SELF-IMPROVING VIDEO-TO-ACTION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Video generative models are an emerging tool for embodied agents: they can serve as scalable visual planners, generating videos that depict how to complete previously unseen tasks. However, executing these plans typically requires a video-to-action or "inverse dynamics" model trained on large, human-collected video-action datasets, severely limiting scalability. We introduce Mirror, a self-supervised approach that learns video-to-action mappings through environment interactions, without prior action data or hand-crafted rewards. Mirror learns by attempting to follow a visual plan by executing its best guess of the underlying actions in the environment. Mirror then trains on the resulting video-action pairs, incrementally self-improving and following the visual planner more and more closely. Across Atari, OGBench, and DeepMind Control, Mirror achieves overall performance comparable to leading learning-from-observation methods with learned video plans, and near-expert performance with oracle plans. On Minecraft BASALT and text-conditioned tasks, it outperforms methods with access to expert actions or hand-engineered task knowledge. Using generated video plans, it surpasses VPT-BC in survival mode, even though VPT-BC was trained on thousands of hours of action-labeled human gameplay. With oracle video plans, it learns to acquire diamonds without task rewards, using an order of magnitude fewer environment interactions than DreamerV3. Our results suggest that advances in video models can translate into increasingly capable agents without a corresponding need for hand-crafted rewards or large collections of expert action data. Please find video and other qualitative results on our anonymous website.

1 INTRODUCTION

Video models are emerging as scalable learners of the physical world. Trained on vast collections of unstructured web video, they absorb knowledge about objects, environments, human behavior, and how tasks unfold over time.

Recent work has begun to harness this knowledge for robot control by using video models as *visual planners* (e.g., Du et al., 2023; Chen et al., 2025; Ye et al., 2026; Rhoda, 2026). Given past frames and optional task conditioning, these models generate a video depicting how the task could be completed. An embodiment-specific video-to-action (V2A) or "inverse dynamics" model then translates the generated visual trajectory into robot actions. This has yielded impressive generalization across unseen tasks, objects, environments, and embodiments.

However, training the V2A model still requires significant amounts of expert-collected, paired video-action data from oracle motion planners or human teleoperation. Moreover, agents cannot easily generate such video-action data by rolling out random policies in the environment, as critical environment transitions—such as executing a robot grasp or crafting a diamond pickaxe in Minecraft—are unreachable by a random policy. Today, video planners thus inherit the same scalability bottleneck as video-language-action models: a reliance on expensive expert teleoperation data.

In this paper, we introduce Mirror, a *fully self-supervised* method for learning to execute the policy encoded by a video model only by interacting with the world, without any expert video-action data. We assume access to an expert video planner that, given the current agent observation, can generate an expert-level visual plan. Under this assumption, we show that no further expert action supervision or hand-crafted rewards are necessary to achieve expert-level performance. Mirror thus represents a bet on the continued scaling of video models: as visual planners improve, their capabilities can be translated into actions without a corresponding increase in expert action data.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Our key idea is to let a randomly initialized video-to-action model self-improve against the expert video planner. Given the current agent state, we query the planner to generate a video plan. The current video-to-action model then predicts the corresponding actions and the agent executes them in the environment. Interactions produce new paired video-action data, which is used to improve the video-to-action model. Critically, this data is not random but concentrated at the frontier of the agent's current capability to follow the video plan. Repeating this process allows the agent to follow progressively more of the plan until it faithfully reproduces the planner's behavior. Mirror does not require rewards, privileged state information, verifiers, or trajectory filtering. Furthermore, we show that our self-improving approach can outperform training on expert actions.

We evaluate Mirror in two settings: (i) the challenging Minecraft BASALT competition tasks (Shah et al., 2021) and Minecraft Survival challenge (Guss et al., 2019); and (ii) standard decision-making benchmarks, including Atari (Bellemare et al., 2013), OGBench (Park et al., 2025a), and DeepMind Control (Tassa et al., 2018). To the best of our knowledge, no prior method without either hand-crafted rewards or expert action labels achieves non-zero success on these Minecraft tasks. Mirror outperforms baselines, including those with privileged information, and further exploits the emergent in-context learning of large video models, improving faster than a baseline that cannot leverage context across attempts.

Our contributions are summarized as follows:

- Mirror, a scalable agent that learns long-horizon tasks by self-improving against a pre-trained video planner, without task rewards or expert video-action data.
- A heavily adapted 14B-parameter NVIDIA Cosmos3 video planner with efficient long-context planning via causal KV-caching. Code and weights will be public.
- The first non-zero success rates on the BASALT Minecraft suite without expert video-action data or hard-coded actions.
- On Minecraft survival, Mirror obtains a wooden pickaxe and cobblestone (approx. 8 minutes of gameplay) after only 13 hours of in-game training. With an oracle video planner, Mirror successfully obtains diamonds.
- On a suite of 24 common RL tasks, Mirror is competitive with baselines using a generative video planner, and, when paired with an oracle video planner, nearly matches expert returns.

2 Related Work

Video planners. A nascent line of work leverages video generative models to generate a visual plan of how a given task should be solved, and then translates it into actions (Du et al., 2023; Ko et al., 2024; Liang et al., 2024a; Luo & Du, 2024; Bharadhwaj et al., 2024). As video models can be pre-trained on web-scale data, robots can inherit impressive generalization capabilities (Chen et al., 2025; Li et al., 2026). However, translating videos into robot actions generally requires training on paired video-action data (Du et al., 2023; Tian et al., 2025; Pai et al., 2025). Our agent learns to act from video and self-play in an interactive environment alone, reaching imagined futures without a single action label.

Learning from demonstrations. We focus on imitation from observation, where an agent learns from expert observations without action labels or task rewards (Torabi et al., 2018). Adversarial approaches match expert observation distributions through learned discriminator rewards (Torabi et al., 2019; Liu et al., 2023; Giammarino et al., 2024), but can be unstable and sensitive to hyperparameters (Huang et al., 2024). Iterative behavioral cloning alternates between training a V2A model on the agent's own experience and imitating V2A-labeled demonstrations (Torabi et al., 2018), but collecting informative transitions remains an exploration challenge on long-horizon tasks (Schmidt & Jiang, 2024). Latent-action approaches such as LAPO and LAPA still ground their latent actions with action labels or rewards (Schmidt & Jiang, 2024; Ye et al., 2025). Mirror instead uses a video planner to guide exploration, collecting V2A training data without action labels, rewards, or adversarial optimization.

Minecraft agents. Minecraft has become a standard benchmark for long-horizon decision making, ranging from the fuzzy objectives of BASALT (Shah et al., 2021) to the canonical MineRL challenges (Guss et al., 2019). Existing agents rely on some combination of expert action demonstrations and hand-designed rewards. For example, VPT (Baker et al., 2022) trains a behavioral cloning

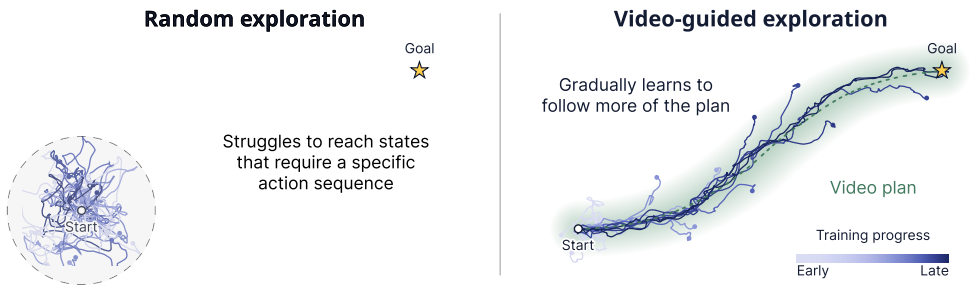

Figure 1: **Enabling long-horizon exploration with a video planner. Left:** Random exploration struggles to reach later stages of long-horizon tasks that require specific sequences of actions. **Right:** The video planner guides exploration. As the V2A model learns from achieved transitions, the agent follows the plan more accurately and reaches progressively later stages of the task. These new transitions further improve the model, eventually enabling the agent to execute the full plan.

(BC) policy using about 2,000 hours of action-labeled gameplay recordings from paid contractors, then fine-tunes the policy with about 230,000 hours of online RL. While DreamerV3 (Hafner et al., 2025a) trains a Minecraft RL agent with only 1,400 hours of online RL, they simplify parts of the Minecraft action space and still require hand-designed reward. DreamerV4 (Hafner et al., 2025b) likewise uses reward, and also requires expert actions to train a video world model. Pan-4B (Pantograph, 2026) trains a goal-conditioned RL policy on 500k hours of internet video, yet still post-trains on 2,000 hours of expert action labels. To make long-horizon tasks more tractable, many works improve upon Minecraft policies with language models. MineDojo (Fan et al., 2022) trains PPO agents on shaped rewards and a learned video-language critic, while STEVE-1 (Lifshitz et al., 2023), DEPS (Wang et al., 2023), and JARVIS-1 (Wang et al., 2024) use language-based planning.

Prior agents that solve long-horizon tasks thus depend on hand-crafted rewards, expert video-action data, or both; Mirror instead only assumes access to an expert-level video planner.

3 Method

We seek a task-conditioned policy $\pi(a_t|o_{\leq t}, l_{\texttt{task}})$, which predicts actions a_t based on past observations $o_{\leq t}$ and a task $l_{\texttt{task}}$. We model π indirectly with two components: a video planner (VP) and a video-to-action (V2A) model[1]. The VP predicts future observations $\hat{o}_{t+1}$ based on historical observations $o_{\leq t}$ and a task $l_{\texttt{task}}$, while the V2A model predicts actions a_t given the combined video sequence $(o_{\leq t}, \hat{o}_{t+1})$. Sequentially sampling from the two models yields the agent's policy:

$$a_t \sim p_{\texttt{V2A}}(\cdot|o_{\leq t}, \hat{o}_{t+1}) \quad \longleftarrow \quad \hat{o}_{t+1} \sim p_{\texttt{VP}}(\cdot \mid o_{\leq t}, l_{\texttt{task}})$$

Only the V2A model requires embodiment-specific action supervision, while the VP can be trained on internet-scale, action-free video. Crucially, every plan is conditioned on the frames the agent *actually* observed, never on its previous plans. The VP therefore always plans from wherever the agent ended up, so deviations from earlier plans can be corrected. Section 3.1 shows how the V2A model can be trained from scratch on the data the agent collects while following these plans, and Section 3.2 describes how we make a generative video model VP practical for an online agent.

3.1 Self-Improving Video-to-Action Models

Our key insight is to use a pretrained video planner to guide data collection for training the V2A model. The model learns from the agent's own rollouts, without expert action labels. As summarized in Algorithm 1, training alternates between data collection and model updates. During data collection, we sample an action from the agent's two-stage policy π and add exploration noise $n_t \sim \pi_{\texttt{explore}}$. We execute the resulting action a_t in the environment and store the transition $(o_{\leq t}, a_t, o_{t+1})$ in a replay buffer. To update the model, we sample batches from the buffer and train the V2A model to predict a_t from $(o_{\leq t}, o_{t+1})$. Because the targets are conditioned on the *achieved* observation o_{t+1} rather than the planned one $\hat{o}_{t+1}$, every transition is a valid training example, even when execution deviates from the plan. Further details can be found in Algorithm 1.

[1] V2A models are also commonly referred to as *inverse dynamics models (IDMs)*.

Algorithm 1 Learning without action labels: VP-guided self-training of the V2A model

Require: video planner $p_{\texttt{VP}}$, environment $\mathcal{E}$, exploration distribution $\pi_{\texttt{explore}}$, horizon T.
Initialize V2A model $p_{\texttt{V2A}}$ with parameters θ from scratch; replay buffer $\mathcal{D} \leftarrow \emptyset$.
Repeat until convergence:

Data collection episode

Reset $\mathcal{E}$; observe o_0; initialize $\tau \leftarrow (o_0)$ and $t \leftarrow 0$
while $t < T$ **do**
$\hat{o}_{t+1} \sim p_{\texttt{VP}}(\cdot \mid o_{\leq t}, l_{\texttt{task}})$
$a_t \sim p_{\texttt{V2A}}(\cdot \mid o_{\leq t}, \hat{o}_{t+1}; \theta)$ and $n_t \sim \pi_{\texttt{explore}}$
$a_t \leftarrow a_t + n_t$
Execute a_t in $\mathcal{E}$ and observe o_{t+1}
$\tau \leftarrow (\tau, a_t, o_{t+1})$ and $t \leftarrow t + 1$
end while
$\mathcal{D} \leftarrow \mathcal{D} \cup \{\tau\}$

V2A model learning

Sample trajectories $\mathcal{B} \sim \mathcal{D}$ and timesteps t from each trajectory
$\theta \leftarrow \theta + \eta \nabla_\theta \sum_{\tau \in \mathcal{B}} \log p_{\texttt{V2A}}(a_t \mid o_{\leq t}, o_{t+1}; \theta)$

Exploration via stochastic V2A. Figure 1 illustrates how this process guides exploration. Early in training, the V2A model often fails to follow the plan, so the realized next observation o_{t+1} may differ substantially from the video plan $\hat{o}_{t+1}$. However, even these failed attempts provide useful training examples: the V2A model still learns to predict the action that produced the observed transition. As the model improves, the agent can follow the planner more closely and reach later stages of the task. In turn, these rollouts provide data to improve the V2A model at the newly reached stages. Repeating this cycle allows the agent to progress further and further through the task. In contrast, purely random exploration often struggles to reach these stages because they require long sequences of consistently coordinated actions.

If the V2A model were trained only on data from its own rollouts, it would collapse to predicting the same action regardless of the observations. We hence make the V2A model stochastic during data collection by adding noise to its predictions. Stochasticity also plays a second, central role: it is what carries the agent past the stages it cannot yet execute (Figure 2). Where the V2A model is already competent, action noise is harmless: while the agent may drift off the currently active video plan, the video planner re-plans from the state actually reached, so error does not accumulate and the agent progresses. At a milestone the V2A model has not yet learned, such as the first time a block has to be mined, the V2A model predicts the wrong action and thus the agent fails to produce the planned transition. The planner, however, can keep re-planning and thus, keep proposing the same transition. The agent is thus held at the milestone. Here, the stochastic V2A model is what overcomes the roadblock: eventually, the correct random action is sampled, the resulting transition enters the replay buffer, the V2A model learns it, and the agent accomplishes the milestone. Random exploration, in contrast, would first have to stumble upon the entire action sequence leading to the milestone before it could even attempt the critical action, see Fig. 1.

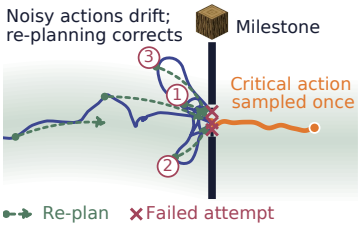

Figure 2: Re-planning holds the agent at the milestone until noise supplies the critical action.

3.2 A Video Planner for Online Learning and Closed-Loop Control

Following plans in closed-loop places three demands on the VP. First, it must re-plan often and cheaply. Second, it must run over long episodes, often well over a minute in our Minecraft benchmarks. Third, it must stay on task even when the agent's failures push it off the distribution of expert videos. To avoid re-encoding the history at every re-plan, we use a block-causal architecture: each observed frame is encoded once and its keys and values are cached, while the plan is denoised

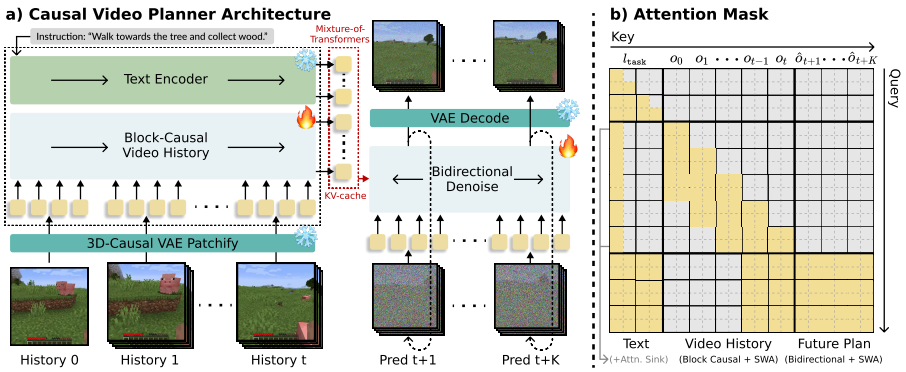

Figure 3: **The causal video planner. (a)** Tokenized frames and text are fed into the 14B-parameter Cosmos3-Nano (NVIDIA, 2026) model, forming a causal stream of contexts whose keys and values are cached and reused. The denoiser generates K future frames while cross-attending to the context KV cache, so each observation and instruction is encoded only once. **(b)** Attention mask. Video history is encoded block-causally. The future plan attends to this history and bidirectionally to itself. Sliding-window attention lets the planner generalize beyond the window while keeping the KV cache within a constant budget.

bidirectionally against this cache. Sliding-window attention over a fixed-size cache keeps the cost per plan constant, so a planner trained on 32 s clips runs on over 290 s of context without degradation. Finally, we propose task-aware guidance, which leverages diffusion guidance (Ho & Salimans, 2022; Karras et al., 2024) to guide the task-conditional video planner with a task-unconditional one, which significantly improved the task adherence of the video planner in our experiments. Figure 3 and Appendix A detail the architecture, training, and inference.

4 Experiments

We empirically validate Mirror's ability to solve tasks across Minecraft (Section 4.1) and traditional RL environments (Section 4.2). In Section 4.1.2, we study what emergent abilities generative video planners can provide with scale and how to leverage them. Section 4.1.3, with an oracle VP, asks whether our self-improvement algorithm can scale to longer horizon tasks, beyond the capabilities of current-day video planners. Finally, in Section 4.2, we investigate how our method generalizes to various environments and control settings. Please find video results on our anonymous website.

4.1 Solving Open-Ended Challenges in Minecraft

On Minecraft, we pre-train a video planner on 2,946 hours of automatically captioned videos from the VPT dataset (Baker et al., 2022), equipping our video planner with general physical and procedural knowledge of Minecraft. The VP is finetuned further for BASALT and Minecraft Survival Mode. With this VP, we are able to train our self-improving V2A model from scratch with self-play. Our agent operates in the MineRL 1.0 (Guss et al., 2019) action space, consisting of over 8,000 keyboard and mouse combinations and does not rely on simplified action spaces or game mechanics.

4.1.1 Learning to Solve Fuzzy Tasks from Human Demonstration

We first evaluate Mirror on the Minecraft BASALT (Benchmark for Agents that Solve Almost-Lifelike Tasks) (Shah et al., 2021) challenge. Unlike traditional RL environments, its task objectives are intentionally fuzzy and hard-to-verify. For instance, consider the MakeWaterfall task:

> *Place a water bucket to make a beautiful waterfall. Then take an aesthetic picture of it by moving to a good location, positioning the player to have a nice view of the waterfall.*

Our method is particularly suited to this kind of task where reward specification is difficult. It learns directly from fuzzy human demonstration videos without any reward supervision. For each task, we finetune our video planner on the corresponding BASALT human video demonstrations. We then train our V2A model from scratch only via self-play. After only 5 hours (360,000 environment steps)

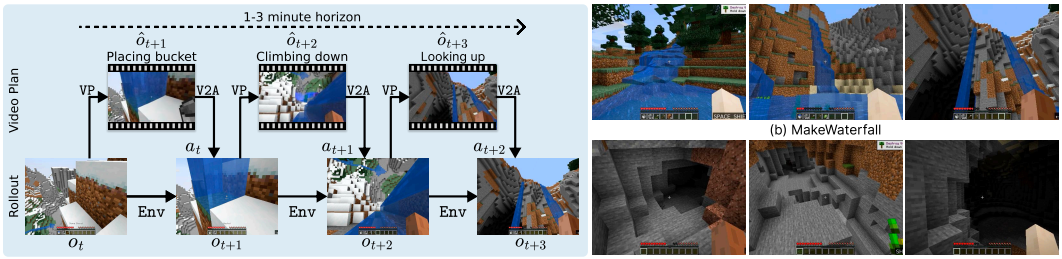

Figure 4: **Learning from human demonstration.** Mirror can complete a variety of open-ended BASALT tasks from self-play without any expert action labels or reward supervision. **(a)** The VP inherits long-horizon planning capabilities learned from video pre-training, while the V2A model learns to ground the video plans to action via online interaction. **(b)** The constructed waterfalls are diverse and similar to human creations. **(c)** Mirror achieves near-human performance on FindCave.

Table 1: **Solving fuzzy tasks from BASALT.** Success rates for methods without action labels (our self-improving method and VIPER (Escontrela et al., 2023)), action-supervised and heuristic BASALT baselines (Milani et al., 2022), vision-language-action (VLA) models, and humans.

	No action labels		**With action labels**				**Vision-language-action (VLA) models**				**Heuristic**	**Human**
Task	Mirror (Ours)	VIPER	Mirror (Ours)	KAIROS	BC (VPT-SFT)	UniTeam	STEVE-1	GROOT	JARVIS-VLA	OpenHA	GoUp	
FindCave	**82%**	0%	**91%**	27%	27%	64%	45%	45%	36%	45%	36%	95%
MakeWaterfall	**36%**	0%	**45%**	0%	0%	0%	0%	0%	0%	0%	36%†	100%
CreateAnimalPen	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	18%†	95%
BuildVillageHouse	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	91%†	100%

†Obtained by pre-programmed action sequences rather than learned behavior.

of in-game play, our method is already able to outperform the baselines. Figure 4 demonstrates how Mirror solves these tasks. Using the same VP, we also evaluate a V2A model pretrained on expert actions from the VPT dataset.

In Table 1, we compare our method against an action free baseline, VIPER (Escontrela et al., 2023), a behavioral cloning (BC) policy fine-tuned from VPT, action-supervised methods and heuristic methods from the BASALT challenge (Milani et al., 2022), and vision-language-action (VLA) Minecraft models: STEVE-1 (Lifshitz et al., 2023), GROOT (Cai et al., 2024), JARVIS-VLA (Li et al., 2025), OpenHA (Wang et al., 2025). Of the BASALT challenge submissions, KAIROS is an inverse RL method based on IQ-Learn (Garg et al., 2021), UniTeam augments the VPT BC policy with search, and GoUp executes pre-programmed scripts (Milani et al., 2022).

Even without expert actions, Mirror outperforms all baselines that do not rely on hand-crafted heuristics. Only 5 hours of self-play was required to achieve the highest success rate on FindCave and MakeWaterfall tasks. To the best of our knowledge, Mirror is the first method to achieve non-zero success on MakeWaterfall from learned behavior alone, without classifiers or hard-coded actions. The baselines fail in different ways (Figure 15): the action-supervised methods (KAIROS, BC and UniTeam) drift out of distribution on multi-step tasks such as MakeWaterfall, the VLA models struggle to follow the task descriptions, and GoUp's hard-coded scripts are brittle, placing, for instance, the water bucket but freezing once underwater, whereas Mirror escapes the water using knowledge learned from human gameplay. CreateAnimalPen and BuildVillageHouse remain unsolved by all learning-based methods. Lastly, while Mirror does not require action labels, it can still leverage them for a further small performance gain.

4.1.2 Emergent Capabilities of the Video Planner

Training the video planner on internet-scale video data unlocks novel capabilities without requiring additional fine-tuning for each task. We explore such emergent behaviors in two settings.

Text-conditioned video planning. As a proof of concept, we evaluate text control with various goals such as "chop down the tree" or "dig into the ground." We evaluate our text-conditioned VP with a V2A model pre-trained on random walks, a V2A model pre-trained on the VPT expert actions,

World	Instruction	No action labels: Mirror (random)	No action labels: Mirror (self-improv.)	No action labels: VIPER	With action labels: BC	With action labels: Mirror (supervised)
Woodland	Swim	70%	80%	0%	10%	80%
	Chop tree	0%	100%	0%	0%	80%
	Ascend hill	100%	90%	0%	0%	90%
	Dig ground	0%	90%	0%	0%	100%
Camp	Ascend hill	70%	90%	0%	0%	90%
	Enter cave	90%	50%	0%	0%	60%
	Light cave	0%	0%	0%	0%	20%
	Mine stone	0%	10%	0%	0%	40%
	Place blocks	70%	30%	20%	20%	70%
	Dig ground	20%	100%	0%	0%	80%
	Total	42%	**64%**	2%	3%	**71%**

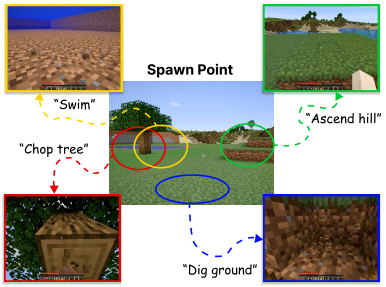

Figure 5: **Text-conditioned video planning.** We evaluate text-conditioned control on 10 tasks across two unseen environments (see Appendix B.1.3). Each task is specified only through a text prompt, with no task-specific fine-tuning. VIPER (Escontrela et al., 2023) is included as a reward-free video-based baseline. Combined with our self-improving V2A approach, the agent achieves high success rates although it has never seen expert demonstrations on these tasks. *Left:* Success rates of each method. *Right:* Text guidance leads the agent to complete different tasks.

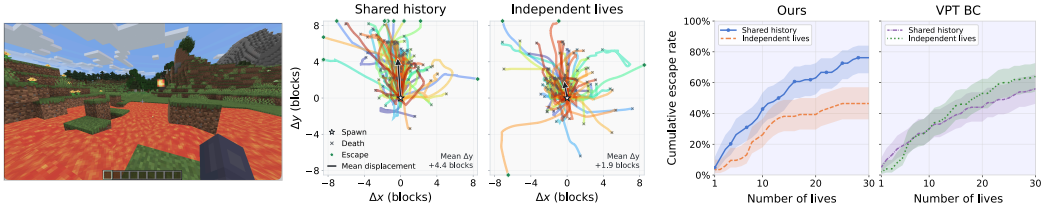

(a) Lava escape challenge

(b) In-context learning causes a common directional bias

(c) Our method leverages context to improve escape rates; VPT BC remains largely unchanged.

Figure 6: **In-context learning from failure.** **(a)** We challenge our agent to escape a lava lake without fine-tuning. **(b)** Retaining context between deaths causes a directional bias in the agent's trajectories. **(c)** Shared history allows the agent to learn from its past failures in-context, leading to higher escape rates. Retaining recurrent state does not provide a comparable benefit for VPT BC.

our self-improving V2A model. We also compare our method to VIPER (Escontrela et al., 2023), which similarly does not use action labels, and a text-conditioned BC policy. Shown in Figure 5, Mirror reliably harvests trees, climbs, and digs underground. Some tasks, however, are difficult across all V2A variants. We see that the V2A model trained with self-improvement generally outperforms the V2A model trained on random walks. Although the self-improving V2A model has never seen expert demonstrations, it is still competitive with the V2A model supervised on expert actions. Appendix B.1.3 provides further information on the evaluation protocol. We interpret this result as evidence that text-conditioned video planning can provide useful task-level control.

In-context learning from failure. We examine if Mirror can utilize its history, $o_{\leq t}$, for in-context learning. In Figure 6, we challenge our agent to escape a lava lake. If the agent falls in the lava, it respawns at the origin. We compare a case where the agent's history is preserved between deaths, and a case when the history is reset after every death. The task is difficult: with fresh histories every life, the escape rate is only 3%. When video context is retained between lives, however, the agent leverages its past failures to improve its escape rate. Figure 6 (b) shows the resulting agent's trajectories. With shared history, the agent exhibits a clear directional bias, demonstrating that prior attempts influence the agent's behavior. With independent history, the agent moves in relatively uniform directions around the starting point. Retaining history does not provide a comparable benefit to the VPT BC policy, perhaps even hindering performance. This ablation shows that our video planning agent exhibits some emergent in-context adaptation.

4.1.3 Self-Improvement in Minecraft Survival Mode

While video generative planners are already highly capable, they remain limited on longer horizon tasks such as Minecraft survival which can last over 10 minutes. To decouple our self-improvement algorithm from the performance of the video planner, we study self-improvement with an *oracle video planner.* The oracle VP is a policy that takes in Mirror's prior observations $o_{\leq t}$, then acts in an

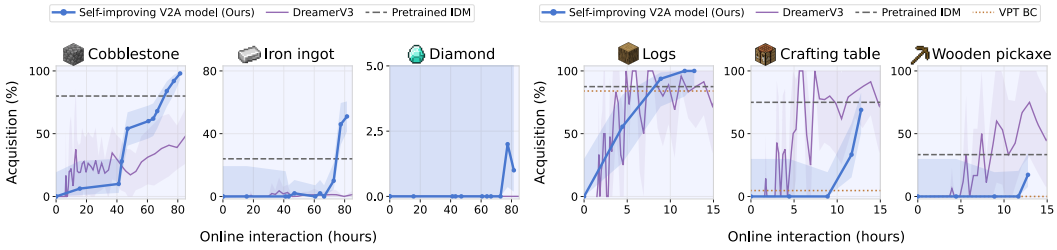

(a) Self-improvement with an oracle VP

(b) Self-improvement with a generative VP

Figure 7: **Self-improvement in Minecraft survival.** Guided by either an oracle or a generative VP, our V2A model improves with online experience, outperforming a V2A model pretrained on 2k hours of action-labeled human gameplay on several acquisition milestones. With the oracle VP, Mirror acquires diamonds within 80 hours of online interaction; with the generative VP, it learns to acquire a wooden pickaxe within 13 hours. These results suggest that Mirror can learn increasingly complex behaviors as video planners improve. Shaded regions denote 95% confidence intervals.

environment duplicated from Mirror to obtain predicted observations $\hat{o}_{t+1}$. This provides the same input-output behavior as a true video planner: $p_{\texttt{VP}}(\hat{o}_{t+1} \mid o_{\leq t})$.

Oracle VP. Using VPT's early game RL policy (Baker et al., 2022) as an oracle VP, we train a V2A model from scratch with self-play. Our results show that the V2A model rapidly improves with experience, acquiring progressively harder items. After only 80 hours of online interaction, the agent is able to acquire diamonds for the first time (Figure 7 (a)). RL agents take significantly longer due to the complexity of exploring the action space. Our approach is an order of magnitude more sample efficient than DreamerV3 (Hafner et al., 2025a), even though DreamerV3 uses simplified crafting. Our approach even outperforms a V2A model we pretrain on 2.5k hours of expert actions from the VPT dataset (Baker et al., 2022). This experiment suggests that the remaining bottleneck is therefore not Mirror's ability to learn from interactions, but the capabilities of the VP itself.

Generative VP. Having established that the self improvement is possible on Minecraft survival, we ask how far we can go with today's VPs. We fine-tune our VP on 1,000 hours of survival gameplay collected from the VPT RL agent and human demonstrations. Because the VAE struggles to reconstruct fine GUI details, sometimes failing to render the cursor entirely, we enlarge the inventory, add cursor trails and add click indicators to compensate for these limitations. After 13 hours of online interaction, the agent is able to acquire a wooden pickaxe, and acquires the crafting table over 50% of the time. The VPT BC policy is unable to craft the wooden pickaxe and only acquires the crafting table in 4.7% of episodes. While video planners are still in their infancy, we believe that they will only improve with time, and Mirror will improve alongside them.

4.2 Solving Robotics Tasks and Video Games

Setup. We evaluate Mirror on standard subsets of key synthetic decision-making benchmarks: OGBench (Park et al., 2025a), DeepMindControl (Tassa et al., 2018), and Atari (Bellemare et al., 2013; Escontrela et al., 2023). We benchmark with BCO (Torabi et al., 2018), which alternates between learning a V2A on policy rollouts and cloning its inferred actions on expert demonstrations, VIPER (Escontrela et al., 2023), which uses video prediction likelihoods as rewards, LAIfO (Giammarino et al., 2024), which learns a policy by adversarially matching expert and agent transitions, and the method proposed by Luo & Du (2024) (which we abbreviate as "Grounding") which learns a goal-conditioned policy to reach video-predicted subgoals. We evaluate Mirror with both a learned video planner and an oracle planner that supplies future frames from an RL expert. Unlike in our Minecraft setting, large-scale human play datasets are not available, and we thus found it essential to augment the generative planner's training data with recovery examples collected using DAgger (Ross et al., 2011). Please find an in-depth discussion of dataset construction, evaluated tasks, planner training, and baseline implementations in Appendix C.

Generative VP results. As shown in Figure 8, among methods evaluated without oracle planners, Mirror achieves the best performance on OGBench and Atari, only trailing LAIfO on DMC. LAIfO performs less consistently across domains, learning less efficiently on OGBench and achieving lim-

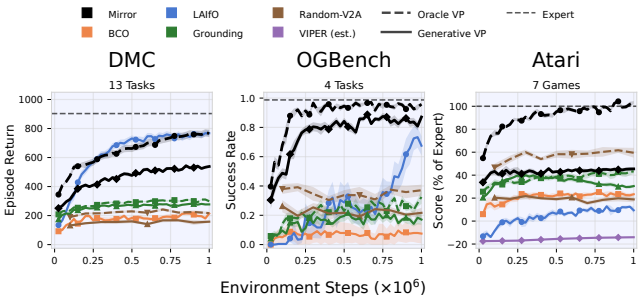

Final @ 1M		DMC return	OGBench success	Atari % expert
Oracle VP	Mirror	**767** ± **25**	**0.96** ± **0.01**	**103** ± **1**
	Grounding	299 ± 5	0.32 ± 0.05	43 ± 2
	Random-V2A	217 ± 9	0.37 ± 0.04	59 ± 4
Generative VP	Mirror	538 ± 7	0.87 ± 0.02	45 ± 1
	Grounding	277 ± 3	0.17 ± 0.04	30 ± 1
	Random-V2A	158 ± 8	0.22 ± 0.03	19 ± 2
Baselines	LAIfO	**767** ± **21**	0.68 ± 0.07	9 ± 3
	BCO	214 ± 10	0.07 ± 0.07	23 ± 2
Expert		903	0.99	100

Figure 8: **Imitation learning performance across three benchmark suites.** *Left:* learning curves on 13 DeepMind Control (DMC) tasks, four OGBench tasks, and seven Atari games. Colors identify methods; dashed method curves use an oracle VP and solid method curves use a generative VP. The horizontal dashed line marks expert performance. Curves average tasks equally, with shaded regions showing ±1 standard error across three seeds per task. The Atari VIPER curve is digitized from Escontrela et al. (2023) over its first 1M decisions and rescaled to % of expert. *Right:* final performance at 1M environment transitions (mean ± SE, equal task weights), separated by planner type.

ited returns on Atari. VIPER also achieves low returns on Atari, suggesting that using video model predictions directly provides better supervision than likelihood-based rewards.

The comparisons with BCO and Grounding help explain these gains. Although BCO also learns a V2A model from self-collected experience, it performs substantially worse, indicating the importance of using the video planner to guide data collection. Grounding improves on BCO, consistent with the benefit of video-guided exploration, but still remains substantially below Mirror. The remaining gap suggests an advantage to learning a V2A model that translates frame-to-frame transitions into actions over learning a goal-conditioned policy to reach predicted subgoals.

Oracle VP results. To evaluate our self-improvement algorithm independently of video-planner quality, we also use an Oracle VP that supplies frames from an RL expert. In this setting, our method matches LAIfO on DMC, and reaches expert performance on OGBench and Atari. This result demonstrates that as video models continue to scale and improve, our method provides a simple framework that is able to solve a wide variety of tasks. In Appendix B.2, we also demonstrate that Mirror learns a V2A model that generalizes better to unseen goals.

Mirror also exceeds the expert's return in Assault and Zaxxon by 38% and 31%, respectively. We observe that the expert issues redundant fire commands that do not produce additional shots. The V2A model can thus follow the plan without copying them. As a result, Mirror makes fewer ineffective firing attempts in Assault, reducing overheating, and releases FIRE more frequently in Zaxxon, enabling subsequent presses to trigger new shots.

5 Conclusion

We have introduced Mirror, a method that learns to translate behaviors encoded in large video planners to policies executable by an agent. Unlike prior work, our method assumes neither paired video-action data nor task-specific rewards, but instead learns to translate video into actions simply by self-play in an interactive environment.

Limitations. When combined with a generative video planner, our method achieves SOTA results on some of the most challenging decision-making benchmarks, such as BASALT. However, we were unable to train a generative video planner that solves the hardest tasks, such as mining diamonds. This would require more data and a higher resolution to depict inventory use, which is beyond our compute budget. In addition, training a capable generative video planner can be data-intensive. In our control benchmarks, the expert data lacked diversity, so we needed additional recovery data collected with DAgger to train a sufficiently robust generative video planner. We expect continued improvements in the quality and robustness of video models to help mitigate these limitations.

We believe Mirror offers a promising path toward general-purpose agents that acquire diverse behaviors from video pre-training and learn to self-improve their execution through interaction.

AI Use Statement. We used generative AI tools to assist with the implementation and editing of experimental code, literature retrieval and discovery, the generation of synthetic data, and language polishing of the manuscript for clarity and readability. Generative AI tools were not used to formulate or prove mathematical claims. All experimental results and conclusions were independently evaluated by the authors. The authors take full responsibility for the final content of this work.

REFERENCES

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025. 15

Bowen Baker, Ilge Akkaya, Peter Zhokhov, Joost Huizinga, Jie Tang, Adrien Ecoffet, Brandon Houghton, Raul Sampedro, and Jeff Clune. Video PreTraining (VPT): Learning to act by watching unlabeled online videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2, 5, 8, 15, 18, 21

M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013. ISSN 1076-9757. doi: 10.1613/jair.3912. URL http://dx.doi.org/10.1613/jair.3912. 2, 8, 24

Homanga Bharadhwaj, Debidatta Dwibedi, Abhinav Gupta, Shubham Tulsiani, Carl Doersch, Ted Xiao, Dhruv Shah, Fei Xia, Dorsa Sadigh, and Sean Kirmani. Gen2Act: Human video generation in novel scenarios enables generalizable robot manipulation. *arXiv preprint arXiv:2409.16283*, 2024. 2

Shaofei Cai, Bowei Zhang, Zihao Wang, Xiaojian Ma, Anji Liu, and Yitao Liang. Groot: Learning to follow instructions by watching gameplay videos. In *International Conference on Learning Representations*, volume 2024, pp. 5523–5554, 2024. 6

Boyuan Chen, Tianyuan Zhang, Haoran Geng, Caiyi Zhang, Peihao Li, Kiwhan Song, William T. Freeman, Jitendra Malik, Pieter Abbeel, Russ Tedrake, Vincent Sitzmann, and Yilun Du. Large video planner enables generalizable robot control. *arXiv preprint arXiv:2512.15840*, 2025. 1, 2

Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016. 17

Juechu Dong, Boyuan Feng, Driss Guessous, Yanbo Liang, and Horace He. Flex attention: A programming model for generating optimized attention kernels. *arXiv preprint arXiv:2412.05496*, 2024. 15

Yilun Du, Sherry Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Joshua B. Tenenbaum, Dale Schuurmans, and Pieter Abbeel. Learning universal policies via text-guided video generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 1, 2

Alejandro Escontrela, Ademi Adeniji, Wilson Yan, Ajay Jain, Xue Bin Peng, Ken Goldberg, Youngwoon Lee, Danijar Hafner, and Pieter Abbeel. Video prediction models as rewards for reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023. 6, 7, 8, 9, 24, 26

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 16

Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, and Anima Anandkumar. MineDojo: Building open-ended embodied agents with internet-scale knowledge. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2022. 3

Divyansh Garg, Shuvam Chakraborty, Chris Cundy, Jiaming Song, and Stefano Ermon. Iq-learn: Inverse soft-q learning for imitation. *Advances in Neural Information Processing Systems*, 34: 4028–4039, 2021. 6

Vittorio Giammarino, James Queeney, and Ioannis Ch. Paschalidis. Adversarial imitation learning from visual observations using latent information, 2024. URL https://arxiv.org/abs/2309.17371. 2, 8, 24, 26

William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. MineRL: A large-scale dataset of Minecraft demonstrations. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. 2, 5, 22

Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 640:647 – 653, 2025a. 3, 8

Danijar Hafner, Wilson Yan, and Timothy Lillicrap. Training agents inside of scalable world models. *arXiv preprint arXiv:2509.24527*, 2025b. 3

Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6185–6194. IEEE, 2023. 15

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 5, 16, 19

Bo-Ruei Huang, Chun-Kai Yang, Chun-Mao Lai, Dai-Jie Wu, and Shao-Hua Sun. Diffusion imitation from observation, 2024. URL https://arxiv.org/abs/2410.05429. 2

Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. Guiding a diffusion model with a bad version of itself. *Advances in Neural Information Processing Systems*, 37:52996–53021, 2024. 5, 19

Po-Chen Ko, Jiayuan Mao, Yilun Du, Shao-Hua Sun, and Joshua B. Tenenbaum. Learning to act from actionless videos through dense correspondences. In *International Conference on Learning Representations (ICLR)*, 2024. 2

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pp. 611–626, 2023. 19

Muyao Li, Zihao Wang, Kaichen He, Xiaojian Ma, and Yitao Liang. Jarvis-vla: Post-training large-scale vision language models to play visual games with keyboards and mouse. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 17878–17899, 2025. 6

Sizhe Lester Li, Evan Kim, Xingjian Bai, Tong Zhao, Tao Pang, Max Simchowitz, and Vincent Sitzmann. Turning video models into generalist robot policies. *arXiv preprint arXiv:2605.27817*, 2026. 2

Junbang Liang, Ruoshi Liu, Ege Ozguroglu, Sruthi Sudhakar, Achal Dave, Pavel Tokmakov, Shuran Song, and Carl Vondrick. Dreamitate: Real-world visuomotor policy learning via video generation. In *Conference on Robot Learning (CoRL)*, 2024a. 2

Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, et al. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *arXiv preprint arXiv:2411.04996*, 2024b. 15

Shalev Lifshitz, Keiran Paster, Harris Chan, Jimmy Ba, and Sheila McIlraith. STEVE-1: A generative model for text-to-behavior in Minecraft. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 3, 6

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 16

Minghuan Liu, Tairan He, Weinan Zhang, Shuicheng Yan, and Zhongwen Xu. Visual imitation learning with patch rewards, 2023. URL https://arxiv.org/abs/2302.00965. 2

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 16

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 18

Yunhao Luo and Yilun Du. Grounding video models to actions through goal conditioned exploration. *ArXiv*, abs/2411.07223, 2024. 2, 8, 26

Sam McCandlish, Jared Kaplan, Dario Amodei, and OpenAI Dota Team. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*, 2018. 17

Stephanie Milani, Anssi Kanervisto, Karolis Ramanauskas, Sander Schulhoff, Brandon Houghton, Sharada Mohanty, Byron Galbraith, Ke Chen, Yan Song, Tianze Zhou, Bingquan Yu, He Liu, Kai Guan, Yujing Hu, Tangjie Lv, Federico Malato, Florian Leopold, Amogh Raut, Ville Hautamäki, Andrew Melnik, Shu Ishida, João Henriques, Robert Klassert, Walter Laurito, Lucas Cazzonelli, Cedric Kulbach, Nicholas Popovic, Marvin Schweizer, Ellen Novoseller, Vinicius Goecks, Nicholas Waytowich, David Watkins, Josh Miller, and Rohin Shah. Towards solving fuzzy tasks with human feedback: A retrospective of the minerl basalt 2022 competition. In Marco Ciccone, Gustavo Stolovitzky, and Jacob Albrecht (eds.), *Proceedings of the NeurIPS 2022 Competitions Track*, volume 220 of *Proceedings of Machine Learning Research*, pp. 171–188. PMLR, 28 Nov–09 Dec 2022. URL https://proceedings.mlr.press/v220/milani23a.html. 6, 21, 23

Stephanie Milani, Anssi Kanervisto, Karolis Ramanauskas, Sander Schulhoff, Brandon Houghton, and Rohin Shah. BEDD: The MineRL BASALT evaluation and demonstrations dataset for training and benchmarking agents that solve fuzzy tasks. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2023. 21

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning, 2013. URL https://arxiv.org/abs/1312.5602. 25

Lorenzo Mur-Labadia, Matthew Muckley, Amir Bar, Mido Assran, Koustuv Sinha, Michael Rabbat, Yann LeCun, Nicolas Ballas, and Adrien Bardes. V-jepa 2.1: Unlocking dense features in video self-supervised learning. *ArXiv*, abs/2603.14482, 2026. URL https://api.semanticscholar.org/CorpusID:286573147. 22

Michal Nauman, Mateusz Ostaszewski, Krzysztof Jankowski, Piotr Miłoś, and Marek Cygan. Bigger, regularized, optimistic: scaling for compute and sample-efficient continuous control, 2024. URL https://arxiv.org/abs/2405.16158. 25

NVIDIA. Cosmos 3: Omnimodal world models for physical ai. *ArXiv*, abs/2606.02800, 2026. 5, 15, 19

Jonas Pai, Liam Achenbach, Victoriano Montesinos, Benedek Forrai, Oier Mees, and Elvis Nava. mimic-video: Video-action models for generalizable robot control beyond VLAs. *arXiv preprint arXiv:2512.15692*, 2025. 2

Pantograph. A general goal-conditioned minecraft model. Pantograph Research, July 2026. URL https://pantograph.com/journal/pan-1. Accessed: 2026-09-04. 3

Seohong Park, Kevin Frans, Benjamin Eysenbach, and Sergey Levine. Ogbench: Benchmarking offline goal-conditioned rl, 2025a. URL https://arxiv.org/abs/2410.20092. 2, 8

Seohong Park, Qiyang Li, and Sergey Levine. Flow q-learning, 2025b. URL https://arxiv.org/abs/2502.02538. 25

Rhoda. Causal video models are data-efficient robot policy learners. Rhoda AI Research blog, 2026. URL https://www.rhoda.ai/research/direct-video-action. 1

Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning, 2011. URL https://arxiv.org/abs/1011.0686. 8

Dominik Schmidt and Minqi Jiang. Learning to act without actions, 2024. URL https://arxiv.org/abs/2312.10812. 2

Rohin Shah, Cody Wild, Steven H. Wang, Neel Alex, Brandon Houghton, William Guss, Sharada Mohanty, Anssi Kanervisto, Stephanie Milani, Nicholay Topin, Pieter Abbeel, Stuart Russell, and Anca Dragan. The minerl basalt competition on learning from human feedback. 2021. URL https://arxiv.org/abs/2107.01969. 2, 5

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 15

Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin Riedmiller. Deepmind control suite, 2018. URL https://arxiv.org/abs/1801.00690. 2, 8, 24

Yang Tian, Sizhe Yang, Jia Zeng, Ping Wang, Dahua Lin, Hao Dong, and Jiangmiao Pang. Predictive inverse dynamics models are scalable learners for robotic manipulation. In *International Conference on Learning Representations (ICLR)*, 2025. 2

Philippe Tillet, Hsiang-Tsung Kung, and David Cox. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, pp. 10–19, 2019. 15

Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation, 2018. URL https://arxiv.org/abs/1805.01954. 2, 8, 25

Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation, 2019. URL https://arxiv.org/abs/1807.06158. 2

Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning, 2015. URL https://arxiv.org/abs/1509.06461. 26

Alibaba Group Wan Team. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 15, 16

Zihao Wang, Shaofei Cai, Guanzhou Chen, Anji Liu, Xiaojian Ma, and Yitao Liang. Describe, explain, plan and select: Interactive planning with large language models enables open-world multi-task agents. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 3

Zihao Wang, Shaofei Cai, Anji Liu, Yonggang Jin, Jinbing Hou, Bowei Zhang, Haowei Lin, Zhaofeng He, Zilong Zheng, Yaodong Yang, Xiaojian Ma, and Yitao Liang. JARVIS-1: Open-world multi-task agents with memory-augmented multimodal language models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. arXiv:2311.05997. 3

Zihao Wang, Muyao Li, Kaichen He, Xiangyu Wang, Zhancun Mu, Anji Liu, and Yitao Liang. Openha: A series of open-source hierarchical agentic models in minecraft. *arXiv preprint arXiv:2509.13347*, 2025. 6

Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. Dueling network architectures for deep reinforcement learning, 2016. URL https://arxiv.org/abs/1511.06581. 26

Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Sejune Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, Lars Liden, Kimin Lee, Jianfeng Gao, Luke Zettlemoyer, Dieter Fox, and Minjoon Seo. Latent action pretraining from videos, 2025. URL https://arxiv.org/abs/2410.11758. 2

Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, et al. World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026. 1

Ted Zadouri, Markus Hoehnerbach, Jay Shah, Vijay Thakkar, and Tri Dao. Flashattention-4: Algorithm and kernel pipelining co-design for asymmetric hardware scaling. *Proceedings of Machine Learning and Systems*, 8:912–926, 2026. 15

Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36:49842–49869, 2023a. 18

Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*, 2023b. 17

A VIDEO PLANNER

Design principle. The central goal behind the design of our video planner is long context. Many of our tasks require memory over long horizons. Long context may also help error recovery in closed-loop planning. Early in training, the imperfect V2A model often deviates from expert behavior, forcing the planner to operate on trajectories outside its training distribution. With a short context, we often saw the planner repeat the same failed behavior. With long context, we instead observed emergent recovery behavior, as the planner can use its previous failures to change strategy (see Sec. 4.1.2) and replan. This section describes how we make such long context affordable.

A.1 ARCHITECTURE

Videos are encoded with the Wan2.2 VAE (Wan Team, 2025), which compresses 4 frames in time and 16×16 pixels in space into 48 latent channels. The zeroth pixel frame is encoded on its own rather than grouped with other frames, so L latent frames correspond to $1 + 4(L - 1)$ pixel frames. The MineRL environment (Baker et al., 2022) renders 640×360 frames at 20 fps. We resize each frame to 320×192, so that every four-frame chunk (0.2 s) is encoded into a $12 \times 20 \times 48$ latent array (height $\times$ width $\times$ channels). A 2×2 patchifier then turns each latent frame into 60 tokens of 192 dimensions, which form the input of the video planner.

Backbone. The planner is initialized from Cosmos3-Nano (NVIDIA, 2026), a mixture of transformers (Liang et al., 2024b) with a text tower and a video tower. The two towers share the same Qwen3-VL (Bai et al., 2025) block design but have separate weights. Each has 7B parameters, 36 layers of width 4096, 32q / 8kv GQA heads of dimension 128, SwiGLU MLPs of width 12,288. For positional embedding, 3D-MROPE (Su et al., 2024; Bai et al., 2025; NVIDIA, 2026) is used in which the 128-dim attention head is split into $[48, 24, 24]$-dims, each encoding the time, height and width of the video. We freeze the text tower and finetune only the video tower. The video tower plays two roles: as a *context encoder*, it encodes the clean history of observed frames, and as a *denoiser*, it denoises the frames of the plan. We refer to the two roles separately for clarity, although they share the same architecture and weights.

Block-sparse attention with FlexAttention. The original Cosmos3 implementation uses NATTEN (Hassani et al., 2023), which cannot efficiently express our block-causal sliding-window attention pattern. We instead use FlexAttention (Dong et al., 2024), which supports efficient attention with arbitrary block-sparse masks. On NVIDIA Hopper and Blackwell GPUs, we use FlexAttention's experimental FlashAttention-4 (Zadouri et al., 2026) backend, which is 1.35–1.55$\times$ faster than the default Triton (Tillet et al., 2019) backend on our workload. On AMD GPUs, we use the Triton backend.

Mixture-of-transformers attention. In the original Cosmos3, all video tokens attend to each other bidirectionally under a single mask. In our implementation, however, the video tower's role is split into a context encoder and a denoiser, which require distinct attention patterns. The video history is block-causally encoded frame by frame, whereas the future plans are jointly denoised and thus require bidirectional attention. Together with the text, this yields three token streams. At every layer, each stream self-attends to its own tokens and cross-attends to the streams it is conditioned on:

- **Text** (text tower): self-attends causally, token by token. It carries the video metadata (resolution, aspect ratio, duration and frame rate), which also serves as an attention sink, and optional task instructions.
- **History** (context encoder): self-attends block-causally, i.e., bidirectionally within a frame and causally across frames. Cross-attends to the text.
- **Plan** (denoiser): self-attends bidirectionally within the plan. Cross-attends to the history and the text.

Handling all these attention patterns in a single joint attention would require an overly complex attention mask. This is not only cumbersome to implement but also substantially slows down the FlexAttention kernel. We instead compute attention with a simpler mask for each source and merge

the results as

$$o = \sum_s \frac{\exp(\mathrm{LSE}_s)}{\sum_{s'} \exp(\mathrm{LSE}_{s'})} o_s, \tag{1}$$

where the attention output o_s and the log-sum-exp LSE_s are computed separately for each source s. This merge is mathematically equivalent to standard scaled dot-product attention (SDPA) computed jointly over all sources. Keeping the sources separate makes the implementation of variable-length training and cached inference straightforward, at negligible overhead.

Attention sink. In video models such as Wan (Wan Team, 2025), the zeroth latent frame serves as the attention sink. In Cosmos3, whose mixture-of-transformers structure lets every video token attend to the text, the text tokens appear to play this role instead. Removing the text tokens greatly slowed down the finetuning. Provided that every video query can properly attend to the text metadata tokens, adapting the pretrained bidirectional model to block-causal attention is fast, taking fewer than 1,000 finetuning steps. On the other hand, removing the zeroth latent frame from the attention context of the denoiser caused only a small drop in performance if the history is long enough. The sliding-window attention described below therefore does not keep the zeroth latent frame; instead, every query always attends to the text.

Text conditioning format. Our video planner is conditioned on two kinds of captions, written in Cosmos3's structured JSON format and wrapped in a chat template. Each caption contains the video metadata (resolution, aspect ratio, duration and frame rate) and, optionally, temporal information and language instructions.

The context encoder takes a *task-level caption*, which describes the long-horizon task sequence. It can optionally state how much time has elapsed since the current task sequence began. We find that providing this resume timestamp helps when the model context must be reset in the middle of a task. The caption may also describe the preceding clip, although we do not use this in our experiments. With all fields present, it reads:

```
{"temporal_caption": "<observations of the preceding 10 s> The clip resumes at 1:23.4. Task
sequence: collect wood; then craft a crafting table", "resolution": {"H": 192, "W": 320}, "
aspect_ratio": "5,3", "duration": "32.0s", "fps": 20}
```

The denoiser takes a *control-level caption*, which describes the desired short-horizon behavior of the plan. The `duration` field is included to help the denoiser handle variable-length prediction. It can optionally state the time range of the plan to generate. An unconditional plan uses the metadata-only caption.

```
{"actions": [{"time": "1:25.0-1:26.6", "description": "Place the crafting table"}], "
resolution": {"H": 192, "W": 320}, "aspect_ratio": "5,3", "duration": "1.6s", "fps": 20}
```

A.2 Training

Objective. We train with the linear interpolant $x_\tau = (1-\tau)x + \tau\epsilon$ and velocity prediction $v = \epsilon - x$ (Lipman et al., 2022; Liu et al., 2022), using a squared error averaged over latent dimensions. Every target latent frame carries equal weight in the loss. Following Esser et al. (2024) and Wan (Wan Team, 2025), τ is drawn from a mixture of a logit-normal distribution shifted towards higher noise scale. As our plan resolution and length are smaller than those of typical video models, we use relatively low shift. We use $\texttt{shift} = 2$ for VPT pretraining and BASALT finetuning, and use 2.5 for the `HumanSurvival` finetuning. We also sample τ from a uniform distribution 10% of the time to ensure better GUI interaction details and likelihood estimation.

Variable plan length. The right length for a plan is unknown in advance. Using too short plans or planning autoregressively with the shortest chunk the VAE allows (one latent frame, 0.2 s) produced poor plans, while 4 s is too long for open-loop execution. We therefore train a single model on a mixture of plan lengths from 0.2 s to 4 s. All frames of a prediction window share one diffusion time, and different windows draw diffusion time and noise independently. We also stochastically drop the text conditions of a subset of the prediction windows for unconditional generation and text-conditional classifier-free guidance (Ho & Salimans, 2022).

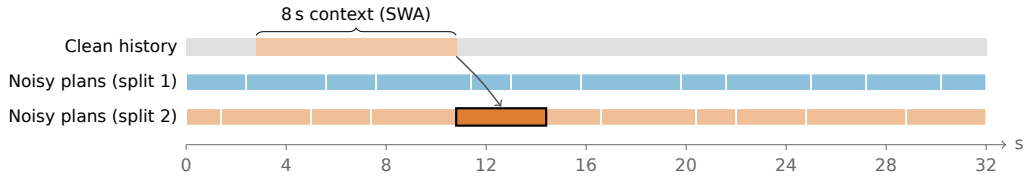

Figure 9: **Variable prediction length training.** Prediction windows of variable length over the same 32 s clip. The clean history is encoded once by the context encoder and shared by all prediction windows. With sliding-window attention, the clean history tokens and the noisy plan windows attend to the preceding 12 s and 8 s clean frames, respectively.

Each training sample is a 32 s video clip, which is block-causally encoded as the clean history. Because this history is shared, we predict many prediction windows from a single video clip at once. This reuse of the context encoding is our key scaling knob. The learning signal comes only from the denoised tokens, so predicting many windows per video clip maximizes the learning signal per token. In our training setup, the prediction windows of each video clip span twice its 32 s length, totaling 64 s. Since the windows differ in position, length, diffusion time and noise, predicting the same video clip twice still provides new learning signal. In our gradient noise scale (McCandlish et al., 2018) measurement at 30k steps, this double pass over a video clip reduced the gradient noise as much as an additional fresh video clip would. As the cost of encoding 32 s of history is amortized over all windows of the video clip, long-context training becomes affordable. Combined with sliding-window attention, this long-context training lets the planner roll out far beyond the training length, to about five minutes in our experiments.

Static shape training for variable-length plans. Compiling every transformer block is essential for training speed, and it requires static tensor shapes. FlexAttention makes this even more important: its block mask is built for a specific sequence shape, so any change of shape forces the mask and the attention kernel to be rebuilt and recompiled. While FlexAttention also supports dynamic-shape compilation, it is far less efficient than static compilation in both throughput and memory. We therefore keep every tensor shape fixed for training and express the variable prediction windows purely through the attention mask. We illustrate this in Figure 9.

In our setup, each GPU processes 6,144 text tokens for the frozen text tower and one 32 s video clip as a fixed-length sequence of 9,600 clean video tokens (160 latent frames $\times$ 60 tokens), followed by 19,200 noisy plan tokens. All three numbers are multiples of FlexAttention's 128-token block size, so the sequences align exactly with the 128×128 attention blocks for best efficiency. The number, lengths and positions of the prediction windows change at every step, but they only change the values of the boolean and integer window/frame id tensors that the mask function reads. The mask function for the video self-attention is:

$$\texttt{attend}(q,k) = \big[\, k \text{ is clean } \wedge\ k \in \texttt{SWA}(q) \,\big] \ \vee\ \big[\, q,k \text{ are noisy } \wedge\ \texttt{plan_id}(q) = \texttt{plan_id}(k) \,\big],$$

where $\texttt{SWA}(q)$ is the set of clean tokens in the sliding window of q. For a clean query, this is the preceding 12 s up to and including its own frame. For a noisy query, this corresponds to the 8 s of clean frames immediately before its prediction window. $\texttt{plan_id}(q)$ is the index of the prediction window containing q. We use $\texttt{plan_id}(q) = 0$ for clean tokens.

Parallelism and memory. We shard the video tower with Hybrid Sharded Data Parallel (HSDP) strategy on multiple nodes using Pytorch FSDP2 library (Zhao et al., 2023b). Parameters are sharded within each 8-GPU node and replicated across nodes. We use activation checkpointing (Chen et al., 2016) for each transformer block. On AMD MI350X GPUs (288 GB of memory) we do not use activation checkpointing.

Precision. We use BF16 precision for the computation. FSDP2 automatically downcasts the FP32 weights to BF16. The gradients are reduced across ranks in FP32. Numerically sensitive parts stay in FP32: the input patch projection, the timestep embedder, the velocity head, RMS/LayerNorm, the rotary embeddings, and the LSE attention merge. Training losses are computed in double precision. The frozen text tower and the VAE run in BF16. For pretraining, we additionally use FP8 mixed

precision training with dynamic tensorwise scaling for the linear layers. Because we use sliding window attention (SWA), which is computed by an efficient block-sparse kernel, the training step is dominated by these matrix multiplications: FP8 speeds them up by about 1.6× and a full training step by about 1.3×. We found no evidence that FP8 training degrades the planning quality. FP8 can be disabled for cooldown or finetuning to reach a lower loss floor.

Optimization and pretraining schedule. We use AdamW (Loshchilov & Hutter, 2017) with $\beta = (0.9, 0.95)$, $\epsilon = 10^{-8}$ and weight decay 0.01, a 1,000-step linear warm-up followed by a constant learning rate, gradient clipping at global norm 3.0, and an EMA of the weights with decay 0.9999. We chose AdamW because Cosmos3 was trained with it. We pretrained on VPT contractor dataset (Baker et al., 2022) for 220k steps. As GPUs became available we grew the global batch from 8 to 32, raising the learning rate roughly with the square root of the batch size (Table 2), and ended with a 5k-step cooldown at a constant 5×10^{-6}. The run took roughly 4,700 H200 GPU-hours of training steps. The final pretrain checkpoint is the step-220k EMA.

Table 2: Pretraining schedule of the video planner on VPT dataset.

Steps	GPUs	Global batch	Learning rate	Wall-clock
0–30k	8 H200	8	1.5×10^{-5}	25 h
30k–80k	16 H200 (2 nodes)	16	2.0×10^{-5}	42 h
80k–215k	32 H200 (4 nodes)	32	2.7×10^{-5}	115 h
215k–220k	8 H200	32	5×10^{-6} (cooldown)	10 h

Pretraining data. The VPT contractor dataset contains 2,945.7 hours of gameplay. We excluded one game-version subset (6.9; 271 recordings, 21.8 h) whose logged cursor coordinates are malformed; the coordinates are needed because the recordings do not show the cursor inside menus, so we draw it back into those frames. We also excluded low-quality data (7.9 h), leaving 55,919 recordings (2,916.0 h) for training. Training samples a recording uniformly and then a random 32 s clip from it. Frames are resized from 640×360 to 320×192 with bilinear interpolation and antialiasing. Over 220k steps, the planner saw about 98,000 hours of noisy video clips. The captions are generated from the videos with the Cosmos3-Nano VLM, as described in Appendix A.4.

Finetuning. The pretrained video planner is finetuned for each downstream task with global batch 8 and an EMA decay of 0.9995. The number of steps varies from 10k to 50k depending on the task difficulty and the amount of available video data. The BASALT planners are trained only on human gameplay videos, without text instructions, with a learning rate of 10^{-5}, a 1k-step warm-up and cosine decay. BASALT episodes end when the player presses the Esc button, which is not visible in the video. We therefore finetune the video planner to predict a gray screen once the player presses the Esc button. The Diamond planners were finetuned on 8 H200 and 8 MI350X GPUs on mixtures of MineDiamond expert episodes and human VPT gameplay. Text conditioning informs the planner of the current stage of the tech tree.

A.3 Inference

Rolling cache with a fixed budget. At inference, the planner keeps the keys and values of only the most recent 12 s of observations in a fixed-size ring-buffer cache. Each newly observed frame is encoded once by the context encoder using this cache. This new frame replaces the oldest frame in the cache, which no query can attend to anymore because of the sliding window. To plan, the denoiser generates the next 1.6 s future by attending to the last 8 s of the cache and optionally to the control-level caption. Because the cache size never grows, every plan costs the same and the model can be compiled with static shape. We illustrate this in Figure 10.

Sampling. Plans are sampled with the second-order UniPC solver (Zhao et al., 2023a) for 20 steps with a time shift of 3. We quantize the linear layers of the video tower to FP8 with per-tensor dynamic scaling, which makes planning 1.43× faster than BF16 on an H200. With 20 solver steps and autoguidance, the planner generates about 4.2 s of video per second of computation on a single H200.

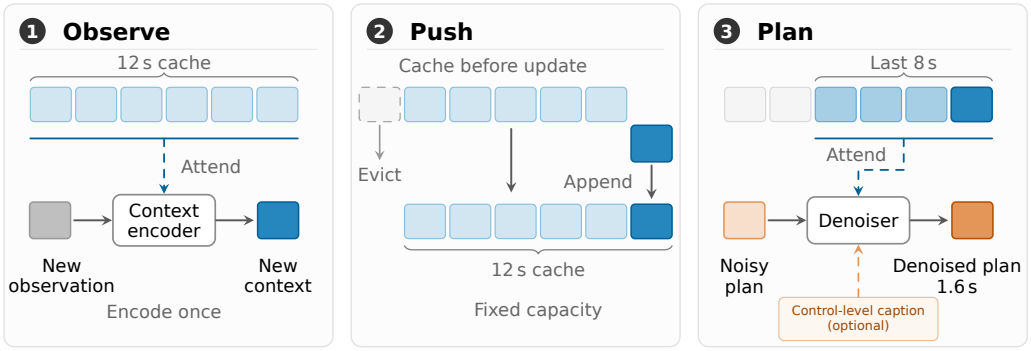

Figure 10: **Rolling cache at inference.** (1) Each new observation is encoded once by the context encoder, which attends to the 12 s cache. (2) The resulting keys and values are appended and the oldest context is evicted, keeping the cache size fixed. (3) The denoiser generates the next 1.6 s of video while attending to the last 8 s of cached context and, optionally, to a control-level caption. Cache cells schematically group frames.

Task-aware guidance. We use two forms of task-aware guidance. For the BASALT tasks, we use a variant of autoguidance (Karras et al., 2024), where the positive model is the finetuned planner and the negative model is a planner trained on broader, less task-specific data, such as its parent before task finetuning; we use a guidance scale of 1.5. For Sec. 4.1.2 and Sec. 4.1.3, we use text classifier-free guidance (Ho & Salimans, 2022) with a scale of 7. For Sec. 4.1.3, the text consists of 15 captions, one per milestone of the `HumanSurvival` tech tree (from log to diamond pickaxe), each stating the current milestone and its target quantity. The conditional and unconditional branches share the cached history, so guidance doubles only the denoiser's computation. We find that task-aware guidance is crucial for goal persistence. Without it, the planner often abandons the current goal after repeated failures. As a result, it may skip necessary steps and move on to the next sub-goal before completing the current one. Task-aware classifier-free guidance effectively suppresses repeated mistakes and strengthens adherence to the intended goal.

A.4 Caption Generation

Pipeline. We captioned the VPT corpus with the vision-language model of Cosmos3-Nano (NVIDIA, 2026), served offline by vLLM (Kwon et al., 2023) with JSON-schema-constrained decoding (temperature 0.7, top-p 0.8, top-k 20). We aim for text that is both *descriptive*, grounded in what is visible, and *instructional*, phrased the way a user would command an agent. We therefore generate captions in three stages. A strictly factual stage comes first, and later stages only rephrase or group its output.

1. *Descriptions.* The model watches overlapping windows of every video at four temporal scales. These are 8 s windows every 4 s at 4 fps, 24 s windows every 12 s at 1 fps, and 64 s windows every 32 s at 0.5 fps, all at 192×320. The fourth scale is short windows at full 640×360 resolution centered on salient moments (attacks, opening a menu, jumps, sharp turns, climbs and descents of at least four blocks). The model is not told why a window was chosen. The model lists up to six atomic, intent-free descriptions per window, each with a start and end time.
2. *Instructions.* Batches of 40 descriptions, together with the video covering their time span, are rewritten one-to-one into imperative instructions. Each instruction inherits the description's time span.
3. *Tasks.* From the video duration and up to 110 descriptions evenly distributed across time, the model partitions the whole recording into 1–8 consecutive task segments and infers the intent that makes each segment's behavior coherent. We run this stage five times per video with different seeds, which yields five alternative task partitions.

The description stage took 4.6 hours on 32 H200 GPUs, plus 1.2 hours for the coarsest scale.

Prompts. The instructions of the three stages, verbatim:

Stage 1: descriptions (system: "You are a helpful assistant."; input: video)

```
This is a clip from a Minecraft gameplay recording; each frame is preceded by its <t seconds>
 timestamp within the full episode. List the distinct things the player OBSERVABLY does, as
separate facts. For each fact give start_s and end_s in those same episode seconds, read from
 the frame timestamps. Rules: describe only what is visible on screen --- never goals,
purposes or attempts (no 'trying to', 'in order to', 'prepares to', 'searching for'); one
atomic action or event per fact; repeated actions separated in time are separate facts; name
blocks/items/mobs/GUI screens concretely when identifiable (hotbar, inventory, crafting table
, crosshair target); 6-20 words per fact. If nothing notable happens, return one fact
describing the ongoing state (movement direction, terrain).
```

Stage 2: instructions (system: "You are a precise Minecraft gameplay annotator."; input: video + indexed descriptions)

```
Below are OBSERVATIONS of what a Minecraft player did, each on its own line with an index.
Rewrite EVERY observation as a CONTROL: the instruction a coach could have given just before
it, which that behaviour then obeys.
Rules for each control:
- Imperative, 3-14 words, concrete Minecraft vocabulary. Command the ACTION, not the scene: '
Mine the coal seam in the left wall', not 'The player mines coal'.
- Say only what the observation says. Add no goals, no reasons, no purpose clauses ('to get
wood', 'so that...'), no outcomes the player could not know in advance.
- A coach speaks before the fact, so drop discoveries and results from the wording: 'Open the
 furnace' rather than 'Open the furnace, revealing an iron ingot'.
- Keep the affirmative form; never write 'Don't', 'Avoid' or 'Stop'.
- Return one control per index, for every index given, in the same order.
```

Stage 3: tasks (system: "You are a precise annotator of MINECRAFT gameplay. Every session, task and control below comes from a recording of a human playing Minecraft, and every word you write must be Minecraft vocabulary."; input: text only)

```
You are splitting one Minecraft recording into its TASKS.

A TASK answers exactly one question:

    What ongoing purpose makes this stretch of the trajectory coherent?

It is NOT a summary of what happened in order.

  BAD  (activity summary): "Walks around, gathers wood, opens the inventory, and mines."
  GOOD (intent hypothesis): "Prepare equipment for underground exploration."

The segments run back to back and together they cover the WHOLE recording, from the
first frame to the last. You give only the END of each segment: the first starts at
0 s, every later one starts where the previous ended, and the last one ends at the
video duration. Give the ends in increasing order.

Rules:

- A segment should normally cover at least 30 seconds.
- Split where the player's PURPOSE changes, not where the scenery changes.
- Name the intent that makes each whole segment coherent, in 6-20 words of concrete
  Minecraft vocabulary, grounded in the facts.
- Do not add narrative detail the evidence does not support.
- Cite the fact ids that support each segment.

VIDEO DURATION: <duration> s
FACTS (intent-neutral observations, video-relative seconds):
  [f12] 3.0-5.5s  <observation>
  ...
Return the JSON object with your segments.
```

Filtering. We clean the generated captions with simple rules:

- decoding is constrained to a JSON schema;
- description times are clamped to their window;
- empty or zero-length items are dropped, 1.5% of descriptions in the rewrite;
- task boundaries are closed into an exact partition on the frame clock.

Statistics. Table 3 summarizes the corpus: 58,121 videos (2,943 h), about 28M captions and 489M tokens. Descriptions and instructions overlap heavily by design. Their total span is about 6× the video duration, so every moment is described several times at several scales. Each task partition tiles its video exactly. 94.6% of task intents and 84.9% of instructions are unique strings.

Table 3: Caption corpus statistics for the VPT dataset. Spans are median [10th, 90th percentile]. Tokens are counted with the Cosmos3-Nano tokenizer.

	Count	Per video (median)	Span (s)	Tokens / caption
Videos	58,121	–	241.0 (duration)	–
Descriptions	13,352,359	244	3.5 [1.0, 8.0]	19.8
Instructions	13,155,310	237	3.5 [1.0, 8.0]	15.2
Task segments	1,506,066	26 (5 partitions)	31.5 [12.0, 62.6]	16.1
Total tokens				489.1M

Table 4: Data scale of different Minecraft agents

Method	**Inputs**	**Outputs**	**Minecraft Data (hours)**		
			Offline	Web	Online
DreamerV3	Video, inventory state	Keyboard, camera, abstract crafting	–	–	1.4K
DreamerV4	Video	Keyboard, mouse	2.5K	–	–
VPT (BC)	Video	Keyboard, mouse	2K	70K	–
VPT (RL)	Video	Keyboard, mouse	2K	70K	230K
Video Planner (Ours)	Video, text prompt (optional)	Video	3K	–	–
V2A Model (Ours)	Video	Keyboard, mouse	–	–	$<$100

B Additional Results

B.1 Minecraft

B.1.1 Data Scale

We compare the data scale of our method with prior Minecraft agents in Table 4. The VPT BC and RL policies (Baker et al., 2022) were trained on significantly greater amounts of data taken from internet videos.

B.1.2 BASALT Challenge

Figure 15 contains an additional comparison to the BASALT challenge methods (Milani et al., 2022). Episodes from each BASALT baseline were taken directly from Milani et al. (2023).

B.1.3 Text-Conditioned Minecraft Evaluation

We evaluate text control in two worlds: Woodland and Camp. The worlds are designed to contain various locations of interest, such as trees, a pool, a hill, and a cave. In each world, the player starts with an iron sword, iron tools, torches, wood planks, bread, and a water bucket. We evaluate 10 episodes per text prompt. In each episode, we first make the player pan the camera left and right to observe their environment. Then it must complete each task given by the text prompt within 30 seconds. The video planner uses a classifier-free guidance scale of 7.0. We include more results on text-conditioned task guidance in Figure 11.

B.2 OGBench Generalization Analysis

Table 5 evaluates transfer to unseen OGBench goals, with all methods using oracle planners for the target tasks. Our V2A model outperforms Random-V2A in three of four environments and BCO-V2A, the V2A model learned by BCO, in all four. These results suggest that video-guided exploration helps the model acquire behaviors that remain useful beyond the tasks used for data collection. Grounding achieves at most 3.8% success, suggesting that learning to translate frame-level transitions into actions supports stronger task generalization than learning a goal-conditioned policy.

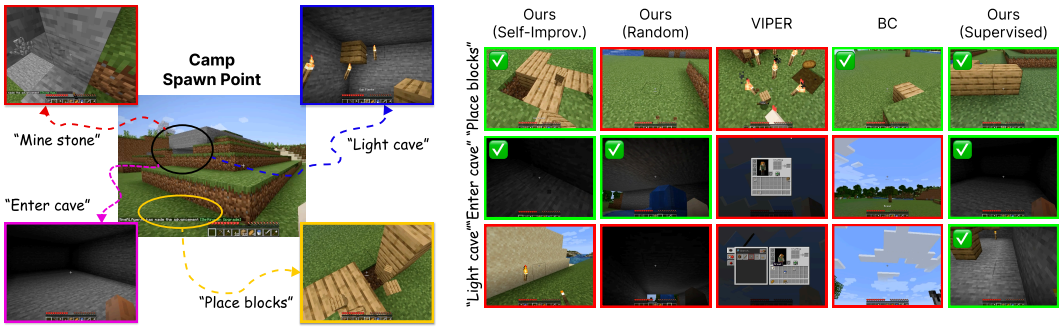

Figure 11: **Text-conditioned planning on Camp.** We demonstrate text-conditioned planning on the camp environment. While our self-improving approach is able to solve many tasks, it struggles on tasks with fine-grained environment interactions such as placing torches in the cave. Meanwhile, VIPER, our action-free baseline, struggles to learn the action space of Minecraft. BC solves some tasks, but remains limited in its generalization ability.

Table 5: **Transfer to unseen OGBench goals.** Success rate (%, mean ± SE across three seeds), averaged over target goals. All methods use target-task oracle planners; BCO-V2A uses BCO's saved V2A model.

Environment	Ours	Random-V2A	BCO-V2A	Grounding
Cube Single	**69.6±1.7**	33.8±5.7	13.8±13.8	3.8±2.2
Cube Double	**36.1±3.1**	20.6±6.4	6.7±6.7	1.1±0.6
Scene	20.8±9.8	**55.0±6.3**	0.0±0.0	3.3±1.7
AntMaze Large	**37.9±3.7**	0.0±0.0	0.0±0.0	0.0±0.0

B.3 Individual Task Learning Curves for Robotics and Atari

We report per-task results for DeepMind Control, OGBench, and Atari in Figures 12, 13, and 14, respectively.

C Implementation Details and Baseline Selection

C.1 Minecraft

C.2 V2A Implementation Details

Architecture. For all Minecraft experiments, the V2A model for Minecraft uses a 80-million parameter VJEPA 2.1 model as a backbone. The model (Mur-Labadia et al., 2026) takes up to 8 history frames and 16 future frames as input, and predicts a chunk of up to 16 actions. The V2A predicts keyboard and mouse actions from the MineRL 1.0 action space (Guss et al., 2019). The button actions are parameterized as categorical distributions, following the factorized button representation of VPT. The x and y offsets for mouse movement are parameterized as a 2D vector.

Exploration Noise. Exploration noise helps the agent learn novel interactions with the environment. However, if the noise is too high frequency, it may harm performance. Real-world actions are typically low frequency—a player will often repeat the same action, such as moving forward, over several time steps. We generate such low frequency exploration noise with random walks. We inject Brownian motion with standard deviation 1.0 for mouse exploration, and add noise corresponding to random walks of button presses for discrete actions.

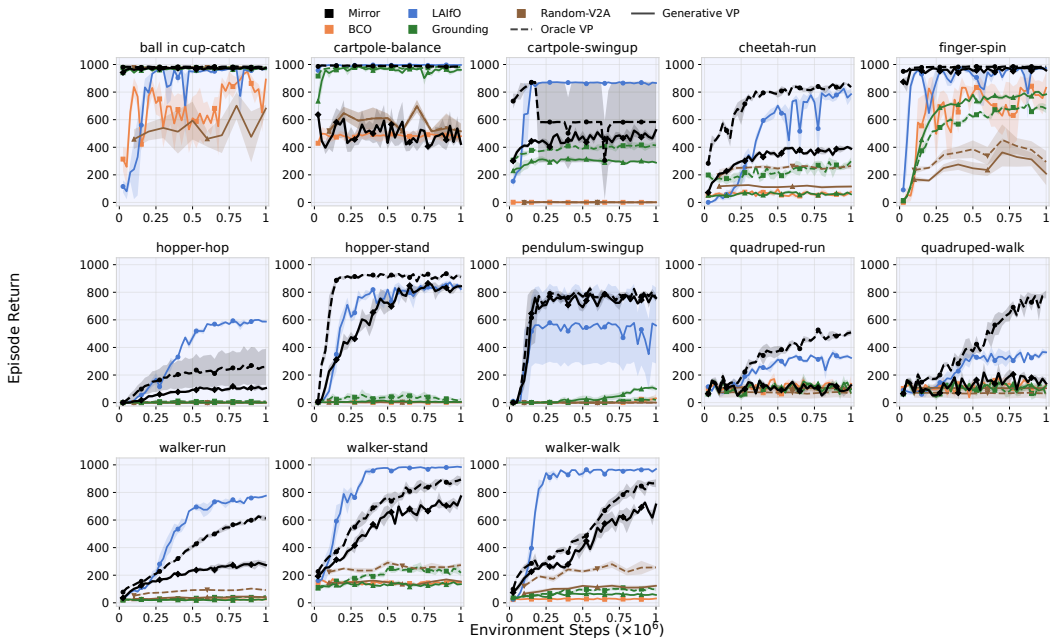

Figure 12: **Per-task performance on DeepMind Control.** Episode return on 13 tasks. Curves show the mean ± one standard error over three seeds.

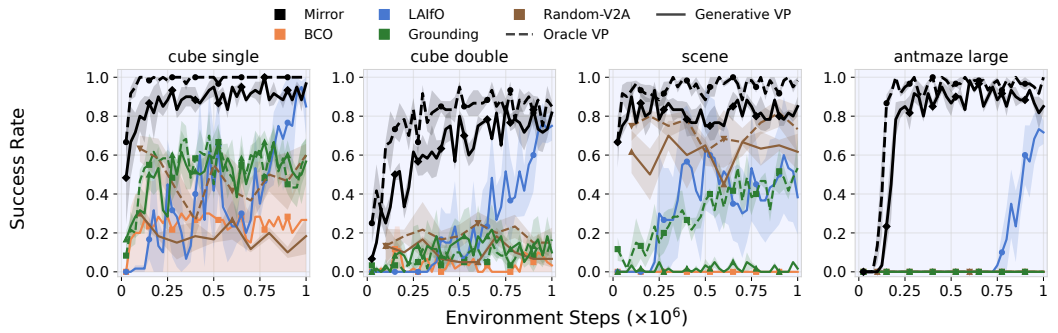

Figure 13: **Per-task performance on OGBench.** Success rate on four single-task environments. Curves show the mean ± one standard error over three seeds.

C.3 BASALT Evaluation

Our method and the BASALT baselines were evaluated using the 11 world seeds from the BASALT 2022 challenge (Milani et al., 2022). Specific task prompts and environment settings can be found on the MineRL website. The BASALT 2022 retrospective report (Milani et al., 2022) contains comprehensive details and code describing the winning solutions to the challenge: GoUp, UniTeam, KAIROS, and others.

While the BASALT criteria are intentionally fuzzy, we evaluate each method on objective criteria for fairness. For FindCave, success is granted when agent entering a cave or ravine. On MakeWaterfall, we report success as when the player successfully places a water bucket, and ends the episode facing the waterfall to "take a picture". The waterfall must be in view during the episode's final frame for success to be granted.

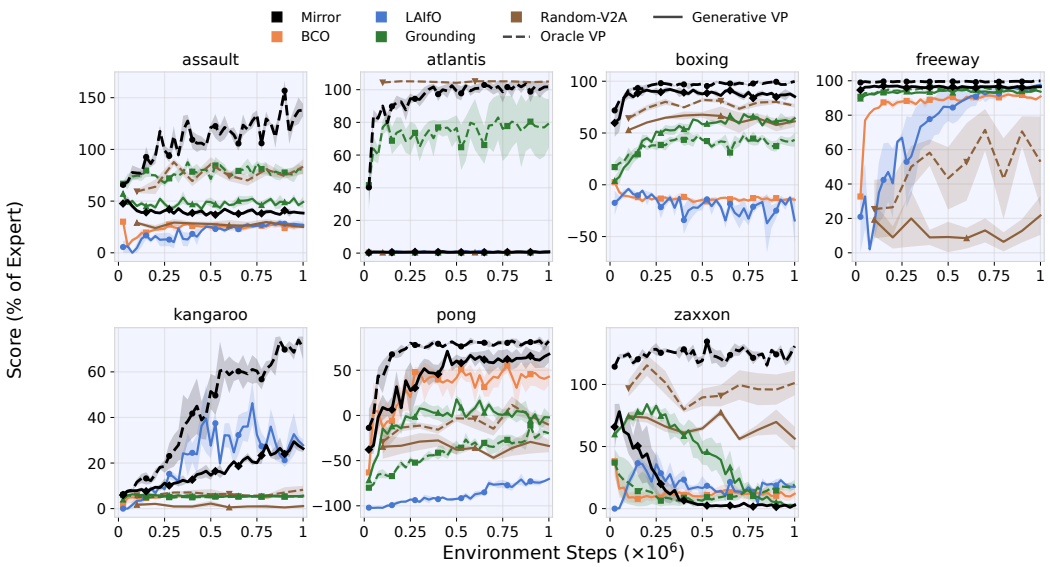

Figure 14: **Per-game performance on Atari.** Scores on seven games, expressed as a percentage of each game's expert score. Curves show the mean ± one standard error over three seeds.

C.4 Minecraft Survival

We train our V2A model with online interaction for 6 minute episodes. We enable exploration noise after the 3 minute mark to encourage exploration. To learn longer horizon behaviors, we occasionally save worlds where the agent has acquired new items, then resume collection episodes from those saves rather than fresh worlds.

C.5 Robotics and Atari

C.5.1 Environments and Evaluation

We evaluate on 13 DeepMind Control (DMC) tasks, four single-task OGBench environments, and seven Atari games. Each online method receives 10^6 environment actions and is evaluated every 25,000 actions over 20 episodes, with three training seeds per task. Random-V2A collects its interaction budget before offline fitting on different-sized subsets. We average tasks equally, reporting native return on DMC, success rate on OGBench, and $100R_{\text{agent}}/R_{\text{expert}}$ on Atari.

OGBench. We use four visual single-task environments: `cube-single` and `cube-double` require rearranging one and two cubes; `scene` requires manipulating a drawer, window, and locking buttons; and `antmaze-large` requires navigating a quadruped robot to a goal. Observations are 64×64 RGB images. We use task 2 for Cube Single, Cube Double, and Scene, and task 1 for AntMaze Large.

DeepMind Control. Following Giammarino et al. (2024), we use 13 tasks from DMC (Tassa et al., 2018): Ball-in-Cup Catch, Finger Spin, Cartpole Balance and Swingup, Pendulum Swingup, Walker Stand/Walk/Run, Hopper Stand/Hop, Quadruped Walk/Run, and Cheetah Run. Agents receive RGB observations with three-frame histories. Oracle-planner experiments use 84×84 images; Mirror and Grounding learned-planner experiments use native 64×64 images for observations, action-model inputs, and predicted frames. For evaluation of Random-V2A with a generative VP, we upsample VP predicted frames to 84×84.

Atari. Following Escontrela et al. (2023), we evaluate on Assault, Atlantis, Boxing, Freeway, Kangaroo, Pong, and Zaxxon from the Arcade Learning Environment (Bellemare et al., 2013). We

use grayscale observations, four-frame histories, the minimal action set, sticky-action probability 0.25, and action repeat four. Oracle-planner experiments use 84×84 images; Mirror and Grounding generative planner experiments use native 64×64 images throughout the student pipeline. One million Atari decisions correspond to approximately four million emulator frames.

C.5.2 Shared Implementation

Action models. Mirror, BCO, Grounding, and Random-V2A use flow-based action models with method-specific conditioning. Continuous-control models use an IMPALA-style visual encoder and four 1024-unit flow layers. Atari uses a smaller encoder, one 512-unit flow layer, and factorized one-hot action codes projected to legal actions. Training uses Adam with learning rate 3×10^{-4}, batch size 256, gradient clipping at 10, and random-shift augmentation. Action sampling uses ten Euler integration steps. Mirror, BCO, and Random-V2A use zero initialization at evaluation; Grounding uses stochastic initialization during both collection and evaluation. LAIfO uses its separate single-action learner.

The action horizon K is five on OGBench; three on DMC except Cartpole Balance (seven), Cartpole Swingup (ten), and Pendulum Swingup (twenty); and two on Atari except Atlantis and Kangaroo (four).

Expert information. Experts are BRO (Nauman et al., 2024) on DMC, FQL (Park et al., 2025b) on OGBench, and DQN (Mnih et al., 2013) on Atari. BCO and LAIfO receive observation-only demonstrations: 50 episodes per Atari game and 1,000/500/500/50 episodes for Cube Single/Cube Double/Scene/AntMaze Large. Their OGBench corpora have matching episode counts. LAIfO uses full-horizon trajectories that continue after success, whereas BCO uses the original shorter demonstrations. DMC uses 50 episodes for BCO and 100 for LAIfO. No method uses expert action labels, task rewards, or proprioceptive state.

Learned video planners. Mirror and Grounding use task-specific latent flow predictors with four residual blocks of width 512. A spatial autoencoder with encoder widths $64, 128, 128, 128$ maps 64×64 images to $4 \times 4 \times 2$ latents. The predictor models normalized latent differences from the current frame. Flow training uses Adam with peak learning rate 3×10^{-4}, cosine decay, batch size 256, gradient clipping at 1, and exponential-moving-average weights with decay 0.999; prediction uses ten Euler steps. Autoencoders are frozen during flow training, and both models remain frozen during downstream learning. Each task shares one planner checkpoint across downstream seeds.

Atari frames are replicated across RGB channels for the planner and converted back to grayscale after decoding. Both Atari planners use three context frames, while action models use four.

C.5.3 Methods and Hyperparameter Selection

Mirror. After 5,000 random actions, Mirror repeatedly obtains K future frames, predicts a K-action sequence from the current visual history and these frames, and executes it before replanning. Achieved observations and executed actions enter a 100,000-transition replay buffer, with 0.2 V2A model updates per environment action. Continuous-control exploration adds Gaussian noise to one randomly selected action coordinate per step, with standard deviation 0.1 on OGBench and task-specific values on DMC. Atari uses stochastic flow sampling. Evaluation uses deterministic V2A model readout without collection noise.

For learned-planner experiments, we train the planner on expert rollouts augmented with expert recoveries from learner-visited states through iterative dataset aggregation. It conditions on three frames and predicts one future frame at a time, generating K-frame plans autoregressively. Checkpoints are selected by rollout performance, with validation flow loss breaking ties. DMC and Atari use one stochastic planner sample per predicted frame during collection and the mean of eight latent samples during evaluation. OGBench uses the mean of eight latent samples during both collection and evaluation.

Behavioral Cloning from Observation. BCO (Torabi et al., 2018) alternates learner interaction, V2A model fitting, and policy fitting. The V2A model infers action sequences from expert observations, and an observation-conditioned policy learns these pseudo-labels and executes action chunks

open loop. Both networks retain their parameters and optimizer state across rounds. Episode-level validation selects each fit, and pending fits finish before final evaluation.

We tune the initial collection size i_{pre} and collection ratio α, giving approximately αi_{pre} actions between refits. Nine pairs are evaluated with three seeds on Cube Single, Finger Spin, and Pong, ranked by the final three evaluations. Selected (i_{pre}, α) values are $(40{,}000, 0.1)$ for OGBench, $(20{,}000, 0.2)$ for DMC, and $(20{,}000, 0.4)$ for Atari.

LAIfO. LAIfO (Giammarino et al., 2024) learns an imitation reward by discriminating expert and learner observation transitions, then trains a single-action off-policy learner. Transition distance n controls discriminator inputs and critic bootstrapping. DMC uses released task-specific settings with our BRO experts; Atari uses a dueling Double-DQN (Wang et al., 2016; van Hasselt et al., 2015) adaptation with greedy evaluation.

We tune exploration decay, transition distance, batch size, and a paired learning-rate/feature-dimension setting on Cube Single and Pong using 24- and 36-configuration grids, respectively, with three seeds per configuration. Selection uses performance at one million actions. Both suites select a 100k exploration decay, batch size 512, learning rate 8×10^{-5}, and feature dimension 100. Transition distance is three for OGBench and one for Atari; discriminator learning rate is 4×10^{-4}.

Grounding. Grounding (Luo & Du, 2024) trains a goal-conditioned action model from random and plan-guided rollouts. Each training segment uses its final achieved image as the goal and executed actions as targets. Following an initial random-only fitting phase, replay sampling uses 30% random and 70% guided data. On OGBench, we follow the released LIBERO implementation, conditioning the model on a single current frame, with each guided collection following a single video plan without replanning. During evaluation, we periodically replan from the latest observation, with environment-specific caps on the number of plans chosen to approximately match the evaluation horizon of other methods. For DMC and Atari, we use three and four frame observation histories, respectively. To accommodate their longer task horizon, we replan during both collection and evaluation after every two subgoals in oracle planner experiments and after every subgoal in generative planner experiments. We round DMC episode limits down to a multiple of the action-chunk length.

We tune random-action hold duration and optimizer-update intervals between guided and random collection using 27-configuration, three-seed sweeps on Cube Single, Finger Spin, and Pong, ranked by the final five evaluations. Selected $(\text{hold}, \text{guided interval}, \text{random interval})$ values are $(5, 100, 500)$, $(10, 100, 500)$, and $(5, 200, 1000)$ for OGBench, DMC, and Atari, respectively. OGBench collection intervals scale with task length.

Grounding's learned planners use the shared architecture with task-specific subgoal targets. DMC and Atari reuse the aggregated expert/recovery corpus and frozen autoencoder, conditioning on three frames to predict an expert-continuation subgoal K actions ahead. OGBench uses 1,000 expert episodes per task with a 950/50 training/validation split, conditions on one frame, and jointly predicts 7/14/20/70 subgoals for Cube Single/Cube Double/Scene/AntMaze Large. Following the original implementation, targets are spaced evenly from a sampled starting frame to shortly after expert success, with the final frame used for unsuccessful episodes. OGBench autoencoders train for 100,000 updates. Predictors train for up to 600,000 updates, with validation every 1,000 updates and early stopping after 200,000 updates without improvement. We select the lowest-validation-loss checkpoint and predict from zero initialization during both collection and evaluation.

Random-V2A. This baseline fits the same V2A model as Mirror independently on cumulative subsets of 100k to 1M transitions from a frozen uniformly random corpus. We hold out 5% of episodes and select checkpoints by validation flow loss. Evaluation translates visual plans into action chunks using deterministic IDM readout. Learned VP evaluations take the mean of eight samples from the video planner.

VIPER. VIPER (Escontrela et al., 2023) learns a generative video model from expert observations and uses observation likelihoods as an RL reward. We include it as an Atari-only likelihood-reward reference, using estimates from its published results.



Figure 15: **MakeWaterfall comparisons.** Mirror is capable of the long horizon planning necessary for MakeWaterfall. The agent must make a waterfall, remember its location, and move to another location to take a picture of it. Even though BC, UniTeam and KAIROS are trained with action labels, they achieve 0% success on this long-horizon task. They often go out of distribution. GoUp, which relies on scripted actions, cannot adapt to unseen scenarios such as when underwater. While it is able to place the water bucket with a pre-programmed instruction, it does not know what to do when underwater. The VLA approaches, STEVE-1, GROOT, JARVIS-VLA, and OpenHA, struggle at following the task instructions. VIPER, our action-free baseline, struggles to learn the Minecraft action space.