

Statler: State-Maintaining Language Models for Embodied Reasoning

Takuma Yoneda^{*1}, Jiading Fang^{*1}, Peng Li^{*2}, Huanyu Zhang^{*3}, Tianchong Jiang³, Shengjie Lin¹,
Ben Picker³, David Yunis¹, Hongyuan Mei¹, and Matthew R. Walter¹

¹Toyota Technological Institute at Chicago
{takuma, fjd, slin, dyunis, hongyuan, mwalter}@ttic.edu

²Fudan University
lip21@m.fudan.edu.cn

³University of Chicago
{huanyu, tianchongj, bpicker}@uchicago.edu

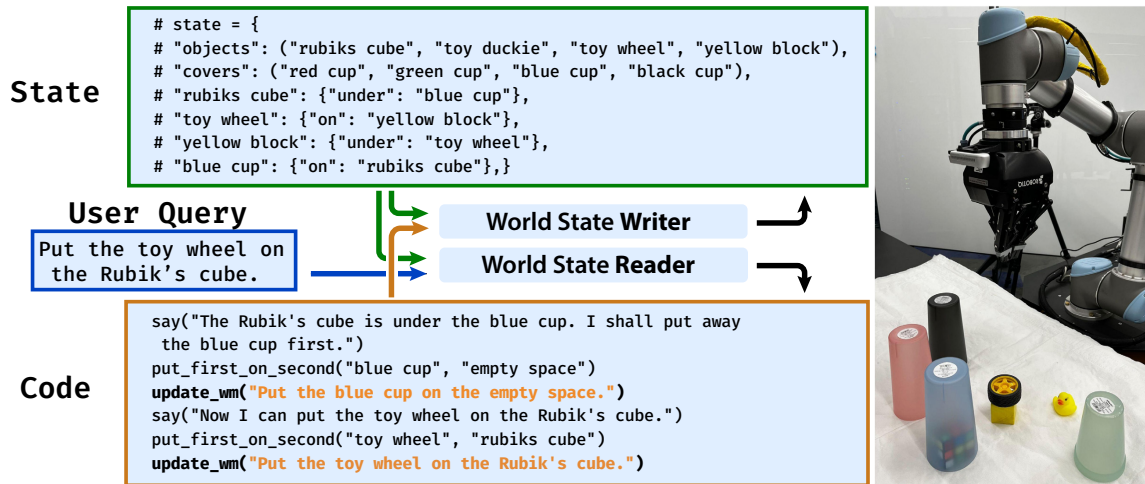


Fig. 1: Our Statler framework enables robots to carry out complex tasks specified in natural language that require reasoning over long time horizons. Integral to our model are its world-state writer and world-state reader, two instances of general LLMs responsible for maintaining the explicit world state and generating code that enables the robot to carry out the task.

Abstract—There has been a significant research interest in employing large language models to empower intelligent robots with complex reasoning. Existing work focuses on harnessing their abilities to reason about the histories of their actions and observations. In this paper, we explore a new dimension in which large language models may benefit robotics planning. In particular, we propose Statler, a framework in which large language models are prompted to maintain an estimate of the world state, which are often unobservable, and track its transition as new actions are taken. Our framework then conditions each action on the estimate of the current world state. Despite being conceptually simple, our Statler framework significantly outperforms strong competing methods (e.g., Code-as-Policies) on several robot planning tasks. Additionally, it has the potential advantage of scaling up to more challenging long-horizon planning tasks. We release our code [here](#).

I. INTRODUCTION

Large language models (LLMs) exhibit strong reasoning capabilities that are harnessed to perform a wide range of downstream tasks such as dialogue and code generation [1–3]. The robotics community has recently seen a significant interest in empowering robots with LLMs, enabling them to

understand natural language commands and perform tasks that require sophisticated reasoning [4–7]. However, existing methods are model-free: they use LLMs as policy functions that generate future actions only conditioned on previous actions and observations.

In this paper, we propose a simple yet effective model-based approach. Our framework—named Statler—maintains a running estimate of the world state by prompting large language models and performs multistep embodied reasoning conditioned on the estimated state. Figure 1 illustrates this framework. In particular, Statler utilizes a pair of prompted LLMs: instructed by a few demonstrations, the **world-state reader** takes as input the user query, reads the estimated world state, and generates an executable action (e.g., a code snippet); instructed by another set of demonstrations, the **world-state writer** updates the world state estimate based on the action. This mechanism resembles how a domain-specific formal language tracks a symbolic world state [8], but enjoys greater flexibility since pretrained large language models are known to be domain-agnostic. As we will see soon in Section IV, the prompts in our experiments are generic and users of our framework will have minimal design workload.

Our Statler framework is primarily inspired by classical

^{*}Denotes equal contribution; Contribution by each author can be found in the appendix.

```

1 # Initial state
2 cups = [False, True, False]
3 Swapping cup 1 with cup 2
4 Swapping cup 0 with cup 2
5 Swapping cup 1 with cup 2
6 cups = [True, False, False]
7
8

```

Prompt 1: The prompt and desired output of a vanilla LLM.

```

1 # Initial state
2 cups = [False, True, False]
3 Swapping cup 1 with cup 2
4 Swapping cup 0 with cup 2
5 Swapping cup 1 with cup 2
6 cups = [False, False, True]
7 cups = [True, False, False]
8 cups = [True, False, False]

```

Prompt 2: The prompt and desired output of LLM+CoT.

```

1 # Initial state
2 cups = [False, True, False]
3 Swapping cup 1 with cup 2
4 cups = [False, False, True]
5 Swapping cup 0 with cup 2
6 cups = [True, False, False]
7 Swapping cup 1 with cup 2
8 cups = [True, False, False]

```

Prompt 3: The prompt and desired output of LLM+State.

model-based reinforcement learning. In a model-based approach, an environment (or world) model learns to capture the dynamics of the environment (e.g., possible outcomes of an action) so that the policy conditioned on the model state will take more informed actions [9]. In our framework, the LLMs have acquired massive amounts of commonsense knowledge from pretraining, and they are elicited—by a few demonstrations—to behave as an environment model, estimating the world state and facilitating decision making.

Another motivation of our design is to handle missing data. In robotics tasks, we often have to cope with latent world dynamics that are not directly observable. In such scenarios, explicitly maintaining an estimated world state improves decision making, although the estimates might be imperfect. This is analogous to graphical models with latent variables: spelling out latent variables and imputing their values is often helpful for reasoning about the target variables, although the imputation may not be perfect [10].

The final motivation of our state-maintaining design is its potential to scale up to long-horizon planning tasks. In multistep reasoning and planning, an LLM has to implicitly maintain the world state in its internal representation, which has been demonstrated to be difficult in previous work [11–17]. By explicitly maintaining an estimated world state, our framework makes it easier to track and consult the world state at any step of reasoning and planning, thus carrying a higher chance of success in long-horizon tasks.

In the following sections, we will show that our framework performs as expected: in Section II, we demonstrate the concept with a pedagogical example; in Section III, we introduce the Statler framework; in Section IV, we present the experiments, in which our framework significantly outperforms strong competing methods such as Code-as-Policies [5].

II. MOTIVATING EXAMPLE

We use *three-cups-and-a-ball*, a simple shell game, to demonstrate the effectiveness of our state-maintaining idea. In this game, a ball is covered under one of three identical cups and the initial position of the ball is known to the player. In each of K rounds, we randomly swap two cups’ positions. Finally, we ask the player to guess the position of the ball.

We present three separate cases of using LLMs to play this game using GPT-3 (precisely, text-davinci-003). Prompt 1 demonstrates the simplest case: We represent the state with Boolean variables with `True` indicating “ball is here”. We feed the initial state and the K rounds of swaps into GPT-3, instructing it to complete the final state. Prompt 2 is an improved way: after reading K rounds of swaps, GPT-3




Fig. 2: Model accuracies on the *three-cups-and-a-ball* shell game. LLM+State is a simplified version of our proposed Statler framework. For each method, the solid line shows how its accuracy $a(n)$ changes with the number of swaps n . The dashed line is the *relative* accuracy: $r(n) = a(n)/a(1)$. Intuitively, it measures how fast the performance decreases from a *hypothetically perfect* one-swap performance. Note that LLM+State indeed achieves $a(1) = 100\%$

is asked to give all the intermediate states over the game. This version is inspired by Chain-of-Thought prompting [18], which improves the performance of an LLM by requesting it to spell out its intermediate reasoning steps. Finally, Prompt 3 is a simple instantiation of our state-maintaining idea: we ask GPT-3 to return the current state immediately after reading each round of swaps, stimulating the model to track and update the state as the game progresses.

We evaluate these methods with a range of K ; for each K and each method, we feed 30 demonstrations with various numbers of swaps to the model, and repeat the experiment 100 times. Figure 2 visualizes the average accuracies. The state-maintaining method significantly outperforms the other methods, with the performance gap increasing with K .

III. METHOD

As exemplified in Section II, the key to our approach is to allow the LLM to describe the next state while responding to each user query. The motivating example is simple in that the next state *is* the response. Instead, we now consider more general scenarios where there is a significant burden on the LLM to track the state updates as well as generate responses. (Fig. 3). For the general cases, we propose to *split* the burden across multiple different prompted LLMs. Precisely, we maintain a separate prompt that includes instructions and demonstrations for each subtask (i.e., state-tracking or query-

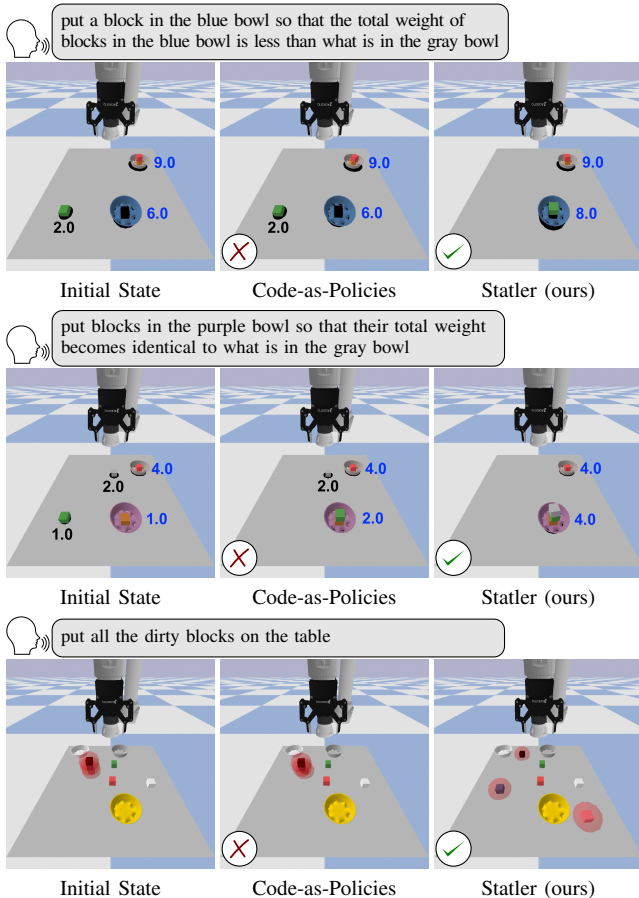


Fig. 3: Examples of simulations that show the result of executing different natural language instructions using Code-as-Policies and our state-maintaining Statler algorithm.

responding) and then use the prompt to elicit an LLM to perform the particular subtask. As we discuss shortly, our framework includes a **world-state reader** that responds to the user query and a **world-state writer** that is responsible for updating the state representation. Our framework (Fig. 1) does not pose any fundamental limitation on which domain it can be applied to. Our approach can be regarded as a model-based extension of Code-as-Policies (CaP) in the sense that it keeps the core capabilities of CaP (e.g., hierarchical code generation) and incorporates a means to explicitly maintain an estimated world state.

It is useful to consider example prompts to understand the operation of the reader and writer models. Prompt 4 is an example of the input passed to the world-state reader. Initially, we initialize a JSON-formatted state with a reference to object-oriented principles. Given a user query “Put the cyan block on the yellow block” (Line 13) and the current state representation (Lines 1–12), the world-state reader should generate the code that responds to the query, taking into account the current state. The expected code to be generated is highlighted in green. After generating the code, our model executes it to complete the query. When the state needs to be updated, the generated code will contain an `update_wm` function that triggers the world-state writer with

```

1 # state = {
2 #   "objects": ["cyan block", "yellow block", "brown block",
3 #             "purple block", "blue block", "green bowl", "red bowl", "
4 #             disinfectant"],
5 #   "relations": [],
6 #   "disinfectant": {"contains": []},
7 #   "cyan block": {"is": ["dirty"]},
8 #   "yellow block": {"is": ["clean"]},
9 #   "brown block": {"is": ["clean"]},
10 #   "purple block": {"is": ["dirty"]},
11 #   "blue block": {"is": ["clean"]},
12 #   "green bowl": {},
13 #   "red bowl": {}
14 # }
15 # query: Put the cyan block on the yellow block
16 put_first_on_second("cyan block", "yellow block")
17 update_wm("Put the cyan block on the yellow block")

```

Prompt 4: World-state reader. The text highlighted in green represents the part that the model is expected to generate.

```

1 # state = {
2 #   "objects": ["cyan block", "yellow block", "brown block",
3 #             "purple block", "blue block", "green bowl", "red bowl", "
4 #             disinfectant"],
5 #   "relations": [],
6 #   "disinfectant": {"contains": []},
7 #   "cyan block": {"is": ["dirty"]},
8 #   "yellow block": {"is": ["clean"]},
9 #   "brown block": {"is": ["clean"]},
10 #   "purple block": {"is": ["dirty"]},
11 #   "blue block": {"is": ["clean"]},
12 #   "green bowl": {},
13 #   "red bowl": {}
14 # }
15 # query: Put the cyan block on the yellow block.
16 # state = {
17 #   "objects": ["cyan block", "yellow block", "brown block",
18 #             "purple block", "blue block", "green bowl", "red bowl",
19 #             "disinfectant"],
20 #   "relations": [{"cyan block is on yellow block"}],
21 #   "disinfectant": {"contains": []},
22 #   "cyan block": {"is": ["dirty"]},
23 #   "yellow block": {"is": ["dirty"]},
24 #   "brown block": {"is": ["clean"]},
25 #   "purple block": {"is": ["dirty"]},
26 #   "blue block": {"is": ["clean"]},
27 #   "green bowl": {},
28 #   "red bowl": {}
29 # }

```

Prompt 5: World-state writer. The text rendered in blue highlights the updated part of the state.

the query specified in its argument. In Prompt 5, we show the corresponding example for the world-state writer. Similar to the reader, we prepend the current state representation before the user query and the model generates the updated state representation (highlighted in green). Whenever the writer updates the state representation, we store it in external memory and refer to it as the current state.

IV. EXPERIMENTS

We evaluate the capabilities of Statler alongside state-of-the-art LLM models on three tabletop manipulation domains (Fig. 4): pick-and-place, block disinfection, and relative weight reasoning. For each domain, we design in-context examples and consider 20 evaluation episodes each of which consists of 5–16 consecutive steps of user queries. Every episode contains at least one query that requires reasoning over the interaction history (i.e., requires “memory” across steps), which makes the task significantly challenging.




Fig. 4: The simulated domains we consider include (a) Pick-and-Place; (b) Block Disinfection, where the translucent sphere around a block represents its dirtiness (this is not visible to the robot); and (c) Relative Weight Reasoning, where the radius of the disk under each block indicates its weight (this is not visible to the robot).

A. Simulated Tabletop Manipulation Domains

The **Pick-and-Place** domain involves scenarios that require a robot arm to sequentially pick up and place a block onto another block, bowl, or the table. The model needs to remember and reason over the block locations. The example user queries are “Put the green block in the red bowl.”, “What is the color of the block under the pink block?”, and “How many blocks are in the green bowl?”.

In the **Block Disinfection** domain, we consider a scenario in which a block can be either *dirty* or *clean*, the state of which is not observable by the robot. When a clean block touches a dirty block (e.g., as a result of stacking one block on another), the clean block becomes dirty. There is a *disinfector* on the table that cleans any block placed inside it. This scenario emulates a clean-up task in which you might ask a robot to put dirty dishes in a dishwasher or dirty clothes in a washing machine. The user query contains pick-and-place commands similar to those in the pick-and-place domain as well as textual utterances that require reasoning over which blocks are clean and dirty, such as “Put all the clean blocks in the green bowl.” This domain presents a particular challenge as the model must track the cleanliness of each block and accurately capture the state mutations that happen when a dirty block comes into contact with another clean block.

Relative Weight Reasoning involves memorizing and reasoning over the relative weights of the blocks. User queries provide information about the weight of blocks (e.g., “The red block is twice the weight of the bronze block.”), which are followed by queries that require reasoning over the weights (e.g., “Put blocks in the purple bowl so that their total weight becomes identical to what is in the gray bowl.”).

We compare our proposed approach, Statler, to two strong competing methods: Code-as-Policies [5] (CaP) and CaP with Chain-of-Thought prompting [18] (CaP+CoT). CaP generates code for the current question at each step based on the past actions, but it does not maintain a state. Following the CoT framework, *at every step*, CaP+CoT deduces the intermediate states based on an *initial state* and past actions, which are considered as its thoughts, to generate the current code. But it leads to redundant reasoning and increases the length of the prompt, which may then exceed the LLM’s context window

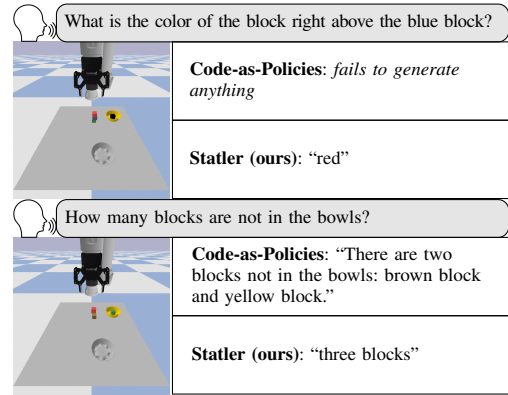


Fig. 5: Examples that show the result of querying LLMs with and without maintained state. In the first scenario, CaP fails to produce an answer, while our Statler model produces the correct response. In the second example, one block is not visible and CaP incorrectly identifies two blocks as not being a bowl. By maintaining a persistent world state, our method is aware of the third block and correctly answers the query.

size limitations. Furthermore, longer reasoning also demands longer, more intricate demo example prompts, contributing to increased developer effort. We ensure that the demonstrations (i.e., in-context examples) given to each of the models are equivalent. Namely, we use the same sequence of user queries and code snippets, except for necessary differences due to their designs such as state representation.

TABLE I: Episode success rates and individual step success rates (in parentheses) for each sequential task. † indicates that the context limit was often exceeded.

	Pick & Place	Disinfection	Weight
CaP	0.00 (0.54)	0.00 (0.68)	0.00 (0.84)
CaP+CoT	0.25 (0.76)	0.00 (0.20) [†]	0.30 (0.88)
Statler (ours)	0.50 (0.88)	0.40 (0.82)	0.55 (0.93)

Table I reports the episode success rates of each method along with the the success rate for individual steps. An episode is considered to be a failure if a model fails to respond to one of the user queries in the episode. While the CaP baseline correctly processes more than half of the individual steps in each domain, it fails to successfully complete any of the episodes. As we show later, CaP correctly processes most queries that do not require reasoning over previous steps (e.g., “Put the red block on the blue block.”), but tends to generate incorrect (or no) code in response to queries that require reasoning over the history (e.g., “Put all the dirty blocks in the pink bowl.” and “What is the color of the block under the purple block?”) (see Figure 5 (top)). CaP+CoT fares slightly better in the Pick-and-Place and Relative Weight Reasoning, but still fails in most episodes. In contrast, Statler successfully handles the majority of these queries, demonstrating strong improvement over the others. It should be noted we explicitly chose queries that were challenging for LLM-based models, which partially accounts

for why our model’s scores show room for improvement.

TABLE II: Success rates of Code-as-Policies (CaP) and Statler for non-temporal and temporal queries.

	Non-temporal		Temporal	
	CaP	Statler (ours)	CaP	Statler (ours)
Pick & Place	1.00 (62/62)	1.00 (68/68)	0.31 (9/29)	0.83 (48/58)
Disinfection	0.99 (148/149)	0.98 (164/168)	0.05 (1/20)	0.65 (15/23)
Weight	1.00 (107/107)	1.00 (107/107)	0.00 (0/20)	0.55 (11/20)

In order to better understand the behavior of Statler in comparison to Code-as-Policies, we analyze the success rates based on the type of user queries. Specifically, we categorize each query as either *temporal* or *non-temporal* depending on whether responding to the query necessitates temporal reasoning. We emphasize that contemporary methods, including the baselines that we consider, use non-temporal queries for evaluation. Table II summarizes the resulting accuracy. The models often fail at different steps in an episode. We note that, when calculating accuracy we only consider the sequence of steps until the model fails to generate the correct code, which explains the mismatch in the denominators.

We see that both models achieve near-perfect performance on commands that do not require temporal reasoning. However, the performance of CaP noticeably decreases for tasks that require reasoning over the past. In contrast, Statler achieves success rates of 83% (vs. 31% for CaP) on Pick-and-Place, 65% (vs. 5% for CaP) on Block Disinfection, and 55% (vs. 0% for CaP) on Relative Weight Reasoning.

Although our method enjoys a better robustness than the baseline methods, it inherits some issues of large language models, which hinders its performance. For example, it hallucinates block conditions (e.g., clean or dirty) or locations when the cleanliness of the block is never explicitly described. Moreover, the model’s reasoning strategy seems to predominantly treat weight as an abstract concept, e.g. light vs. heavy, rather than executing mathematical computations. This weakness is evident when asking the model to accumulate blocks in a bowl until their total weight surpasses that of another bowl, yet the model underfills the bowl. In the disinfection domain, our model struggles to comprehend ambiguous terms like “other” in queries such as “the other blocks are clean.” It can also wrongly infer from the training prompt that a block *at the bottom* becomes dirty when a block is placed on top of it, irrespective of the latter’s cleanliness.

B. Real Robot Experiments

In order to validate Statler on a real robot, we implement it on a UR5 arm in a similar tabletop domain as the simulated experiments. We use MDETR [19], an open-vocabulary segmentation model, to obtain segmentation masks for objects within an RGB image captured by a RealSense camera mounted on the gripper. Using the predicted object mask and the depth image, the object point-cloud can be obtained, from which we estimate its center position and bounding box. All of the primitive functions are identical to those used in simulation. In this domain, the robot is asked to stack objects

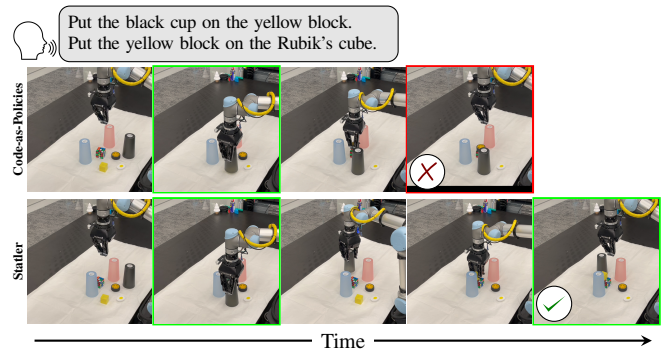


Fig. 6: A comparison of the resulting behavior for (top) Code-as-Policies and (bottom) our Statler model for the real robot experiments for the given multi-sentence instruction. Frames correspond to instances when the robot has placed an object. In order to successfully carry out the instruction, the robot must subsequently remove the black cup immediately after using it to cover the yellow block so that it can place the yellow block on the Rubik’s cube. However, the the baseline Code-as-Policies (top row, third frame) fails to move the black cup aside, leaving the yellow block covered. It then places a wrong object on top of the Rubik’s cube.

- Your task is to maintain the status of these items using a JSON dictionary and update the status of the corresponding items after a new query.
- This JSON dictionary will be commented, meaning that the starting character of each line is #.

Prompt 6: Portion of Statler-Auto prompt.

and cover objects with different colored cups. At any point, an object is only permitted to be covered by at most a single object or cover. If the robot is asked to manipulate the bottom object, it must put away the top one. If a new cover or object is to be stacked on it, the existing one must be removed.

We evaluate the performance of Statler vs. CaP in the real robot domain using 10 episodes. Statler achieves episode and step (in parentheses) success rate of 40% (70%), where 67% of the failure cases are due to LLM reasoning while others are caused by either perception or manipulation issues. The success rate for CaP is 20% (46%), where LLM reasoning accounts for 88% of failures. In Figure 6, we also provide a short example where the CaP baseline fails. The difficulty is in recognizing that yellow block is hidden under the black cup, which must be removed before picking up the yellow block as Statler correctly spots. Instead, CaP is not aware of this and tries to pick up the yellow block nonetheless, which leads MDETR to incorrectly detect the toy wheel that has yellow color in it as the yellow block.

C. State-Maintenance Ablations

To better understand Statler’s state-maintenance strategy, we consider three different approaches to tracking the state.

The first (Statler-Unified) employs a single LLM as both the world-state reader and writer using a prompt that interleaves Statler’s reader and writer prompts. At each step, the LLM first generates the action and then predicts the state that results

from executing that action. The LLM then uses the resulting state when reasoning over the next query. Using a single LLM is conceptually simple, but it incurs an added burden for reasoning and generalization. Inspired by InstructGPT [20], the second (Statler-Auto) does not receive any in-context state-update examples for the world-state writer. Instead, we provide a natural language description of how the state should be maintained. Prompt 6 shows the relevant portion of the prompt. With an instruction and no in-context state-update examples, we ran our model on the same set of tasks. The third (Statler w/o State) ablates the world-state maintenance components of Statler entirely, resulting in a model that reduces to Code-as-Policies.

TABLE III: Ablation episode (individual step) success rates.

	Pick & Place	Disinfection	Weight
Statler w/o State	0.00 (0.54)	0.00 (0.68)	0.00 (0.84)
Statler-Unified	0.40 (0.85)	0.35 (0.79)	0.50 (0.92)
Statler-Auto	0.75 (0.88)	0.45 (0.82)	0.40 (0.90)
Statler (ours)	0.50 (0.88)	0.40 (0.82)	0.55 (0.93)

Table III compares the performance of Statler to the three variations in terms of both their full-episode completion rates (using 20 episodes for each domain) as well their individual step success rates. Without maintaining the world-state, Statler w/o State fails to complete any episodes (recall that an episode is considered to be a failure if the model fails to respond to one of the user queries during the episode) and results in individual step success rates that are significantly lower than Statler. Meanwhile, we see that Statler’s use of separate LLMs for the world-state reader and world-state writer results in consistently higher episode success rates compared with the use of a unified reader and writer (Statler-Unified). The individual step success rates are higher for Pick-and-Place and Block Disinfection, and comparable for Relative Weight Reasoning. With regards to Statler’s use of separate LLMs for the world-state writer and reader, we note that in-context learning has been shown to be sensitive to variations in prompt templates, the order of examples, and the examples used [21, 22]. In light of this, it is plausible that the performance gains that we achieve by dividing our reader and writer may be attributed in part to this sensitivity, allowing the models to, in effect, become specialized at their respective tasks. Interestingly, Statler-Auto performs noticeably better than Statler and Statler-Unified with regards to the episode success rate on the Pick-and-Place and Block Disinfection domains, but comparable to Statler in terms of the individual success rates, and worse for Relative Weight Reasoning.

V. RELATED WORK

Language Understanding for Robotics A common approach for language understanding for robotic agents involves symbol grounding [23], whereby phrases are mapped to symbols in the robot’s world model. Early work [24, 25] relies upon hand-engineered rules to perform this mapping. More recent methods replace these rules with statistical

models the parameters of which are trained on annotated corpora [26–37]. Other methods use neural network-based architectures to jointly reason over natural language utterances and the agent’s (visual) observations of the scene [38–42].

LLMs for Robotics Since LLMs are trained with Internet-scale corpora, their infused common sense have shown to help in the domain of robotics in terms of high-level planning from natural language instructions [4, 5, 43] for both object manipulation [44, 45] and navigation tasks [46–49]. Combining LLMs with expressive visual-language embeddings also enables impressive capabilities [50]. This has led to efforts to push for general multi-modality embodied models [51, 52].

Code Generation with LLMs Code generation has been one of the most successful use cases for LLMs [2, 3, 53–56]. Since code can connect with executable APIs for tasks including computation, vision and manipulation, a large chunk of work has focused on code generation with different tools [57–59]. In particular, Code-as-policies [5] is one of the first to use code generation within a robotics context.

State Representation in Reasoning The use of state representations have been shown to help in algorithmic reasoning tasks [60, 61]. Instead of using one forward pass to predict the execution result for the entire code snippet, Nye et al. [60] proposes to spell out step-by-step intermediate outputs to help infer the final execution results. Also relevant are research efforts that aim to enhance language modeling by rolling out possible future tokens [62].

Language Models and Planning Recent work shows that vanilla and instruction-tuned LLMs plan poorly [14, 16, 17]. Some works propose using the LLM as an intermediary between natural language and a domain-specific programming language, and then uses a traditional planner [15, 16, 63]. Silver et al. [17] employ Chain-of-Thought and iterative reprompting with feedback on generated plans, but require GPT-4 for good performance. Xiang et al. [64] use parameter-efficient finetuning of LLMs on top of traces from a world-model and show improved performance on related tasks.

VI. CONCLUSION

We presented Statler, a language model that maintains an explicit representation of state to support longer-horizon robot reasoning tasks. Integral to Statler are a world-state reader that responds to a user query taking into account the current internal state, and a world-state writer that maintains the world state. Evaluations on various simulated and real robot manipulation tasks reveal that Statler significantly outperforms contemporary models on non-trivial tasks that require reasoning over the past. Ablations demonstrate the contributions of our world-state reader and writer, and suggest Statler’s flexibility to the state representation.

VII. ACKNOWLEDGEMENTS

We are grateful to National Science Foundation for enabling this work under HDR TRIPODS (No. 2216899), and to Adobe for supporting Hongyuan Mei through an Adobe Research gift. We thank Luzhe Sun and Richard Xu for their help at the early stage of the project.

REFERENCES

- [1] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," *arXiv preprint arXiv:2205.11916*, 2022.
- [2] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, and W. Zaremba, "Evaluating large language models trained on code," *arXiv preprint arXiv:2107.03374*, 2021.
- [3] OpenAI, "GPT-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [4] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, E. Jang, R. J. Ruano, K. Jeffrey, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, K.-H. Lee, S. Levine, Y. Lu, L. Luu, C. Parada, P. Pastor, J. Quiambao, K. Rao, J. Rettinghouse, D. Reyes, P. Sermanet, N. Sievers, C. Tan, A. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, S. Xu, M. Yan, and A. Zeng, "Do as I can, not as I say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.
- [5] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. R. Florence, and A. Zeng, "Code as policies: Language model programs for embodied control," *arXiv preprint arXiv:2209.07753*, 2022.
- [6] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. R. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar, P. Sermanet, N. Brown, T. Jackson, L. Luu, S. Levine, K. Hausman, and B. Ichter, "Inner monologue: Embodied reasoning through planning with language models," in *Proceedings of the Conference on Robot Learning (CoRL)*, 2022.
- [7] S. Yang, O. Nachum, Y. Du, J. Wei, P. Abbeel, and D. Schuurmans, "Foundation models for decision making: Problems, methods, and opportunities," *arXiv preprint arXiv:2303.04129*, 2023.
- [8] A. Nordmann, N. Hochgeschwender, and S. B. Wrede, "A survey on domain-specific languages in robotics," in *Simulation, Modeling, and Programming for Autonomous Robots*, 2014.
- [9] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- [10] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*, 2009.
- [11] "Anthropic introducing 100k Context windows," <https://www.anthropic.com/index/100k-context-windows>, accessed: 2023-05-11.
- [12] N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, and P. Liang, "Lost in the middle: How language models use long contexts," *arXiv preprint arXiv:2307.03172*, 2023.
- [13] S. Sun, K. Krishna, A. Mattarella-Micke, and M. Iyyer, "Do long-range language models actually use long-range context?" *arXiv preprint arXiv:2109.09115*, 2021.
- [14] K. Valmeekam, S. Sreedharan, M. Marquez, A. Olmo, and S. Kambhampati, "On the planning abilities of large language models (a critical investigation with a proposed benchmark)," *arXiv preprint arXiv:2302.06706*, 2023.
- [15] L. Guan, K. Valmeekam, S. Sreedharan, and S. Kambhampati, "Leveraging pre-trained large language models to construct and utilize world models for model-based task planning," *arXiv preprint arXiv:2305.14909*, 2023.
- [16] B. Liu, Y. Jiang, X. Zhang, Q. Liu, S. Zhang, J. Biswas, and P. Stone, "Llm+ p: Empowering large language models with optimal planning proficiency," *arXiv preprint arXiv:2304.11477*, 2023.
- [17] T. Silver, S. Dan, K. Srinivas, J. B. Tenenbaum, L. P. Kaelbling, and M. Katz, "Generalized planning in PDDL domains with pretrained large language models," *arXiv preprint arXiv:2305.11014*, 2023.
- [18] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [19] A. Kamath, M. Singh, Y. LeCun, G. Synnaeve, I. Misra, and N. Carion, "MDETR - Modulated detection for end-to-end multi-modal understanding," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [20] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe, "Training language models to follow instructions with human feedback," 2022.
- [21] Q. Dong, L. Li, D. Dai, C. Zheng, Z. Wu, B. Chang, X. Sun, J. Xu, and Z. Sui, "A survey for in-context learning," *arXiv preprint arXiv:2301.00234*, 2022.
- [22] Z. Zhao, E. Wallace, S. Feng, D. Klein, and S. Singh, "Calibrate before use: Improving few-shot performance of language models," in *International Conference on Machine Learning*. PMLR, 2021, pp. 12 697–12 706.
- [23] S. Harnad, "The symbol grounding problem," *Physica D*, vol. 42, pp. 335–346, 1990.
- [24] T. Winograd, "Procedures as a representation for data in a computer program for understanding natural language," Ph.D. dissertation, Massachusetts Institute of Technology, 1971.
- [25] M. MacMahon, B. Stankiewicz, and B. Kuipers, "Walk the talk: Connecting language, knowledge, and action in route instructions," in *Proceedings of the National*

- Conference on Artificial Intelligence (AAAI)*, 2006.
- [26] T. Kollar, S. Tellex, D. Roy, and N. Roy, “Toward understanding natural language directions,” in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2010.
- [27] C. Matuszek, D. Fox, and K. Koscher, “Following directions using statistical machine translation,” in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2010.
- [28] D. L. Chen and R. J. Mooney, “Learning to interpret natural language navigation instructions from observations,” in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2011.
- [29] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. Teller, and N. Roy, “Understanding natural language commands for robotic navigation and mobile manipulation,” in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2011.
- [30] C. Matuszek, E. Herbst, L. Zettlemoyer, and D. Fox, “Learning to parse natural language commands to a robot control system,” in *Proceedings of the International Symposium on Experimental Robotics (ISER)*, 2012.
- [31] J. Thomason, S. Zhang, R. J. Mooney, and P. Stone, “Learning to interpret natural language commands through human-robot dialog,” in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [32] T. M. Howard, S. Tellex, and N. Roy, “A natural language planner interface for mobile manipulators,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [33] D. K. Misra, J. Sung, K. Lee, and A. Saxena, “Tell me Dave: Context-sensitive grounding of natural language to manipulation instructions,” *International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 281–300, January 2016.
- [34] J. Thomason, J. Sinapov, M. Svetlik, P. Stone, and R. J. Mooney, “Learning multi-modal grounded linguistic semantics by playing “I spy”,” in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- [35] J. Thomason, J. Sinapov, R. J. Mooney, and P. Stone, “Guiding exploratory behaviors for multi-modal grounding of linguistic descriptions,” in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2018.
- [36] M. Shridhar and D. Hsu, “Interactive visual grounding of referring expressions for human-robot interaction,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- [37] R. Paul, J. Arkin, D. Aksaray, N. Roy, and T. M. Howard, “Efficient grounding of abstract spatial concepts for natural language interaction with robot platforms,” *International Journal of Robotics Research*, vol. 37, no. 10, pp. 1269–1299, June 2018.
- [38] H. Mei, M. Bansal, and M. Walter, “Listen, attend, and walk: Neural mapping of navigational instructions to action sequences,” in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2016.
- [39] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. D. Reid, S. Gould, and A. van den Hengel, “Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [40] D. Fried, R. Hu, V. Cirik, A. Rohrbach, J. Andreas, L.-P. Morency, T. Berg-Kirkpatrick, K. Saenko, D. Klein, and T. Darrell, “Speaker-follower models for vision-and-language navigation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2018.
- [41] F. Zhu, Y. Zhu, X. Chang, and X. Liang, “Vision-language navigation with self-supervised auxiliary reasoning tasks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020.
- [42] S. Y. Min, D. S. Chaplot, P. Ravikumar, Y. Bisk, and R. Salakhutdinov, “FILM: Following instructions in language with modular methods,” *arXiv preprint arXiv:2110.07342*, 2021.
- [43] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, “Language models as zero-shot planners: Extracting actionable knowledge for embodied agents,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2022.
- [44] R. Wang, J. Mao, J. Hsu, H. Zhao, J. Wu, and Y. Gao, “Programmatically grounded, compositionally generalizable robotic manipulation,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [45] A. Z. Ren, B. Govil, T.-Y. Yang, K. R. Narasimhan, and A. Majumdar, “Leveraging language for accelerated learning of tool manipulation,” in *Proceedings of the Conference on Robot Learning (CoRL)*, 2023.
- [46] A. Majumdar, A. Shrivastava, S. Lee, P. Anderson, D. Parikh, and D. Batra, “Improving vision-and-language navigation with image-text pairs from the Web,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [47] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song, “Cows on pasture: Baselines and benchmarks for language-driven zero-shot object navigation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [48] D. Shah, B. Osinski, S. Levine *et al.*, “LM-Nav: Robotic navigation with large pre-trained models of language, vision, and action,” in *Proceedings of the Conference on Robot Learning (CoRL)*, 2023.
- [49] C. Huang, O. Mees, A. Zeng, and W. Burgard, “Visual language maps for robot navigation,” *arXiv preprint arXiv:2210.05714*, 2022.
- [50] M. Shridhar, L. Manuelli, and D. Fox, “CLIPort: What and where pathways for robotic manipulation,” *arXiv preprint arXiv:2109.12098*, 2021.

- [51] A. Zeng, A. S. Wong, S. Welker, K. Choromanski, F. Tombari, A. Purohit, M. S. Ryoo, V. Sindhvani, J. Lee, V. Vanhoucke, and P. R. Florence, “Socratic models: Composing zero-shot multimodal reasoning with language,” *arXiv preprint arXiv:2204.00598*, 2022.
- [52] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu *et al.*, “PaLM-E: An embodied multimodal language model,” *arXiv preprint arXiv:2303.03378*, 2023.
- [53] D. Hendrycks, S. Basart, S. Kadavath, M. Mazeika, A. Arora, E. Guo, C. Burns, S. Puranik, H. He, D. X. Song, and J. Steinhardt, “Measuring coding challenge competence with APPS,” *arXiv preprint arXiv:2105.09938*, 2021.
- [54] Y. Li, D. H. Choi, J. Chung, N. Kushman, J. Schrittwieser, R. Leblond, Tom, Eccles, J. Keeling, F. Gimeno, A. D. Lago, T. Hubert, P. Choy, C. de, M. d’Autume, I. Babuschkin, X. Chen, P.-S. Huang, J. Welbl, S. Goyal, Alexey, Cherepanov, J. Molloy, D. J. Mankowitz, E. S. Robson, P. Kohli, N. de, Freitas, K. Kavukcuoglu, and O. Vinyals, “Competition-level code generation with AlphaCode,” *Science*, vol. 378, pp. 1092–1097, 2022.
- [55] B. Chen, F. Zhang, A. Nguyen, D. Zan, Z. Lin, J.-G. Lou, and W. Chen, “CodeT: Code generation with generated tests,” *arXiv preprint arXiv:2207.10397*, 2022.
- [56] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. J. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” *arXiv preprint arXiv:2005.14165*, 2020.
- [57] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, and T. Scialom, “Toolformer: Language models can teach themselves to use tools,” *arXiv preprint arXiv:2302.04761*, 2023.
- [58] D. Sur’is, S. Menon, and C. Vondrick, “ViperGPT: Visual inference via Python execution for reasoning,” *arXiv preprint arXiv:2303.08128*, 2023.
- [59] S. G. Patil, T. Zhang, X. Wang, and J. E. Gonzalez, “Gorilla: Large language model connected with massive APIs,” *arXiv preprint arXiv:2305.15334*, 2023.
- [60] M. Nye, A. Andreassen, G. Gur-Ari, H. Michalewski, J. Austin, D. Bieber, D. Dohan, A. Lewkowycz, M. Bosma, D. Luan, C. Sutton, and A. Odena, “Show your work: Scratchpads for intermediate computation with language models,” *arXiv preprint arXiv:2112.00114*, 2021.
- [61] A. J. H. Nam, M. Ren, C. Finn, and J. L. McClelland, “Learning to reason with relational abstractions,” *arXiv preprint arXiv:2210.02615*, 2022.
- [62] L. Du, H. Mei, and J. Eisner, “Autoregressive modeling with lookahead attention,” *arXiv preprint arXiv:2305.12272*, 2023.
- [63] L. Wong, G. Grand, A. K. Lew, N. D. Goodman, V. K. Mansinghka, J. Andreas, and J. B. Tenenbaum, “From word models to world models: Translating from natural language to the probabilistic language of thought,” *arXiv preprint arXiv:2306.12672*, 2023.
- [64] J. Xiang, T. Tao, Y. Gu, T. Shu, Z. Wang, Z. Yang, and Z. Hu, “Language models meet world models: Embodied experiences enhance language models,” *arXiv preprint arXiv:2305.10626*, 2023.