

Ground Manipulator Primitive Tasks to Executable Actions Using Large Language Models

Yue Cao and C.S. George Lee

Elmore Family School of Electrical and Computer Engineering, Purdue University
{yuecao, csgelee}@purdue.edu

Abstract

Layered architectures have been widely used in robot systems. The majority of them implement planning and execution functions in separate layers. However, there still lacks a straightforward way to transit high-level tasks in the planning layer to the low-level motor commands in the execution layer. In order to tackle this challenge, we propose a novel approach to ground the manipulator primitive tasks to robot low-level actions using large language models (LLMs). We designed a program-function-like prompt based on the task frame formalism. In this way, we enable LLMs to generate position/force set-points for hybrid control. Evaluations over several state-of-the-art LLMs are provided.

Introduction

In robotics, the layered architectures break down the system design into multiple layers, then separately configure each layer. The development of layered architectures can be traced back to the Stanford Research Institute’s *Shakey* mobile robot (Nilsson 1969). Its system was programmed by high-level functions solving planning problems in first-order logic and low-level functions specifying motor commands. By the early 1990s, a number of layered architectures have been developed and summarized as the sense-model-plan-act (SMPA) paradigm in (Brooks 1991a). Nowadays, the layered architectures typically consisting of a planning layer and an execution player have been widely used in many robot systems (Kortenkamp, Simmons, and Brugali 2016).

However, there is a long-lasting challenge in the layered system – the transition from planning to execution. Back in 1987, Rodney Brooks had pointed out the drawbacks of separating planning and execution layers (Brooks 1987). To deal with this issue, Brooks proposed a biological-system-inspired subsumption architecture (Brooks 1991b). But after years, the subsumption architecture didn’t achieve much progress in modern robot control. Most robot systems stay

with the layered architectures and one common solution for the planning-to-execution transition problem is using primitive tasks. As early as in the *Shakey* robot, it introduced a library of intermediate-level actions such as “push object” or “go through door” to communicate between the high-level logic and low-level motor commands (Nilsson 1984). A classic 3-layer architecture in (Firby 1989) was separated into high-level planning, middle-level execution, and low-level hardware control layers. Its middle-level layer implements a reactive-action-package (RAP) system (Firby 1992) consisting of primitive tasks to bridge between the symbolic plans and low-level control actions. Nevertheless, the existing primitive-task-based approaches require laborious manual specification. They usually yield in some libraries of a small number of primitive tasks, thereby limiting the task adaptability and generalizability of robots.

To seek a new method for the planning-to-execution transition problem, we turn our attention to the state of the art – large language models (LLMs). The large language models are neural-network-based language models with a huge number of parameters and trained over massive amount of data. They significantly outperform other natural language processing (NLP) models and exhibit remarkable capabilities in generating textual responses to a broad range of questions. Since the release of OpenAI ChatGPT, the emergent abilities of LLMs have drawn the interests of researchers from a variety of disciplines, including robotics. In robotics, pilot studies have primarily focused on the planning layer (Huang et al. 2022; Ahn et al. 2022; Cao and Lee 2023; Xiang et al. 2023; Lin et al. 2023). They utilized the large language models to generate high-level task plans in a zero shot or few shot manner. However, the generated abstract plans are described in natural languages, making them uninterpretable by the execution layer.

In this paper, we focus on the manipulator tasks and study the ability of LLMs to solve the planning-to-execution transition problem. Specifically, given a manipulator primitive task described in the natural language domain, we seek a LLM-based solution to translate it to low-level motor commands as shown in Figure 1. We propose a novel approach to ground textual manipulator primitive tasks to executable actions using LLMs. Specifically, we design a program-like prompt based on the task frame formalism (TFF), an object-centric specification that facilitates task transfer across dif-

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[†]This work was supported in part by the National Science Foundation under Grant IIS-1813935. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.




Figure 1: (a) Previous work in LLM & robotics mainly studied task generation in the planning layer. (b) Our proposed approach targets the planning-to-execution transition problem.

ferent manipulators. The prompt takes the text of a primitive task as input and outputs a set of position/force set-points in the task frame. These position/force set-points will allow manipulators to compute their low-level motor commands.

Major contributions of this paper are:

- (1) We propose a LLM-based approach that enables language-described manipulator primitive tasks to be converted to the set-points for robot position/force hybrid control. It offers a new solution for the planning-to-execution transition problem in layered architectures of manipulators.
- (2) The remarkable generalizability of LLMs makes it possible to apply many manipulator primitive tasks described in natural language within our approach. Hence, it has the potential to reduce the burden on end-users when it comes to designing a complex library of primitive tasks.

Related Work

Language Grounding to Robot Actions

Language grounding was introduced to robotics in (Roy 2005) to establish connections between natural languages and robot capabilities, including perception and action. Language grounding to action is generally referred to as associating high-level language commands with the low-level robot control system (Chai et al. 2016). In the area of language grounding to action, the study of grounding schema follows one step behind the development of NLP techniques.

The early work in (Kress-Gazit, Fainekos, and Pappas 2008) enforced strict grammar rules to transform task specifications to linear temporal logic form. Later on, grammar parsers (Tellex et al. 2011; Matuszek et al. 2012) such as the Stanford parser (de Marneffe, MacCartney, and Manning 2006) were applied to decompose sentence instructions into different grammatical elements, mainly involving verbs and objects. Subsequently, word-based statistical machine translation models were investigated in (Squire et al. 2015) to enable generalization to new environments. In the beginning of the deep-learning era, word embeddings and recurrent neural networks (RNNs) became the main steam and were later introduced in the robot grounding schema (Arumugam et al. 2019; Toyoda et al. 2021).

Since the 2020s, large language models have achieved tremendous progress and become the primary focus of NLP. Using LLMs, one can bypass analyzing grammatical structure or word semantics, as was used in previous work.

Task Frame Formalism

The concept of the task frame formalism originated from the study of compliant motions (Mason 1981). It was then explicitly formulated as an intermediate transfer between task planning and force/position control (Bruyninckx and De Schutter 1996). The task frame, also known as the compliance frame, is a local coordinate frame attached to the object being manipulated. Its translational and rotational directions are configured to be either force controlled or position controlled. Such formalism allows end-users to implement hybrid position/force control strategy (Raibert and Craig 1981) alongside separate frame directions.

Based on the task frame formalism, a variety of manipulator primitive tasks have been explored (Ballard and Hartman 1986; Morrow and Khosla 1997; Kroger, Finkemeyer, and Wahl 2004; Vuong, Pham, and Pham 2021). Nevertheless, they mainly focused on the taxonomy of low-level action representation and provided little consideration to the association with high-level task nomenclature. For example, manipulator primitive tasks were specified as “rotate about y 5 deg” or “translate along x until next contact” in (Vuong, Pham, and Pham 2021). But there is no straightforward way to translate natural-language-described tasks like “insert key” or “open bottle” to such low-level movements.

Proposed Approach

Preprocess: Identify Manipulator Primitive Tasks

To start with, we need to ensure that the given primitive task conforms to the task frame formalism. For instance, we generally consider tasks like “insert peg” or “open bottle” as primitive tasks because they can be accomplished by a single control strategy without changing the coordinate setting. Tasks such as “assembly GPU” or “make coffee” are examples of non-primitive tasks because they must be further decomposed into multiple steps that use different control strategies and coordinates.

We use LLMs to assess the primitiveness of tasks. In fact, current LLMs including OpenAI ChatGPT, Google Bard, and Meta LLaMA-2, have some knowledge of the task frame formalism. For example, they can respond to the following prompt with reasonable answers.

Do you understand task frame formalism in robotics?

But when it comes to identifying the manipulator primitive tasks, their criteria become inconsistent. Therefore, we write a prompt including detailed explanation to query the LLM, as shown below.

We specify primitive tasks based on the concept of task frame formalism in robotics. For example, “open door”, “slice bread”, “insert card” are primitive tasks because they use a single control strategy and coordinate setting. While “assembly com-

puter”, “make coffee”, “drive car” are not primitive tasks because they must be decomposed into several low-level tasks in single control strategies and coordinate settings.
So, is “open bottle” a primitive task? Just answer yes or no.

With such prompt, the LLMs can effectively determine whether a given manipulator task is primitive or not. If a certain task is not primitive, we need to perform task decomposition. Since the task decomposition is out of the scope of this work, we refer to other literature (Huang et al. 2022; Ahn et al. 2022; Cao and Lee 2023).

TFF-based Prompt Design

After obtaining a manipulator primitive task, we use LLMs to transit it to low-level execution in a TFF fashion. Since LLMs tend to produce lengthy answers and detailed explanations to casual questions, we need to design a specific prompt to regulate their output into a structured form. Inspired by the work (Singh et al. 2022) that uses code blocks for robot-plan generation, we design our prompt in a function format in computer programming. The details of our prompt design are listed below, with a Python example in Figure 2.

```
# Source function 1
def turn_screw(translational_x, translational_y, translational_z,
              angular_x, angular_y, angular_z):
    # Coordinate setting: Z axis as the direction of screw
    translational_x=0
    translational_y=0
    translational_z=-5 # N

    angular_x=0
    angular_y=0
    angular_z=5 # rad/sec
    return translational_x, translational_y, translational_z, \
           angular_x, angular_y, angular_z

# Source function 2
def wipe_table(...)
    ...
    return

# Source function 3
def open_door_from_doorknob(...)
    ...
    return

# Target function
def turn_steering_wheel(translational_x, translational_y, translational_z,
                       angular_x, angular_y, angular_z):
    # Coordinate setting: Z axis as the direction of the steering wheel
    translational_x=0
    translational_y=0
    translational_z=0

    angular_x=0
    angular_y=0
    angular_z=5 # rad/sec
    return translational_x, translational_y, translational_z, \
           angular_x, angular_y, angular_z
```

Figure 2: An exemplary 3-shot generation in Python syntax. The 3-shot prompt is shown in the grey background. Details in the source function 2 and 3 are skipped. The text generated by LLM is displayed in the green background. Note that the prompt can also be written in other common programming languages. More source functions are preferred.

- Outside of the function, we use program comments to indicate whether a function is set for source specification or target generation. We recommend using multiple source functions with different TFF configurations to provide better guidance for the LLMs.
- The function name is configured as the text of the primitive task. For example, the task “turn screw” is written as “def turn_screw()”. Stop words such as “a” and “the” can be removed from the task name. In addition, the function parameters are specified as six-directional motions in the task frame formalism.
- In the function, we first set the coordinates in the function comment. Then, we designate position/force set-points in a TFF manner. If one direction is inactive, we simply assign 0 as its value. For the active directions, we use different units to distinguish whether they are position- or force-controlled. For the example in Figure 2, we activate two out of the six directions using “translational_z=-5 # N” and “angular_z=5 # rad/sec”. This specification implies that we rotate the screw about the z axis at a velocity of 5 rad/sec while simultaneously applying a downward force of 5 newtons.
- The target function sets the desired manipulator primitive task that we want to generate for. In this function, the target task name is the only element that needs to be filled.

Once we input the whole prompt to a LLM, it will automatically generate the coordinate setting and specification of six-directional motions for the given target task. Using the generated TFF-based coordinates and set-points, robots can compute control commands based on their own manipulator configurations.

Evaluation

In this section, we evaluate our TFF-based prompt with 4 state-of-the-art large language models – OpenAI GPT-3.5-turbo, OpenAI GPT-4, Google Bard, and Meta LLaMA-2-70B.

Evaluation Setting

For the GPT-3.5-turbo and GPT-4 models, we applied the same setting. The temperature hyperparameter was set to be 0 to minimize the randomness. The top P, frequency penalty, and presence penalty hyperparameters were configured to 1, 0, and 0, respectively.

For the Google Bard model, there is no hyperparameter configuration available towards end-users. Thereby we cannot guarantee the deterministicness of the generated outcomes. We entered the prompt in its web GUI and only recorded the initial results generated.

The Meta LLaMA-2 model is the most recently released LLM among them. We chose its 70B (70 billion) version, the one with the most parameters. We also set its temperature hyperparameter to the minimum value 0.01. The repetition penalty was set to 1 to grant repeated words with no penalty. Another hyperparameter top P was set to 1.

We assessed the proposed prompt in a set of 30 manipulator primitive tasks as listed in Table 1. We conducted

1. cut pizza	11. rasp wood	21. open door from hinge
2. scrub desk with bench brush	12. scrape substance from surface	22. slide block over vertical surface
3. spear cake with fork	13. peel potato	23. turn steering wheel
4. fasten screw with screwdriver	14. slice cucumber	24. shake cocktail bottle
5. loosen screw with screwdriver	15. flip bread	25. cut banana
6. unlock lock with key	16. shave object	26. crack egg
7. fasten nut with wrench	17. use roller to roll out dough	27. press button
8. loosen nut with wrench	18. insert peg into pegboard	28. insert GPU into socket
9. spread paint with brush	19. brush across tray	29. open bottle
10. hammer in nail	20. insert straw through cup lid	30. open childproof bottle

Table 1: Manipulator primitive tasks used in our evaluation. The tasks in the first two columns were selected from the Daily Interactive Manipulation dataset (Paulius, Eales, and Sun 2020) with minor modification, while the tasks listed in the third column were created by us.

tests in 0-shot, 1-shot, 3-shot and 5-shot manners. The 5-shot prompt was consisted of 5 source tasks as outlined in the block below.

```
# Source function 1
def turn_screw(...)
...
# Source function 2
def wipe_table(...)
...
# Source function 3
def open_door_from_doorknob(...)
...
# Source function 4
def cut_sandwich(...)
...
# Source function 5
def slide_box_upward_on_wall(...)
...
```

For simplicity, we only show the task indicators and function names in this block while omitting other details. The 3-shot prompt utilized the first three source tasks from the 5-shot prompt, whereas the 1-shot prompt only used the first one from it.

Results

The evaluations for GPT and Bard were conducted in early July, 2023. The test for LLAMA-2 was carried out in late July, 2023, shortly after its release. In order to validate the generated TFF specifications, We manually set up metrics for all 30 manipulator primitive tasks. The details of our metrics can be found in Appendix. The overall correct rates are reported in Table 2.

	0-shot	1-shot	3-shot	5-shot
GPT-3.5-turbo	0	0.30	0.70	0.67
GPT-4	0	0.47	0.63	0.83
Bard	0	0	0.47	0.70
LLaMA-2-70B	0	0	0.07	0.13

Table 2: Correct rates in the evaluation over 30 manipulator primitive tasks.

From the results, we notice that none of LLMs were able to produce a valid response in the 0-shot test. When we increased the number of source tasks to 5, GPT-3.5-turbo, GPT-4, and Bard generated plenty of correct answers. Notably, the GPT-4 in the 5-shot setting achieved the highest correct rate of 0.83. Meanwhile, the LLaMA-2-70B failed the majority of tests. For the 5-shot tests, the correctness of each manipulator primitive task is shown in Table 3.

	Task No.	Correctness
GPT-3.5-turbo	1-10	●●●●●○●●●●
	11-20	○○○●●○○●●●
	21-30	●○○●○○●●○○
GPT-4	1-10	●●●●●●●●●●
	11-20	●●○○●●●●○○
	21-30	●●●●○○●●○○
Bard	1-10	●○●●●●○○●●
	11-20	○○●●●●○○●●
	21-30	●●●●○○●●●●
LLaMA-2-70B	1-10	●○○○●○○○○○
	11-20	○○○○●○○○○○
	21-30	○○○○●○○○○○

Table 3: The correctness of each manipulator primitive task in the 5-shot tests. Solid circles indicate the correct generation for tasks, while hollow circles represent incorrect ones. Task No. corresponds to the numbering in Table 1, with every 10 tasks sequentially lined-up in a row.

For GPT-3.5-turbo, GPT-4, and Bard models, they strictly followed the desired “# *Coordinate setting*; translational ...; angular ... ” format even since the one-shot prompt. Both GPT models can also figure out some task specification just depending on the 1-shot prompt. For the Bard in the 1-shot setting, most of its target task specifications essentially duplicated the direction setting from the source task “turn screw.” The only improvisational part was the new object regarding the z axis in the coordinate setting.

Next, we take a close look into the GPT-4 5-shot case, the one with the best performance. In total, its five failed cases were “peel potato”, “shave object”, “insert straw through cup lid”, “crack egg”, and “open childproof bottle.” These failed tasks are quite divergent from the source tasks and

pose more challenges to LLMs. Take the “insert straw through cup lid” task as an example, the GPT-4 generated a single active direction “translational. $z = -5 \# N$ ”, which implied using force control in the straw-insertion direction. However, this task should be position controlled because the straw has no interaction with any outside rigid object and there is only minor friction between the straw and the cup lid. For another failed task “open childproof bottle”, it is supposed to have the opposite force direction compared to the normal “open bottle” task. However, the GPT-4 model generated an identical answer for these two tasks.

Regarding the LLaMA-2-70B, our current prompt has difficulty in guiding it to the desired specification. Even up to the 5-shot prompt, it barely generated any suitable response. One typical response is shown below. Note that some redundant details of the function parameters are skipped using “...”.

```
# Combine the five source functions
turn_screw(...)
wipe_table(...)
open_door_from_doorknob(...)
cut_sandwich(...)
slide_box_upward_on_wall(...)
```

The LLaMA-2-70B just repeatedly called the source functions. It was not capable of generating any new coordinate setting or position/force set-points with our prompt.

In addition to focusing our attention on the direction settings, we also inspected the numerical values that were generated. We used values such as 5 (rad/sec), 5 (N) and 2 (cm) in the prompt. Most of the generated values were close to these ones and sometimes reached 10. Among all generated values in the 5-shot tests, we noticed two unusual values, both produced by Bard – 1) “hammer in nail” task: “translational. $z = 50 \# N$ ”, 2) “cut banana” task: “angular. $z = 100 \# rad/sec$ ”. In the hammer case, it is reasonable to set a large force. However, the direction setting the banana cutting case is incorrect, and more notably, the generated 100 (rad/s) set-point is excessively high for common manipulators to execute. This raises an additional concern: if the LLMs generate abnormal set-point values, might it potentially result in safety issues in robot execution?

Conclusions and Discussions

This paper presented a LLM-based approach to convert textual manipulator primitive tasks to TFF-based position/force set-points. By integrating the generalizability of LLMs and the cross-task transferability of TFF together, we have provided a new solution for the planning-to-execution transition problem in layered architectures. The evaluations showed the effectiveness of our prompt in three LLMs – OpenAI GPT-3.5-turbo, OpenAI GPT-4, and Google Bard.

Since the correct rate increases as more source tasks are offered, we suggest building a more comprehensive prompt that includes tasks in all contact and movement types. In addition, further study can be conducted to guide the LLaMA-2 model in producing correct specifications.

Furthermore, the evaluation on the LLM-generated content remains a challenging problem. We believe that there

exists a necessity to design a comprehensive and standardized evaluation set for assessing the language grounding to actions. The recent work of Google RT-1 (Brohan et al. 2023) established a large-scale real-world dataset connecting images and natural language instructions with manipulator actions but is limited to a few verbs. In the future, we will establish an evaluation set that includes a broader range of verbs that can better capture the contact-rich characteristics of manipulators.

Appendix: Evaluation Metrics

In this appendix, we provide detailed metrics used in the Evaluation section. In certain cases, the x , y , and z axes can be interchangeable depending on the coordinate setting.

1. cut pizza: 1 translational direction activated.
2. scrub desk with bench brush: 1 or 2 translational directions activated, must apply force on plane.
3. spear cake with fork: only z translational direction activated.
4. fasten screw with screwdriver: z angular direction activated.
5. loosen screw with screwdriver: opposite as Task 4.
6. unlock lock with key: z angular direction activated.
7. fasten nut with wrench: z angular direction activated.
8. loosen nut with wrench: opposite as Task 7.
9. spread paint with brush: translational direction on $x - y$ plane activated.
10. hammer in nail: only z translational direction activated.
11. rasp wood: same as Task 2.
12. scrape substance from surface: same as Task 2.
13. peel potato: 1 translational direction activated, must apply force on plane.
14. slice cucumber: same as Task 1.
15. flip bread: 1 angular direction activated.
16. shave object: same as Task 2.
17. use roller to roll out dough: same as Task 2.
18. insert peg into pegboard: same as Task 3.
19. brush across tray: same as Task 2 or Task 9.
20. insert straw through cup lid: 1 translational direction activated, position control.
21. open door from hinge: only z angular direction activated.
22. slide block over vertical surface: z translational direction activated, must apply force on $x - y$ plane.
23. turn steering wheel: same as Task 6.
24. shake cocktail bottle: angular direction activated, position control.
25. cut banana: same as Task 1.
26. crack egg: same as Task 20.
27. press button: same as Task 3.
28. insert GPU into socket: same as Task 3.
29. open bottle: same as Task 4.
29. open childproof bottle: add downside force versus Task 29.

References

- Ahn, M.; Brohan, A.; Brown, N.; Chebotar, Y.; Cortes, O.; David, B.; Finn, C.; Gopalakrishnan, K.; Hausman, K.; Herzog, A.; et al. 2022. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. *arXiv preprint arXiv:2204.01691*.
- Arumugam, D.; Karamcheti, S.; Gopalan, N.; Williams, E. C.; Rhee, M.; Wong, L. L. S.; and Tellex, S. 2019. Grounding Natural Language Instructions to Semantic Goal Representations for Abstraction and Generalization. *Auton. Robots*, 43(2): 449–468.
- Ballard, D. H.; and Hartman, L. 1986. Task Frames: Primitives for Sensory-Motor Coordination. *Comput. Vis. Graph. Image Process.*, 36(2-3): 274–297.
- Brohan, A.; Brown, N.; Carbajal, J.; et al. 2023. RT-1: Robotics Transformer for Real-World Control at Scale. In *Robotics: Science and Systems*.
- Brooks, R. A. 1987. Planning is Just a Way of Avoiding Figuring Out What To Do Next. Working Papers WP-303, MIT Artificial Intelligence Laboratory.
- Brooks, R. A. 1991a. Intelligence Without Reason. In *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, 569–595.
- Brooks, R. A. 1991b. New Approaches to Robotics. *Science*, 253(5025): 1227–1232.
- Bruyninckx, H.; and De Schutter, J. 1996. Specification of Force-Controlled Actions in the “Task Frame Formalism”-A Synthesis. *IEEE Trans. Robotics Autom.*, 12(4): 581–589.
- Cao, Y.; and Lee, C. S. G. 2023. Robot Behavior-Tree-Based Task Generation with Large Language Models. In *Proc. AAAI Spring Symp. Challenges Requiring the Combination of Machine Learning and Knowledge Engineering*.
- Chai, J. Y.; Fang, R.; Liu, C.; and She, L. 2016. Collaborative Language Grounding Toward Situated Human-Robot Dialogue. *AI Mag.*, 37(4): 32–45.
- de Marneffe, M.; MacCartney, B.; and Manning, C. D. 2006. Generating Typed Dependency Parses from Phrase Structure Parses. In *Proc. Int. Conf. Language Resources and Evaluation, (LREC)*, 449–454.
- Firby, R. J. 1989. Adaptive Execution in Complex Dynamic Worlds. *PhD Thesis, Yale University*.
- Firby, R. J. 1992. Building Symbolic Primitives with Continuous Control Routines. In *Artificial Intelligence Planning Systems*, 62–69. Elsevier.
- Huang, W.; Abbeel, P.; Pathak, D.; and Mordatch, I. 2022. Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents. In *Proc. Int. Conf. Machine Learning (ICML)*, 9118–9147.
- Kortenkamp, D.; Simmons, R. G.; and Brugali, D. 2016. Robotic Systems Architectures and Programming. In Siciliano, B.; and Khatib, O., eds., *Springer Handbook of Robotics*, 283–306. Springer.
- Kress-Gazit, H.; Fainekos, G. E.; and Pappas, G. J. 2008. Translating Structured English to Robot Controllers. *Adv. Robotics*, 22(12): 1343–1359.
- Kroger, T.; Finkemeyer, B.; and Wahl, F. M. 2004. A Task Frame Formalism for Practical Implementations. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, volume 5, 5218–5223.
- Lin, K.; Agia, C.; Migimatsu, T.; Pavone, M.; and Bohg, J. 2023. Text2Motion: From Natural Language Instructions to Feasible Plans. *arXiv preprint arXiv:2303.12153*.
- Mason, M. T. 1981. Compliance and Force Control for Computer Controlled Manipulators. *IEEE Trans. Syst. Man Cybern.*, 11(6): 418–432.
- Matuszek, C.; Herbst, E.; Zettlemoyer, L.; and Fox, D. 2012. Learning to Parse Natural Language Commands to a Robot Control System. In *Int. Symp. Experimental Robotics*, volume 88, 403–415. Springer.
- Morrow, J. D.; and Khosla, P. K. 1997. Manipulation Task Primitives for Composing Robot Skills. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, volume 4, 3354–3359.
- Nilsson, N. J. 1969. A Mobile Automaton: An Application of Artificial Intelligence Techniques. In *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, 509–520.
- Nilsson, N. J. 1984. Shakey the Robot. Technical Report 323, SRI AI Center.
- Paulius, D.; Eales, N.; and Sun, Y. 2020. A Motion Taxonomy for Manipulation Embedding. In *Robotics: Science and Systems*.
- Raibert, M. H.; and Craig, J. J. 1981. Hybrid Position/Force Control of Manipulators. *ASME. J. Dyn. Sys., Meas., Control.*, 103(2): 126–133.
- Roy, D. 2005. Semiotic Schemas: A Framework for Grounding Language in Action and Perception. *Artificial Intelligence*, 167(1-2): 170–205.
- Singh, I.; Blukis, V.; Mousavian, A.; Goyal, A.; Xu, D.; Tremblay, J.; Fox, D.; Thomason, J.; and Garg, A. 2022. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. *arXiv preprint arXiv:2209.11302*.
- Squire, S.; Tellex, S.; Arumugam, D.; and Yang, L. 2015. Grounding English Commands to Reward Functions. In *Robotics: Science and Systems*.
- Tellex, S.; Kollar, T.; Dickerson, S.; Walter, M. R.; Banerjee, A. G.; Teller, S. J.; and Roy, N. 2011. Understanding Natural Language Commands for Robotic Navigation and Mobile Manipulation. In *Proc. AAAI Conf. Artificial Intelligence*, volume 25, 1507–1514.
- Toyoda, M.; Suzuki, K.; Mori, H.; Hayashi, Y.; and Ogata, T. 2021. Embodying Pre-Trained Word Embeddings Through Robot Actions. *IEEE Robotics Autom. Lett.*, 6(2): 4225–4232.
- Vuong, N.; Pham, H.; and Pham, Q.-C. 2021. Learning Sequences of Manipulation Primitives for Robotic Assembly. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 4086–4092.
- Xiang, J.; Tao, T.; Gu, Y.; Shu, T.; Wang, Z.; Yang, Z.; and Hu, Z. 2023. Language Models Meet World Models: Embodied Experiences Enhance Language Models. *arXiv preprint arXiv:2305.10626*.