

Notes-to-Self: Scratchpad Augmented VLAs for Memory Dependent Manipulation Tasks

Sanjay Haresh¹, Daniel Dijkman¹, Apratim Bhattacharyya¹, Roland Memisevic¹

Abstract—Many dexterous manipulation tasks are non-Markovian in nature, yet little attention has been paid to this fact in the recent upsurge of the vision-language-action (VLA) paradigm. Although they are successful in bringing internet-scale semantic understanding to robotics, existing VLAs are primarily “stateless” and struggle with memory-dependent long horizon tasks. In this work, we explore a way to impart both spatial and temporal memory to a VLA by incorporating a language scratchpad. The scratchpad makes it possible to memorize task-specific information, such as object positions, and it allows the model to keep track of a plan and progress towards subgoals within that plan. We evaluate this approach on a split of memory-dependent tasks from the ClevrSkills environment, on MemoryBench, as well as on a challenging real-world pick-and-place task. We show that incorporating a language scratchpad significantly improves generalization on these tasks for both non-recurrent and recurrent models.

I. INTRODUCTION

The advent of LLMs and VLMs has started the paradigm of scaling learning to large data, leading to human-like generalization on a myriad of tasks. Vision-language-action (VLA) [1], [2], [3], [4] models are the instantiation of that paradigm in Robotics. VLAs leverage the rich internet-scale representations learned by VLMs and adapt them to predict actions in the robotics domain. These models are trained end-to-end to take language instructions and camera images as inputs and to predict low-level actions to be performed by the robot.

VLAs trained on a large amount of robotics data [5] have shown great promise in generalization to new tasks, robot morphologies, objects and lighting conditions. However, despite this progress, the models are still largely “reactive” in nature, as they are defined to map a visual input at any given timestep to an action independently of past observations or actions. This makes the policies fail in memory-dependent tasks, where the policy either needs to make sense of the evolving state of the world over a long time-horizon or its own past actions.

Accordingly, the impressive recent progress in vision-language-action modeling has been confined almost exclusively to Markovian tasks, where the current state of the environment suffices to predict the next action. However, many tasks in the real world require temporal or spatial memory to be solved. Consider as a simple example of the task of swapping two objects: To swap the positions of the objects, a model first needs to remember their initial positions. Then, once it has moved one of them, it needs to

remember which object it moved and which one it needs to move next. Contemporary VLAs fail on this task, as they are a) unable to remember the object positions, and b) unable to remember which sub-task they have already completed and what to do next.

In this work, we investigate how VLAs can be endowed with a memory of their past interactions to succeed in complex, multi-stage tasks. We introduce a simple yet effective strategy of training VLAs with a language scratchpad. The core idea is to leverage the innate language capabilities of the underlying VLM. Before executing an action, the model first predicts a textual description of what it intends to do. This “thought” is then stored and provided as part of the input context for all subsequent steps, creating an explicit, evolving memory trace of its own behavior. This allows the model to condition future actions not just on the current visual input and the original instruction, but also on a rich, semantic history of its own actions. Since the thoughts generated by VLMs can be very expressive descriptions of the environment, the language scratchpad not only imparts temporal memory but also spatial memory.

We evaluate our approach on three kinds of memory-dependent tasks. First, we introduce ClevrSkills-Mem, a split of memory-dependent tasks from the ClevrSkills [6] environment aimed at evaluating both temporal and spatial memory capabilities of robotic policies. We show that equipping either recurrent or non-recurrent policies with language scratchpad can lead to significant improvements in performance across all these tasks. Second, we evaluate the approach on MemoryBench [7] and show that a scratchpad allows task-agnostic generalist VLAs to achieve performance similar to highly specialized, task-specific methods. Third, we show results on a challenging *real-world* memory-dependent pick-and-place task, demonstrating that a scratchpad can enable a VLA to perform this task successfully. On all tasks, our experiments show significant performance gains with incorporation of a language scratchpad in VLAs.

In summary, our contributions are as follows:

- We present a simple memory mechanism using language based scratchpad, that allows us to train a VLA to create a plan of actions before performing them and to use it to condition all next actions.
- To test this approach, we present a memory dependent split of the ClevrSkills benchmark including five tasks (Touch-Reset-Pick, Place-Next-to-and-Restore, Swap, Stack-and-Topple and Rotate-Restore) where the first two are new tasks, and the last three are taken from [6].
- We show that both recurrent and non-recurrent poli-

¹Qualcomm AI Research

*Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

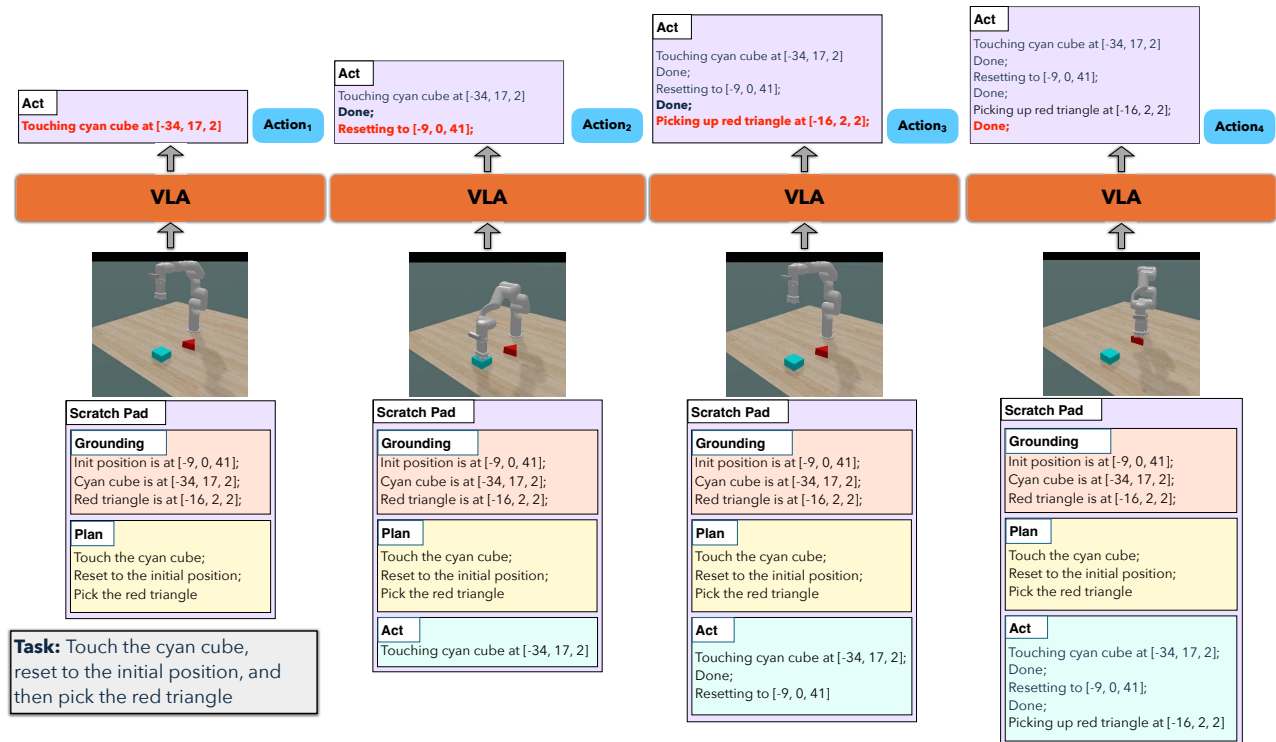


Fig. 1. Scratchpad-augmented VLAs. The VLA generates and updates scratchpad which is stored and provided as part of the input context for all subsequent steps, creating an explicit, evolving memory trace of its own behavior.

cies show significant improvement in performance on memory dependent tasks ($\approx 48\%$ and $\approx 11\%$ for non-recurrent and recurrent VLAs respectively).

II. RELATED WORK

a) Vision Language Action models.: Recent efforts in robotics have focused on creating general-purpose policies by adapting large-scale models. Early approaches involved training large transformer architectures on extensive robotics datasets, pioneered by models like Gato [8], which handled a mix of vision, language and robotics tasks, and RT-1 [1], which focused on robotics data. This trend continued with open-source policies like Octo [9] and JAT [10], which are designed for efficient fine-tuning.

More recently, the trend has shifted to fine-tuning pre-trained Vision-Language Models (VLMs) to leverage their powerful generalization capabilities. This is exemplified by the closed-source RT-2 [2] and similar open-source efforts like OpenVLA [3] and π_0 and $\pi_{0.5}$ [4], [11]. More recent work has explored various ways to incorporate spatial awareness [12], 3D perception [13], streaming and recurrence [14], [15], [16], [6], and embodied chain-of-thought reasoning [17], [18] in these VLAs to improve generalization and efficiency. However, incorporating memory into VLAs has been largely missing from the literature. In this work, we explore a way to incorporate memory in these models to enable them to solve memory-dependent tasks.

b) Memory in robotics.: Memory is crucial in enabling robots to achieve complex tasks [19]. A large portion of

the navigation literature has explored incorporating spatial memory by building explicit maps of the environment [20], [21], [22] or attention based memory modules [23], [24]. [25] also explored emergence of maps in recurrent navigation architectures. Recent works leveraging vision-language models use 3D representations like voxel-maps, neural-fields and gaussian splats to encode spatial features of the environment [26], [27]. Alternative frameworks have explored video tracking as a mechanism to reason about unobserved, disappearing and reappearing objects in the scene [28]. Very recent work SAM2Act [7] explores incorporating a memory bank and a memory attention mechanism in a multi-view transformer architecture enabling it to perform tasks requiring spatial memory. In contrast, we leverage the flexibility of language to encode both temporal and spatial context and the inherent capability of VLAs to predict language to endow them with memory. While being very flexible, our framework is also compatible with state-of-the-art reasoning models in both robotics [17], [29] and vision-language modeling [30].

c) Recurrent policies in robotics.: Recent Vision-Language-Action (VLA) models, such as Octo [9], RoboVLMs [31], and Interleave-VLA [32], have successfully processed robotic video data using an interleaved image-text format. A key limitation of these approaches, however, is their reliance on a fixed-size context window, which restricts their memory to only the most recent observations. While recurrent architectures present a natural alternative for maintaining long-term memory, prior works like RoboMamba [15], RoboFlamingo [14], and ClevrSkills [6] have

primarily explored them as an efficient substitute for transformers rather than as a solution for long-horizon memory. Our work builds on efforts to train recurrent policies like XLSTM [16] on large-scale robotics data. We demonstrate that as task horizons lengthen, these recurrent policies still struggle to retain crucial task-relevant information. Critically, we find that the performance of recurrent networks can be substantially improved by incorporating an explicit separate memory mechanism that aids in information retention over temporally extended tasks. Our findings echo more general observations made in the recurrent network community about the necessity of memory-augmentation in tasks that go beyond simple state tracking or regular language modeling (see, for example, [33]).

III. METHOD

In this section, we introduce our scratchpad-augmented VLAs and how to generate data to train them.

Preliminaries. We leverage Vision-Language-Action (VLA) models as our base models. VLAs generally start from a pre-trained VLM and are finetuned on robotics data to produce robot actions auto-regressively. More concretely: given current observation o_t , and a language instruction l , VLAs model the distribution $p(a_t|o_t, l)$ where a_t is the action at a given timestep. Actions are often discretized into 256 bins per dimension and mapped to tokens in the VLMs vocabulary [3].

More recent VLAs, however, are also able to reason about tasks in language [17]. More specifically, along with each action, these VLAs are able to predict a language description d of the task and the environment leading to a model $p(a_t, d_t|o_t, l)$. The description can be as simple as stating the current task or as descriptive as including object positions and descriptions of low level actions, such as the direction of the end-effector [17].

A. Scratchpad-augmented VLAs.

While they are powerful, current state of the art VLAs [3], [4], [11] are stateless i.e. they look at the current observation and predict the next action without taking any history into account. This renders them ineffective in memory-dependent tasks, where the policy needs to remember its own actions or the past state of the environment. In this work, we propose the use of an external scratchpad of language descriptions $S_t = \{d_1, d_2, \dots, d_n\}$ as a means to enable an existing VLA to solve such tasks. Formally, along with the current observation o_t and the language instruction l , the policy takes in the scratchpad S_t as the input, and it produces an action and a description d_t as the output, leading to a model of the conditional distribution $p(a_t, d_t|o_t, S_t, l)$.

Next, we describe our proposed scratchpad in more detail. The scratchpad is updated by the VLA when necessary, i.e. when a sub-task is successfully completed by the VLA. Specifically, the VLA predicts a special update token that triggers the next description to be included in the scratchpad as seen in Alg. 1. In practice, the model is trained to recognize the completion of the current sub-task and produce

Algorithm 1 Scratchpad Update Logic

Require:

- p : The policy.
- o_t : The observation from the environment at timestep t .
- S_t : The scratchpad at timestep t .
- l : The high-level language instruction or goal.

```

1: function FORWARDPASS( $p, o_t, S_t, l$ )
2:   ( $a_t, d_t$ )  $\leftarrow$   $p(o_t, S_t, l)$   $\triangleright$  Generate action and
   description
3:   if  $d_t \in d^{\text{update}}$  then  $\triangleright$  Check if the description
   triggers an update
4:      $S_{t+1} \leftarrow S_t \cup d_t$   $\triangleright$  Update scratchpad by adding
   the description
5:   else
6:      $S_{t+1} \leftarrow S_t$   $\triangleright$  Keep scratchpad unchanged
7:   end if
8:   return ( $a_t, S_{t+1}$ )  $\triangleright$  Return the action and the new
   scratchpad
9: end function

```

a special $\langle done \rangle$ token which triggers the description to be updated in the scratchpad. Thus, conditioning our proposed scratchpad lends the model a language defined memory. Since the descriptions can be a detailed description of the environment state, the scratchpad acts a flexible and extendable memory, allowing the model to perform memory dependent tasks. An example instantiation of a scratchpad augmented VLA can be seen in Fig. 1.

B. Generating scratchpad data

Memory dependent tasks mainly require two types of memory that we shall refer to as *Spatial Memory* and *Temporal Memory*, respectively. Here, Spatial memory refers to where objects are placed in the environment and how they change positions during the execution of the task, whereas Temporal memory refers to what subtasks are being carried out by the robot. We generally define the scratchpad to include **grounding**, **plan** and **act** sections. The **grounding** section includes the initialization conditions of the task, which usually include the object positions and the end-effector position (spatial memory), the **plan** section includes the sub-tasks that the policy needs to carry out to complete the task, and the **act** section is used to track current progress by holding the sub-tasks that have been completed by the policy (temporal memory). An example instantiation of a scratchpad for a Touch-Reset-Pick task can be seen in Fig. 1. Note that we design our scratchpad to be flexible, in that, it does not assume a specific type of position coordinates or specific type of semantic action labels. The positions can be either 3D coordinates in the environment or 2D pointers in the image space. Similarly, action labels can but do not have to carry semantic meaning for the scratchpad to be effective.

C. Extending scratchpad to Recurrent VLAs

In addition to an explicit language defined memory via a scratchpad, we can further impart implicit memory to VLAs by using a recurrent backbone which is largely missing from the state-of-the-art literature. Therefore, we also experiment with VLAs with recurrent backbones to highlight the effectiveness of using scratchpad with such models. As recurrent methods inherently include memory, at first glance, it is not clear if they require an explicit scratchpad. However, in case of tasks with long time horizons, it becomes increasingly difficult to train recurrent models on full-length trajectories. Training of recurrent networks on long trajectories is challenging not only due to GPU memory requirements, but also due to training instability and vanishing gradients. To alleviate the first challenge, recurrent networks are commonly trained on sub-sampled trajectories with lengths small enough to fit the GPU memory which, however, makes it impossible to retain initial environment conditions in long-horizon memory-dependent tasks. To adapt scratchpad to recurrent models, we convert our ground truth trajectories into interleaved sequences of language instructions l , observations o , action a and descriptions d . Concretely, our recurrent models are trained on sequences $Q \in [l, d_1, o_1, a_1, o_2, a_2, \dots, d_k, o_n, a_n]$ all represented in text with a simple next token prediction task. We use special $\langle image \rangle$ tokens to represent observation positions in the sequences and replace these with the image embeddings from a vision head after tokenization.

IV. CLEVRSKILLS-MEM BENCHMARK

While there is no dearth of robotic simulators and benchmarks, memory dependent tasks are largely missing from these. Memory dependent tasks violate the Markov assumption prevalent in robotics, according to which the current state of the environment and action are sufficient to determine the next state. To rigorously evaluate our approach, we leverage the ClevrSkills [6] benchmark and propose the ClevrSkills-Mem split, which includes tasks that require varying levels of information retention to be solved successfully. ClevrSkills is based on the ManiSkill2 [34] simulator, and it includes oracle solvers to allow generation of ground truth trajectories. It adopts a vacuum gripper and uses simple objects, which allows us to focus on the memory aspect of the tasks rather than manipulation complexity of objects.

ClevrSkills-Mem consists of 5 simple memory-dependent tasks:

- 1) **Touch-reset-pick**: the environment is initialized with 2-3 objects and the agent is required to touch a specific object, reset to the initial position and then pick another object. The task requires simple temporal memory, as the agent, once reset, does not know if it needs to do the first sub-task of touching an object or the second of picking an object.
- 2) **Place-next-to and restore**: the environment is initialized with 2 or more objects and the agent is required to place a specific object near another object. The

task evaluates spatial memory, as the agent needs to remember the initial position of the object to be able to restore it to the correct position.

- 3) **Stack-and-topple**: the environment is initialized with 2-4 objects and the agent is required to stack them in a specific order and then topple the stack. The long-horizon task requires temporal memory to discern which part of the sub-task is being solved at any instance.
- 4) **Swap**: the environment is initialized with 2 or more objects and the agent is required to swap the positions of 2 specific objects. The task evaluates both spatial and temporal memory, as the agent needs to remember initial positions of the objects to be able to swap them correctly, and it furthermore needs to remember which phase of the task it is in to carry out the next sub-task.
- 5) **Rotate-restore**: the environment is initialized with 2 or more objects, and the agent is required to rotate a specific object by a pre-defined amount and subsequently restore it to the original position. The task requires very fine-grained temporal memory, as once the agent starts rotating the object, it needs to keep track of how much it has rotated it relative to the initial position.

Fig. 2 shows some examples of these tasks.

V. EXPERIMENTS

We evaluate scratchpad augmented VLAs in two simulation environments (ClevrSkills [6] and MemoryBench [7]) and on a real-world task using a UFACTORY xArm 6.

A. Experimental setup

We use language description for task specification across all models and benchmarks. We resize all images to 224×224 before feeding them to the model. We use end-effector pose deltas as our action space. Actions are 7-dimensional vectors that include 6DOF pose deltas and 1D gripper state. We experiment with two types of VLA model: a transformer based VLA (called **T-VLA** from here on) which takes a single time-slice as input and a recurrent VLA (called **R-VLA** from here on) which autoregressively generates the whole trajectory.

For the T-VLA, we linearize the scratchpad and append it with the prompt during training. We denote different parts of the scratchpad using special tokens. The initialization and plan part of the scratchpad are enclosed in $\langle plan \rangle \dots \langle /plan \rangle$ tokens. The model is trained to predict the current sub-activity with $\langle think \rangle$ tokens followed by the action where the start is denoted by the $\langle act \rangle$ token. End points of a specific subtask are denoted by the $\langle done \rangle$ token. During evaluation, we let the model predict the $\langle plan \rangle \dots \langle /plan \rangle$ on the first observation and then follow the algorithm in Alg. 1 to progressively update the scratchpad every time a $\langle done \rangle$ token is predicted which act as the updated context for the next action.

For R-VLA, we interleave the scratchpad and actions in a single long text sequence which follows the format, $[obs_0; plan_0; think_0; act_0; obs_1; act_1; \dots; think_1; act_1 \dots]$.




Fig. 2. Example memory dependent tasks from ClevrSkills-Mem we evaluate on.

Here, observation denotes the image tokens, plan denotes the grounding and plan part of the scratchpad, followed by observation, thought and action triples at each timestamp. The model is trained on a simple next token prediction task and also evaluated autoregressively in a similar manner. Practically, we found R-VLA to struggle with discretized actions and therefore use a simple MLP as an action head where we represent the action with a $\langle action \rangle$ token in text and use the embeddings from the action token as input to the action head which then produces the 7-dimensional real valued vectors for actions.

For T-VLA, we start from the PaliGemma-2 VLM model with 3B parameters [35], which is based on the Gemma-2 LLM [36] and on the SigLIP vision encoder [37]. We perform full-finetuning of the model with a batch size of 8 per GPU and a learning rate of 2×10^{-5} , using the Adam optimizer [38]. For R-VLA, we start from Mamba [39] language model with 130M parameters and use a ViT [40] as an image backbone. We also experimented with 1B and 3B versions of Mamba model but found little to no performance difference and therefore selected the 130M version as the base language model. We attribute this to the fact that Mamba is not a vision-language model and therefore bigger models do not have better pre-trained visual representations as in

other vision-language models. We perform full-finetuning of the model with a batch size of 4 per GPU and a learning rate of 10^{-5} using the Adam optimizer [38]. All experiments were performed using PyTorch [41] on 4 Nvidia A100 GPUs.

B. Results on ClevrSkills-Mem

The ClevrSkills [6] benchmark is based on the ManiSkill2 [34] manipulation environment, and it includes an oracle solver to generate data including actions, object positions, and solver traces which are used to get sub-task segmentation to create scratchpad. We evaluate the models on a range of tasks that require varying levels of spatial and temporal memory. Specifically, Touch-reset-pick and Stack-and-topple require strong temporal memory, Swap and Place-next-to-restore require strong spatial memory and Rotate-restore require fine-grained low-level temporal memory. Note that tasks requiring spatial memory have a success criterion of $5cm$ delta, i.e. the final object position should not be more than $5cm$ away from the required placement. Also note that the objects can be placed anywhere in the workspace and are not located at fixed positions as common in other benchmarks.

We evaluate all models on 50 rollouts of unseen starting positions of the objects seen during training. The main results on ClevrSkills can be seen in Fig. 3. As we can see, due to




Fig. 3. **Results on the ClevrSkills-Mem benchmark.** We show results for all the models considered with and without scratchpad on 5 tasks of ClevrSkills-Mem benchmark. We report success rate of each model on 50 rollouts on unseen starting positions objects. On the right we show the average performance across all tasks.

the challenging nature of the memory-dependent ClevrSkills tasks, T-VLA struggles on all the tasks, and adding the scratchpad significantly improves the performance across all the tasks. More specifically, adding the scratchpad improves the performance by 68%, 72%, 68% and 30% on Touch-reset-pick, Swap, Place-next-to-restore and Stack-and-topple, respectively, leading to an average performance gain of 48.8% across all the tasks. We note that T-VLA+Scratchpad still fails to achieve any success on rotate task as it requires low-level fine-grained temporal memory.

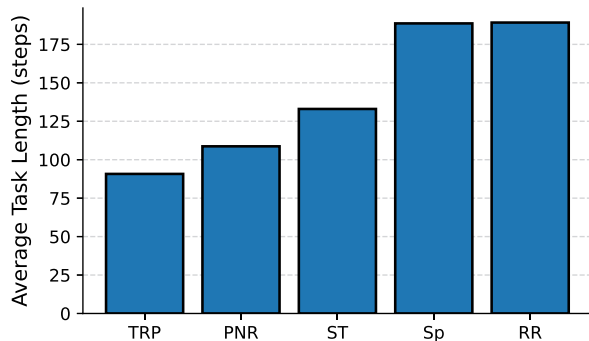


Fig. 4. Average trajectory length of tasks in ClevrSkills-Mem. Here, TRP denotes Touch-Reset-Pick, PNR denotes Place-Next-to-Restore, ST denotes Stack-and-Topple, Sp denotes Swap and RR denotes Rotate-Restore task respectively.

We also experiment with a recurrent VLA (R-VLA) with Mamba as backbone to compare against a model with inherent memory. While R-VLA usually performs better than T-VLA, we note that T-VLA+Scratchpad is able to match R-VLA performance on average even surpassing it on some tasks. This shows that language scratchpad is flexible enough to endow strong memory capabilities to stateless VLAs. We also note that recurrent VLAs despite having inherent memory capabilities also see an average improvement of 11% when trained with scratchpad. In Fig. 4, we show the average trajectory length of the tasks in ClevrSkills. As we can see, generally increasing task length leads to a higher performance improvement with scratchpad in R-VLAs with exception of Rotate-Restore which see no performance

TABLE I
EVALUATION ON PUT-BLOCK-BACK TASK FROM MEMORYBENCH [7].

Models	Put Block Back
RVT-2 [43]	50.0
SAM2Act [7]	35.0
SAM2Act+ [7]	100.0
T-VLA	0
T-VLA (sim-eval)	0
T-VLA w/ SP (Our)	40.0
T-VLA w/ SP (Our) (sim-eval)	100.0

improvement as the task requires low-level fine-grained temporal memory which can not be aided by a scratchpad.

C. Results on MemoryBench

To show the flexibility and ability to generalize, we also evaluate the models on the recently proposed MemoryBench [7]. MemoryBench is built on top of RLbench [42], and it includes 3 tasks: Put-Block-Back, Rearrange-Block, and Reopen-Drawer. MemoryBench differs from ClevrSkills in that objects are placed at fixed position areas and variations are generated based on which area the objects are placed on. For example in Put-Block-Back, the cube can start at 1 of the 4 starting positions and the task is to move the cube to the middle and then press a button and move the cube back to the initial position. This task requires both temporal and spatial memory, as a policy needs to remember if it has pressed the button or not (temporal) and what the initial position of the object was (spatial). Note that along with memory, MemoryBench also requires very precise manipulation capabilities, especially for pressing the button, as it is very small compared to the scene and therefore very easy to miss. We focus on the Put-Block-Back task from MemoryBench as it encompasses both types of memory.

To evaluate on MemoryBench, we use the open-sourced 100 trajectories to train our model and evaluate on the 25 trajectories as done in prior work [7]. Since MemoryBench requires very precise manipulation, we also use the wrist image as input during training. We follow the prior work to predict the key-frame poses of the end-effector as actions.

We show performance of our T-VLA and T-VLA+Scratchpad on MemoryBench in Table I. Prior




Fig. 5. Key frames of the real world task: Pick the tomato, place it in the bowl and then restore it to the initial position.

baselines evaluated on MemoryBench all take scene point clouds as input and are driven towards very precise manipulation whereas we just take front RGB and wrist RGB images as input. As we can see, scratchpad significantly improves performance over the stateless baseline. We also note that since we take RGB images as input as opposed to 3D point clouds, our model struggles to press the button and majority of the failure cases stem from this failure mode. Therefore, we also report a relaxed evaluation criteria (denoted by “sim-eval” in Table I) where if the VLA gets close to the button, we execute the ground-truth action to press that button instead of taking the action predicted by the VLA and continue the rest of the trajectory. Doing so leads to our model getting a 100% success rate on the Put-Block-Back task showing that our model is able to perform both temporal and spatial recall perfectly.

D. Results in the real world

We also evaluate scratchpad augmented VLAs on a real world task, Pick-Place-Restore, where we task the robot to pick an object (a tomato), place it in a bowl and then restore its position to the original position. We show an example trajectory of the task in Fig. 5. Note that we do not assign fixed starting positions for the tomato and the bowl and therefore both of them can be located anywhere in the workspace. To be successful at the task, the policy needs to remember what sub-task it is currently performing (temporal memory), as placing the object in the bowl and restoring it to the original position are both “reverse” of each other. This leads to a temporal ambiguity when looking at a single time-slice at a time. The policy also needs to remember the starting position of the object in the workspace (spatial memory) for it to be successful at restoring it to the original position.

We collected a dataset comprising 200 trajectories from the Pick-Place-Restore task by tele-operating a UFactory xArm 6 with a flexible two-fingered gripper, operating on a white tabletop. The agent observes the environment through RGB images captured by a RealSense D435 camera, positioned at a corner of the table. We compare a pretrained OpenVLA model with 7B parameters [3] to the scratchpad-augmented version of the same model. We perform LoRA finetuning with rank 32 [44], a batch size of 8 and a learning rate of 5×10^{-4} , using the Adam optimizer [38].

We present the results in Table II. The table shows the average success rate, average sub-task completion rate

TABLE II
EVALUATION ON THE REAL WORLD TASK.

Models	Avg. Success	Sub-task CR	Avg. Dist.
Human	100%	3	3.2cm
OpenVLA	0%	0.9	N/A
OpenVLA w/ SP (Our)	65%	2.4	10.62cm

(denoted by Sub-task CR) and the average distance between the starting position of the object and the final restored position (denoted by Avg. Dist.). We evaluate each model on 20 rollouts, where the starting position of the objects are randomized every time. As we can see, the VLA without scratchpad fails on this task, as it has no memory of what it is doing. Qualitatively, we note that it generally is able to pick the object but restores it back without putting it in the bowl as it can not discern between the start and end of the trajectory as both are roughly identical in the dataset. In contrast, the scratchpad augmented model is able to complete all the subtasks 65% of the time demonstrating the memory capacity of the scratchpad. We note that the average replacement distance for the model is relatively high as compared to the ground truth trajectories. This can be attributed to the limited amount of training data (200 examples) given the large variability in the starting positions.

VI. CONCLUSION

In this work, we present a simple language scratchpad mechanism that allows us to endow existing VLAs with a flexible and strong memory mechanism enabling them to solve memory-dependent tasks. We evaluate the approach in both simulation and in the real world. As the number of existing simulation environments for testing long-range memory is limited, we also introduce the ClevrSkills-Mem split within the ClevrSkills environment, consisting of existing and novel tasks. Together these tasks test for a spectrum of different memory-capabilities with respect to spatial and temporal information. Experiments show that scratchpad-augmented VLAs significantly outperform contemporary stateless VLAs for memory-dependent tasks.

REFERENCES

- [1] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.

- [2] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control. arxiv,” *arXiv preprint arXiv:2307.15818*, 2023.
- [3] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, *et al.*, “Openvla: An open-source vision-language-action model,” *arXiv preprint arXiv:2406.09246*, 2024.
- [4] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, *et al.*, “ π_0 : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.
- [5] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. Jain, *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6892–6903.
- [6] S. Haresh, D. Dijkman, A. Bhattacharyya, and R. Memisevic, “Clevrskills: Compositional language and visual reasoning in robotics,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 38 235–38 266, 2024.
- [7] H. Fang, M. Grotz, W. Pumacay, Y. R. Wang, D. Fox, R. Krishna, and J. Duan, “Sam2act: Integrating visual foundation model with a memory architecture for robotic manipulation,” *arXiv preprint arXiv:2501.18564*, 2025.
- [8] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, *et al.*, “A generalist agent,” *arXiv preprint arXiv:2205.06175*, 2022.
- [9] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, C. Xu, J. Luo, T. Kreiman, Y. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine, “Octo: An open-source generalist robot policy,” in *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024.
- [10] Q. Gallouédec, E. Beeching, C. Romac, and E. Dellandréa, “Jack of all trades, master of some, a multi-purpose transformer agent,” *arXiv preprint arXiv:2402.09844*, 2024.
- [11] K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, M. Y. Galliker, *et al.*, “ $\pi_0.5$: a vision-language-action model with open-world generalization,” *arXiv preprint arXiv:2504.16054*, 2025.
- [12] D. Qu, H. Song, Q. Chen, Y. Yao, X. Ye, Y. Ding, Z. Wang, J. Gu, *et al.*, “Spatialvla: Exploring spatial representations for visual-language-action model,” *arXiv preprint arXiv:2501.15830*, 2025.
- [13] H. Zhen, X. Qiu, P. Chen, J. Yang, X. Yan, Y. Du, Y. Hong, and C. Gan, “3d-vla: A 3d vision-language-action generative world model,” *arXiv preprint arXiv:2403.09631*, 2024.
- [14] X. Li, M. Liu, H. Zhang, C. Yu, J. Xu, *et al.*, “Vision-language foundation models as effective robot imitators,” *arXiv preprint arXiv:2311.01378*, 2023.
- [15] J. Liu, M. Liu, Z. Wang, P. An, X. Li, K. Zhou, S. Yang, R. Zhang, Y. Guo, and S. Zhang, “Robomamba: Efficient vision-language-action model for robotic reasoning and manipulation,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 40 085–40 110, 2024.
- [16] T. Schmied, T. Adler, V. Patil, M. Beck, K. Pöppel, J. Brandstetter, G. Klambauer, R. Pascanu, and S. Hochreiter, “A large recurrent action model: x lstm enables fast inference for robotics tasks,” *arXiv preprint arXiv:2410.22391*, 2024.
- [17] M. Zawalski, W. Chen, K. Pertsch, O. Mees, C. Finn, and S. Levine, “Robotic control via embodied chain-of-thought reasoning,” *arXiv preprint arXiv:2407.08693*, 2024.
- [18] X. Li, C. Mata, J. Park, K. Kahatapitiya, Y. S. Jang, J. Shang, K. Ranasinghe, R. Burgert, M. Cai, Y. J. Lee, *et al.*, “Llara: Supercharging robot learning data for vision-language policy,” *arXiv preprint arXiv:2406.20095*, 2024.
- [19] S. Jockel, M. Weser, D. Westhoff, and J. Zhang, “Towards an episodic memory for cognitive robots,” in *Proc. of 6th Cognitive Robotics workshop at 18th European Conf. on Artificial Intelligence (ECAI)*, 2008, pp. 68–74.
- [20] T. Bailey and H. Durrant-Whyte, “Simultaneous localization and mapping (slam): Part ii,” *IEEE robotics & automation magazine*, vol. 13, no. 3, pp. 108–117, 2006.
- [21] D. S. Chaplot, D. P. Gandhi, A. Gupta, and R. R. Salakhutdinov, “Object goal navigation using goal-oriented semantic exploration,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 4247–4258, 2020.
- [22] J. Fuentes-Pacheco, J. Ruiz-Ascencio, and J. M. Rendón-Mancha, “Visual simultaneous localization and mapping: a survey,” *Artificial intelligence review*, vol. 43, no. 1, pp. 55–81, 2015.
- [23] L. Mezghan, S. Sukhbaatar, T. Lavril, O. Maksymets, D. Batra, P. Bojanowski, and K. Alahari, “Memory-augmented reinforcement learning for image-goal navigation,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3316–3323.
- [24] K. Fang, A. Toshev, L. Fei-Fei, and S. Savarese, “Scene memory transformer for embodied agents in long-horizon tasks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 538–547.
- [25] E. Wijmans, M. Savva, I. Essa, S. Lee, A. S. Morcos, and D. Batra, “Emergence of maps in the memories of blind navigation agents,” *AI Matters*, vol. 9, no. 2, pp. 8–14, 2023.
- [26] H. Huang, F. Lin, Y. Hu, S. Wang, and Y. Gao, “Copa: General robotic manipulation through spatial constraints of parts with foundation models,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 9488–9495.
- [27] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “Voxposer: Composable 3d value maps for robotic manipulation with language models,” *arXiv preprint arXiv:2307.05973*, 2023.
- [28] Y. Huang, J. Yuan, C. Kim, P. Pradhan, B. Chen, L. Fuxin, and T. Hermans, “Out of sight, still in mind: Reasoning and planning about unobserved objects with video tracking enabled memory models,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3108–3115.
- [29] D. Kim, S. Park, H. Jang, J. Shin, J. Kim, and Y. Seo, “Robot-r1: Reinforcement learning for enhanced embodied reasoning in robotics,” *arXiv preprint arXiv:2506.00070*, 2025.
- [30] M. Nye, A. J. Andreassen, G. Gur-Ari, H. Michalewski, J. Austin, *et al.*, “Show your work: Scratchpads for intermediate computation with language models,” *arXiv preprint arXiv:2112.00114*, 2021.
- [31] H. Liu, X. Li, P. Li, M. Liu, D. Wang, J. Liu, B. Kang, X. Ma, T. Kong, and H. Zhang, “Towards generalist robot policies: What matters in building vision-language-action models,” *arXiv preprint arXiv:2412.14058*, 2025.
- [32] C. Fan, X. Jia, Y. Sun, Y. Wang, J. Wei, Z. Gong, X. Zhao, M. Tomizuka, X. Yang, J. Yan, *et al.*, “Interleave-vla: Enhancing robot manipulation with interleaved image-text instructions,” *arXiv preprint arXiv:2505.02152*, 2025.
- [33] G. Deletang, A. Ruoss, J. Grau-Moya, T. Genewein, L. K. Wenliang, E. Catt, C. Cundy, M. Hutter, S. Legg, J. Veness, and P. A. Ortega, “Neural networks and the chomsky hierarchy,” in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=WbxHAzkeQcn>
- [34] J. Gu, F. Xiang, X. Li, Z. Ling, X. Liu, T. Mu, Y. Tang, S. Tao, X. Wei, Y. Yao, X. Yuan, P. Xie, Z. Huang, R. Chen, and H. Su, “Maniskill2: A unified benchmark for generalizable manipulation skills,” in *International Conference on Learning Representations*, 2023.
- [35] A. Steiner, A. S. Pinto, M. Tschannen, D. Keysers, X. Wang, *et al.*, “PaliGemma 2: A family of versatile vlms for transfer,” *arXiv preprint arXiv:2412.03555*, 2024.
- [36] G. Team, M. Riviere, S. Pathak, P. G. Sessa, C. Hardin, *et al.*, “Gemma 2: Improving open language models at a practical size,” *arXiv preprint arXiv:2408.00118*, 2024.
- [37] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, “Sigmoid loss for language image pre-training,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 11 975–11 986.
- [38] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [39] A. Gu and T. Dao, “Mamba: Linear-time sequence modeling with selective state spaces,” *arXiv preprint arXiv:2312.00752*, 2023.
- [40] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [41] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, 2019.
- [42] S. James, Z. Ma, D. Rovick Arrojo, and A. J. Davison, “Rlbench: The robot learning benchmark & learning environment,” *IEEE Robotics and Automation Letters*, 2020.
- [43] A. Goyal, V. Blukis, J. Xu, Y. Guo, Y.-W. Chao, and D. Fox, “Rvt2: Learning precise manipulation from few demonstrations,” *RSS*, 2024.
- [44] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=nZeVKeeFYf9>