N/A	LTL formula	LTL-MDP planner, Low-level controllers	Navigation (Real)
[Xie <i>et al.</i> , 2023]	Formal	N/A	PDDL goal	Symbolic planner, Low-level controllers	Navigation, Mobile manipulation (Sim)
LLM+P [Liu <i>et al.</i> , 2023a]	Formal	N/A	PDDL	Symbolic planner, Low-level controllers	Manipulation (Real)
SayCan [Ahn <i>et al.</i> , 2022]	Formal	N/A	predefined skill	Low-level controllers	Mobile manipulation (Real)
Code as Policies [Liang <i>et al.</i> , 2022]	Formal	N/A	Python code	Python interpreter, Low-level controllers	Pick-and-place, Mobile manipulation Drawing (Real)
ProgPrompt [Singh <i>et al.</i> , 2023]	Formal	N/A	Python code	Python interpreter, Low-level controllers	Mobile manipulation (Sim) Manipulation (Real)
Voyager [Wang <i>et al.</i> , 2023]	Formal	N/A	Javascript Code	Javascript interpreter, Low-level controllers	Minecraft survival tasks (Sim)
ITP [Li <i>et al.</i> , 2023a]	Intermediate	High-level action descriptions	Python code	Python interpreter, Low-level controllers	Manipulation (Real)
UniSim [Yang <i>et al.</i> , 2023]	Intermediate	Image subgoals	predefined skill	Low-level controllers	Manipulation (Real)
SuSIE [Black <i>et al.</i> , 2023]	Intermediate	Image subgoals	End-effector pose	Low-level controllers	Manipulation (Real)
VLP [Du <i>et al.</i> , 2023]	Intermediate	Image subgoals	End-effector pose	Low-level controllers	Manipulation (Real)
PaLM-E [Driess <i>et al.</i> , 2023]	End-to-End	N/A	predefined skill	Low-level controllers	Pick-and-place, Manipulation (Real)
VIMA [Jiang <i>et al.</i> , 2023]	End-to-End	N/A	predefined skill	Low-level controllers	Manipulation (Sim)
PerAct [Shridhar <i>et al.</i> , 2023]	End-to-End	N/A	End-effector pose	Low-level controllers	Manipulation (Real)
RT-1 [Brohan <i>et al.</i> , 2023a]	End-to-End	N/A	End-effector pose	Low-level controllers	Manipulation (Sim, Real)
RT-2 [Brohan <i>et al.</i> , 2023b]	End-to-End	N/A	End-effector pose	Low-level controllers	Mobile Manipulation (Real)
RT-X [O’Neill <i>et al.</i> , 2024]	End-to-End	N/A	End-effector pose	Low-level controllers	Manipulation (Real)

Table 1: Method Comparison

as $\neg red_room$ while an action-based expression might be $North; North; West; West; North; North$ (depending on the specific geometry of the environment). We order our review from high-level, abstract representations to low-level, concrete, fine-grained representations.

2.1 Logics

Logics are mathematically precise goal-based representations that specify robotic goals and provide guarantees for robot behaviors. Temporal logics can concisely represent long-horizon, temporally extended tasks. [Kress-Gazit *et al.*, 2018] surveyed the uses of several temporal logics as task specifications for the formal synthesis of robot controllers. We consider logical expressions at the most formal end of the spectrum because they map to goals, describing abstractly the state of the world corresponding to the language, leaving the plan to achieve this state of the world to other modules.

To train their language grounding system on diverse natural language commands, [Pan *et al.*, 2023] used LLMs to paraphrase structured English commands generated from algorithmically produced LTL formulas. [Wang *et al.*, 2021] and [Patel *et al.*, 2020] trained a semantic parser to map language commands to LTL formulas using weak supervision of execution trajectories without any LTL annotations. Lang2LTL [Liu *et al.*, 2023b] is a modular system that uses LLMs to ground navigation commands to linear temporal logic (LTL) formulas and their propositions to physical landmarks in a given semantic map. The same system solved navigational tasks in indoor and outdoor environments without retraining on language data by harnessing pre-trained LLMs. Similarly, [Hsiung *et al.*, 2022] first translated commands to lifted LTL formulas and then grounded them to specific domains for better generalization. Other approaches [Fuggitti and Chakraborti, 2023; Chen *et al.*, 2023] also used LLMs to translate natural language commands to logical representations but did not ground the formulas to a robot domain.

AutoTAMP [Chen *et al.*, 2024] uses LLMs to translate task and state descriptions to signal temporal logic (STL) formulas [Maler and Nickovic, 2004] and correct syntax errors if detected. It then uses an STL planner to generate trajectories. By using an intermediate formal task specification, AutoTAMP outperforms LLM planners on tasks with geometric and temporal constraints in 2D domains.

Recent work also leveraged LLMs for grounding natural language commands to first-order logic (FOL) formulas. LEFT [Hsu *et al.*, 2023] used an LLM to translate natural language queries to FOL programs, which a differentiable FOL executor executed. At the same time, a domain-specific grounding model grounded the symbols of the FOL program to various input modalities, e.g., 2D images and point clouds. The advantage of using first-order logic over LTL is capturing commands that use quantifiers and predicates for generalization. On the other hand, LTL provides a natural way to concisely represent temporally extended commands. Different logic representations can be formally composed to form a new, more expressive logic representation. These representations can enable direct mapping between abstract concepts in a language such as “avoid” and a precise, formal logical expression that guarantees task satisfaction.

2.2 Planning Domain Definition Language (PDDL)

The Planning Domain Definition Language (PDDL) [McDermott *et al.*, 1998; Fox and Long, 2003; Edelkamp and Hoffmann, 2004; Kovacs, 2011] is a structured representation that defines a planning problem. It consists of a domain, which defines objects, predicates, and actions that govern the world’s rules, and a problem, a grounded problem instance with an initial state and a desired goal state. A symbolic planner takes as input the PDDL domain and problem and then outputs a plan [Helmert, 2006], i.e., a sequence of actions, to reach the goal state from the initial state. In this sense, it is a goal-based representation, but because it encodes actions and

their effects, it can also be used imperatively. Recent works used LLMs to translate natural language descriptions of the world and the task to a PDDL representation of the planning problem, which a symbolic planner can then use. Compared to logical representations, PDDL provides a language of predicate states and transitions that allow goals to be translated into lower-level actions and skills but requires a full domain specification. People have defined PDDL domains for a wide variety of real-world applications. [Konidaris *et al.*, 2018] showed how a robot can learn symbols for low-level skills that are both necessary and sufficient to enable planning in PDDL.

[Xie *et al.*, 2023] prompted LLMs with a PDDL domain description, an initial state, and examples to translate a natural language goal description to a PDDL goal specification. Results in simulation showed that LLMs could translate unambiguous goal descriptions and fill in missing details for under-specified goals but struggled with numerical and spatial reasoning. [Collins *et al.*, 2022] also showed that translating natural language goal descriptions to PDDL goals and then solving them using a symbolic planner outperformed directly using an LLM as a planner.

[Guan *et al.*, 2023] and [Liu *et al.*, 2023a] prompted an LLM to translate a natural language description of a problem into a complete PDDL problem definition, which was then fed into a symbolic planner together with a PDDL domain definition to find an optimal plan. These approaches outperform LLM planners and provide strong correctness guarantees from using a symbolic planner. Recent work has also created benchmarks for evaluating LLM’s ability to make plans concerning PDDL and other AI baselines [Valmeekam *et al.*, 2023a; Valmeekam *et al.*, 2023b], finding that LLMs achieve low success rates across domains in isolation. Still, they can improve the search process for underlying sound planners, leading to higher performance.

Unlike other works that use LLMs for AI planning, [Silver *et al.*, 2024] developed a generalized planner to synthesize Python programs to solve novel tasks by prompting LLMs with a domain specification and a few training tasks emphasizing satisfaction and efficiency. Their system also automatically detects planning errors and then re-prompts the LLM with feedback.

2.3 Code

Code is a very flexible form of formal representation that can be used in a goal-based or an action-based manner. Indeed, a Python program can be generated that directly outputs motor torques for robots or implements an arbitrary deep-learned function. A formal language is typically specified as a subset of the language at hand using a manually defined programming API consisting of functions and their arguments analogous to robot skills. Skills or functions are not simply linearly called but can be embedded in more complex logic, like conditionals and loops.

Code as Policies [Liang *et al.*, 2022] prompted an LLM with import statements, example code, and code comments that describe the desired policy to generate Python code executable directly on the robot. It solved manipulation and mobile manipulation tasks with API calls to Python li-

braries and predefined perception and control modules. Prog-Prompt [Singh *et al.*, 2023] applied a similar approach with additional assertion statements to recover from errors when reliable state tracking is available. [Varley *et al.*, 2024] developed a modular bi-arm system that employs an LLM to map natural language instructions to high-level API calls to manipulation skills powered by VLMs and a control module to solve three tabletop bimanual manipulation tasks. Their experiment results highlighted the benefits of modularity for ensuring safety, interpreting failures, and identifying modules to improve. Socratic Models [Zeng *et al.*, 2023a] also demonstrated the code generation capabilities of LLMs for simulated pick-and-place tasks, provided with pre-trained perception modules and robot policies. ITP [Li *et al.*, 2023a] used LLMs to generate high-level action sequences and then translated the action descriptions to predefined function calls using APIs for an object detector and robot policies to solve tabletop manipulation tasks. By storing completed high-level actions, ITP can replan given new commands at any step during the execution. Voyager [Wang *et al.*, 2023] extended the code generation capabilities of LLMs to build a lifelong learning agent in Minecraft that continuously explores and learns a skill library of executable code.

2.4 Predefined Skills

Recent methods have used LLMs as a planner to map language commands to a sequence of predefined skills. Skills can be learned from data or manually specified. These predefined skills are a formal representation since they are discrete and manually specified. (Even if skills are learned, they are frequently learned from human-provided demonstrations, which provide the discrete structure for the skills.) They are also action-based representations specifying actions to take rather than goals or resultant states. For example, a robot might be provided with skills such as “pick up,” “put down,” “drive to refrigerator,” “drive to microwave,” and “clean table.” Then, the challenge is to map a natural language instruction such as “Get me the apple” and “Clean the table with the sponge” to a sequence of skills. [Huang *et al.*, 2022a] iteratively prompted an LLM to decompose a high-level task specification in natural language to a sequence of action descriptions. They used sentence similarities to map the proposed action descriptions to actions available in a simulated environment where their corresponding predefined low-level controllers can be executed. SayCan [Ahn *et al.*, 2022] iteratively prompted an LLM to sequence pre-trained skills described by predefined verb phrases based on their probabilities of success from the current state to solve mobile manipulation tasks specified by natural language on a physical robot. Their approach relies on both formal representations and high-dimensional, end-to-end differentiable representations. Despite using a fixed ontology of skills, these are implemented using a deep approach. SayCan learns these skills using a multi-task value function conditioned on language. CAPE [Raman *et al.*, 2024] also used an LLM planner to sequence predefined skills. When a plan does not meet the precondition of some skill, it re-prompted the LLM with corrective feedback. In addition to using an LLM to ground language to predefined skills, Inner Monologue [Huang *et al.*,

2022b] also enabled language feedback based on pre-trained perception models for describing the scene and detecting successful execution of skills. In these methods, the specific skills are often given relatively little attention, yet having the right set of skills at the right level of granularity is critical to success. We categorize this method on the formal side of the spectrum because the skills impose significant structure. That said, it imposes less structure than goal-based methods such as logical representations, leaving the LLM to keep track of higher-level constructs such as constraints, sequencing, and conditionals.

3 Mapping from Natural Language to High-Dimensional Vectors to Actions

On the other end of the spectrum are approaches that map natural language instructions to high-dimensional, non-formal representations. These deep approaches are largely defined by data and learning, whereas formal, symbolic representations are largely human-crafted and invoke manually provided structure. The large models that enable deep and many symbolic approaches are created by self-supervised pretraining on large datasets scraped from the web. These models thereby learn a generative representation of images, text, and other modalities that capture background and task knowledge that is useful for robotic applications, for example, as in [Driess *et al.*, 2023]. By learning to model a large and varied distribution, the models also learn to perform tasks useful for robotic language grounding, including translation, semantic parsing, and broadly useful ones like in-context learning [Brown *et al.*, 2020].

Pretrained models can represent high-level and low-level semantics, so they can represent high-level instructions and low-level actions. This semantic understanding can extend across modalities while retaining these task-learning abilities, forming combined representations of language and vision useful to robotics problems. Beyond their representation capacity, a unified interface for training and inference makes these methods particularly powerful. Multiple models across modalities can be connected and trained end-to-end through gradient descent, and pretrained representations can be improved through increases in model and data scale [Brown *et al.*, 2020; Henighan *et al.*, 2020; Hoffmann *et al.*, 2022; Wei *et al.*, 2022]. Provided with a method for generating and standardizing data, a single pipeline with minimal human intervention can be used for continued improvement and deployment. One limitation of deep methods is that they require abundant data, which limits their applicability to domains where data is easy to collect. However, large-scale pretraining can significantly reduce the amount of task-specific training data required [Brown *et al.*, 2020].

In robotic language grounding, deep learning approaches map input language commands to one of the following representation types: actions represented as low-level controls (policies) or high-level goals represented as natural language, images, or neural network activations. These high-level goals are used to condition low-level actions from planners or policies to drive robot actions. At the lowest level, the approaches map language to sequences of joint states, end-effector poses,

or motor torques. High levels of abstraction are easier for humans to interpret versus low-level joint states, making it difficult to discern the outcome of the action. We order from high-level to low-level to indicate the levels of abstraction as they have evolved in the field and to create a spectrum for interpretability.

3.1 Image and Language Subgoals

These methods use images and natural language to express subgoals. [Black *et al.*, 2023] perform pick-and-place tasks on a robot by alternating between generating images representing sub-goals using a text-guided image-editing model and executing a low-level policy conditioned on the goal images. They map language to subgoals represented as images and then map these images to end-effector positions. VLP [Du *et al.*, 2023] transforms the input image and command into a sequence of language and image subgoals and then generates plan rollouts with tree search using a language-conditioned video generation model. It was deployed on a robot to perform tabletop arrangement tasks.

UniSim [Yang *et al.*, 2023] trains a generative video model to predict the outcome of both high and low-level actions. Using this model, the authors train high-level planners and low-level policies in simulation and demonstrate zero-shot transfer to real-world autonomous driving tasks. Language and images are mapped to high-level skills and low-level controls, e.g., “move the gripper to (x, y)”. GAIA-1 [Hu *et al.*, 2023] similarly leverages real-world driving data to develop a world model that generates life-like driving behavior by predicting outcomes in the video space. Coarse symbols, such as words in English, are not as rich of a representation as images. Incorporating multimodality increases the generalizability and expressiveness of the systems’ states. This allows for directly representing the visual world state without introducing an intermediate representation like natural language, scene graphs, or maps. The downside to using images is that it requires more pretraining and fine-tuning data to learn this multimodal representation.

3.2 End-Effector and Joint-State Goals

These approaches use low-level skills and parameterize actions using joint states or end-effector position, in contrast to high-level skills like “pick up” that we classify on the formal side of the spectrum. Often, they are not learned because they are sufficiently low-level, and the built-in controllers on the robot are sufficient. For example, the RT-* family of models [Brohan *et al.*, 2023a; Brohan *et al.*, 2023b; O’Neill *et al.*, 2024] use an action space that specifies the arm, base, and end-of-episode token. For the arm, they use the end-effector pose of the gripper (vs. lower-level joint states or joint torques) and rely on inverse kinematics to find joint states to move to the end-effector pose.

RT-1, RT-2, and RT-X [Brohan *et al.*, 2023a; Brohan *et al.*, 2023b; O’Neill *et al.*, 2024] perform various language-conditioned tasks on 22 distinct robotic platforms. They collected a large and diverse dataset of demonstrations and trained a large multimodal transformer that maps language and state observations to low-level discrete end-effector controls. These methods attempt to demonstrate positive transfer,

the transfer of knowledge between tasks that allows a system to perform better at both tasks simultaneously. The Open-X Embodiment dataset [O’Neill *et al.*, 2024] release came with baselines showing this phenomenon, whereas GATO [Reed *et al.*, 2022] did not. Open-X Embodiment targeted end-effector positions and used models from 35M to 55B parameters, whereas GATO targeted joint states directly and used 1.2B parameters. Octo [Octo Model Team *et al.*, 2023] also shows better performance across tasks by leveraging the Open-X Embodiment dataset while simultaneously outputting end-effector positions and joint-state control. PaLM-E [Driess *et al.*, 2023] is an image and text multimodal model finetuned end-to-end for various multimodal reasoning tasks, including robot manipulation. This model maps visual-language input describing a scene to a limited vocabulary of actions, which are then turned into robot actions. VIMA [Jiang *et al.*, 2023] learns a multimodal model that maps tasks specified with images and language to high-level actions parameterized by low-level end-effector controls to perform tabletop arrangement tasks on a simulated robot. One consequence of this level of abstraction is that what is learned is often specific to a robotic platform and cannot be generalized from one environment and robot to another without more data. Generating data automatically in simulation allowed it to overcome limitations on the data required to learn a meaningful representation. Future work is focused on increasing the size and diversity of the datasets to enable more generalization.

Recent works in foundation models have shown tremendous improvement in generalization by leveraging large amounts of data. These methods often use symbols in the form of vector-quantized image patches to reduce the number of tokens required [Yan *et al.*, 2023]. ALOHA [Zhao *et al.*, 2023] seeks to address the limited data for training these systems by utilizing accessible hardware for many research institutions. They then use an end-to-end foundation model that targets joint-state control of robots to execute policies on the bimanual robot platform. [Li *et al.*, 2023b] addresses domain limitations by leveraging vision and text models that are larger and trained on more data. [Shridhar *et al.*, 2023] develop a language-conditioned behavior-cloning model that encodes the inputs of a natural language command and a voxel grid to predict a collision-free pose that a motion planner can execute. Similarly, VPT [Baker *et al.*, 2022] utilizes real demonstrations of users playing Minecraft to train better policies that target a discrete set of actions for the agent to follow. They expand their dataset by training an inverse dynamics model to label data from YouTube. Joint-state and end-effector goals are useful for simultaneously targeting multiple robotic platforms but can lack the interpretability of their output.

4 Discussion and Future Directions

Each end of the spectrum has tradeoffs and advantages, and the two approaches are complementary in many ways. As shown, many recent works combine aspects of symbolic and end-to-end methods. A number of open problems and exciting future research directions are to leverage the best of both worlds.

4.1 Representation

Table 1 summarizes the grounding representations used by various systems we review and how they execute their output on a robot in various simulated and real-world domains. On the formal side of the spectrum, system output ranges from more structured logic formulas and code to less structured predefined skills. On the deep side of the spectrum, methods ground language input to high-dimensional representations of joint states or end-effector poses. Some deep approaches first ground language to an intermediate representation of image or language subgoals, then low-level controls. Formal approaches to executing the system output on a robot use a planner or a code interpreter together with low-level controllers. In contrast, deep approaches use low-level controllers to produce desired robot motion. A key takeaway is using intermediate representations at different system levels, which enables the use of “off-the-self” robot modules such as SLAM and object detectors. When using these modules, there is a crisp abstraction barrier between the learning model and the robot, limiting what needs to be learned and the flexibility of the learned system. Automatically learning an extendable set of symbols and grounding those symbols [Gopalan *et al.*, 2020] can address this limitation of formal approaches. Another key research question for formal representations is what formal language to use. Yet a formal language that is powerful and flexible enough to capture all English or other natural languages is unknown. Approaches such as RLANG [Rodriguez-Sanchez *et al.*, 2023] expand the scope of what language can talk about (from goals and actions to observations, states, and transition functions). Integrating these formal, manually specified languages with open-class learned models is critical, for example, jointly learning skills and symbols [Konidaris *et al.*, 2018; Konidaris, 2020]. This survey only focuses on natural language, yet using multimodal input, e.g., text, audio, RGB-D images, videos, and joint trajectories, etc., to complement language [Reed *et al.*, 2022] in solving robotic tasks is also a promising future research direction. Current state-of-the-art approaches combine multiple modalities in a single neural architecture capable of running in real-time, including speech, image, and text [Reid *et al.*, 2024]. This promises to enable more audio interactivity between humans and robots.

4.2 Situated Natural Language Commands

Table 2 reviews various domains and natural language commands used to evaluate each system we review. Domains vary from tabletop manipulation tasks, usually focusing on pick-and-place, to mobile manipulation, usually involving chaining pick-and-place actions and navigation, in both simulation and the physical world. Table 2 shows examples of natural language commands that specify temporally extended tasks and spatial relations among objects in various domains. Natural language was created based on an assumption about human reaction time. The methods described in this work have control loops in the space of 1-3 Hz, which cannot close the loop fast enough to react to sentences such as “move to the left” followed by “okay, stop.” For real-time robotics applications, faster language processing methods need to be created. These could be through the advent of faster neural net-

Command Types	Domains	Command Examples	Implementation Examples
Temporal	Navigation	“Move back and forth between the table and the countertop twice.”	[Liang <i>et al.</i> , 2022]
		“Go to the store on Main Street, but only after visiting the bank”	[Liu <i>et al.</i> , 2023b]
	Manipulation	“Remain on the first floor and navigate to the red room.”	[Pan <i>et al.</i> , 2023]
		“Scan for blocks and insert any found into the bin.”	[Gopalan <i>et al.</i> , 2018]
	Mobile Manipulation	“Create a stack that contains two blocks.”	[Xie <i>et al.</i> , 2023]
		“Get Glass of Milk”	[Huang <i>et al.</i> , 2022a]
		“Always take the pear and go to the tree and stay there.”	[Wang <i>et al.</i> , 2021]
		“I spilled my coke. Can you bring me something to clean it up?”	[Ahn <i>et al.</i> , 2022]
		“Can you bring me the drink from the table?”	[Huang <i>et al.</i> , 2022b]
		“Bring the plate to the kitchen”	[Hsiung <i>et al.</i> , 2022]
“Move KeyChain1 to the box with more books”	[Xie <i>et al.</i> , 2023]		
Spatiotemporal	Navigation	“Go to Landmark A then Landmark B, in addition, avoid Landmark C;” “Move back and forth between Object A in Room B and Object C in Room D”	[Liu <i>et al.</i> , 2023b; Pan <i>et al.</i> , 2023]
	Manipulation	“Put the red block to the left of the rightmost bowl”	[Liang <i>et al.</i> , 2022]
		“Pack the ring into the brown box”	[Hsu <i>et al.</i> , 2023]
		“Sort fruits on the plate and bottles in the box”	[Singh <i>et al.</i> , 2023]
		“Move all the blocks to different corners.”	[Zeng <i>et al.</i> , 2023a]
		“May I have a cup of milk with taro?”	[Li <i>et al.</i> , 2023a]
		“Stack Objects X, Y, and Z” “Push Object X to Object Y”	[Black <i>et al.</i> , 2023; Brohan <i>et al.</i> , 2023a; Driess <i>et al.</i> , 2023; Black <i>et al.</i> , 2023; Yang <i>et al.</i> , 2023; Jiang <i>et al.</i> , 2023; Liang <i>et al.</i> , 2022; Du <i>et al.</i> , 2023; Brohan <i>et al.</i> , 2023b; O’Neill <i>et al.</i> , 2024]
	Mobile Manipulation	“Take the Coca-Cola can from the desk and put it in the middle of fruits on the table”	[Liang <i>et al.</i> , 2022]
		“Microwave salmon”	[Singh <i>et al.</i> , 2023]
		“Bring Object A to Location B;” “Open the door and proceed to the kitchen”	[Ahn <i>et al.</i> , 2022; Driess <i>et al.</i> , 2023; Huang <i>et al.</i> , 2022b; Hsiung <i>et al.</i> , 2022]
		“Build a house out of dirt blocks”	[Wang <i>et al.</i> , 2023]
		“Put away groceries in the pantry”	[Huang <i>et al.</i> , 2022a; Li <i>et al.</i> , 2022]

Table 2: Situated Natural Language Commands

work hardware or methods that learn hierarchical abstraction at different frame rates to support the cross-cutting ability of language to talk about any part of the robot system.

4.3 Datasets

Both approaches require datasets for learning and evaluation. Formal methods generally require parallel datasets that map natural language to structures in the formal representation or at least the ability to test if a formal representation is correct. These parallel corpora can be expensive at scale, but new LLM methods enable good performance even with much smaller datasets. A common approach is to show a trajectory of robot behavior generated from an underlying formal representation and then ask annotators to describe the trajectory in language [Gopalan *et al.*, 2018; Patel *et al.*, 2020]. This approach can enable untrained annotators to provide parallel data. Still, it often leads to ambiguous situations because the trajectory does not overtly show the constraints present in the underlying representation. Thus, it remains an open problem to collect diverse and unambiguous language commands and their corresponding labels.

To compensate for this problem, approaches, especially less formal approaches, try to learn from large unannotated datasets or from datasets obtained from simulation. LLMs can also enable robust learning to map from natural language to formal representation using few-shot learning or fine-tuning [Liu *et al.*, 2023b; Hsu *et al.*, 2023; Liu *et al.*, 2023a; Liang *et al.*, 2022; Ahn *et al.*, 2022]. Methods that use less constrained representations, such as VIMA [Jiang *et al.*, 2023], collect data in simulation across many scenarios and eschew a formal representation that requires annotation of any kind, instead directly train from joint trajectories, images, and video.

4.4 Generalization and Bias

Models that map language to a structured representation can exploit the regularities of that structure to train from smaller datasets and can generalize by porting the formal representation to different domains. Depending on the domain, this generalization can be significant by relying on state-of-the-art robot algorithms such as SLAM and motion planning rather than trying to learn everything end-to-end. This ability is possible because of the strict, modular compositionality introduced by the formal model.

On the other hand, end-to-end models have the potential to be applied in domains given large enough training datasets and computational resources. Looking at the power of LLMs, one predicts that a multimodal dataset of video, other robot sensor data, and joint states may be able to generalize as flexibly and powerfully as an LLM does with language. However, training this model requires significantly higher dimensions than language data and potentially much larger dataset sizes. However, if provided with general-purpose robot data, foundation models have the potential to understand language in a very general way, analogous to the success of LLMs. Key research challenges to address this problem include acquiring semantically diverse datasets covering desired robot behaviors, developing sample-efficient model architectures and training algorithms for multimodal datasets many times larger than those used to train LLMs, and translating learned models to robot action.

Many existing works that use foundation models also leverage symbols but do not speak to the symbols’ role in building practical systems. We observe that in many recent RoboNLP papers including [Ahn *et al.*, 2022; Jiang *et al.*, 2023; Brohan *et al.*, 2023a; Brohan *et al.*, 2023b; O’Neill *et al.*, 2024] all leverage symbolic representations, either in the form of discretized action spaces or predefined skills. We would like to see more authors acknowledge that their work

is made better because of the symbolic representation their model uses or propose how other symbolic representations can improve the model.

In the future, we expect models to better represent the world through these finer-grained symbolic representations. After a certain point, combining text, audio, color images, depth images, point clouds, etc., becomes functionally continuous. While the POMDP and PDDL methods are theoretically sound, in practice, they lack representational capacity outside of their domain due to the brittleness of out-of-domain representation. We believe that future successful methods will leverage deep learning methods with smarter representations of the input and output data via transformations of the intermediate representations. This includes representing color images as point clouds, maps, or meshes as derivative intermediates. We can then augment these with state representations of the environment as text or other labels. There is an eventuality where the bitter lesson [Sutton, 2019] takes over, and throwing even more parameters and data at the problem will win out, but that is predicated on an evergrowing data supply. It remains to be seen where additional data will come from and if methods that leverage only partial observations from one modality can synergize multiple modalities simultaneously.

4.5 Limitations of Natural Language

As robot operators, we want to be able to express the state of our robotic system in a way that is auditable by us and useful to the robot in accomplishing tasks. Many of the methods we discussed also focus on English as the representation of natural language encoded in LLMs and multimodal models. While English works well for much of the world, it does not precisely describe objects’ physical locality. Saying that an apple is to the left of an orange can put that apple in various positions with just a language description alone. While adding geometric priors could enhance the specificity of the language, we should, as a field, look to building benchmarks to evaluate the physical plausibility of these language models as they operate in different languages [Whorf, 1956; Lucy, 1992]. Robot morphology and human morphology are very different. It is likely that expressing tasks as being completed by one’s arm, such as “clean the dishes with your left hand” may be completely moot to a robot that uses soft-body mechanics to interact with the world. Human language evolved in a biological domain that necessitates a more survival vocabulary. These actions are not necessarily helpful for describing robot actions and could contribute to skewing the success rates of robots in accomplishing tasks.

4.6 Safety and Interpretability

Interpretable and explainable robotic systems provide transparency in decision-making and gain trust in human-robot interaction [Gunning and Aha, 2019; Anjomshoae *et al.*, 2019; Silva *et al.*, 2023]. Moreover, verifiable safe operation is essential for deployments that satisfy worldwide standards such as ISO 61508 [International Organization for Standardization, 2010], which defines standards for safely deploying robots in industrial factory environments worldwide.

These standards require that robotic systems be mathematically proven to have a failure rate lower than 10^{-5} dangerous failures per hour. Robots should also provide feedback to unsatisfiable task specifications [Raman *et al.*, 2013] or explain their actions when execution fails [Das *et al.*, 2021]. Safety can be classified into semantic and kinematic/physical safety [Varley *et al.*, 2024]. Examples of semantic safety are never entering the nursery and always transporting a cup full of coffee in an upright orientation. Examples of kinematic safety are avoiding collisions with humans and objects and avoiding reaching joint/velocity/torque limits. Formal approaches using LLMs mostly focus on enforcing semantic safety. Many works on formal methods and safe control do not use LLMs and encode kinematic safety in trajectory optimization [Dawson *et al.*, 2023]. For example, linear temporal logic (LTL) has been used extensively to develop provably safe controllers [Lignos *et al.*, 2015; Chinchali *et al.*, 2012]. Given a logical specification and the system model, formal methods can synthesize correct-by-construction robot controllers or provide counterexamples to explain when the task is unsatisfiable with respect to how well the model captures reality [Kress-Gazit *et al.*, 2018]. A key challenge for any safety framework of this kind is grounding predicates in real-world noisy and partial perceptual data.

Existing deep learning approaches do not have a clear pathway for achieving this level of interpretability and safety. If we look at human behavior, defining a set of safety guidelines to assess adherence has been challenging. While we can continue to apply formal representations to these models to varying levels of success, there remains a need for a system that, at a meta-level, operates on symbols but whose abilities are as general and flexible as humanity’s. For example, one of our recent papers uses a combined approach with LLMs and a formal representation, showing that we can exploit the generalization ability of LLMs with the formal safety guarantees provided by LTL to create a safe yet robust and flexible system for following commands [Yang *et al.*, 2024]. Yet this approach only scratches the surface. Much more needs to be done to integrate these systems, especially to study the interpretability and safety guarantees inherent in perceptual systems and provide formal guarantees and bounds about the behavior of deep networks.

5 Conclusion

Our review characterizes the literature in robotic language grounding along a spectrum from using more formal, biased, discrete representations to less formal, less biased, higher-dimensional continuous representations. There are benefits and tradeoffs to each approach. More formal methods induce structure that can limit the size of the learning problem and provide interpretability and formal safety guarantees. However, they also constrain the output space, limiting the flexibility and expressive power of what can be learned. Less formal methods impose fewer constraints but require more data and possibly more structured neural networks to be learned. Methods such as SayCan [Ahn *et al.*, 2022], traditionally considered less structured, use a formal representation of the robot’s skills. Seen in this context, it is clear that a limitation

for all methods is the lack of physical capability of existing robots: a key area of future work is to enable robots to perform a larger variety of tasks in a larger variety of environments. Hence, they can physically perform the tasks people ask them to do.

Acknowledgments

The work done by Vanya Cohen and Raymond Mooney is supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001122C0007. The work done by Jason Xinyu Liu and Stefanie Tellex is supported by the Office of Naval Research (ONR) under grant number N00014-22-1-2592 and with support from Amazon Robotics. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies. The authors thank Ankit Shah and Jessica Hodgins for their insightful feedback.

Contribution Statement

All authors contributed equally. Listed in alphabetical order.

References

- [Ahn *et al.*, 2022] Michael Ahn, Anthony Brohan, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, 2022.
- [Anjomshoae *et al.*, 2019] Sule Anjomshoae, Amro Najjar, Davide Calvaresi, and Kary Främling. Explainable agents and robots: Results from a systematic literature review. In *18th International Conference on Autonomous Agents and Multiagent Systems*, pages 1078–1088. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- [Baier and Katoen, 2008] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT press, 2008.
- [Baker *et al.*, 2022] Bowen Baker, Ilge Akkaya, Peter Zhokhov, Joost Huizinga, Jie Tang, Adrien Ecoffet, Brandon Houghton, Raul Sampedro, and Jeff Clune. Video pretraining (VPT): Learning to act by watching unlabeled online videos, 2022.
- [Black *et al.*, 2023] Kevin Black, Mitsuhiro Nakamoto, Pranav Atreya, Homer Walke, Chelsea Finn, Aviral Kumar, and Sergey Levine. Zero-shot robotic manipulation with pretrained image-editing diffusion models, 2023.
- [Brohan *et al.*, 2023a] Anthony Brohan, Noah Brown, et al. RT-1: Robotics transformer for real-world control at scale. *Robotics: Science and Systems*, 2023.
- [Brohan *et al.*, 2023b] Anthony Brohan, Noah Brown, et al. RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, 2020.
- [Chen *et al.*, 2023] Yongchao Chen, Rujul Gandhi, Yang Zhang, and Chuchu Fan. NL2TL: Transforming natural languages to temporal logics using large language models. *arXiv:2305.07766*, 2023.
- [Chen *et al.*, 2024] Yongchao Chen, Jacob Arkin, Yang Zhang, Nicholas Roy, and Chuchu Fan. AutoTAMP: Autoregressive task and motion planning with llms as translators and checkers. In *IEEE International Conference on Robotics and Automation*, 2024.
- [Chinchali *et al.*, 2012] Sandeep Chinchali, Scott C Livingston, Ufuk Topcu, Joel W Burdick, and Richard M Murray. Towards formal synthesis of reactive controllers for dexterous robotic manipulation. In *IEEE International Conference on Robotics and Automation*, 2012.
- [Collins *et al.*, 2022] Katherine M. Collins, Catherine Wong, Jiahui Feng, Megan Wei, and Joshua B. Tenenbaum. Structured, flexible, and robust: benchmarking and improving large language models towards more human-like behavior in out-of-distribution reasoning tasks. In *Proceedings of the 44th Annual Conference of the Cognitive Science Society*, 2022.
- [Das *et al.*, 2021] Devleena Das, Siddhartha Banerjee, and Sonia Chernova. Explainable AI for robot failures: Generating explanations that improve user assistance in fault recovery. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pages 351–360, 2021.
- [Dawson *et al.*, 2023] Charles Dawson, Sicun Gao, and Chuchu Fan. Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods for robotics and control. *IEEE Transactions on Robotics*, 39(3):1749–1767, 2023.
- [Driess *et al.*, 2023] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. PaLM-E: An embodied multimodal language model. *International Conference on Machine Learning*, 2023.
- [Du *et al.*, 2023] Yilun Du, Mengjiao Yang, Pete Florence, Fei Xia, Ayzan Wahid, Brian Ichter, Pierre Sermanet, Tianhe Yu, Pieter Abbeel, Joshua B Tenenbaum, et al. Video language planning. *arXiv:2310.10625*, 2023.
- [Edelkamp and Hoffmann, 2004] Stefan Edelkamp and Jörg Hoffmann. PDDL2. 2: The language for the classical part of the 4th international planning competition. Technical report, Technical Report 195, University of Freiburg, 2004.
- [Fox and Long, 2003] Maria Fox and Derek Long. PDDL1: An extension to pddl for expressing temporal planning domains. *Journal of artificial intelligence research*, 20:61–124, 2003.
- [Fuggitti and Chakraborti, 2023] Francesco Fuggitti and Tathagata Chakraborti. NL2LTL – a python package for converting natural language (NL) instructions to linear temporal logic (LTL) formulas. *AAAI Conference on Artificial Intelligence*, 2023. System Demonstration.
- [Gopalan *et al.*, 2018] Nakul Gopalan, Dilip Arumugam, Lawson LS Wong, and Stefanie Tellex. Sequence-to-sequence language grounding of non-markovian task specifications. In *Robotics: Science and Systems*, volume 2018, 2018.
- [Gopalan *et al.*, 2020] Nakul Gopalan, Eric Rosen, George Konidaris, and Stefanie Tellex. Simultaneously learning transferable symbols and language groundings from perceptual data for instruction following. In *Robotics: Science and Systems*, 2020.
- [Guan *et al.*, 2023] Lin Guan, Karthik Valmeekam, Sarath Sreedharan, and Subbarao Kambhampati. Leveraging pre-trained large language models to construct and utilize world models for model-based task planning, 2023.
- [Gunning and Aha, 2019] David Gunning and David Aha. Darpa’s explainable artificial intelligence (XAI) program. *AI magazine*, 40(2):44–58, 2019.

- [Harnad, 1990] Stevan Harnad. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1):335–346, 1990.
- [Harnad, 2007] Stevan Harnad. Symbol grounding problem. *Scholarpedia*, 2(7):2373, 2007. revision #73220.
- [Helmert, 2006] Malte Helmert. The fast downward planning system. *JAIR*, 26:191–246, 2006.
- [Henighan *et al.*, 2020] Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B. Brown, Prafulla Dhariwal, Scott Gray, Chris Hallacy, Benjamin Mann, Alec Radford, Aditya Ramesh, Nick Ryder, Daniel M. Ziegler, John Schulman, Dario Amodei, and Sam McCandlish. Scaling laws for autoregressive generative modeling. *ArXiv*, abs/2010.14701, 2020.
- [Hoffmann *et al.*, 2022] Jordan Hoffmann, Sebastian Borgeaud, et al. An empirical analysis of compute-optimal large language model training. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [Hsiung *et al.*, 2022] Eric Hsiung, Hiloni Mehta, Junchi Chu, Xinyu Liu, Roma Patel, Stefanie Tellex, and George Konidaris. Generalizing to new domains by mapping natural language to lifted ltl. In *IEEE International Conference on Robotics and Automation*, 2022.
- [Hsu *et al.*, 2023] Joy Hsu, Jiayuan Mao, Joshua B. Tenenbaum, and Jiajun Wu. What’s left? concept grounding with logic-enhanced foundation models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Hu *et al.*, 2023] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. GAIA-1: A generative world model for autonomous driving, 2023.
- [Huang *et al.*, 2022a] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pages 9118–9147. PMLR, 2022.
- [Huang *et al.*, 2022b] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. *arXiv:2207.05608*, 2022.
- [International Organization for Standardization, 2010] International Organization for Standardization. International Electrotechnical Commission International Organization for Standardization. Functional safety of electrical/electronic/programmable electronic safety-related systems - part 1: General requirements, 2010.
- [Jiang *et al.*, 2023] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. VIMA: General robot manipulation with multimodal prompts. In *Fortieth International Conference on Machine Learning*, 2023.
- [Konidaris *et al.*, 2018] George Konidaris, Leslie Pack Kaelbling, and Tomas Lozano-Perez. From skills to symbols: Learning symbolic representations for abstract high-level planning. *Journal of Artificial Intelligence Research*, 61:215–289, 2018.
- [Konidaris, 2020] GD Konidaris. Simultaneously learning transferable symbols and language groundings from perceptual data for instruction following. *Robotics: Science and Systems*, 2020.
- [Kovacs, 2011] Daniel L Kovacs. BNF definition of PDDL 3.1. *Unpublished manuscript from the IPC-2011 website*, 15, 2011.
- [Kress-Gazit *et al.*, 2018] Hadas Kress-Gazit, Morteza Lahijanian, and Vasumathi Raman. Synthesis for robots: Guarantees and feedback for robot behavior. *Annual Review of Control, Robotics, and Autonomous Systems*, 1:211–236, 2018.
- [Li *et al.*, 2022] 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:2202.01771*, 2022.
- [Li *et al.*, 2023a] Boyi Li, Philipp Wu, Pieter Abbeel, and Jitendra Malik. Interactive task planning with language models. In *2nd Workshop on Language and Robot Learning: Language as Grounding*, 2023.
- [Li *et al.*, 2023b] Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, Hang Li, and Tao Kong. Vision-language foundation models as effective robot imitators. *arXiv:2311.01378*, 2023.
- [Liang *et al.*, 2022] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *arXiv:2209.07753*, 2022.
- [Lignos *et al.*, 2015] Constantine Lignos, Vasumathi Raman, Cameron Finucane, Mitchell Marcus, and Hadas Kress-Gazit. Provably correct reactive control from natural language. *Autonomous Robots*, 38:89–105, 2015.
- [Liu *et al.*, 2023a] Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. LLM+P: Empowering large language models with optimal planning proficiency. *arXiv:2304.11477*, 2023.
- [Liu *et al.*, 2023b] Jason Xinyu Liu, Ziyi Yang, Ifrah Idrees, Sam Liang, Benjamin Schornstein, Stefanie Tellex, and Ankit Shah. Grounding complex natural language commands for temporal tasks in unseen environments. In *Conference on Robot Learning*, 2023.
- [Lucy, 1992] John A Lucy. *Language diversity and thought: A reformulation of the linguistic relativity hypothesis*. Cambridge University Press, 1992.
- [Maler and Nickovic, 2004] Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In *International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*, pages 152–166, 2004.
- [McDermott *et al.*, 1998] Drew McDermott, Malik Ghallab, Adele E. Howe, Craig A. Knoblock, Ashwin Ram, Manuela M. Veloso, Daniel S. Weld, and David E. Wilkins. PDDL-the planning domain definition language. In *Technical Report, Tech. Rep.*, 1998.
- [Ng and Jordan, 2001] Andrew Ng and Michael Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. *Advances in neural information processing systems*, 14, 2001.
- [Octo Model Team *et al.*, 2023] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. <https://octo-models.github.io>, 2023.
- [O’Neill *et al.*, 2024] Abby O’Neill, Abdul Rehman, et al. Open X-Embodiment: Robotic learning datasets and RT-X models. In *IEEE International Conference on Robotics and Automation*, 2024.

- [Pan *et al.*, 2023] Jiayi Pan, Glen Chou, and Dmitry Berenson. Data-efficient learning of natural language to linear temporal logic translators for robot task specification. In *IEEE International Conference on Robotics and Automation*, 2023.
- [Patel *et al.*, 2020] Roma Patel, Ellie Pavlick, and Stefanie Tellex. Grounding language to non-markovian tasks with no supervision of task specifications. In *Robotics: Science and Systems*, 2020.
- [Raman *et al.*, 2013] Vasumathi Raman, Constantine Lignos, Cameron Finucane, Kenton CT Lee, Mitchell P Marcus, and Hadas Kress-Gazit. Sorry dave, i'm afraid i can't do that: Explaining unachievable robot tasks using natural language. *Robotics: Science and Systems*, 2(1):2–1, 2013.
- [Raman *et al.*, 2024] Shreyas Sundara Raman, Vanya Cohen, Eric Rosen, Ifrah Idrees, David Paulius, and Stefanie Tellex. Planning with large language models via corrective re-prompting. In *IEEE International Conference on Robotics and Automation*, 2024.
- [Reed *et al.*, 2022] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gómez Colmenarejo, Alexander Novikov, Gabriel Barthmaron, Mai Giménez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent. *Transactions on Machine Learning Research*, 2022. Featured Certification, Outstanding Certification.
- [Reid *et al.*, 2024] Machel Reid, Nikolay Savinov, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv:2403.05530*, 2024.
- [Rodriguez-Sanchez *et al.*, 2023] Rafael Rodriguez-Sanchez, Benjamin Adin Spiegel, Jennifer Wang, Roma Patel, Stefanie Tellex, and George Konidaris. RLang: a declarative language for describing partial world knowledge to reinforcement learning agents. In *International Conference on Machine Learning*, pages 29161–29178. PMLR, 2023.
- [Shridhar *et al.*, 2023] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, 2023.
- [Silva *et al.*, 2023] Andrew Silva, Mariah Schrum, Erin Hedlund-Botti, Nakul Gopalan, and Matthew Gombolay. Explainable artificial intelligence: Evaluating the objective and subjective impacts of xai on human-agent interaction. *International Journal of Human-Computer Interaction*, 39(7):1390–1404, 2023.
- [Silver *et al.*, 2024] Tom Silver, Soham Dan, Kavitha Srinivas, Josh Tenenbaum, Leslie Kaelbling, and Michael Katz. Generalized planning in PDDL domains with pretrained large language models. In *AAAI Conference on Artificial Intelligence*, 2024.
- [Singh *et al.*, 2023] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. ProgPrompt: Generating situated robot task plans using large language models. In *IEEE International Conference on Robotics and Automation*, 2023.
- [Sutton, 2019] Richard Sutton. The bitter lesson. *Incomplete Ideas (blog)*, 13(1), 2019.
- [Tellex *et al.*, 2020] Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. Robots that use language. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:25–55, 2020.
- [Valmeekam *et al.*, 2023a] Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. PlanBench: An extensible benchmark for evaluating large language models on planning and reasoning about change, 2023.
- [Valmeekam *et al.*, 2023b] Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. On the planning abilities of large language models: A critical investigation, 2023.
- [Varley *et al.*, 2024] Jake Varley, Sumeet Singh, Deepali Jain, Krzysztof Choromanski, Andy Zeng, Somnath Basu Roy Chowdhury, Avinava Dubey, and Vikas Sindhwani. Embodied ai with two arms: Zero-shot learning, safety and modularity. *arXiv:2404.03570*, 2024.
- [Wahlström *et al.*, 2015] Niklas Wahlström, Thomas B Schön, and Marc Peter Deisenroth. From pixels to torques: Policy learning with deep dynamical models. *arXiv:1502.02251*, 2015.
- [Wang *et al.*, 2021] Christopher Wang, Candace Ross, Yen-Ling Kuo, Boris Katz, and Andrei Barbu. Learning a natural-language to ltl executable semantic parser for grounded robotics. In *Conference on Robot Learning*, 2021.
- [Wang *et al.*, 2023] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv:2305.16291*, 2023.
- [Wang *et al.*, 2024] Jiaqi Wang, Zihao Wu, Yiwei Li, Hanqi Jiang, Peng Shu, Enze Shi, Huawei Hu, Chong Ma, Yiheng Liu, Xuhui Wang, et al. Large language models for robotics: Opportunities, challenges, and perspectives. *arXiv:2401.04334*, 2024.
- [Wei *et al.*, 2022] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large language models, 2022.
- [Whorf, 1956] Benjamin Lee Whorf. *Language, thought, and reality: Selected writings of Benjamin Lee Whorf*. MIT press, 1956.
- [Xie *et al.*, 2023] Yaqi Xie, Chen Yu, Tongyao Zhu, Jinbin Bai, Ze Gong, and Harold Soh. Translating natural language to planning goals with large-language models. *arXiv:2302.05128*, 2023.
- [Yan *et al.*, 2023] Wilson Yan, Danijar Hafner, Stephen James, and Pieter Abbeel. Temporally consistent transformers for video generation, 2023.
- [Yang *et al.*, 2023] Mengjiao Yang, Yilun Du, et al. Learning interactive real-world simulators. *arXiv:2310.06114*, 2023.
- [Yang *et al.*, 2024] Ziyi Yang, Shreyas S Raman, Ankit Shah, and Stefanie Tellex. Plug in the safety chip: Enforcing constraints for LLM-driven robot agents. In *IEEE International Conference on Robotics and Automation*, 2024.
- [Zeng *et al.*, 2023a] Andy Zeng, Maria Attarian, brian ichter, Krzysztof Marcin Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aavek Purohit, Michael S Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, and Pete Florence. Socratic models: Composing zero-shot multimodal reasoning with language. In *The Eleventh International Conference on Learning Representations*, 2023.
- [Zeng *et al.*, 2023b] Fanlong Zeng, Wensheng Gan, Yongheng Wang, Ning Liu, and Philip S Yu. Large language models for robotics: A survey. *arXiv:2311.07226*, 2023.
- [Zhang *et al.*, 2023] Ceng Zhang, Junxin Chen, Jiatong Li, Yanhong Peng, and Zebing Mao. Large language models for human-robot interaction: A review. *Biomimetic Intelligence and Robotics*, page 100131, 2023.
- [Zhao *et al.*, 2023] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv:2304.13705*, 2023.