

History-Aware Visuomotor Policy Learning via Point Tracking

Jingjing Chen^{*1}, Hongjie Fang^{*1}, Chenxi Wang², Shiquan Wang^{2,3,†} and Cewu Lu^{1,2,3,4,†}




Fig. 1: Consider the task of adding some salt to a dish, where the decision depends on the question “Have I added the salt?” since the dish looks similar before and after adding the salt. (Left) Humans form object-centric memories from past interactions to answer such questions. (Right) [orange] Policies using only the current observation cannot decide and perform random guesses. [blue] Policies that look back at a fixed number of past frames provide partial history-awareness but remain limited by their short temporal horizon and computational inefficiency, as they must repeatedly process many redundant images. [green] Our history-aware policies instead use object-centric point tracks — akin to human memory — to compactly encode history, enabling efficient and correct decision making.

Abstract—Many manipulation tasks require memory beyond the current observation, yet most visuomotor policies rely on the Markov assumption and thus struggle with repeated states or long-horizon dependencies. Existing methods attempt to extend observation horizons but remain insufficient for diverse memory requirements. To this end, we propose an object-centric history representation based on point tracking, which abstracts past observations into a compact and structured form that retains only essential task-relevant information. Tracked points are encoded and aggregated at the object level, yielding a compact history representation that can be seamlessly integrated into various visuomotor policies. Our design provides full history-awareness with high computational efficiency, leading to improved overall task performance and decision accuracy. Through extensive evaluations on diverse manipulation tasks, we show that our method addresses multiple facets of memory requirements — such as task stage identification, spatial memorization, and action counting, as well as longer-term demands like continuous and pre-loaded memory — and consistently outperforms both Markovian baselines and prior history-based approaches. Project website: <https://tonyfang.net/history/>.

I. INTRODUCTION

In recent years, visuomotor policies have achieved significant advancements in robotic manipulation using the behavior cloning paradigm [5, 6, 11, 17, 28, 41, 57, 70], which learns an end-to-end mapping from visual observation to robot action [44]. However, these policies typically rely on the Markov assumption, which presumes that the current observation alone provides sufficient information for optimal

decision-making, so that future actions depend only on the present state while ignoring historical information. In practical manipulation scenarios, such an assumption often breaks down [29, 38], especially in tasks that involve repeated states or similar observations [8], or in long-horizon tasks that require memorization of previous actions [35, 49, 54].

For example, in a cooking scenario where a robot is required to *add some salt to a dish* as shown in Fig. 1, the visual appearance of the dish before and after adding salt may look nearly identical, yet the correct action depends on whether salt has already been added in previous steps. Relying solely on the current observation is therefore insufficient, as the robot must recall its past actions to avoid either over-salting or forgetting the seasoning. Similar issues can arise in other manipulation tasks: in object swapping, the robot must remember the original positions of the objects to complete the exchange, while in button pressing, the scene may appear unchanged after the action. These examples illustrate that many practical tasks inherently require memory of previous behaviors and thus violate the Markov assumption.

Faced with this challenge, many visuomotor policies attempt to relax the Markov assumption by extending the observation horizon to a fixed temporal window [11, 31, 33, 54, 60, 70, 71]. Nonetheless, these approaches still suffer from two major limitations: (1) repeated states in long-horizon tasks cannot be reliably distinguished using a limited observation window; (2) a fixed horizon cannot adapt to the diverse memory requirements of different tasks. These shortcomings underscore the need for history-aware visuomotor policies that can fully exploit past information

^{*}Equal Contribution. [†]Corresponding Authors.

¹Shanghai Jiao Tong University. ²Noematrix.

³Flexiv. ⁴Shanghai Innovation Institute.

and move beyond the constraints of the Markov assumption.

However, leveraging complete historical information introduces new challenges, as it requires an effective representation of history in robotic manipulation scenarios. A straightforward approach is to feed all historical frames into the policy [31, 72], but this quickly incurs substantial computational cost. One possible solution is to treat the history as a video and apply the video encoders for representation learning [3, 36, 51, 53, 69]. Yet, video encoders are typically designed around a keyframe-motion scheme [23, 68], which can overlook subtle but critical information necessary for precise action execution in manipulation tasks [48, 52].

Humans do not explicitly store every visual frame of a process or memorize all fine-grained details in the history sequence. Instead, we abstract past experiences into compact representations that capture only the essential information for decision-making [1, 47]. Recall the previous salt-adding example, when cooking, we humans do not retain the entire visual history of the dish; instead, we maintain a high-level object-centric memory such as “*the salt has been added*”. Such abstraction enables efficient use of past information without being overwhelmed by redundant details.

Building on this insight, we introduce **an object-centric history representation realized through point tracking**, which provides a compact and structured encoding of historical information. As illustrated in Fig. 1, this design bridges the human-inspired abstraction of past experiences with a practical mechanism for visuomotor policies, enabling them to leverage *infinite* history at minimal cost. By encoding tracked points with a dedicated track encoder and performing object-level aggregation, the resulting history representation can be efficiently integrated into many visuomotor policies [11, 57]. Experimental results across multiple manipulation tasks, with diverse memory requirements in both type and length, demonstrate that our history-aware visuomotor policies significantly outperform both Markovian baselines and prior history-based methods.

II. RELATED WORKS

A. History Representation

Researchers have explored history representations to encode past observations into compact forms that capture temporal dependencies for manipulation. Existing approaches mainly fall into three categories: (1) video representations process history as videos using video encoders [2, 53, 58, 64], effectively modeling long horizons but introducing redundancy and high computational cost; (2) keyframe-motion representations compress history by selecting salient frames [16, 52, 65], reducing redundancy yet risking the loss of subtle transitions and causal cues. (3) hybrid representations combine multi-level abstractions for richer temporal modeling [25, 49, 50, 63], but in manipulation tasks, they often emphasize irrelevant details over task-critical features.

A recent work [71] employs visual prompting [4, 24], overlaying past frames as point traces on the current observation to encode history. Although intuitive, this design scales poorly along the temporal axis due to accumulated trace

overlays on the image, and conveys limited historical cues within the image domain (§IV-B). These limitations highlight the need for a history representation that is both scalable and informative for visuomotor policies in manipulation.

B. History-Aware Visuomotor Policy

Incorporating historical information into visuomotor policies remains a fundamental challenge in robot learning. Conventional policies typically rely on fixed-length observation windows [11, 57, 60, 70], while LongDP [54] extends this by freezing encoders trained on short-context policies and using past action tokens as auxiliary training signals and inference guidance. UniVLA [7] incorporates temporal context by integrating past actions into input prompts. CronusVLA [31] extracts features from historical observations to guide decisions, though it incurs substantial inference cost and produces redundant features. TraceVLA [71] represents history by overlaying past frames as point traces on the current observation, but it conveys limited information and scales poorly to long context windows due to clutter.

Other approaches aim to leverage complete historical information. Several works adopt keyframe-based representations to select representative moments from past observations [9, 16, 20]. While computationally efficient, they risk missing critical temporal details not captured by keyframes. MTIL [72] uses Mamba [19] to compress temporal sequences via recurrent hidden states, but suffers from training instability and limited performance. Consequently, effectively compressing history representations of arbitrary length for use in visuomotor policies remains a critical open question.

C. Point Representation in Robotics

Point representations are widely used in robotic manipulation as a high-level abstraction of high-dimensional data. Several works leverage point trajectories (flows) to represent and abstract robot actions [18, 21, 42, 46, 59, 62], enabling policies to focus on essential motion patterns while reducing computational complexity and improving generalization across tasks. Points can also serve as affordances, providing external knowledge relevant to robot actions [32, 66]. Beyond action abstraction, object-centric keypoints are used to encode observations [21, 30, 39, 42], enhancing the generalization of visuomotor policies to novel environments. Building on these ideas, we adopt object-centric point track representations as an efficient history representation to equip visuomotor policies with history awareness.

D. Point Tracking

Point tracking is fundamental for maintaining spatial awareness of objects over time in manipulation tasks. Early methods [13, 22] introduce iterative refinement and two-stage designs, and CoTracker [26] further improved robustness through joint tracking. Tracking performance is further advanced through self-supervised real-world training [14, 27]. However, 2D trackers struggle with occlusions and depth ambiguities. To address this, 3D trackers [40, 56, 61, 67] operate directly in world space: SpatialTracker [61]

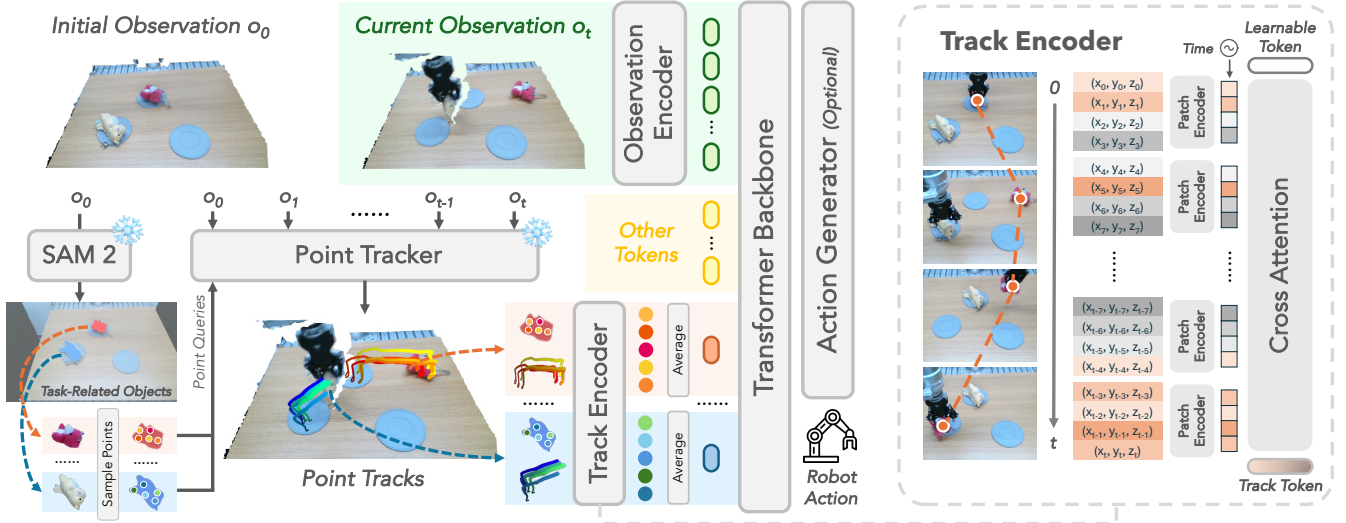


Fig. 2: History-Aware Visuomotor Policy Architecture. (left) We first identify and segment task-related objects via SAM [45]. From each object, several points are sampled and tracked using an off-the-shelf point tracker [26, 67], producing point tracks as our history representations. Each track is then encoded into a track token, and the tokens from all points of an object are aggregated into a single history-aware object track token. These object track tokens, together with the original observation tokens and other tokens, are fed into the transformer backbone of visuomotor policies [11, 57, 70] to generate history-aware robot actions. (right) The track encoder first encodes each track patch using an MLP, then applies a cross-attention module with temporal positional encodings to produce a compact point track token from point tracks of arbitrary length, effectively compressing and preserving historical information.

combines 2D tracking with monocular depth estimation, and TAPI3D [67] enables robust dense tracking with stabilized 3D representations. Given the frequent occlusions and complex spatial interactions in robotic manipulation scenarios, we adopt 3D point tracking as a more reliable solution.

III. METHOD

A. Overview

Consider a dataset \mathcal{D} of manipulation demonstrations, where each trajectory $(o_1, a_1, \dots, o_T, a_T)$ pairs observations o_t with actions a_t . Current Markovian and semi-Markovian policies π_M predict actions from observations within a short history horizon h , i.e., $\hat{a}_t = \pi_M(o_{t-h:t})$, but many tasks require adaptive reasoning over longer horizons. In this work, we aim to learn a history-aware policy π_H that predicts actions based on the full observation history $o_{1:t}$ with an optional pre-loaded memory $o_{-h_p:0}$ (please refer to §IV-C for details). Formally,

$$\hat{a}_t = \pi_H(o_{1:t}, o_{-h_p:0}) = \pi_H(o_{-h_p:t}) = \pi_H(o_t, o_{-h_p:t-1}).$$

Directly using raw observations from the full history is computationally expensive and redundant. To address this, we first convert the history into an **object-centric point tracking representation** $\tau(o_{-h_p:t-1})$, capturing the motion and state of task-relevant objects (§III-B). Although more structured and abstract than raw image observations, this representation remains unbounded and has an arbitrary length. Hence, we apply a **compression module** $\phi(\cdot)$ to obtain a compact summary of the history (§III-C). This compact history representation can then be integrated into existing visuomotor policies π_M , yielding **history-aware policies** (§III-D) that leverage long-horizon contextual information:

$$\pi_H(o_t, o_{-h_p:t-1}) = \pi_M(o_t, \phi(\tau(o_{-h_p:t-1}))).$$

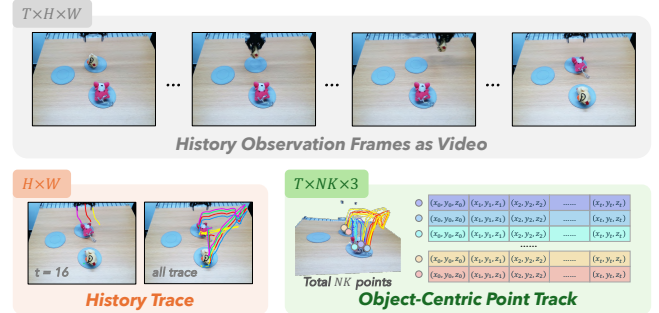


Fig. 3: Different History Representations Comparisons. Suppose the full history horizon is t and the observation image size is $H \times W$. (Gray) Treating history as video is redundant and computationally expensive. (Orange) History overlays are ineffective for short horizons and cluttered for long horizons. (Green) Our object-centric point track representation is efficient while preserving motion patterns and object-centric dynamics across all horizons.

In other words, our approach can equip various visuomotor policies [11, 57, 70] with enhanced history-awareness, enabling them to tackle non-Markovian manipulation tasks and achieve improved performance. An overview of our history-aware visuomotor policy architecture is illustrated in Fig. 2.

B. History Representation via Object Point Tracking

As discussed in §III-B, current history representations suffer from redundancy, limited temporal scalability, and a lack of task-relevant cues, making them difficult to integrate effectively into visuomotor policies. To overcome these issues, we propose an **object-centric point track representation** that compactly encodes historical information as trajectories of points on task-related objects, as illustrated in Fig. 5.

Formally, we first identify and segment N task-related objects in the initial observation o_0 using SAM-2 [45]. For

each object, we randomly sample K points within its mask and track them across frames using TAPIP3D [67], an off-the-shelf 3D tracker. This process yields $N \times K$ point tracks of length $(h_p + t)$, which constitute our history representation:

$$\tau(o_{-h_p:t-1}) \in \mathbb{R}^{N \times K \times (h_p + t) \times 3}.$$

In our representation, each point track captures the spatio-temporal trajectory of a location on a task-relevant object, preserving motion patterns and object-centric dynamics across arbitrary horizons. As shown in Fig. 5, unlike video-based representations that are highly redundant and computationally expensive, or trace overlays [71] that become cluttered and make it difficult to distinguish task-relevant targets, our approach scales naturally with temporal length while highlighting essential object motions and interactions. In this way, it offers a principled and scalable solution to equip visuomotor policies with history awareness efficiently.

C. Compressing Unbounded Tracks

Although we have abstracted history into task-relevant object point tracks, the resulting tracks are unbounded and can vary arbitrarily in length at different stages of task execution. Such variability not only increases storage and computational costs but also complicates learning and inference, making it challenging to incorporate long-term history efficiently. Therefore, it is essential to compress the tracks of arbitrary length into a compact yet informative feature representation.

For a single point track $\mathcal{T} \in \mathbb{R}^{(h_p + t) \times 3}$, we first divide it into $\lceil (h_p + t)/P \rceil$ consecutive patches of size P . Each patch is encoded using a shared patch encoder to capture local spatio-temporal patterns. We add temporal positional encodings to the patch embeddings to inject temporal order information, giving the model a clear notion of time within the trajectory. Inspired by [60], we then introduce a learnable token that directly cross-attends to the patch encodings, condensing the sequence into a compact track token that summarizes the entire trajectory. This patching and attention design not only preserves critical motion information but also ensures scalability, allowing long point tracks to be encoded without incurring excessive computational or memory costs.

After encoding each point track into a compact feature, we aggregate the features of all point tracks belonging to the same object into a single object track token. This consolidates object-level features and reduces the randomness introduced by point sampling, yielding a more stable and informative representation. Repeating this process for all task-relevant objects produces a set of object track tokens that collectively form the compact history representation $\phi(\tau(o_{-h_p:t-1})) \in \mathbb{R}^{N \times D}$ where D is the feature dimension, which can be directly consumed by the visuomotor policy.

D. History-Aware Policy Learning

After compressing the object-centric point track representation into N feature tokens, we can integrate the history tokens into various modern visuomotor policies [11, 57, 70]. By appending the history tokens alongside the current observation and other tokens and adding potential

token type encodings, the Transformer backbone can jointly attend to past and present information. This enables the policy to leverage compact yet informative historical cues during decision making, effectively incorporating temporal context into action prediction while preserving the original policy architecture as much as possible. Formally, given a Markovian policy $\pi_M(o_t)$, we incorporate the encoded history tokens to obtain the history-aware visuomotor policy $\pi_H(o_{-h_p:t}) = \pi_M(o_t, \phi(\tau(o_{-h_p:t-1})))$.

Point trackers operate considerably slower than visuomotor policies. To accelerate inference during deployment, we run the tracker asynchronously with the policy, inevitably introducing discrepancies in the observed point tracks. To address this, we apply random dropping augmentation during training, which simulates asynchronous deployment conditions, mitigates the observation mismatch, and enables the policy to remain robust under asynchronous tracking.

IV. EXPERIMENTS

In this section, we aim to answer the following research questions: **(Q1)** Is our object-centric point track history representation effective in improving policy performance across tasks with varying demands for history awareness? **(Q2)** Do our history-aware policies better leverage historical information for decision making compared to existing history representations and policies? **(Q3)** Can our representation be integrated with pre-loaded memory before manipulation? **(Q4)** How does tracking quality affect policy performance, and how can we ensure efficiency while maintaining performance? **(Q5)** What points should be selected to enhance both history awareness and policy performance?

A. Setup

Platform. The experimental platform consists of a Flexiv Rizon 4 arm with a Dahuan AG-95 gripper, one global Intel RealSense D415 RGB-D camera, and one wrist-mounted Intel RealSense D435 RGB-D camera. The global camera serves as the primary observation source, while the wrist-mounted camera provides auxiliary inputs for 2D policies.

Tasks. We carefully designed 7 real-world manipulation tasks that feature multiple repeated or difficult-to-distinguish states across different horizons, aiming to evaluate both the history-awareness and overall performance of visuomotor policies. The tasks collectively evaluate several aspects of history-awareness, including “counting”, “spatial memorization”, “task stage identification”, “pre-loaded memory” and “continuous memory”. Please refer to Fig. 4 for task details and the explanation of evaluation aspects.

Baselines. We select three representative base policies: 2D image-based *ACT* [70], *Diffusion Policy (DP)* [11], and 3D point-cloud-based *RISE* [57]. To incorporate history information, we use trace images [71] as an additional representation, constructing 2D policies *TraceACT* and *TraceDP* for ACT and DP, respectively. We further compare with *LongDP* [54], which optimizes history-aware learning based DP.

Implementations. We build our history-aware visuomotor policies on three representative baselines [11, 57, 70],

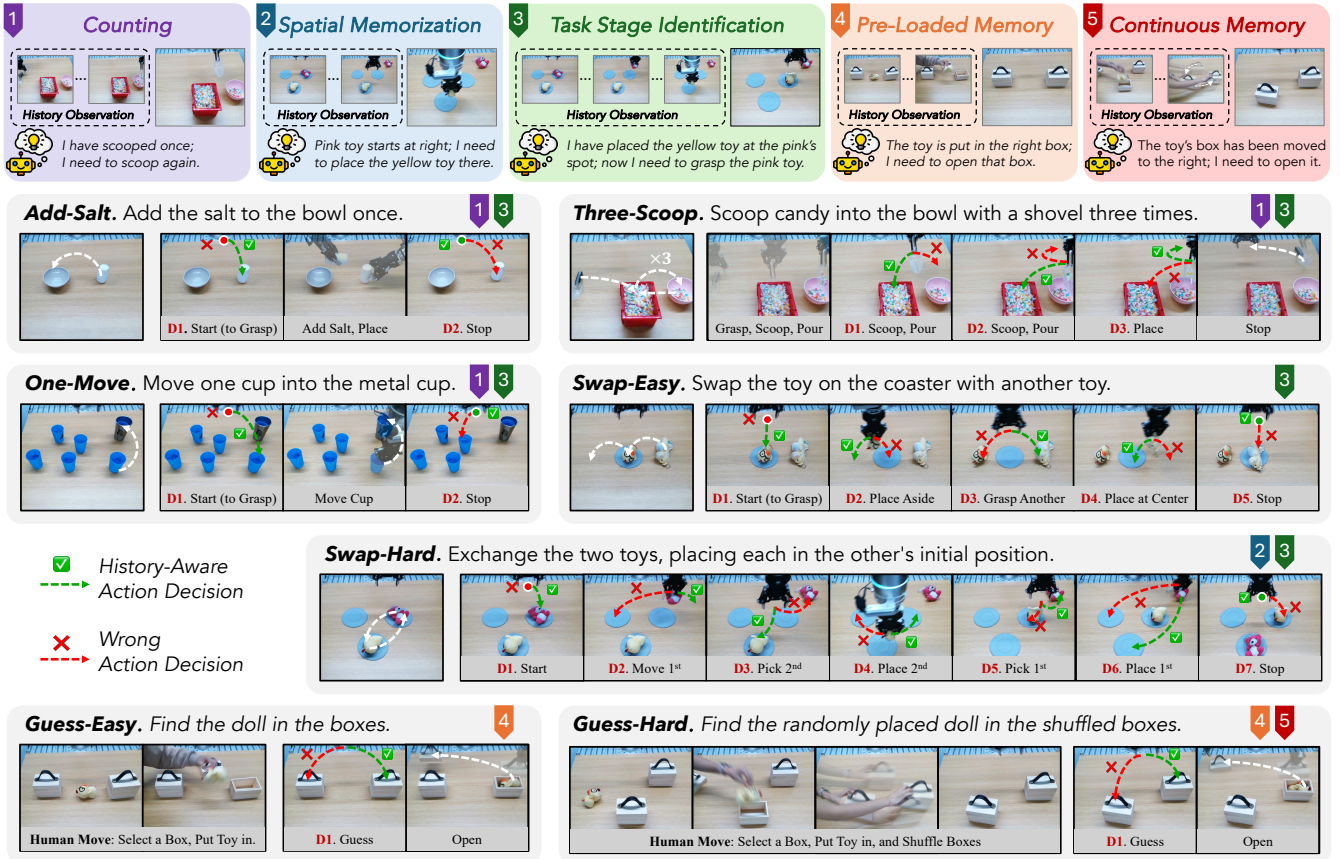


Fig. 4: Tasks and Evaluation Aspects. (Top) Five evaluation aspects of history-awareness: (1) *counting* evaluates the policy’s ability to track repeated actions; (2) *spatial memorization* tests remembering spatial locations during manipulation; (3) *task stage identification* requires inferring the correct phase in long-horizon tasks; (4) *pre-loaded memory* assesses use of pre-loaded information for decision-making; and (5) *continuous memory* examines the ability to retain continuous history for critical decisions. (Bottom) Detailed descriptions show the full process of each task, specifying every decision phase (D1, D2, ...) where history is required for correct action decisions (green arrow). The top-right bookmarks of each task indicate which aspects of history-awareness are required for this task.

denoted as *HistACT*, *HistDP*, and *HistRISE*. For *HistACT* and *HistRISE*, history-aware point-track tokens are directly injected into the Transformer backbone. For *HistDP*, a Transformer fuses observation tokens with point-track tokens before passing them to the diffusion transformer [43] to condition action denoising. For point selection, we sample $K = 5$ points per task-relevant object across all tasks.

Metrics. We evaluate each task using *Success Rate (SR)* as the primary metric, complemented by *Decision Accuracy (DA)*. Tasks are divided into action and decision phases, and *DA* evaluates whether the policy makes correct decisions to advance through the next phase of the task. For tasks involving multiple decision stages (>3), we also report *Average Length (AL)*, the average number of decision stages reached in policy rollouts [34, 37].

Protocols. We collect 50 expert demonstrations for each task using haptic teleoperation [15] as training data, and deploy all policies on a workstation equipped with an NVIDIA RTX 3090 GPU. Following [12, 17, 60], we adopt a consistent evaluation protocol to reduce performance variability and ensure reproducibility. For each task, the workspace is configured identically across all policies, and test positions are randomly sampled within the workspace beforehand.

Each policy is evaluated over 20 trials per task, and the metrics are computed from these trials.

B. Basic Results

Our object-centric point track history representation enables effective history-aware decision-making (Q1). As shown in Tab. I, our history representations substantially boost decision-making performance across diverse manipulation tasks among various base visuomotor policies, lifting decision accuracy over 90% across nearly all tasks. Even in challenging long-horizon tasks that involve multiple repeated states, such as *Swap-Easy* and *Swap-Hard*, *HistRISE* still demonstrates superior performance, outperforming previous policies by a large margin. It is worth noticing that while our method significantly aids history-aware decision-making, it does not directly improve action accuracy. Action accuracy, such as whether the robot can successfully grasp an object, fundamentally depends on the base policy. Since task success requires both correct decisions and accurate execution, methods like *HistACT* and *HistDP* show lower success rates compared to *HistRISE*, as *RISE* exhibits higher action accuracy than *ACT* and *DP* [57].

Our history representation demonstrates strong effectiveness across all five history-awareness evaluation

Method	Obs. Len.	Add-Salt		One-Move		Three-Scoop		Swap-Easy			Swap-Hard		
		SR	DA	SR	DA	SR	DA	SR	DA	AL	SR	DA	AL
ACT [70]	1	20%	36%	25%	89%	5%	68%	0%	61%	0.90	-	-	-
TraceACT [71]	16	15%	29%	15%	82%	5%	67%	0%	63%	1.05	-	-	-
<i>HistACT (ours)</i>	∞	50%	97%	35%	100%	95%	100%	25%	87%	2.05	-	-	-
DP [11]	1	45%	73%	30%	76%	5%	73%	10%	71%	1.85	10%	76%	2.85
DP [11]	2	35%	73%	35%	96%	10%	73%	0%	62%	0.95	15%	78%	3.10
LongDP [54]	16	50%	81%	50%	100%	10%	70%	10%	68%	1.95	10%	77%	3.05
TraceDP [71]	16	40%	76%	50%	84%	0%	64%	40%	86%	3.10	20%	80%	3.20
<i>HistDP (ours)</i>	∞	85%	97%	65%	100%	80%	95%	70%	96%	3.80	40%	97%	4.55
RISE [57]	1	45%	74%	65%	83%	15%	74%	15%	71%	2.05	10%	78%	3.10
<i>HistRISE w/ 2D tracker</i>	∞	85%	100%	95%	100%	65%	87%	80%	95%	4.00	70%	95%	6.30
<i>HistRISE w/ EE points</i>	∞	80%	100%	85%	100%	55%	98%	60%	98%	3.60	65%	94%	5.70
<i>HistRISE (ours)</i>	∞	85%	100%	95%	100%	90%	100%	90%	98%	4.50	80%	97%	6.50




TABLE I: Basic Evaluation Results and Ablation Results. (Left) Our history-aware policies outperform both base policies and other history-aware methods in overall success rates and decision-making performance. ACT [70] struggles with long-horizon manipulation tasks involving complex object motions; thus we report its results on *Swap-Easy* but omit them for *Swap-Hard*, as all ACT-based policies can hardly complete that task. We also report *HistRISE* variants with 2D trackers and with end-effector (EE) points for ablation comparisons. (Right) Integration with pre-loaded memory allows visuomotor policies to effectively utilize pre-loaded memory in both *Guess* tasks.




Fig. 5: Evaluation Results on Five History-Awareness Aspects. Average success rate and decision accuracy across five history-awareness aspects are reported, with each dimension representing the average performance of the policies on the corresponding tasks.

aspects (Q2). As shown in Fig. 5 and Tab. I, built upon the same base policy DP [11], *HistDP* consistently outperforms other history-aware variants, highlighting the performance advantage of our object-centric point track representation. Notably, *HistRISE* further surpasses *HistDP*, achieving state-of-the-art performance across all five aspects. On the contrary, fixed-window approaches such as *TraceDP* and *LongDP* often struggle in long-sequence tasks. As shown in Tab. I, both methods obtain low success rates due to truncation in *Three-Scoop* and *Swap-Hard*. This limitation arises because they restrict the accessible history to a predefined length, inevitably discarding critical early-stage information. Beyond this, the object-centric design enables continuous memory, whereas keyframe-based methods capture only discrete snapshots and fail to preserve temporal continuity. These results collectively confirm that representing history as object-centric point tracks provides a robust and generalizable foundation for enhancing visuomotor decision-making.

C. Incorporating Pre-Loaded Memory

Introducing history into visuomotor policies offers the additional benefit of a pre-loaded memory, allowing the policy to access past observations or human actions before deployment. This capability not only enhances the policy’s contextual understanding but also opens up potential applications such as in-context learning [55], task adaptation,

and long-horizon planning [10]. To specifically evaluate this ability, we design the *Guess-Easy* and *Guess-Hard* tasks, which test whether the policy can effectively exploit pre-loaded memory for correct task execution.

Our history representation can be seamlessly integrated with pre-loaded memory, enabling policies to utilize prior observations before manipulation (Q3). As shown in Tab. I (right), our history-aware policy *HistRISE* not only supports incorporating pre-loaded memory but also effectively leverages it to assist the robot in finding the correct box with a toy. This demonstrates that pre-loaded memory, which provides access to human demonstrations in this task, supplies valuable contextual information about the environment and task dynamics. By harnessing this information effectively, our approach enables the robot to make more informed decisions, anticipate changes, and respond appropriately. Looking forward, pre-loaded memory represents a promising mechanism for further enhancing robot decision-making by incorporating information beyond its own direct interactions.

Interestingly, *HistDP* shows relatively low accuracy in *Guess-Easy* and *Guess-Hard* tasks compared to *HistRISE*. Nevertheless, *HistDP* does effectively incorporate historical information in other tasks as illustrated in Tab. I. Hence, we hypothesize that the base policy *DP* may overfit to implicit behavioral patterns in the dataset, *e.g.*, humans may show unintended preferences for which box they place the toys in, and consequently overemphasize visual observation features while neglecting the history point track modalities.

D. Ablations

The following ablations are based on the *HistRISE* policy, our best-performing policy in previous evaluations.

3D trackers offer better point tracking quality and superior policy performance compared to 2D trackers (Q4). We compare CoTracker [26] and TAPIP3D [67], the representative 2D and 3D trackers, in terms of tracking quality and their impact on history-aware policy performance.




Fig. 6: Comparisons of Different Trackers on the *Guess-Hard* Task. 2D trackers suffer from tracking failures under occlusions, which frequently happen during manipulations, while 3D trackers maintain more robust point tracking performance.

Results in Tab. I show that *HistRISE* consistently outperforms its variant with 2D tracker across all tasks. In scenarios requiring spatial memorization (*Guess-Easy*, *Swap-Hard*) or involving occlusions (*Guess-Hard*), 2D trackers suffer from depth ambiguities and tracking discontinuities, producing low-quality tracks, as illustrated in Fig. 6 and Tab. II. By contrast, 3D trackers maintain accurate spatial relationships and handle occlusions effectively, yielding substantially better performance than 2D trackers, making them better suited for robotic manipulations.

Asynchronous tracking with train-time augmentation improves efficiency while preserving policy performance (Q4). Although 3D trackers are slower due to heavy processing, performing tracking asynchronously reduces latency but inevitably lowers performance compared to synchronous tracking and policy inference; this tradeoff is mitigated by randomly dropping recent point tracks as an augmentation during training, yielding a policy that maintains performance while achieving higher efficiency, as shown in Table III.

Selecting object-centric points improves history-aware policy performance, whereas end-effector points degrade it (Q5). Although end-effector (EE) points are easy to compute from proprioception and camera extrinsics, including them in our history representations consistently degrades performance across all tasks, as depicted in Tab. I. After analyzing the performance, we hypothesize that the resulting policy overrelies on EE point tracks while neglecting critical visual cues. When initial actions such as grasping fail, the policy struggles to recover and often continues subsequent steps based on incomplete EE information, leading to eventual failure. In contrast, object-centric points do not suffer from this issue, as task success primarily depends on object state rather than EE motion, enabling the policy to recover and complete tasks reliably.

V. CONCLUSION

In this paper, we introduce history-awareness into visuomotor policies, enabling more informed decision-making across diverse manipulation tasks. Inspired by how humans rely on past interactions, we propose to represent history as object-centric point tracks, a compact yet expressive form that integrates seamlessly with different policy backbones. We instantiate this design in various visuomotor policies, and extensive experiments show that these history-aware policies consistently outperform both their base counterparts and prior history-aware methods, achieving over 90% decision accuracy on nearly all tasks. Among them, *HistRISE* achieves state-of-the-art performance across the board, including on

Tracker	Tracking Accuracy	<i>Guess-Hard</i>	
		SR	DA
(2D) CoTracker [26]	56%	55%	70%
(3D) TAPI3D [67]	93%	85%	95%

TABLE II: Tracking Quality and Policy Performance. Better tracking quality directly leads to better history-aware decision accuracy and overall policy performance.

Tracking Type	w/wo Aug.	<i>Swap-Hard</i>			Avg. Freq.
		SR	DA	AL	
Synchronous		80%	97%	6.55	0.93 Hz
Asynchronous		0%	65%	1.85	5.44 Hz
Asynchronous	✓	80%	97%	6.50	5.44 Hz

TABLE III: Performance-Efficiency Tradeoff and Optimization. Asynchronous tracking with random-dropping augmentation preserves performance and enhances efficiency. Frequencies are reported for the full tracking and policy inference process.

long-horizon tasks such as *Swap-Hard*. We further demonstrate that our object-centric point track history representation naturally supports pre-loaded memory, as validated in the *Guess-Easy* and *Guess-Hard* tasks, broadening the applicability of visuomotor policies. Finally, we analyze the impact of tracking quality on history-awareness and overall policy performance, and introduce asynchronous tracking with point-drop augmentation to maintain performance while improving efficiency by $5.8\times$.

Together, these results establish history-aware visuomotor policies as a general and powerful framework for robotic manipulation. Future work will explore (1) maintaining history-awareness while preserving the policies' failure recovery ability, and (2) extending our framework with pre-loaded memory to broader applications, like in-context learning.

ACKNOWLEDGEMENT

This work was supported in part by the National Natural Science Foundation of China (No. 62595774), Science and Technology Major Project of Jiangsu Province (No. BG2024041), Shanghai Artificial Intelligence Laboratory, XPLOER PRIZE grants.

REFERENCES

- [1] Brian P Ackerman. "Fuzzy-Trace Theory: A Grand Theory". In: *Learning and Individual Differences* 7.2 (1995), pp. 77–81.
- [2] Anurag Arnab et al. "ViViT: A Video Vision Transformer". In: *ICCV*. 2021, pp. 6836–6846.
- [3] Max Bain et al. "Frozen in Time: A Joint Video and Image Encoder for End-To-End Retrieval". In: *ICCV*. 2021, pp. 1728–1738.
- [4] Amir Bar et al. "Visual Prompting via Image Inpainting". In: *NeurIPS*. 2022, pp. 25005–25017.
- [5] Kevin Black et al. " π_0 : A Vision-Language-Action Flow Model for General Robot Control". In: *RSS*. 2025.
- [6] Anthony Brohan et al. "RT-1: Robotics Transformer for Real-World Control at Scale". In: *RSS*. 2023.
- [7] Qingwen Bu et al. "Learning to Act Anywhere with Task-centric Latent Actions". In: *RSS*. 2025.
- [8] Boyuan Chen et al. "Diffusion Forcing: Next-token Prediction Meets Full-Sequence Diffusion". In: *NeurIPS*. 2024, pp. 24081–24125.
- [9] Shizhe Chen et al. "History Aware Multimodal Transformer for Vision-and-Language Navigation". In: *NeurIPS*. 2021, pp. 5834–5847.

- [10] Yuanpei Chen et al. “Sequential Dexterity: Chaining Dexterous Policies for Long-Horizon Manipulation”. In: *CoRL*. 2023, pp. 3809–3829.
- [11] Cheng Chi et al. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion”. In: *RSS*. 2023.
- [12] Cheng Chi et al. “Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots”. In: *RSS*. 2024.
- [13] Carl Doersch et al. “TAPIR: Tracking Any Point with Per-Frame Initialization and Temporal Refinement”. In: *ICCV*. 2023, pp. 10061–10072.
- [14] Carl Doersch et al. “Bootstap: Bootstrapped training for tracking-any-point”. In: *ACCV*. 2024, pp. 3257–3274.
- [15] Hao-Shu Fang et al. “RH20T: A Comprehensive Robotic Dataset for Learning Diverse Skills in One-Shot”. In: *ICRA*. 2024, pp. 653–660.
- [16] Haoquan Fang et al. “SAM2Act: Integrating Visual Foundation Model with a Memory Architecture for Robotic Manipulation”. In: *ICML*. 2025.
- [17] Hongjie Fang et al. “AirExo-2: Scaling up Generalizable Robotic Imitation Learning with Low-Cost Exoskeletons”. In: *CoRL*. 2025.
- [18] Xiaolin Fang et al. “KALM: Keypoint Abstraction using Large Models for Object-Relative Imitation Learning”. In: *ICRA*. 2025, pp. 8307–8314.
- [19] Albert Gu and Tri Dao. “Mamba: Linear-Time Sequence Modeling with Selective State Spaces”. In: *COLM*. 2024.
- [20] Pierre-Louis Guhur et al. “Instruction-Driven History-Aware Policies for Robotic Manipulations”. In: *CoRL*. 2023, pp. 175–187.
- [21] Siddhant Haldar and Lerrel Pinto. “Point Policy: Unifying Observations and Actions with Key Points for Robot Manipulation”. In: *CoRL*. 2025.
- [22] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. “Particle video revisited: Tracking through occlusions using point trajectories”. In: *ECCV*. 2022, pp. 59–75.
- [23] Bo He et al. “MA-LMM: Memory-Augmented Large Multimodal Model for Long-Term Video Understanding”. In: *CVPR*. 2024, pp. 13504–13514.
- [24] Menglin Jia et al. “Visual Prompt Tuning”. In: *ECCV*. 2022, pp. 709–727.
- [25] Yang Jin et al. “Video-LaVIT: Unified Video-Language Pre-training with Decoupled Visual-Motional Tokenization”. In: *ICML*. 2024, pp. 22185–22209.
- [26] Nikita Karaev et al. “CoTracker: It Is Better to Track Together”. In: *ECCV*. 2024, pp. 18–35.
- [27] Nikita Karaev et al. “Cotracker3: Simpler and better point tracking by pseudo-labelling real videos”. In: *arXiv preprint arXiv:2410.11831* (2024).
- [28] Moo Jin Kim et al. “OpenVLA: An Open-Source Vision-Language-Action Model”. In: *CoRL*. 2024, pp. 2679–2713.
- [29] Hanna Kurniawati. “Partially observable markov decision processes and robotics”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 5.1 (2022), pp. 253–277.
- [30] Mara Levy et al. “P3-PO: Prescriptive Point Priors for Visuo-Spatial Generalization of Robot Policies”. In: *ICRA*. 2025, pp. 4167–4174.
- [31] Hao Li et al. “CronusVLA: Transferring Latent Motion Across Time for Multi-Frame Prediction in Manipulation”. In: *arXiv preprint arXiv:2506.19816* (2025).
- [32] Xiaoqi Li et al. “Object-Centric Prompt-Driven Vision-Language-Action Model for Robotic Manipulation”. In: *CVPR*. 2025, pp. 27638–27648.
- [33] Xinghang Li et al. “Towards Generalist Robot Policies: What Matters in Building Vision-Language-Action Models”. In: *arXiv preprint arXiv:2412.14058* (2024).
- [34] Xinghang Li et al. “Vision-Language Foundation Models as Effective Robot Imitators”. In: *ICLR*. 2024.
- [35] Yishu Li et al. “Learn from What We HAVE: History-Aware VErifier that Reasons about Past Interactions Online”. In: *CoRL*. 2025.
- [36] Ruyang Liu et al. “ST-LLM: Large Language Models Are Effective Temporal Learners”. In: *ECCV*. 2024, pp. 1–18.
- [37] Oier Mees et al. “CALVIN: A Benchmark for Language-Conditioned Policy Learning for Long-Horizon Robot Manipulation Tasks”. In: *IEEE Robotics Autom. Lett.* 7.3 (2022), pp. 7327–7334.
- [38] Nicolas Meuleau et al. “Learning finite-state controllers for partially observable environments”. In: *UAI*. 1999, pp. 427–436.
- [39] Zhuochen Miao et al. “Knowledge-Driven Imitation Learning: Enabling Generalization Across Diverse Conditions”. In: *IROS*. 2025.
- [40] Tuan Duc Ngo et al. “DELTA: Dense Efficient Long-Range 3D Tracking for Any Video”. In: *ICLR*. 2025.
- [41] Octo Model Team et al. “Octo: An Open-Source Generalist Robot Policy”. In: *RSS*. 2024.
- [42] Norman Di Palo and Edward Johns. “Keypoint Action Tokens Enable In-Context Imitation Learning in Robotics”. In: *RSS*. 2024.
- [43] William Peebles and Saining Xie. “Scalable Diffusion Models with Transformers”. In: *ICCV*. 2023, pp. 4172–4182.
- [44] Dean Pomerleau. “ALVINN: An Autonomous Land Vehicle in a Neural Network”. In: *NeurIPS*. 1988, pp. 305–313.
- [45] Nikhila Ravi et al. “SAM 2: Segment Anything in Images and Videos”. In: *ICLR*. 2025.
- [46] Juntao Ren et al. “Motion Tracks: A Unified Representation for Human-Robot Transfer in Few-Shot Imitation Learning”. In: *ICRA*. 2025, pp. 8802–8810.
- [47] Valerie F Reyna and Charles J Brainerd. “Fuzzy-Trace Theory: An Interim Synthesis”. In: *Learning and Individual Differences* 7.1 (1995), pp. 1–75.
- [48] Richard Roberts et al. “Optimal and Interactive Keyframe Selection for Motion Capture”. In: *SIGGRAPH Asia*. 2018, pp. 1–4.
- [49] Hao Shi et al. “MemoryVLA: Perceptual-Cognitive Memory in Vision-Language-Action Models for Robotic Manipulation”. In: *arXiv preprint arXiv:2508.19236* (2025).
- [50] Min Shi et al. “Slow-Fast Architecture for Video Multi-Modal Large Language Models”. In: *arXiv preprint arXiv:2504.01328* (2025).
- [51] Enxin Song et al. “MovieChat: From Dense Token to Sparse Memory for Long Video Understanding”. In: *CVPR*. 2024, pp. 18221–18232.
- [52] Xi Tang et al. “Adaptive Keyframe Sampling for Long Video Understanding”. In: *CVPR*. 2025, pp. 29118–29128.
- [53] Peter Tong et al. “Cambrian-1: A Fully Open, Vision-Centric Exploration of Multimodal LLMs”. In: *NeurIPS*. 2024.
- [54] Marcel Torne et al. “Learning Long-Context Diffusion Policies via Past-Token Prediction”. In: *CoRL*. 2025.
- [55] Vitalis Vosylius and Edward Johns. “Instant Policy: In-Context Imitation Learning via Graph Diffusion”. In: *ICLR*. 2025.
- [56] Bo Wang et al. “SceneTracker: Long-Term Scene Flow Estimation Network”. In: *TPAMI* 47.9 (2025), pp. 8305–8312.
- [57] Chenxi Wang et al. “RISE: 3D Perception Makes Real-World Robot Imitation Simple and Effective”. In: *IROS*. 2024, pp. 2870–2877.
- [58] Jianyuan Wang et al. “VGGT: Visual Geometry Grounded Transformer”. In: *CVPR*. 2025, pp. 5294–5306.
- [59] Chuan Wen et al. “Any-point Trajectory Modeling for Policy Learning”. In: *RSS*. 2024.
- [60] Shangning Xia et al. “CAGE: Causal Attention Enables Data-Efficient Generalizable Robotic Manipulation”. In: *ICRA*. 2025, pp. 13242–13249.
- [61] Yuxi Xiao et al. “SpatialTracker: Tracking Any 2D Pixels in 3D Space”. In: *CVPR*. 2024, pp. 20406–20417.
- [62] Mengda Xu et al. “Flow as the Cross-domain Manipulation Interface”. In: *CoRL*. 2024, pp. 2475–2499.
- [63] Mingze Xu et al. “SlowFast-LLaVA-1.5: A Family of Token-Efficient Video Large Language Models for Long-Form Video Understanding”. In: *COLM*. 2025.
- [64] Ziyu Yao et al. “Countllm: Towards generalizable repetitive action counting via large language model”. In: *CVPR*. 2025, pp. 19143–19153.
- [65] Sicheng Yu et al. “Frame-Voyager: Learning to Query Frames for Video Large Language Models”. In: *ICLR*. 2025.
- [66] Wentao Yuan et al. “RoboPoint: A Vision-Language Model for Spatial Affordance Prediction in Robotics”. In: *CoRL*. 2024, pp. 4005–4020.
- [67] Bowei Zhang et al. “TAPI3D: Tracking Any Point in Persistent 3D Geometry”. In: *arXiv preprint arXiv:2504.14717* (2025).
- [68] Haoji Zhang et al. “Flash-VStream: Efficient Real-Time Understanding for Long Video Streams”. In: *ICCV*. 2025.
- [69] Jiazhao Zhang et al. “NaVid: Video-based VLM Plans the Next Step for Vision-and-Language Navigation”. In: *RSS*. 2024.
- [70] Tony Z Zhao et al. “Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware”. In: *RSS*. 2023.
- [71] Ruijie Zheng et al. “TraceVLA: Visual Trace Prompting Enhances Spatial-Temporal Awareness for Generalist Robotic Policies”. In: *ICLR*. 2025.
- [72] Yulin Zhou et al. “MTIL: Encoding Full History with Mamba for Temporal Imitation Learning”. In: *arXiv preprint arXiv:2505.12410* (2025).