

Long-Context Robot Imitation Learning by Focusing on Key History Frames

Max Sobol Mark^{1,2}, Jacky Liang², Maria Attarian², Chuyuan Fu², Debidatta Dwivedi², Aviral Kumar^{1†}, Dhruv Shah^{2,3†}

¹Carnegie Mellon University ²Google DeepMind ³Princeton University [†]Equal advising.

<https://bigpicturepolicies.github.io/>




Fig. 1. Big Picture Policies (BPP). Across a suite of non-Markovian tasks, BPP (left) uses an off-the-shelf vision-language model (VLM) to identify task-relevant keyframes from history, enabling robust performance compared to memoryless and naïve history-conditioned baselines (right).

Abstract—Many useful robot tasks require attending to the history of past observations. For example, finding an item in a room requires remembering which places have already been searched. However, the best-performing robot policies typically condition only on the current observation, limiting their applicability to such tasks. Naïvely conditioning on past observations often fails due to spurious correlations: policies latch onto incidental features of training histories that do not generalize to out-of-distribution trajectories upon deployment. In this paper, we analyze why policies latch onto these spurious correlations. We find that this problem arises because of limited coverage over the space of possible histories during training, which grows exponentially with horizon. Existing regularization techniques provide inconsistent benefits across tasks, as they do not fundamentally address this coverage problem. Motivated by these findings, we propose Big Picture Policies (BPP), an approach that conditions on a minimal set of meaningful keyframes detected by a vision-language model. By projecting diverse rollouts onto a compact representation of task-relevant events, BPP substantially reduces distribution shift between training and deployment, without sacrificing expressivity. We evaluate BPP on four challenging real-world manipulation tasks and three simulation tasks, all requiring history conditioning. BPP achieves 70% higher success rates than the best comparison on real-world evaluations.

I. INTRODUCTION

Imitation learning has become a central paradigm for training robotic policies. However, the most successful approaches today typically lack *memory*, conditioning only on the current observation [1, 2, 3, 4, 5] or a short temporal window [6, 7]. While effective on many benchmarks, such policies struggle in real-world settings where the robot must remember what it has already done, seen, or what remains to be completed.

Consider tasks such as searching for an object in a cluttered scene, adding a precise number of ingredients, or executing multi-stage procedures where early mistakes are not easily observed later. In these settings, the current observation is insufficient to determine the correct action; success instead requires reasoning over past events. Motivated by this, we study training policies to effectively leverage history.

A natural approach for learning memory-based policies is to condition policies directly on sequences of past observations. Yet, despite its appeal, prior work has consistently found that naïvely adding history often *hurts* performance [8, 9, 10, 11, 12], a phenomenon we also observe. Rather than improving robustness, history-conditioned policies frequently fail even on simple tasks. These failures arise because policies latch onto incidental features (a.k.a. “spurious correlations”) of training histories that explain expert actions but do not generalize. Once the policy deviates from the expert distribution, it encounters unseen histories and errors compound.

In theory, this problem stems from a fundamental challenge of *coverage*: the space of possible observation histories grows exponentially with horizon. Coverage is especially limited under current robotic data collection paradigms, which rely primarily on human tele-operators gathering “near-expert” demonstrations. As a result, although history-conditioned policies can predict the underlying *oracle* task state accurately on the training set, they fail catastrophically on histories encountered during deployment. Existing approaches attempt to mitigate the spurious correlations through regularization [13], auxiliary objectives [12], or architectural constraints [11].

While these techniques can help in some cases, our experiments show their benefits are inconsistent and often task-dependent. In essence, they reshape the action prediction loss without addressing the underlying lack of coverage over histories.

These observations suggest that effective history conditioning requires directly addressing the challenge of coverage. But *how can we improve coverage without collecting more data?* Our core idea is to train policies in a different input space than raw observation histories. Rather than reasoning over full observation histories, we focus on a small number of past moments that are most relevant for predicting actions. In many non-Markovian tasks, successful behavior depends on identifying a few behaviorally salient events, or *keyframes*, such as whether an object has been grasped, a subgoal has been completed, or a failure has occurred. Representing history through these semantic keyframes increases overlap between training and test-time policy inputs by reducing the effective history space, while preserving task-relevant information needed for action prediction. This enables robust generalization despite limited coverage over all histories in the data.

In this paper, we instantiate this idea in a concrete approach we call **Big Picture Policies (BPP)**, which trains robotic policies conditioned on semantic keyframes. BPP leverages an off-the-shelf vision-language model (VLM) to identify task-relevant keyframes via simple question answering, projecting long and diverse trajectories into a compact set of informative inputs for policy learning. We evaluate BPP on four challenging real-world bimanual manipulation tasks and three simulated tasks, all explicitly designed to require memory. Across all settings, BPP substantially outperforms strong history-conditioned and memoryless baselines, achieving up to 70% higher success rates in real-world evaluations. Our contributions are twofold: (i) we show that the central challenge in history-conditioned policy learning is coverage over histories, and that regularization-based approaches provide inconsistent, task-dependent improvements without resolving the core underlying issue; and (ii) we introduce **BPP**, a simple yet effective method that leverages VLMs to identify task-relevant keyframes, achieving state-of-the-art performance with improved data efficiency and reduced training time.

II. RELATED WORK

History in robot learning policies. Most high-performing robot policies predict actions conditioned only on the current observation [1, 3, 5, 14, 2, 4]. This raises a natural question: is history conditioning actually necessary? Much of robot learning work evaluates on tasks that do not require memory, though some prior work studies settings where history is required [12, 15] or beneficial [16, 7, 11]. In contrast, on-policy deep RL routinely uses memory, via frame stacking or recurrent networks, to handle partial observability [17, 18]. History is less common in robot imitation learning because adding observation histories often hurts performance [8, 9, 10, 11, 12], increasing spurious correlations and causal confusion

under rollout distribution shift [19, 20, 21]. Our goal in this paper is to address these spurious correlations.

Addressing spurious correlations. A well-studied failure mode in imitation from observation histories is the *copycat* problem [8], also known as causal confusion [19], where policies overfit to past actions implicitly encoded in observation sequences rather than responding to the current state. Prior work addressed this issue by upweighting loss on keyframes where actions change significantly [9], discovering keyframe structure for planning [22], or applying information bottlenecks to suppress past-action leakage [13]. However, copycat is not the only source of spurious correlations in long-context policies: even without explicit action histories, extended observation sequences create opportunities for policies to latch onto irrelevant temporal patterns.

Beyond loss reweighting, prior work has proposed auxiliary tasks to address spurious correlations. Torne et al. [12] introduce past-token prediction as a self-supervised objective and employ a test-time selection mechanism to enforce temporal consistency, which we find is effective only under certain task conditions. Another line of work avoids conditioning the low-level policy on raw observation histories altogether. TraceVLA [11] augments observations with a visual trace of the robot’s past trajectory, while MemER [15] retrieves relevant past experiences via a language-based memory and delegates action prediction to a memoryless policy. We argue that such approaches are limited: TraceVLA’s visual trace cannot capture object-level or scene history, and MemER’s language-only memory limits expressivity and assumes a clean sub-task decomposition, *e.g.*, if memory is required within a single sub-task, the low-level policy cannot complete it. By contrast, our method conditions the low-level policy directly on semantically meaningful frames from history, preserving expressivity while mitigating spurious correlations.

Pre-trained models in robotics. Vision-language models (VLMs) pre-trained on internet-scale data have been used for robotics in several ways. A dominant paradigm fine-tunes VLMs directly for action prediction, yielding vision-language-action (VLA) models [1, 5, 3, 6]. Other approaches keep the VLM frozen and use it as a high-level planner issuing language commands to a low-level policy [23, 24, 25], or as a reward model, success detector, or value function for reinforcement learning [26, 27, 28, 29]. Our work falls into a distinct category: using a VLM to identify semantically meaningful moments in the observation history, guiding what the policy should attend to rather than what actions it should take. While some prior work (*e.g.*, [30]) uses VLMs to propose “important” states, these are typically treated as future subgoals. To our knowledge, prior work has not used VLMs to filter observation histories to improve learning signals.

III. WHAT MAKES HISTORY-CONDITIONED IMITATION LEARNING PARTICULARLY CHALLENGING?

We first examine why history-conditioned imitation learning is difficult, and whether practices now standard in robot imitation learning, such as expressive policy classes (*e.g.*,




Fig. 2: Benchmark tasks. We evaluate on 4 real-world tasks (A–D) and 3 simulation tasks (E–G). All tasks require history conditioning for success.

diffusion [10] or flow policies [31]) with action chunking [32] can mitigate these difficulties. Conventional wisdom attributes these challenges to *spurious correlations*: when conditioned on histories, policies may latch onto features that explain training actions but fail to generalize on test-time histories. To assess this issue, we first study standard (naïve) history conditioning with a diffusion policy and action chunking, focusing on the impact of spurious correlations.

Notation. We consider partially observable decision-making problems with latent states $s_t \in \mathcal{S}$, observations $o_t \in \mathcal{O}$, and actions in $\mathcal{A}$. Because tasks are often non-Markovian with respect to o_t , we define the observation history as $h_t = (o_{t-k}, \dots, o_t)$. Given a dataset of N expert demonstrations $\mathcal{D} = \{\tau_i\}_{i=1}^N$, our objective is to learn a history-conditioned policy $\pi_\theta : \mathcal{O}^{k+1} \rightarrow \Delta(\mathcal{A})$ via imitation learning.

A. History with Action-Chunking Diffusion Policies

We first benchmark action-chunking diffusion policies on a suite of simulation tasks explicitly designed to be unsolvable without leveraging history (Figure 2):

- **Ingredient-Insertion:** Place two lemons inside a red bowl, then press a button. Once a lemon is in the bowl, it becomes invisible. This mirrors tasks like adding spoonfuls of sugar to coffee, where previous actions must be remembered as they are no longer observable in the current state.
- **Fixed-Password:** Enter a 7-digit password (1112134) on a 3×3 grid of buttons. While the password remains the same, button positions are randomized each episode. This task needs “dynamic memory”: the number of relevant history frames varies depending on the digit being entered.
- **Variable-Password:** Given a 4-digit password (represented as a one-hot vector), enter it using three available buttons. Each password contains at least one repeated digit, requiring the policy to track its progress.

To establish the lower-bound performance, we run imitation using only the current observation. We then evaluate naïve history conditioning, where policy inputs contain a fixed-length, uniformly-strided window of past observations. We chose the history length and stride so that they provide sufficient representation for the optimal policy (see Appendix D). Finally, we establish an upper bound using an *oracle* policy that has access to privileged state information, making the task solvable without history.

Results in simulation (Figure 7) reveal a dichotomy: while there is a substantial performance gap between naïve history conditioning and the oracle, the naïve method does not suffer from the catastrophic performance drop seen in prior work [19, 8, 12]. In the following section, we investigate both what prevents catastrophic failure for naïve history conditioning and the fundamental factors that limit its performance.

B. Avoiding Catastrophic Failure from Naïve Conditioning

We now ablate various design components of our policy to understand why naïve history conditioning avoids catastrophic failure, despite conventional wisdom suggesting it should fail. We hypothesize that two components play a key role: **1)** the use of action chunking, and **2)** the policy architecture. We explain our hypotheses below and perform controlled experiments to understand the role of each component.

(a) Action chunking. We hypothesize that predicting temporally extended action chunks suppresses reliance on spurious correlations in observation histories, even though these segments correspond only to future timesteps and not past ones as in Torne et al. [12]. To test this, we evaluate how well policy features encode history state (e.g., the number of buttons pressed in Fixed-Password) by training a classifier on top of the learned representations. In Figure 3 (left), policies with short action chunks accurately predict history state on i.i.d. validation trajectories, but they fail under rollout-induced distribution shift, achieving near-random accuracy. By contrast, longer action chunks yield more robust features that generalize better to policy rollouts ($2.9\times$ vs $7.2\times$ error increase). However, even with longer chunks, a substantial gap remains between expert and rollout performance, indicating that action chunking mitigates but does not eliminate the issue.

These representational differences directly translate into performance. Figure 3 (right) shows that while shortening the predicted chunk length has little effect on the oracle policy, it causes a substantial performance decline in the naïve history-conditioned policy. This indicates that action chunking plays a crucial role in mitigating spurious correlations: by enforcing longer-horizon action prediction, it forces the policy to learn features that are more robust to history shift.

(b) Policy architecture. A common strategy for reducing the cost of long-history training is a multi-stage recipe: the visual encoder is trained with short contexts and then frozen [12]. In




Fig. 3: Predicting longer action chunks significantly decreases spurious correlations in naïve history-conditioned policies. Left: We compare how well the features learned by the policy predict the history state (in Fixed-Password, this is the number of buttons pressed so far). Both policies trained with shorter (10) and longer (50) action chunks learn features that predict history state well on in-distribution expert trajectories. However, policies trained with shorter action chunks exhibit a much larger generalization gap to out-of-distribution policy rollouts ($7.2\times$ vs. $2.9\times$ decrease), indicating greater reliance on spurious correlations. Right: These representational differences translate into significant performance gaps. Shorter action chunks drastically hurt naïve history-conditioned policy performance, while having little effect on oracle policies. Notably, training with action chunk 50 but executing with chunk 10 still significantly outperforms training with chunk 10, confirming that longer chunk prediction improves feature learning for history conditioning.

Appendix A, we find that this substantially degrades history conditioning, suggesting that encoder inflexibility contributes to the poor performance observed by prior work [12]. By contrast, our approach trains a single encoder jointly across all timesteps, allowing representations to adapt to extended histories and yielding better performance.

C. Why Does the Performance Gap Persist?

While appropriate architectural choices prevent naïve history conditioning from failing catastrophically, a large performance gap remains compared to the oracle (Figure 7). Recall that we found that even with long action chunking, the policy’s understanding of the underlying state degrades significantly on its own rollouts. This raises a critical question: can we solve this problem by simply adding a better auxiliary regularizer?

To test this, we construct a “golden” auxiliary task: explicitly regularizing the visual encoder to predict the ground-truth history state (e.g., the number of buttons pressed so far in Fixed-Password). Since this loss directly forces the encoder to extract the exact information needed to solve the task, it provides a strong stress-test for auxiliary regularizers. As shown in Figure 4, adding this auxiliary loss increases accuracy on in-distribution validation data. However, surprisingly, it leads to a decrease in both the understanding of history state on out-of-distribution rollouts (from 86% to 18%) and the policy success rate on deployment (from 56% to 19%).

This negative result highlights a critical limitation of history conditioning on expert demonstrations: even when the learned features can make correct state predictions on the training distribution, these features become spurious under distribution shift if the training data lacks sufficient coverage of corrective behaviors. This points to *coverage*, rather than feature learning architecture, as the fundamental bottleneck. In the next section,




Fig. 4: History state prediction regularization hurts history understanding. We evaluate whether regularizing the history encoder to predict ground-truth state information (number of buttons pressed so far in Fixed-Password) improves generalization. While this auxiliary task improves accuracy on unseen expert trajectories, it leads to worse performance on out-of-distribution rollouts, indicating it increases reliance on spurious correlations. This regularization also degrades success rate: $55.5\% \pm 3.3\%$ to $19.0\% \pm 3.3\%$.

we introduce **Big Picture Policies (BPP)**, our method that bypasses this limitation. Observe in Figure 5, while naïve and PTP suffer significant degradation in history understanding when evaluated on policy rollouts, BPP does not.

IV. A SIMPLE RECIPE FOR HISTORY-CONDITIONING IN IMITATION LEARNING

In the previous section, we showed that naïve history conditioning fails due to insufficient coverage over observation histories, so policies learn to rely on spurious correlations. Existing approaches based on regularization do not resolve this issue because they do not address the underlying mismatch between the histories observed during training and those




Fig. 5: BPP is more robust to distribution shifts from policy rollouts. We evaluate accuracy of history state prediction, averaged between fixed-password-entering and ingredient-insertion. Both Naive History Conditioning and PTP suffer big performance degradations when testing on policy rollouts. BPP (Ours), on the other hand, does not.

encountered at deployment. This motivates a simple question: *can we improve history conditioning without changing data collection protocols or training objectives?* In this section, we introduce *Big Picture Policies (BPP)*, which address the coverage challenge by changing the *representation* of history rather than the learning algorithm itself (i.e., standard imitation learning without auxiliary regularization). Instead of conditioning policies on full observation sequences, BPP conditions on a *minimal* set of task-relevant *keyframes*. By compressing long and diverse histories into a small number of keyframes, BPP increases the overlap between policy inputs at training and test time, reducing the exposure to spurious correlations.

A. Big Picture Policies (BPP)

To build intuition, consider the task of adding two spoons of sugar to a coffee. During execution, a robot may miss the spoon several times, producing observation histories that differ substantially from expert demonstrations, even though the task state remains the same. A policy conditioned on the full history must therefore implicitly learn to look for accurate information in a test-time history that is poorly covered by the data. BPP instead conditions the policy on keyframes that correspond to behaviorally salient events, in this case the moment when sugar is successfully dropped into the coffee. With keyframes, trajectories that differ in incidental details collapse to a single, task-relevant representation. Thus, the policy is trained and evaluated on a much narrower and better-covered input distribution, despite variation in raw history.

Concretely, BPP first defines keyframes using concise, task-specific criteria that capture salient events. The policy is then conditioned only on these keyframes, rather than the full history. Formally, let $\phi(o_t) \in \{0, 1\}$ denote a binary keyframe detector (e.g., a VLM) that identifies whether observation o_t corresponds to a keyframe. To avoid redundant keyframes, we define the set of keyframes $\mathcal{K}$ as the timesteps corresponding to the onset of a detected event (a “rising edge”):

$$\mathcal{K} = \{t : \phi(o_t) = 1 \wedge \phi(o_{t-1}) = 0\}. \quad (1)$$

TABLE 1: Real-Robot Experiments. BPP achieves nearly 70% higher average success rate than the best prior approach (PTP).

Task	Current Obs	Naïve History	PTP	BPP
Drawer Search	11.1%	0.0%	0.0%	33.3%
Marshmallows	40.0%	25.0%	35.0%	65.0%
Mug Replacement	0.0%	5.0%	40.0%	60.0%
Stacking Puzzle	6.5%	21.0%	52.0%	56.0%
Average	14.4%	12.8%	31.8%	53.6%

This ensures that if multiple consecutive detections occur, only the first frame is retained, and distinct events must be separated by at least one negative detection. Let $\mathcal{K}_t \subseteq \mathcal{K}$ denote the keyframes detected by time t . During policy learning, we account for the system latency, denoted by Δ , inherent to the detection mechanism (e.g., 3 seconds for VLM queries). We define the *latency-masked* keyframes available at time t as:

$$\mathcal{K}_t^\Delta = \{k \in \mathcal{K}_t : k \leq t - \Delta\}. \quad (2)$$

The policy $\pi_\theta(a_t|o_t, \{o_k\}_{k \in \mathcal{K}_t})$ conditions on the current observation and the available keyframe observations. During training, we use $\mathcal{K}_t^\Delta$ to simulate the latency present at inference time, ensuring the policy learns to act under realistic delays. At inference, the policy simply conditions on all keyframes that have been detected so far.

B. Systems Design for BPP

For real-world tasks, we implement the keyframe detector ϕ using a vision language model (VLM) with a simple prompt as a binary classifier. Both for labeling demonstrations and during evaluation rollouts, we query the VLM at 1 Hz using the current wrist camera image and the image from the previous query. Because keyframe definitions are often ambiguous in time, the rising edge condition in Eq. 1 ensures that consecutive positive responses produce a single keyframe, avoiding duplicate detections. We use Gemini 3 Pro for detection, which incurs an average latency of 3 to 5 seconds per query, and set $\Delta = 3$ seconds during training to construct $\mathcal{K}_t^\Delta$. Processing fewer context frames also reduces training time. On Drawer Search, BPP reduces training time by 41% compared to naïve history conditioning (see Appendix C).

V. EXPERIMENTS

We now evaluate BPP on a suite of challenging manipulation tasks that require history conditioning. We first describe the real-world tasks used in our primary evaluation, then outline the baselines and methods we compare against, and finally analyze results across both real-world and simulation settings. Our experiments address four questions: **1)** Does keyframe-based history conditioning via BPP outperform methods with access to full histories on complex real-world tasks? **2)** When do representation auxiliary losses (e.g., PTP) help or hurt? **3)** How does BPP compare to prior approaches in data efficiency? **4)** How sensitive is BPP to VLM detection errors?




Fig. 6: BPP system architecture. We condition a diffusion transformer policy on a small set of history keyframes. These keyframes are defined by simple, task-specific criteria and are detected using a VLM. We also mask recent histories to account for detection latency.




Fig. 7: Simulation results. BPP outperforms all non-oracle methods, and even surpasses the oracle in one task, suggesting keyframe-based conditioning provides a representation more amenable to learning.

A. Real-World Benchmarks

We evaluate our approach on four real-world manipulation tasks using a bimanual ALOHA 2 platform [33]. The robot observes four RGB camera views (top, worm’s-eye, and two wrist-mounted) along with proprioceptive state, and outputs target joint positions and gripper commands for both arms. The control loop runs at 50 Hz for data collection and policy execution. All policies run locally on a workstation with an NVIDIA RTX 4090 GPU, while VLM inference for BPP is performed via the Vertex AI platform on Google Cloud. Data is collected via teleoperation by multiple operators with varying strategies and expertise to capture diverse demonstration quality. Appendix D provides detailed task specifications, including demonstrations, instructions, and success criteria.

- Mug Replacement: Swap one of two mugs from a coffee machine for another (adapted from PTP [12]).
- Marshmallows: Put two handfuls of marshmallows into a bowl, and then signal completion by pressing a red button.
- Drawer Search: Search for a key in two cabinets with three drawers each, then place the key in the center of the table.
- Stacking Puzzle: Given three different color pieces stacked in random order, sort them into locations of the same color

and reassemble in the original configuration.

The difficulty of Mug Replacement and Marshmallows is comparable to the Mug Swap and Two Scoops tasks from prior work [12]. A notable aspect of our data collection is that multiple teleoperators contributed demonstrations using different strategies, introducing additional behavioral diversity. By contrast, Drawer Search and Stacking Puzzle represent a significant increase in complexity. These long-horizon tasks require up to 90 seconds for successful execution and involve 4-5 sequential subtasks, demanding more complex history reasoning, such as tracking which drawers have been checked or recalling the initial puzzle configuration and task phase.

B. Baselines and Comparisons

All comparisons use the same policy architecture (Diffusion Transformer) and training objective (DDPM [34] with action chunking of 50). We use separate image encoders per camera view, shared across all timesteps within each view. In addition to the **Current Observation** and **Naïve History** baselines from Section III, we compare BPP to **Past-Token Prediction (PTP, Torne et al. [12])**. PTP uses the same inputs as Naïve History but adds an auxiliary objective to predict actions at past timesteps. Unlike the official implementation, we train

the image encoder using gradients from all history steps. We also evaluate the *Oracle* baseline, which has access to ground-truth state and thus requires no history conditioning (see Appendix F) for obtaining insights in simulation.

C. Real-World Results and Analysis

Our real-world results are summarized in Table I. BPP significantly outperforms all baselines, achieving nearly **70% higher** performance than the best prior approach (PTP) on average. Beyond aggregate numbers, it is also very informative to analyze the behavioral differences between policies. We make the following observations:

(i) *Policies that only utilize the current observation fail to track progress.* Without history, policies learn locally reasonable behaviors but cannot maintain task state. This manifests as repetitive loops (e.g. opening and closing the same drawer repeatedly, shuffling the same puzzle piece back and forth) and premature termination (e.g. pressing the completion button after a single handful of marshmallows instead of the required two). These failures are unsurprising: the tasks require history to succeed, and a memoryless policy has no mechanism to distinguish “first handful completed” from “task complete.”

(ii) *Naïve history conditioning introduces new failure modes rather than solving old ones.* Unlike simulation results in Section V-D, history conditioning *hurts* performance relative to the current-observation baseline in two of four real-world tasks, with only modest gains in the others. We hypothesize that reversal in trends stems from additional sources of spurious correlations in the real-world data that are absent from simulation, such as background variation and sensor or actuator noise. The failure modes here differ qualitatively from memoryless policies, exhibiting degraded behaviors such as frequent failed grasps, stalling after opening drawers, and dropping objects outside target locations (Figures 14, 12, 15). Most informatively, on Marshmallows, a history-conditioned policy sometimes executes drop motions with an empty grasp or retries after a successful grasp, effectively replaying demonstration sequences regardless of the current observation. These behaviors indicate overfitting to training histories.

(iii) *PTP helps when predicting past actions helps infer the task state, but not otherwise.* PTP shows significant improvements over the naïve history conditioning approach on Mug Replacement and Stacking Puzzle, but offers little benefit on Drawer Search and Marshmallows. As we shall see later, this pattern of task-dependent performance is repeated in our simulated benchmarks. To explain this pattern, let us consider what the auxiliary loss actually requires. In Stacking Puzzle, correctly predicting past joint positions requires knowing which pieces have been placed and whether the policy is in the unstacking or restacking phase, which are also essential for correctly choosing subsequent actions. Similarly in Mug Replacement, predicting past actions requires knowing whether each mug has been placed in the machine or returned to its plate. However, *this heuristic breaks down when past actions do not reliably indicate task state.* In Drawer Search, the robot may pull away from a drawer

either after discovering no key *or* after a failed grasp attempt (see Fig. 9). Similarly, in Marshmallows, pulling away from the pile (expressed as target joint positions) appears identical whether or not the gripper successfully grasped any marshmallows. In both cases, the same action sequence is consistent with multiple task states, so the auxiliary prediction no longer allows task state disambiguation. In this scenario, PTP may not offer representational benefits. We observe similar results in simulation (Section V-D).

(iv) *BPP produces robust progress tracking.* BPP performs best across all four tasks by conditioning on a compact, semantically meaningful history summary rather than raw observations. Behaviorally, keyframe conditioning gives the policy a clearer sense of task progress: it retries more reliably after failed grasps and, in Drawer Search, executes systematic searches that avoid revisiting checked drawers (Fig. 8). BPP failures fall into three categories. Most stem from data limitations: in Drawer Search, failures arise from dropping keys inside drawers, which never appears in training, or repeated drawer-opening failures. Rarer errors come from VLM false positives, for example misclassifying an empty-handed drop as a successful scoop in Marshmallows, leading to a premature button press. Finally, some errors are due to VLM latency. While usually irrelevant, in Mug Replacement trials where the top mug starts in the machine, the grasp event and subsequent decision occur close enough in time that delayed keyframe detection degrades performance.

D. Analysis in Simulation

Next, we compare our approach to the **Oracle** baseline, where the policy has access to privileged ground-truth state information in simulation in addition to the current observation. We evaluate this on the simulation tasks from Section III. For instance, in Ingredient-Insertion, the policy receives a one-hot vector indicating the number of lemons already inserted (see Appendix F). This comparison estimates the achievable performance if challenges in inferring and reasoning over history were fully resolved. Figure 7 summarizes the results and leads to several key observations.

(i) *PTP performs comparably to Naïve History in simulation tasks*, except on Ingredient-Insertion. In Section V-C, we argued that alignment between past-action prediction and task state inference determined PTP’s effectiveness. Here, we apply the same argument. In the password-entering tasks, predicting target joint positions for each past observation is well-aligned with the task state, as it requires knowing which buttons were pressed and in which order. In Ingredient-Insertion, the relevant task state is simple: how many lemons have been inserted so far. Predicting past actions does encode this information, but it also encodes the specific spatial locations of the lemons that the demonstrator happened to pick. Since lemon positions were randomized across episodes and demonstrators chose arbitrarily, predicting these specific locations will lead to overfitting. Finally, it may seem contradictory that explicit ground-truth state prediction hurts performance (Section III) while PTP does not, since predicting past tokens




Fig. 8: Qualitative behavior on Drawer Search. BPP systematically searches through drawers without revisiting them, demonstrating robust long-horizon progress tracking.




Fig. 9: PTP failure on Drawer Search. PTP fails to track progress after failed grasps or empty drawers, often getting stuck in loops.

should impose a similar regularization. However, PTP does not explicitly optimize a state prediction loss; instead, it applies a weaker signal of past token prediction, which may only encourage useful structure in the learned features without causing them to overfit to the training distribution.

(ii) **BPP performs best; can sometimes outperform oracle.** BPP significantly outperforms all non-oracle comparisons across all tasks. In addition, BPP achieves a higher performance than *Oracle* on Variable-Password. We attribute this surprising result to the possibility that BPP provides a representation of history that is more amenable to learning a generalizable policy than the oracle one-hot encoding.

E. Ablation Studies in Simulation

To further understand the sources of the performance gap between methods, we conduct two complementary ablations.

a) **Data efficiency:** If Naïve History underperforms by latching onto spurious temporal correlations, then providing more demonstrations should eventually allow the model to discover structures that generalize. We test this hypothesis on Mug Replacement by training with varying dataset sizes (Figure 10). Performance evolves non-monotonically: adding demonstrations initially hurts Naïve History, likely due to the increased diversity of behaviors that expose more failure modes. However, with enough data, Naïve History matches the success rate of BPP. BPP achieves strong performance with far less data by reducing exposure to spurious history cues.

b) **VLM keyframes versus ground-truth keyframes:** We also quantify how much performance drops due to imperfect keyframe detection. For this comparison, we construct an “oracle” keyframe baseline on the Mug Replacement task by supplying the first observation of each episode as an additional context frame, which contains enough information to disambiguate the task. This oracle variant achieves a 70% success rate, compared to 60% with VLM-selected keyframes,




Fig. 10: Data efficiency on the Mug task. With sufficient demonstrations, naïve history conditioning catches up with BPP.

indicating that keyframe recognition errors explain only a small portion of the performance gap (Table II).

Method	Success Rate (%)
BPP (w/ VLM keyframes)	60%
BPP w/ Oracle keyframes	70%
Naïve History	5%
PTP	40%

TABLE II: BPP achieves close to oracle performance and vastly outperforms naïve history conditioning and PTP baselines.

VI. CONCLUSION

We show that the primary failure mode of history-conditioned imitation learning stems not from architectural limits, but from insufficient coverage over the exponentially large space of possible histories. To address this, we introduce Big Picture Policies (BPP), which use off-the-shelf vision-language models to compress observation histories into a compact sequence of task-relevant keyframes. By conditioning policies on these semantic events rather than raw temporal sequences, BPP reduces distribution shift between training and deployment. Across three simulation and four real-world tasks, BPP enables robust progress tracking and outperforms state-of-the-art history-conditioned approaches by up to 70%.

Limitations. BPP relies on the inference speed of an external VLM; while we mitigate this during training via latency masking, real-time deployment remains constrained by VLM query latency, which can be critical for highly dynamic tasks (e.g., parts of Mug Replacement). BPP also

relies on accurate keyframe detection. Although our ablations show robustness to some errors, systematic false positives, such as misclassifying an empty grasp, can trigger premature state transitions. Future work could address these issues by distilling keyframe detection into low-latency models. More broadly, BPP points toward event-based robot learning, where abstracting continuous control into semantic milestones offers a promising path for scaling to multi-stage, long-horizon tasks using diverse, non-expert data.

Future Directions. While BPP currently relies on task-specific VLM prompting, the salient event definitions used for keyframe detection are simple enough that large language models could feasibly generate them at scale for new tasks. This opens a path toward automatically labeling keyframes across large, diverse datasets – a key ingredient for training history-conditioned generalist robot policies. Beyond tasks that explicitly require memory, the core insight of BPP, that conditioning on semantically meaningful keyframes rather than raw observation histories yields more robust and efficient learning, extends naturally to other settings. In particular, tasks that demand fast adaptation or forms of in-context learning stand to benefit from similar abstractions. Consider, for example, a policy that must learn from past failed attempts: a single point observation (i.e., an image and state from one timestep) may not capture enough information to understand why a grasp failed. In such cases, the notion of a keyframe can be generalized to *key segments*, or short sequences of consecutive frames that capture the relevant event, such as a failed grasp, providing richer context for downstream learning. Extending BPP to these settings is an important direction for future work.

ACKNOWLEDGEMENTS

We thank the robot operators from GDM Robotics for their efforts in collecting demonstrations and conducting policy evaluations. We are also grateful to the members of the CMU AIRE lab and GDM Robotics for their valuable discussions and feedback on this work. This work was done when MSM was a student researcher at Google DeepMind.

REFERENCES

- [1] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, et al. Openvla: An open-source vision-language-action model. In *Conference on Robot Learning*, pages 2679–2713. PMLR, 2025.
- [2] Gemini Robotics Team, Abbas Abdolmaleki, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Ashwin Balakrishna, Nathan Batchelor, Alex Bewley, Jeff Bingham, et al. Gemini robotics 1.5: Pushing the frontier of generalist robots with advanced embodied reasoning, thinking, and motion transfer. *arXiv preprint arXiv:2510.03342*, 2025.
- [3] Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Es-mail, Michael Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$: A vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- [4] Moo Jin Kim, Yihuai Gao, Tsung-Yi Lin, Yen-Chen Lin, Yunhao Ge, Grace Lam, Percy Liang, Shuran Song, Ming-Yu Liu, Chelsea Finn, et al. Cosmos policy: Fine-tuning video models for visuomotor control and planning. *arXiv preprint arXiv:2601.16163*, 2026.
- [5] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.
- [6] Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, et al. Octo: An open-source generalist robot policy. In *Robotics: Science and Systems*, 2024.
- [7] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [8] Chuan Wen, Jierui Lin, Trevor Darrell, Dinesh Jayaraman, and Yang Gao. Fighting copycat agents in behavioral cloning from observation histories. *Advances in Neural Information Processing Systems*, 33:2564–2575, 2020.
- [9] Chuan Wen, Jierui Lin, Jianing Qian, Yang Gao, and Dinesh Jayaraman. Keyframe-focused visual imitation learning. In *International Conference on Machine Learning*, pages 11123–11133. PMLR, 2021.
- [10] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Robotics: Science and Systems (RSS)*, 2023.
- [11] Ruijie Zheng, Yongyuan Liang, Shuaiyi Huang, Jianfeng Gao, Hal Daumé III, Andrey Kolobov, Furong Huang, and Jianwei Yang. Tracevla: Visual trace prompting enhances spatial-temporal awareness for generalist robotic policies. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [12] Marcel Torne, Andy Tang, Yuejiang Liu, and Chelsea Finn. Learning long-context diffusion policies via past-token prediction. *arXiv preprint arXiv:2505.09561*, 2025.
- [13] Seokin Seo, HyeongJoo Hwang, Hongseok Yang, and Kee-Eung Kim. Regularized behavior cloning for blocking the leakage of past action information. *Advances in Neural Information Processing Systems*, 36:2128–2153, 2023.
- [14] Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauzá, Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *CoRR*,

- 2025.
- [15] Ajay Sridhar, Jennifer Pan, Satvik Sharma, and Chelsea Finn. Memer: Scaling up memory for robot control via experience retrieval. *arXiv preprint arXiv:2510.20328*, 2025.
 - [16] Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Xi Chen, Ken Goldberg, and Pieter Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 5628–5635. Ieee, 2018.
 - [17] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. Pmlr, 2016.
 - [18] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39 (1):3–20, 2020.
 - [19] Pim De Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. *Advances in neural information processing systems*, 32, 2019.
 - [20] Mayank Bansal, Alex Krizhevsky, and Abhijit Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.
 - [21] Gokul Swamy, Sanjiban Choudhury, J Bagnell, and Steven Z Wu. Sequence model imitation learning with unobserved contexts. *Advances in Neural Information Processing Systems*, 35:17665–17676, 2022.
 - [22] Karl Pertsch, Oleh Rybkin, Jingyun Yang, Shenghao Zhou, Konstantinos Derpanis, Kostas Daniilidis, Joseph Lim, and Andrew Jaegle. Keyframing the future: Keyframe discovery for visual prediction and planning. In *Learning for Dynamics and Control*, pages 969–979. PMLR, 2020.
 - [23] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
 - [24] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
 - [25] Beichen Wang, Juexiao Zhang, Shuwen Dong, Irving Fang, and Chen Feng. Vlm see, robot do: Human demo video to robot action plan via vision language model. *CoRR*, 2024.
 - [26] Yuqing Du, Ksenia Konyushkova, Misha Denil, Akhil Raju, Jessica Landon, Felix Hill, Nando de Freitas, and Serkan Cabi. Vision-language models as success detectors. In *Conference on Lifelong Learning Agents*, pages 120–136. PMLR, 2023.
 - [27] Yecheng Jason Ma, Joey Hejna, Chuyuan Fu, Dhruv Shah, Jacky Liang, Zhuo Xu, Sean Kirmani, Peng Xu, Danny Driess, Ted Xiao, et al. Vision language models are in-context value learners. In *The Thirteenth International Conference on Learning Representations*, 2024.
 - [28] Kate Baumli, Satinder Baveja, Feryal MP Behbahani, Harris Chan, Gheorghe Comanici, Sebastian Flennerhag, Maxime Gazeau, Kristian Holsheimer, Dan Horgan, Michael Laskin, et al. Vision-language models as a source of rewards. *CoRR*, 2023.
 - [29] Jingyun Yang, Max Sobol Mark, Brandon Vu, Archit Sharma, Jeannette Bohg, and Chelsea Finn. Robot fine-tuning made easy: Pre-training rewards and policies for autonomous real-world reinforcement learning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4804–4811. IEEE, 2024.
 - [30] Kevin Black, Mitsuhiko Nakamoto, Pranav Atreya, Homer Walke, Chelsea Finn, Aviral Kumar, and Sergey Levine. Zero-shot robotic manipulation with pre-trained image-editing diffusion models. *arXiv preprint arXiv:2310.10639*, 2023.
 - [31] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023.
 - [32] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
 - [33] Jorge Aldaco, Travis Armstrong, Robert Baruch, Jeff Bingham, Sanky Chan, Kenneth Draper, Debiddatta Dwibedi, Chelsea Finn, Pete Florence, Spencer Goodrich, et al. Aloha 2: An enhanced low-cost hardware for bimanual teleoperation. *arXiv preprint arXiv:2405.02292*, 2024.
 - [34] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.