
Looking Back to Move Forward: Temporal Verification for Generative Robot Policies

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Generative policies have emerged as a promising paradigm for robot learning,
2 combining expressive generative action modeling with scalable imitation learning
3 from large demonstration corpora. Yet when trained on heterogeneous demonstra-
4 tions, their chunked generation process can produce suboptimal actions, causing
5 errors to compound during execution and eventually driving the robot into out-of-
6 distribution failure states. Action verification has recently emerged as a test-time
7 scaling approach to mitigate this issue, sampling multiple action candidates and
8 using an external verifier to select the best one for execution. However, existing
9 approaches remain temporally myopic and impractical in design: they evaluate
10 candidates using only static-timestep observations and often rely on large-scale ver-
11 ifiers and additional expert demonstrations. In this paper, we introduce **Temporal**
12 **Verification (TeV)**, an efficient temporally aware action verification framework
13 for flow-matching VLAs. TeV first learns a *temporal token* that summarizes
14 recent observation-action history, allowing candidate chunks to be evaluated as
15 continuations of the robot’s execution trajectory rather than isolated predictions.
16 Conditioned on this token, TeV constructs positive–negative sample pairs without
17 additional expert data or preference annotations, and trains an energy-based verifier
18 with a contrastive objective to favor higher-quality chunks that are temporally
19 compatible with recent execution. Beyond post-hoc ranking, TeV uses the learned
20 energy score to steer intermediate samples toward lower-energy regions, potentially
21 improving the quality of generated candidates. Extensive experiments in simulation
22 and real-world settings demonstrate that TeV provides reliable action candidate
23 ranking, improves task success, and yields smoother execution trajectories.

24 1 Introduction

25 Generative policies based on diffusion and flow models have rapidly become a central paradigm in
26 robot learning [8, 23, 2]. By scaling imitation learning to large collections of human demonstrations,
27 these models have enabled robots to perform increasingly complex skills, including household
28 cleaning [16], cloth folding [47], and contact-rich assembly such as screw driving [45]. In particular,
29 Vision-Language-Action (VLA) policies trained with flow-matching objectives [23] have emerged
30 as a promising foundation for generalist robot control [16, 32, 29, 42, 7]. These models couple
31 the semantic reasoning capacity of vision-language backbones with continuous action generation,
32 allowing language instructions, visual observations, and robot states to be mapped into temporally
33 extended action chunks. Flow-matching further provides an efficient generative objective for action
34 sampling, making it especially attractive for scalable robot policy learning.

35 Despite this promise, the generative flexibility that makes flow-matching policies expressive also
36 exposes a critical weakness. When trained on heterogeneous demonstrations, such a generation
37 pipeline can produce suboptimal actions. Small errors may compound over time, gradually driving

38 the robot away from the demonstrated state-action manifold and into out-of-distribution states that
39 culminate in failure. Previous methods address this issue by retraining or fine-tuning the base
40 policy through reinforcement learning [25, 44, 48] or representation regularization [49, 18]. While
41 effective, these approaches can compromise the base policy’s generalist prior and incur substantial
42 computational cost. Verification-based test-time scaling [1, 46] therefore offers a practical alternative,
43 which improves deployment-time action selection without modifying the base policy. Rather than
44 committing to a single sampled action chunk, these methods allocate additional inference-time
45 compute to sample multiple candidate chunks from the underlying policy. An external verifier is then
46 used to select the candidate most aligned with the expert distribution or preference signal [27, 46].

47 Yet existing verification-based methods face two limitations. First, they suffer from temporal myopia,
48 where candidate actions are generated and evaluated fully based on observations from the nearest static
49 timestep. For chunked generative policies, successful execution depends not only on selecting high-
50 quality candidates under the current observation, but also on ensuring that consecutive action chunks
51 form a coherent trajectory. Without this cross-chunk consistency, the policy may oscillate between
52 individually plausible modes, resulting in jerky motion, unstable execution, and task failure [24, 40].
53 Thus, effective verification must evaluate both *local action plausibility* and *temporal consistency*
54 with the robot’s recent execution history. The second limitation arise from the practical burden of
55 their verifier choice. Existing approaches often rely on large VLM-based verifiers [19, 1], which
56 require careful adaptation through prompt engineering and abundant task-specific supervision. Such
57 supervision is expensive to obtain, since reward labels, trajectory preferences, and failure traces
58 require additional expert annotation, costly rollout collection, or physical interaction [31, 22]. Large-
59 scale verifiers also hinder real-world robot deployment because they are costly to run in latency-
60 sensitive control loops [3, 32], especially under limited onboard compute [17].

61 To address these limitations, we propose **Temporal Verification (TeV)**, an efficient energy-based
62 framework for flow-matching VLAs. TeV improves both the temporal consistency and quality of
63 generated action chunks through two complementary functions: it evaluates completed candidates for
64 execution and steers intermediate samples during test-time generation. Specifically, we first construct
65 a *temporal token* that summarizes the robot’s recent observation-action history. Produced by an
66 efficient temporal encoder and regularized with a dynamics-aware objective, this token provides
67 the verifier with execution context beyond the current observation. Conditioned on the temporal
68 token, a lightweight energy model scores sampled action chunks and ranks them for candidate
69 selection. Since verification only requires relative preference among candidates rather than calibrated
70 reward estimation, we train the verifier contrastively: expert demonstration chunks serve as positives,
71 while temporally mismatched chunks and policy-generated deviations serve as informative negatives.
72 Finally, TeV extends verification beyond post-hoc ranking by using the differentiable energy score as
73 guidance during flow integration, steering intermediate samples toward lower-energy regions that
74 may yield better candidates before final selection.

75 Temporal Verification is a lightweight add-on module for flow-matching VLAs. Its verifier introduces
76 less than 0.15% additional parameters relative to the base policy and requires no additional expert
77 demonstrations. Extensive experiments in simulation and real-world settings show that TeV pro-
78 vides reliable energy-based ranking, yields consistent task-success gains of 6%–18%, and produces
79 smoother robot dynamics, highlighting its effectiveness for generative robot control.

80 2 Related Work

81 **Test-time Scaling via Action Verification.** Test-time verification improves robot execution by
82 allocating additional inference-time computation to evaluate and rank candidate outputs [33, 20, 19].
83 In robotics, this paradigm has been explored through several verifier designs. External-verifier
84 methods use vision-language models to evaluate candidate actions or plans [19, 1], either with task-
85 specific training or zero-shot visual reasoning. RL-based approaches support verification through
86 online policy adaptation [22] or learned value functions for action ranking [27]. Other methods use
87 verification to improve instruction following by checking whether generated actions remain aligned
88 with language goals [20]. The most relevant verification approach to ours is TACO [46], which
89 uses a pseudo-count estimator to favor candidate actions near high-density successful modes of the
90 training distribution. However, existing methods rely on static timestep observations and may use
91 abundant external supervision for action selection, whereas our method conditions verification on
92 recent execution history to encourage temporally coherent and smooth behavior.




Figure 1: **Paradigm comparison.** (a) Conventional action verification selects among multiple sampled candidates at each decision step, relying only on the current observation. (b) Temporal-consistency methods refine the current prediction by using previous action chunks to promote cross-chunk coherence. (c) Temporal Verification integrates action verification with temporal consistency modeling. Conditioned on past observation-action context, the learned temporal verifier steers and selects action chunk candidates that are more consistent with recent execution.

93 **Temporal Consistency for Chunked Policies.** The chunked generation paradigm of diffusion
 94 and flow-matching policies can produce action chunks that are locally plausible but temporally
 95 inconsistent across consecutive predictions, leading to mode oscillation and jerky behaviors during
 96 deployment. Existing methods improve temporal consistency by modifying how action chunks are
 97 executed. ACT [50] smooths execution by averaging overlapping action chunks, while BID [24]
 98 samples multiple candidates and selects the chunk most consistent with the previous prediction over
 99 overlapping timesteps. Real-time control methods [3, 36, 40] study temporal consistency under
 100 asynchronous execution, where stale observations further complicate the problem. In contrast, our
 101 method learns a temporally aware verification signal from recent observation-action history, enabling
 102 an energy-based verifier to select candidates that are compatible with ongoing execution.

103 **Policy Steering.** Inference-time steering improves pretrained generative policies by modifying
 104 the sampling or refinement process while keeping the base policy fixed [27]. Existing methods
 105 typically use external objectives or auxiliary models to guide generated actions toward desired task,
 106 safety, or user-specified constraints. Model-predictive refinement methods [28, 34] guide policy
 107 outputs using learned dynamics models to satisfy safety constraints or improve task performance,
 108 while human-in-the-loop approaches [4, 43, 30] refine actions based on user-provided subgoals,
 109 corrections, or preferences. Classifier- and dynamics-guided methods [10, 35] use latent visual
 110 dynamics models [15, 41] to steer generated actions toward desired outcomes. Our method extends
 111 contrastive energy verification beyond post-hoc candidate ranking to enable inference-time steering.
 112 Figure 1 summarizes the relationship between our method and related works.

113 3 Preliminaries

114 **Flow-matching-based VLAs.** Flow-matching policies [23] generate actions by learning a continuous-
 115 time velocity field that transports samples from a simple prior, such as Gaussian noise, to the action
 116 distribution. In flow-matching VLAs, the current observation context is first encoded by the VLM
 117 backbone into a conditional representation \mathcal{K} . The action expert then conditions on \mathcal{K} and transforms
 118 an initial noise sample $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ into an action chunk by integrating the learned velocity field
 119 \mathbf{v}_π over normalized flow time $s \in [0, 1]$. With N Euler integration steps, the update is

$$\mathbf{x}_{s+\frac{1}{N}} = \mathbf{x}_s + \frac{1}{N} \mathbf{v}_\pi(\mathbf{x}_s, s \mid \mathcal{K}), \quad (1)$$

120 where integration starts from \mathbf{x}_0 and terminates at \mathbf{x}_1 . The final output $\mathbf{x}_1 \in \mathbb{R}^{H \times D}$ is an action
 121 chunk of horizon H , with each action having dimension D .

122 **Contrastive learning and energy-based models.** Contrastive learning trains a scoring function to
 123 distinguish positive samples from negative ones according to relative compatibility. Energy-based
 124 models (EBMs) [21, 11] provide a natural formulation for such scoring: an input \mathbf{z} is assigned a
 125 scalar energy $E_\theta(\mathbf{z}) \in \mathbb{R}$, where lower energy indicates higher compatibility. In conditional settings,
 126 the energy also depends on a context variable \mathbf{c} , written as $E_\theta(\mathbf{z} \mid \mathbf{c})$. Thus, a standard contrastive
 127 objective for EBMs is the margin-based ranking loss [13, 14]:

$$\mathcal{L}_{\text{margin}} = [m + E_\theta(\mathbf{z}^+ \mid \mathbf{c}) - E_\theta(\mathbf{z}^- \mid \mathbf{c})]_+, \quad (2)$$

128 where \mathbf{z}^+ and \mathbf{z}^- denote positive and negative samples under the same context \mathbf{c} , $m > 0$ is a margin,
 129 and $[\cdot]_+ = \max(0, \cdot)$ is the hinge operator. The loss encourages positives to have lower energy

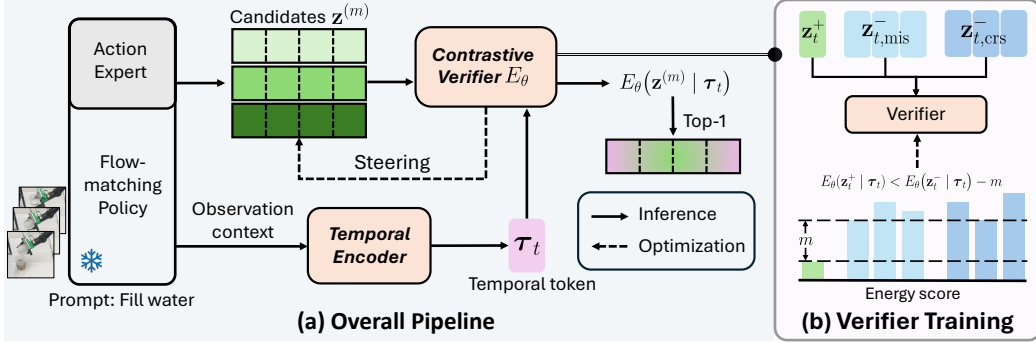


Figure 2: **Illustration of Temporal Verification.** (a) Overall pipeline. We encode observation-action history into a temporal token and use it to condition the energy-based contrastive verifier. At each flow-matching integration step, candidate representations are scored by the verifier. The resulting energy scores are used to steer intermediate samples during integration and rank completed chunks after integration. (b) Verifier training process. Expert chunks serve as positives, while temporally mismatched chunks and policy-generated deviations serve as negatives. The verifier is trained to assign positives lower energy than negatives by at least a specified margin.

130 than negatives by at least m . We observe that verification requires only relative ordering rather than
 131 calibrated absolute scores, which makes contrastive energy modeling well suited to verifier learning.

132 4 Methodology

133 **Overview.** Consider a pretrained flow-matching VLA parameterized by a velocity field \mathbf{v}_π . Our
 134 objective is to learn a lightweight verifier that is also aware of cross-chunk temporal coherence, while
 135 requiring no additional expert demonstrations. To this end, we propose Temporal Verification, whose
 136 overall pipeline is illustrated in Figure 2(a). It first learns a temporal token through a dynamics-aware
 137 loss that compactly summarizes recent observation-action history. Conditioned on this token, we train
 138 an energy-based verifier with a contrastive objective (Figure 2(b)), treating the corresponding expert
 139 demonstration as the positive sample and constructing two negative families: temporally mismatched
 140 chunks and policy-generated deviations. Finally, we use the learned verifier for gradient-based
 141 guidance, steering intermediate candidates toward regions that may yield improved action chunks.

142 4.1 Temporal Token Learning

143 Conditioning stateless flow-matching VLAs [2, 16] on dense observation histories can be compu-
 144 tationally expensive [37]. We address this by compressing recent history into a compact *temporal*
 145 *token*, providing the verifier with temporal context while maintaining inference efficiency.

146 Given a history window of length W , we first construct an observation summary for each timestep
 147 $i \in \{t - W + 1, \dots, t\}$. Specifically, we pass observation \mathbf{o}_i through the pretrained VLM backbone
 148 and extract its token representations. To summarize the observation at timestep i , we project these
 149 tokens into a latent space of dimension d_h and append a learnable readout token shared across
 150 timesteps. A lightweight transformer processes the projected tokens with bidirectional attention,
 151 and the output at the readout position is used as the compressed observation summary, denoted
 152 by $\mathbf{h}_i \in \mathbb{R}^{d_h}$. This step produces compact per-timestep summaries, but each \mathbf{h}_i still represents its
 153 observation independently and does not yet encode temporal evolution.

154 To incorporate execution history, we pair each observation summary \mathbf{h}_i with the immediately pre-
 155 ceding executed action \mathbf{a}_{i-1} . We concatenate \mathbf{h}_i and \mathbf{a}_{i-1} , and feed the resulting sequence into a
 156 temporal transformer \mathcal{T}_ψ with a causal attention mask [39, 5]:

$$(\boldsymbol{\tau}_{t-W+1}, \dots, \boldsymbol{\tau}_t) = \mathcal{T}_\psi([\mathbf{h}_{t-W+1}; \mathbf{a}_{t-W}], \dots, [\mathbf{h}_t; \mathbf{a}_{t-1}]). \quad (3)$$

157 The final output $\boldsymbol{\tau}_t$ causally aggregates information from the recent observation-action history. We
 158 define $\boldsymbol{\tau}_t$ as the **temporal token**, a compact representation of the recent execution trajectory used to
 159 condition the verifier for candidate chunks generated at time t .

160 To encourage the temporal token to capture coherent state-action evolution, we train it with an
 161 auxiliary future-action prediction objective. Let $g(\cdot)$ denote a lightweight prediction head that maps
 162 τ_t to F future actions. We define the auxiliary loss as

$$\mathcal{L}_\tau = \|g(\tau_t) - [\mathbf{a}_t, \dots, \mathbf{a}_{t+F-1}]\|^2, \quad (4)$$

163 where $[\mathbf{a}_t, \dots, \mathbf{a}_{t+F-1}]$ denotes the ground-truth future action sequence from the offline trajectory.
 164 This auxiliary objective acts as a dynamics-aware regularizer: it encourages τ_t to retain information
 165 that links recent observations and executed actions to plausible future motion, making it better suited
 166 for temporally conditioned action ranking and steering.

167 4.2 Contrastive Verifier Training

168 We next introduce a lightweight energy model E_θ that serves as a temporally aware verifier. Conditioned
 169 on the temporal token τ_t , it assigns an energy score to each sampled action chunk candidate
 170 under the same recent execution context. Because verification only requires relative preference among
 171 candidates rather than calibrated reward estimation, we formulate verifier learning as a contrastive
 172 ranking problem, encouraging expert chunks to receive lower energy than inferior alternatives.

173 For each training timestep t , let \mathbf{z}_t^+ denote the latent representation of the expert action chunk. We
 174 use \mathbf{z}_t^+ as the positive sample, reflecting the supervision available from offline imitation data. We
 175 then construct two classes of negative samples, each designed to expose a distinct failure mode of
 176 generative action prediction. Contrasting positives against these negatives trains the verifier to capture
 177 the compatibility criteria required for temporally aware candidate ranking and guidance.

178 **① Temporal-mismatch negatives.** Diffusion- and flow-matching-based policies can model mul-
 179 timodal action distributions [8]. However, this expressiveness can also induce mode switching
 180 during execution: consecutive chunks may appear individually plausible while remaining mutually
 181 inconsistent, resulting in jittery trajectories, mode oscillation, and task failure [24, 3]. We therefore
 182 construct *temporal-mismatch negatives* $\{\mathbf{z}_{t,\text{mis},k}^-\}_{k=1}^{K_{\text{mis}}}$ to train the verifier to recognize temporal
 183 compatibility. For a positive chunk \mathbf{z}_t^+ , we sample expert chunks from other timesteps j within
 184 the same episode, subject to a minimum temporal distance $|j - t| \geq \delta$. Although these chunks are
 185 often high quality, they are not the correct continuation of the history encoded by τ_t . Contrasting
 186 them against \mathbf{z}_t^+ therefore encourages the verifier to distinguish temporally compatible chunks from
 187 plausible but mismatched alternatives. This construction is particularly useful under latent-state
 188 ambiguity, where different physical states induce similar representations. In such cases, action
 189 chunks from other timesteps in the same episode may appear semantically plausible while remaining
 190 temporally misaligned with the robot’s recent execution.

191 **② Coarse-integration negatives.** Temporal-mismatch negatives capture expert-like chunks that
 192 are incompatible with the current history, but they do not expose the verifier to errors produced
 193 by the policy itself. We therefore construct *coarse-integration negatives* $\{\mathbf{z}_{t,\text{crs},k}^-\}_{k=1}^{K_{\text{crs}}}$ from policy-
 194 generated chunks produced under the same observation context but with fewer Euler steps $n < N$.
 195 These negatives arise from a coarser numerical approximation of the learned flow and are therefore
 196 likely to be lower quality than full-step samples. We draw negatives across different values of n :
 197 smaller n yields easier negatives, while larger n produces harder negatives that more closely resemble
 198 standard policy outputs. This encourages the verifier to distinguish high-quality generated chunks
 199 from suboptimal policy samples.

200 Having defined the negative families, we train the verifier with a margin-based contrastive objective.
 201 Let $\mathcal{R} := \{\text{mis}, \text{crs}\}$ denote the set of negative types. For each family $r \in \mathcal{R}$, let $\{\mathbf{z}_{t,r,k}^-\}_{k=1}^{K_r}$ denote
 202 the negative samples constructed for training timestep t , and let $m_r > 0$ denote the corresponding
 203 margin. The contrastive verifier loss is defined as

$$\mathcal{L}_c = \sum_{r \in \mathcal{R}} \frac{1}{K_r} \sum_{k=1}^{K_r} \left[m_r + E_\theta(\mathbf{z}_t^+ | \tau_t) - E_\theta(\mathbf{z}_{t,r,k}^- | \tau_t) \right]_+, \quad (5)$$

204 where $[x]_+ := \max\{x, 0\}$. Since lower energy indicates higher compatibility, minimizing \mathcal{L}_c
 205 encourages the expert chunk \mathbf{z}_t^+ to receive lower energy than each negative sample by at least the
 206 corresponding margin m_r . Although these negative constructions are not exhaustive, they provide
 207 simple and effective supervision for test-time verification.

Table 1: **Performance (%) comparison on LIBERO-Plus.** Best performance in **bold**.

Method	Camera	Robot	Language	Light	Background	Noise	Layout	Average
Base [16]	41.8	72.3	79.9	82.8	84.4	76.2	84.9	73.2
Random	43.9	70.2	86.7	81.8	85.5	79.3	81.1	74.3
TE [50]	44.4	72.3	82.2	76.3	85.1	74.8	84.0	73.0
BID [24]	44.2	73.3	76.2	81.4	80.3	76.6	81.7	72.2
TACO [46]	44.9	77.1	90.1	77.4	86.2	77.1	87.2	76.0
TACO* [46]	42.7	74.8	83.8	76.6	76.8	72.2	80.1	71.5
Ours	48.4	75.6	91.9	84.7	95.8	80.0	92.6	79.8

208 Finally, we jointly optimize the temporal encoder and verifier using

$$\mathcal{L} = \mathcal{L}_c + \lambda_\tau \mathcal{L}_\tau, \quad (6)$$

209 where λ_τ controls the weight of the auxiliary temporal-token prediction loss. The prediction head
210 used for \mathcal{L}_τ is discarded after training.

211 4.3 Verification as Guidance

212 The learned verifier defines a differentiable energy function over action chunk representations
213 conditioned on recent execution history. This enables the verifier to serve not only as a post-
214 hoc ranking module, but also as a guidance signal during flow integration [9]. Let \mathbf{z}_s denote the
215 differentiable action representation extracted at flow time s . Since \mathbf{z}_s depends on the intermediate
216 action sample \mathbf{x}_s , the verifier induces a gradient $\nabla_{\mathbf{x}_s} E_\theta(\mathbf{z}_s | \tau_t)$ through the action expert. For each
217 sampled candidate m , we initialize an independent Gaussian noise sample $\mathbf{x}_0^{(m)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The
218 VLA prefix key-value cache \mathcal{K}_t is computed once from the current observation and shared across
219 candidates. We then augment the standard Euler update in Eq. 1 with an energy-descent term:

$$\mathbf{x}_{s+\frac{1}{N}}^{(m)} = \mathbf{x}_s^{(m)} + \frac{1}{N} \mathbf{v}_\pi(\mathbf{x}_s^{(m)}, s | \mathcal{K}_t) - \frac{\eta_s}{N} \nabla_{\mathbf{x}_s^{(m)}} E_\theta(\mathbf{z}_s^{(m)} | \tau_t), \quad (7)$$

220 where $\eta_s \geq 0$ controls the guidance strength at flow time s . We linearly decay the guidance strength
221 from $\eta_0 = \eta$ to zero over the integration process, making guidance stronger in early steps and weaker
222 near the end. This schedule allows the verifier to provide coarse global correction while the sample is
223 still noisy, while preserving the base policy’s fine-grained refinement in later Euler steps.

224 After integration, each candidate yields a final action chunk $\mathbf{x}_1^{(m)}$ and final representation $\mathbf{z}_1^{(m)}$. We
225 select the executed chunk by minimum verifier energy:

$$m^* = \arg \min_{m \in \{1, \dots, M\}} E_\theta(\mathbf{z}_1^{(m)} | \tau_t), \quad (8)$$

226 and execute $\mathbf{x}_1^{(m^*)}$. The verifier therefore plays a dual role: during integration, it steers intermediate
227 samples toward lower-energy regions of the action space; after integration, it ranks completed
228 candidates for execution. Algorithm 1 in the Appendix summarizes the full procedure.

229 5 Experiments

230 We evaluate TeV in both simulation and real-world settings across diverse perturbations and manipula-
231 tion tasks. In simulation, we follow the standard evaluation protocol for each benchmark. In addition
232 to comparisons with baseline methods, we conduct diagnostic analyses to isolate the contribution of
233 each component and better understand the mechanisms behind the observed performance gains.

234 5.1 Simulation Environment

235 We evaluate on two simulation benchmarks: LIBERO-Plus [12] and RoboTwin 2.0 [6]. LIBERO-Plus
236 contains 10,030 tasks spanning seven perturbation types: object layout (*Layout*), camera viewpoint
237 (*Camera*), robot initial state (*Robot*), language instruction (*Language*), lighting condition (*Light*),
238 background texture (*Background*), and sensor noise (*Noise*). These variations allow us to assess

Table 2: **Performance (%) comparison on RoboTwin2.0 tasks.** Best performance in **bold**.

Method	Adjust Bottle	Pick Bottles	Place Container	Stack Bowls	Place Cup	Open Laptop	Press Stapler	Average
Base [16]	86.0	48.0	86.0	89.0	78.0	75.0	44.0	72.3
Random	85.0	52.0	90.0	88.0	79.0	72.0	50.0	73.7
TE [50]	87.0	43.0	86.0	88.0	82.0	74.0	47.0	72.4
BID [24]	87.0	59.0	87.0	91.0	76.0	77.0	47.0	74.9
TACO [46]	81.0	46.0	87.0	85.0	80.0	62.0	49.0	70.0
TACO* [46]	85.0	32.0	85.0	87.0	80.0	67.0	52.0	69.7
Ours	95.0	64.0	94.0	93.0	88.0	81.0	54.0	81.3

239 policy robustness under diverse visual, physical, and linguistic shifts. For RoboTwin 2.0, we select
 240 seven representative manipulation tasks: *adjust bottle position*, *pick dual bottles*, *place container*
 241 *onto plate*, *stack two bowls*, *place empty cup on coaster*, *open laptop*, and *press stapler*. For each
 242 task, we collect 50 expert demonstrations for policy training. During evaluation, each task is repeated
 243 for 100 trials, and we report task success rate as the primary metric.

244 **Baselines.** We compare TeV against verification-based and temporal-consistency baselines, all using
 245 the same $\pi_{0.5}$ [16] policy backbone and checkpoint. Unless otherwise specified, all test-time scaling
 246 methods use the same candidate budget, sampling $M = 4$ action chunks per observation.

- 247 • **Random selection:** This baseline randomly selects one candidate chunk from the sampled set for
 248 execution, isolating the effect of multi-candidate sampling without any verification signal.
- 249 • **Temporal Ensembling (TE) [50]:** This baseline maintains a buffer of previously predicted action
 250 chunks and executes a weighted average of the actions assigned to the current timestep, smoothing
 251 execution across overlapping chunks and enhances temporal consistency.
- 252 • **Bidirectional Decoding (BID) [24]:** BID samples multiple candidate chunks and selects the one
 253 most similar with the previously executed chunk over overlapping timesteps.
- 254 • **TACO [46]:** TACO uses a lightweight Coin-Flipping Network to estimate pseudo-counts over
 255 internal policy representations, and uses these estimates to prefer action chunks closer to high-
 256 density regions of the training distribution. We also report **TACO***, which uses $M = 50$ candidates
 257 per observation, matching the default candidate budget in the original paper.

258 **Results.** Tables 1 and 2 report the performance comparisons. Temporal Ensembling and
 259 LIBERO-Plus and BID show mixed effects, often performing similarly to the base policy and random selection.
 260 This suggests that directly applying existing temporal-consistency baselines to these simulation
 261 benchmarks is insufficient. TACO improves performance, whereas TACO* underperforms by a large margin, indicating that increasing the candidate
 262 budget does not necessarily lead to higher success rates. By contrast, TeV delivers more
 263 consistent improvements across the evaluated task suites, achieving average gains of 3.8% over the
 264 strongest baseline on LIBERO-Plus and 6.4% on RoboTwin 2.0 subtasks. These results indicate that
 265 verifier design is critical for translating additional test-time samples into performance gains.

270 The limited gains of temporal-consistency baselines may stem from the relatively static task setups
 271 and simplified physical constraints in simulation. Both TE and BID refine the current prediction using
 272 only consecutive action chunks. In these benchmarks, however, the predicted chunks are often already
 273 clean and high quality due to well-collected expert demonstrations. This leaves limited room for
 274 smoothing-based improvement and similarity-based chunk selection. TACO performs less effectively
 275 on RoboTwin 2.0, possibly due to the limited number of expert demonstrations used to train the base
 276 policy. With sparse training coverage, pseudo-count-based estimates of training-distribution support
 277 may become less reliable. In contrast, TeV incorporates recent observation-action history and trains
 278 the verifier contrastively, providing a richer context for candidate evaluation and broader supervision
 279 over plausible but mismatched action chunks.

280 **Ablations.** We conduct the following ablation studies to isolate the contribution of individual
 281 components and assess the effect of key hyperparameters.

- 282 • **Guidance and verification.** Figure 3 compares three variants: verification-only, guidance-only,
 283 and the full method, while also varying the guidance strength η . Verification-only consistently




Figure 3: Effect of guidance and verification.




Figure 4: Effect of temporal token variants.

284 improves task success, showing that the learned energy model provides a reliable signal for
 285 candidate ranking. Guidance-only degrades performance on both benchmarks, suggesting that
 286 gradient-based steering is unstable without final candidate selection. If guidance pushes a candidate
 287 toward an unfavorable region of the action space, the resulting deviation cannot be rejected and
 288 may compound during execution. Combining guidance with final candidate ranking yields the
 289 strongest overall performance. Across the tested values of η , the full method generally outperforms
 290 the baselines and remains competitive in the remaining cases, suggesting that Temporal Verification
 291 is relatively robust to the guidance strength value.

- 292 • **Effect of negative construction.** We evaluate the role of each negative family by training verifiers
 293 with different negative combinations on LIBERO-Plus. The base policy achieves a success rate
 294 of 72.3%. Training with only coarse-integration negatives improves success to 74.3%, while
 295 training with only temporal-mismatch negatives yields a larger gain, reaching 79.2%. These results
 296 indicate that both negative types are useful, with temporal-mismatch negatives providing the larger
 297 contribution. Combining the two families further improves success to 79.8%, suggesting that
 298 coarse-integration negatives add complementary supervision beyond temporal compatibility.
- 299 • **Effect of the number of candidates.** Table 3 examines how the candidate budget affects perfor-
 300 mance. Task success does not increase monotonically with the number of sampled candidates,
 301 indicating that additional samples are not automatically beneficial. This is expected, since an
 302 imperfect verifier may face harder ranking decisions as larger candidate sets include more out-of-
 303 distribution or low-quality chunks.
- 304 • **Effect of the temporal token.** Figure 4 evaluates the learned temporal token against three variants:
 305 **Zero**, which replaces it with 0; **Random**, which replaces it with a Gaussian random vector; and
 306 **Frozen**, which reuses the first temporal token computed in each episode for all subsequent steps.
 307 The learned temporal token consistently outperforms all three variants. Improvements over the
 308 **Zero** and **Random** baselines confirm the effectiveness of temporal-token conditioning, while the
 309 gain over **Frozen** shows that dynamically updating the token with recent observation-action history
 310 provides meaningful context for verifier-based ranking and steering.

311 5.2 Real-world Environment

312 **Setup.** We consider three tasks: (1) put water bottle into the bag (*Bottle to Bag*), (2) move the water
 313 cup from the shelf to the table (*Cup Transport*), and (3) pour water from the right cup into the left cup
 314 (*Water Filling*). We collect 180 trajectories in total for policy training. During execution, the robot
 315 runs at 15Hz with a fixed execution horizon. We compare TeV against the base policy, Bidirectional
 316 Decoding [24], and TACO [46]. For execution safety, all methods adopt Temporal Ensembling [50].
 317 Test-time scaling methods sample $M = 4$ candidates per observation. We evaluate performance
 318 using a task completion score, assigning each episode an integer score according to the successfully
 319 completed substeps. Each setting is evaluated over 30 trials with varied object positions.

320 **Results.** Table 4 reports the average task completion scores, including mean and standard deviation,
 321 across methods and tasks. Different from simulation, these real-world tasks require stable execution
 322 and smooth movement: the bottle in *Bottle to Bag* is deformable, while the other two tasks involve
 323 transporting or pouring liquid. Across the baselines, we observe occasional trajectory oscillations
 324 and unstable execution, resulting in failures such as dropped bottles or spilled water. In contrast, TeV
 325 produces more stable and consistent trajectories, reducing failures associated with jerky motion and
 326 achieving stronger overall performance. Baseline methods also sometimes terminate water pouring
 327 prematurely, before fully emptying the source cup. This may be due to partial observability during
 328 cup tilting, as neither the wrist camera nor the third-view camera reliably captures the liquid state

Table 4: **Performance (%) on real-world tasks.**
Best results marked in **bold**.

Method	Bottle to Bag	Cup Transport	Water Filling
Base [16]	60.7 ± 31.9	79.8 ± 27.5	40.5 ± 32.4
BID [24]	59.8 ± 27.8	71.5 ± 25.9	51.3 ± 30.5
TACO [46]	64.0 ± 30.5	75.2 ± 21.2	51.5 ± 29.7
Ours	69.3 ± 26.7	88.3 ± 15.9	68.0 ± 12.8




Figure 5: **Energy score reliability.**

329 inside the cups throughout the motion. By conditioning verification on encoded temporal history, TeV
 330 can use recent execution context as an additional signal of task progress, which may help maintain a
 331 consistent trajectory when the liquid state is only partially observable.

332 **Analysis.** Figure 5 evaluates the reliability of the predicted energy ranking. We compare the verifier-
 333 induced ranking of sampled candidates with their latent-space distance ranking to the corresponding
 334 expert action chunk. The verifier selects one of the two nearest candidates in more than 70% of
 335 cases, indicating that lower-energy candidates generally align more closely with expert behavior. This
 336 supports the effectiveness of the contrastive verifier learning objective.

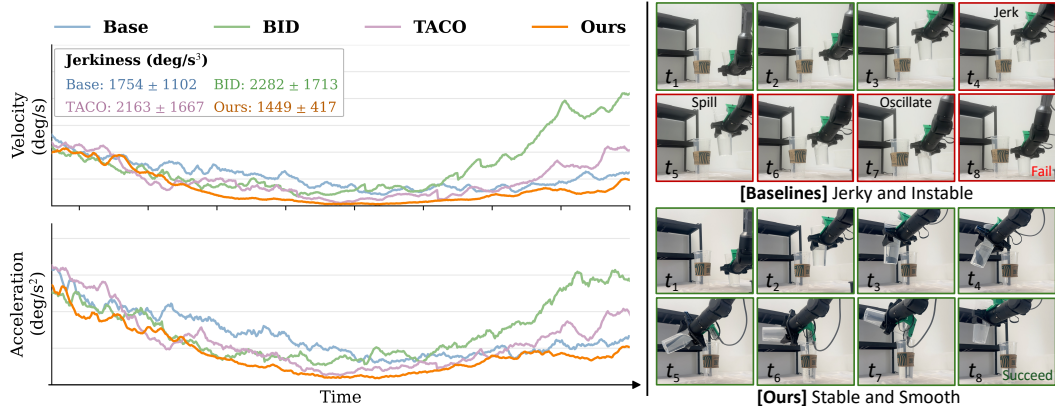


Figure 6: **Illustration of robot kinematics and task progress over time.** **Left:** Average joint velocity and acceleration during interaction with the water cup. **Right:** Real-world rollouts. The baselines often exhibit abrupt motions and oscillations, leading to water spilling or task failure, whereas our method maintains a steadier trajectory and completes the task without spilling.

337 Figure 6 analyzes average robot kinematics and task progress rate. On the left, we illustrate metrics
 338 including joint velocity and acceleration, and quantify jerkiness as the per-trial mean magnitude of
 339 the third time derivative of joint positions, where lower values indicate fewer abrupt motion changes.
 340 The results show that TeV produces smoother kinematic profiles with lower jerkiness, which is
 341 associated with fewer failures such as water spilling. This supports its effectiveness in improving
 342 execution smoothness and temporal consistency. On the right, we provide example rollouts showing
 343 baseline failure cases, such as oscillatory motion and water spilling, which are not observed in the
 344 corresponding TeV rollouts. Additional rollout videos are provided in the supplementary.

345 6 Conclusion

346 Verification-based test-time scaling methods often suffer from temporal myopia, overlooking the
 347 consistency required across consecutive action chunks. In this paper, we introduced Temporal Veri-
 348 fication (TeV), an efficient energy-based framework that makes action verification temporally aware.
 349 TeV equips a lightweight energy-based verifier with a compact temporal token, enabling candidate
 350 actions to be evaluated against recent observation-action history rather than a static timestep. Since
 351 verification requires relative preference among candidates rather than calibrated reward estimation,
 352 we train the verifier with a contrastive objective. The resulting verifier serves two roles: it ranks
 353 completed candidates for selection and actively steers intermediate policy outputs during gener-
 354 ation. Extensive experiments in simulation and real-world manipulation tasks show that TeV is
 355 sample-efficient, improves success rates, and yields smoother robot dynamics.

356 References

- 357 [1] Yusuf Ali, Gryphon Patlin, Karthik Kothuri, Muhammad Zubair Irshad, Wuwei Liang, and Zsolt
358 Kira. Eve: A generator-verifier system for generative policies, 2025.
- 359 [2] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo
360 Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming
361 Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang
362 Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0 : A
363 vision-language-action flow model for general robot control, 2026.
- 364 [3] Kevin Black, Manuel Y. Galliker, and Sergey Levine. Real-time execution of action chunking
365 flow policies, 2025.
- 366 [4] Haoyuan Cai, Zhenghao Peng, and Bolei Zhou. Predictive preference learning from human
367 interventions, 2025.
- 368 [5] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter
369 Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning
370 via sequence modeling, 2021.
- 371 [6] Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang,
372 Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie,
373 Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan ang Gao, Kaixuan
374 Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. Robotwin 2.0: A
375 scalable data generator and benchmark with strong domain randomization for robust bimanual
376 robotic manipulation, 2025.
- 377 [7] Xinyi Chen, Yilun Chen, Yanwei Fu, Ning Gao, Jiaya Jia, Weiyang Jin, Hao Li, Yao Mu,
378 Jiangmiao Pang, Yu Qiao, Yang Tian, Bin Wang, Bolun Wang, Fangjing Wang, Hanqing
379 Wang, Tai Wang, Ziqin Wang, Xueyuan Wei, Chao Wu, Shuai Yang, Jinhui Ye, Junqiu Yu, Jia
380 Zeng, Jingjing Zhang, Jinyu Zhang, Shi Zhang, Feng Zheng, Bowen Zhou, and Yangkun Zhu.
381 Internvla-m1: A spatially guided vision-language-action framework for generalist robot policy,
382 2025.
- 383 [8] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ
384 Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion,
385 2024.
- 386 [9] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021.
- 387 [10] Maximilian Du and Shuran Song. Dynaguide: Steering diffusion polices with active dynamic
388 guidance, 2025.
- 389 [11] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models.
390 *Advances in neural information processing systems*, 32, 2019.
- 391 [12] Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He,
392 Shiduo Zhang, Zhaoye Fei, Jinlan Fu, Jingjing Gong, and Xipeng Qiu. Libero-plus: In-depth
393 robustness analysis of vision-language-action models, 2025.
- 394 [13] Raia Hadsell, Sumit Chopra, and Yann Lecun. Dimensionality reduction by learning an invariant
395 mapping. pages 1735 – 1742, 02 2006.
- 396 [14] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an
397 invariant mapping. In *2006 IEEE computer society conference on computer vision and pattern
398 recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006.
- 399 [15] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and
400 James Davidson. Learning latent dynamics for planning from pixels, 2019.

- 401 [16] Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny
402 Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya
403 Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming
404 Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch,
405 Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz,
406 James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury
407 Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization, 2025.
- 408 [17] Wenqi Jiang, Jason Clemons, Karu Sankaralingam, and Christos Kozyrakis. How fast can i run
409 my vla? demystifying vla inference performance with vla-perf, 2026.
- 410 [18] Taeyoung Kim, Jimin Lee, Myungkyu Koo, Dongyoung Kim, Kyungmin Lee, Changyeon Kim,
411 Younggyo Seo, and Jinwoo Shin. Contrastive representation regularization for vision-language-
412 action models, 2025.
- 413 [19] Jacky Kwok, Christopher Agia, Rohan Sinha, Matt Foutter, Shulu Li, Ion Stoica, Azalia
414 Mirhoseini, and Marco Pavone. Robomonkey: Scaling test-time sampling and verification for
415 vision-language-action models, 2025.
- 416 [20] Jacky Kwok, Xilun Zhang, Mengdi Xu, Yuejiang Liu, Azalia Mirhoseini, Chelsea Finn, and
417 Marco Pavone. Scaling verification can be more effective than scaling policy learning for
418 vision-language-action alignment, 2026.
- 419 [21] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, Fufie Huang, et al. A tutorial on
420 energy-based learning. *Predicting structured data*, 1(0), 2006.
- 421 [22] Yishu Li, Xinyi Mao, Ying Yuan, Kyutae Sim, Ben Eisner, and David Held. Learn from what
422 we have: History-aware verifier that reasons about past interactions online, 2025.
- 423 [23] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow
424 matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- 425 [24] Yuejiang Liu, Jubayer Ibn Hamid, Annie Xie, Yoonho Lee, Max Du, and Chelsea Finn. Bidi-
426 rectional decoding: Improving action chunking via guided test-time sampling. *International
427 Conference on Learning Representations (ICLR)*, 2025.
- 428 [25] Mingyang Lyu, Yinqian Sun, Erliang Lin, Huangrui Li, Ruolin Chen, Feifei Zhao, and Yi Zeng.
429 Reinforcement fine-tuning of flow-matching policies for vision-language-action models, 2025.
- 430 [26] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization
431 for generative adversarial networks, 2018.
- 432 [27] Mitsuhiko Nakamoto, Oier Mees, Aviral Kumar, and Sergey Levine. Steering your generalists:
433 Improving robotic foundation models via value guidance, 2025.
- 434 [28] Kensuke Nakamura, Lasse Peters, and Andrea Bajcsy. Generalizing safety beyond collision-
435 avoidance via latent-space reachability analysis. In *Robotics: Science and Systems XXI*.
436 Robotics: Science and Systems Foundation, June 2025.
- 437 [29] NVIDIA, :, Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding,
438 Linxi "Jim" Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, Joel Jang, Zhenyu Jiang,
439 Jan Kautz, Kaushil Kundalia, Lawrence Lao, Zhiqi Li, Zongyu Lin, Kevin Lin, Guilin Liu,
440 Edith Llonet, Loic Magne, Ajay Mandlekar, Avnish Narayan, Soroush Nasiriany, Scott Reed,
441 You Liang Tan, Guanzhi Wang, Zu Wang, Jing Wang, Qi Wang, Jiannan Xiang, Yuqi Xie,
442 Yinzhen Xu, Zhenjia Xu, Seonghyeon Ye, Zhiding Yu, Ao Zhang, Hao Zhang, Yizhou Zhao,
443 Ruijie Zheng, and Yuke Zhu. Gr00t n1: An open foundation model for generalist humanoid
444 robots, 2025.
- 445 [30] Zhenghao Peng, Wenjie Mo, Chenda Duan, Quanyi Li, and Bolei Zhou. Learning from active
446 human involvement through proxy value propagation, 2025.
- 447 [31] Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A reduction of imitation learning
448 and structured prediction to no-regret online learning, 2011.

- 449 [32] Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil
450 Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, Simon Alibert,
451 Matthieu Cord, Thomas Wolf, and Remi Cadene. Smolvla: A vision-language-action model for
452 affordable and efficient robotics, 2025.
- 453 [33] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute
454 optimally can be more effective than scaling model parameters, 2024.
- 455 [34] Yunzhou Song, Long Le, Yong-Hyun Park, Jie Wang, Junyao Shi, Lingjie Liu, Jiatao Gu, Eric
456 Eaton, Dinesh Jayaraman, and Kostas Daniilidis. Omniguide: Universal guidance fields for
457 enhancing generalist robot policies, 2026.
- 458 [35] Zhanyi Sun and Shuran Song. Latent policy barrier: Learning robust visuomotor policies by
459 staying in-distribution, 2025.
- 460 [36] Jiaming Tang, Yufei Sun, Yilong Zhao, Shang Yang, Yujun Lin, Zhuoyang Zhang, James Hou,
461 Yao Lu, Zhijian Liu, and Song Han. Vlash: Real-time vlas via future-state-aware asynchronous
462 inference, 2025.
- 463 [37] Marcel Torne, Karl Pertsch, Homer Walke, Kyle Vedder, Suraj Nair, Brian Ichter, Allen Z. Ren,
464 Haohuan Wang, Jiaming Tang, Kyle Stachowicz, Karan Dhabalia, Michael Equi, Quan Vuong,
465 Jost Tobias Springenberg, Sergey Levine, Chelsea Finn, and Danny Driess. Mem: Multi-scale
466 embodied memory for vision language action models, 2026.
- 467 [38] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive
468 predictive coding, 2019.
- 469 [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
470 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- 471 [40] Haoxuan Wang, Gengyu Zhang, Yan Yan, Yuzhang Shang, Ramana Rao Kompella, and Gaowen
472 Liu. Real-time robot execution with masked action chunking, 2026.
- 473 [41] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Ken Goldberg, and Pieter Abbeel. Day-
474 dreamer: World models for physical robot learning, 2022.
- 475 [42] Wei Wu, Fan Lu, Yunnan Wang, Shuai Yang, Shi Liu, Fangjing Wang, Qian Zhu, He Sun, Yong
476 Wang, Shuailei Ma, Yiyu Ren, Kejia Zhang, Hui Yu, Jingmei Zhao, Shuai Zhou, Zhenqi Qiu,
477 Houlong Xiong, Ziyu Wang, Zechen Wang, Ran Cheng, Yong-Lu Li, Yongtao Huang, Xing
478 Zhu, Yujun Shen, and Kecheng Zheng. A pragmatic vla foundation model, 2026.
- 479 [43] Yilin Wu, Ran Tian, Gokul Swamy, and Andrea Bajcsy. From foresight to forethought: Vlm-in-
480 the-loop policy steering via latent alignment, 2025.
- 481 [44] Wenli Xiao, Haotian Lin, Andy Peng, Haoru Xue, Tairan He, Yuqi Xie, Fengyuan Hu, Jimmy
482 Wu, Zhengyi Luo, Linxi "Jim" Fan, Guanya Shi, and Yuke Zhu. Self-improving vision-language-
483 action models with data generation via residual rl, 2025.
- 484 [45] Charles Xu, Jost Tobias Springenberg, Michael Equi, Ali Amin, Adnan Esmail, Sergey Levine,
485 and Liyiming Ke. RL token: Bootstrapping online RL with vision-language-action models.
486 Technical report, Physical Intelligence, 2025.
- 487 [46] Siyuan Yang, Yang Zhang, Haoran He, Ling Pan, Xiu Li, Chenjia Bai, and Xuelong Li. Steering
488 vision-language-action models as anti-exploration: A test-time scaling approach, 2025.
- 489 [47] Checheng Yu, Chonghao Sima, Gangcheng Jiang, Hai Zhang, Haoguang Mai, Hongyang Li,
490 Huijie Wang, Jin Chen, Kaiyang Wu, Li Chen, Lirui Zhao, Modi Shi, Ping Luo, Qingwen Bu,
491 Shijia Peng, Tianyu Li, and Yibo Yuan. χ_0 : Resource-aware robust manipulation via taming
492 distributional inconsistencies, 2026.
- 493 [48] Hongyin Zhang, Shuo Zhang, Junxi Jin, Qixin Zeng, Runze Li, and Donglin Wang. Robustvla:
494 Robustness-aware reinforcement post-training for vision-language-action models, 2025.
- 495 [49] Wanpeng Zhang, Ye Wang, Hao Luo, Haoqi Yuan, Yicheng Feng, Sipeng Zheng, Qin Jin, and
496 Zongqing Lu. Dig-flow: Discrepancy-guided flow matching for robust vla models, 2025.

497 [50] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual
498 manipulation with low-cost hardware, 2023.

499 A Implementation Details

500 **Training pipeline.** We implement the temporal encoder as a lightweight transformer-style module
501 with 3.0M parameters. The energy verifier is a shallow feedforward network with 0.9M parameters,
502 equipped with spectral normalization [26] that constrains the verifier’s Lipschitz constant and stabi-
503 lizes energy gradients. Together, the temporal encoder and verifier add less than 0.15% parameters
504 relative to the base model. For training-data construction, we uniformly sample from both negative
505 families. The base model remains frozen throughout training, while only the added modules are
506 optimized from scratch.

507 **Deployment framework.** During deployment, we maintain a rolling buffer of recent observations
508 and executed actions to compute the temporal token τ_t . At the first decision step, when no action
509 history is available, we fall back to standard flow-matching generation; at later steps, any remaining
510 missing action entries are zero-padded. To reduce test-time scaling latency, we use two optimizations.
511 First, we compute the VLA key-value cache \mathcal{K}_t once from the current observation and share it across
512 all M candidates, avoiding repeated backbone computation. Second, we use an efficient inference
513 variant with *guidance cutoff*. Verifier guidance is applied only during the first $\lfloor N/2 \rfloor$ Euler steps, after
514 which candidates are scored and the $\lceil M/2 \rceil$ lowest-energy candidates are retained. The remaining
515 integration steps are completed without guidance. This reduces gradient computation by half and
516 decreases later-stage candidate updates with negligible accuracy degradation.

517 B Experimental Details

518 **Inference acceleration.** Test-time scaling with multiple sampled candidates increases inference
519 latency, which can make real-world execution jerky and instable, leading to abrupt movements that
520 degrade task performance. To reduce this overhead, we use two acceleration strategies. First, we
521 reuse the VLA key-value cache across candidates, since all candidates are conditioned on the same
522 visual-language context, and compile the action-sampling function to reduce execution overhead.
523 Second, we apply progressive pruning: after the early guided integration steps, we retain only the
524 lowest-energy candidates and complete the remaining integration with a smaller candidate set.

525 Following [32], we measure latency as the elapsed time between observation acquisition and the
526 availability of the corresponding action chunk in the control queue. On an RTX 5080 inference device
527 with $M = 4$ candidates, KV-cache reuse reduces latency by approximately 20–30 ms, compilation
528 reduces latency by approximately 10 ms, and progressive pruning reduces latency by approximately
529 30–40 ms. For fair comparison, we apply the same applicable acceleration techniques to all baselines
530 and to TeV; progressive pruning is only used when the method supports candidate scoring during
531 generation.

532 **Computational cost.** We train the verifier on RTX A6000 Ada GPUs. The training cost depends
533 on the dataset size because negative construction is performed from the available offline trajectories.
534 For the real-world tasks, verifier training takes approximately 0.6 GPU hours, which is significantly
535 smaller compared to methods that rely on large-scaled datasets and models.

536 **Hardware setup.** Figure 7 illustrates the real-world hardware setup. We use a master–slave arm
537 framework for demonstration collection and deploy the learned policy on the slave arm. The system
538 uses two cameras for observation and action prediction: a third-view camera that captures the overall
539 workspace and a wrist-mounted first-view camera that provides close-range visual cues for precise
540 manipulation. During execution, the mobile base is fixed to ensure stable and repeatable operation.

541 **Task execution settings.** We evaluate on an AgileX Piper 6-DoF arm equipped with a wrist-mounted
542 Intel RealSense D435i camera and a third-view ZED camera. During real-world execution, we use a
543 fixed execution horizon of 8 for all tasks. The smoothing coefficient for temporal ensembling is set to
544 0.05. We adopt a synchronous execution framework to ensure compatibility across baseline methods,
545 since asynchronous execution may introduce additional latency- and scheduling-related factors that
546 confound performance evaluation.

547 For each task, we evaluate progress using a 100-point completion score defined as follows:

- 548 • *Bottle to Bag*: grasp the water bottle (20), transport the bottle to the target position (50), and
549 successfully drop the bottle into the bag (30).



Figure 7: Robot platform setup.

- 550 • *Water Transport*: grasp the cup of water (30), move the cup stably away from the shelf (35), and
 551 place the cup on the table (35). We subtract 10 points each time water is spilled or the cup collides
 552 with the shelf.
- 553 • *Pour Water*: grasp the right cup of water (10), move it stably to the target position (40), pour water
 554 into the left cup (40), and place the emptied cup back on the table (10). We subtract 10 points if
 555 water is spilled during execution.

556 **Kinematics calculation.** For the kinematic analysis on the *Water Filling* task, we compute the
 557 metrics using joints J1–J5. We exclude J6 because its motion is dominated by the cup-tilting action
 558 required for pouring and is therefore less informative about the stability of the arm trajectory. In
 559 contrast, joints J1–J5 primarily determine how steadily the cup is positioned and transported during
 560 manipulation. For visualization, we smooth the per-timestep velocity and acceleration magnitudes;
 561 all quantitative jerkiness metrics are computed from the original unsmoothed joint trajectories.

562 C Algorithm

563 We show the inference process of our method in Algorithm 1.

564 D Additional Results

565 **Contrastive Loss Comparison** Table 5 compares different contrastive objectives for verifier training.
 566 In addition to the margin-based contrastive loss used in the main experiments, we also evaluate
 567 InfoNCE [38]. Although InfoNCE performs worse than the margin-based objective, it is used
 568 without hyperparameter tuning and still yields substantial gains over the base policy and all baselines.
 569 This suggests that the performance improvement is not tied to a specific contrastive loss, while the
 570 margin-based formulation currently provides the strongest results in our setting.

571 **Energy reliability analysis.** Figure 8 reports the energy-ranking reliability analysis on the LIBERO-
 572 Plus simulation benchmark. The training dataset is built from large-scale, high-quality demonstrations,
 573 and the base policy already produces strong action candidates in many settings. Even in this relatively
 574 saturated regime, our verifier provides a meaningful ranking signal: predicted energy is positively

Algorithm 1 Temporal Verification

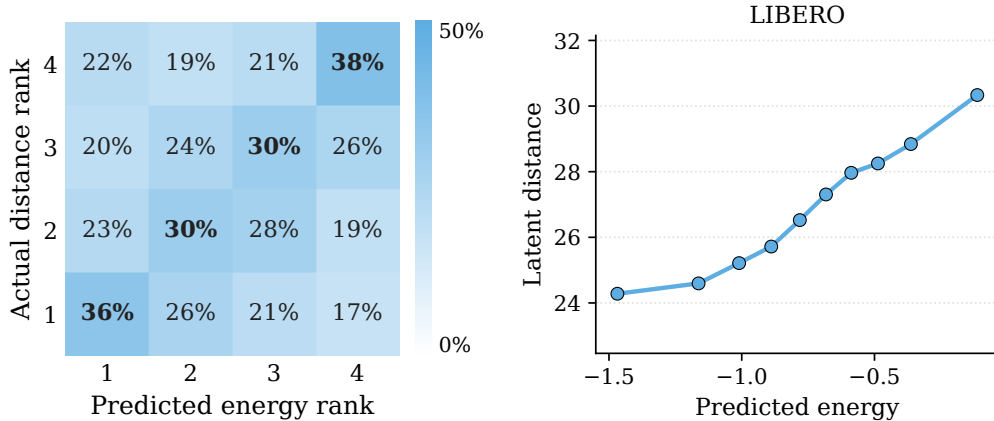
Require: Velocity field \mathbf{v}_π , verifier E_θ , temporal encoder \mathcal{T}_ψ **Require:** Observation history $\{\mathbf{o}_i\}_{i=t-W+1}^t$, action history $\{\mathbf{a}_i\}_{i=t-W}^{t-1}$ **Require:** History length W , candidates M , Euler steps N , guidance schedule $\{\gamma_s\}$ **Ensure:** Selected action chunk $\mathbf{x}_1^{(m^*)}$

```
1: // Encode temporal context
2: Compute temporal token  $\tau_t$  via Eq. (3)
3: Compute VLA context cache  $\mathcal{K}_t$  from  $\mathbf{o}_t$ 
4: // Verifier-guided candidate generation
5: for  $m = 1, \dots, M$  do
6:   Sample  $\mathbf{x}_0^{(m)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
7:   for  $j = 0, \dots, N - 1$  do
8:      $s \leftarrow j/N$ 
9:     Extract intermediate representation  $\mathbf{z}_s^{(m)}$ 
10:    Update  $\mathbf{x}_{s+1/N}^{(m)}$  via Eq. (7)
11:   end for
12:   Compute final representation  $\mathbf{z}^{(m)}$ 
13: end for
14: // Select best candidate
15:  $m^* \leftarrow \arg \min_m E_\theta(\mathbf{z}^{(m)} \mid \tau_t)$  ▷ Eq. (8)
16: return  $\mathbf{x}_1^{(m^*)}$ 
```

Table 5: Performance (%) comparison on LIBERO-Plus under different contrastive losses.

Method	Camera	Robot	Language	Light	Background	Noise	Layout	Average
InfoNCE	45.3	76.8	88.5	83.6	87.5	77.7	87.8	76.9
Margin-based	48.4	75.6	91.9	84.7	95.8	80.0	92.6	79.8

575 correlated with latent-space distance to the corresponding expert action chunk, with lower-energy
576 candidates generally closer to expert behavior.

Figure 8: Energy reliability analysis on LIBERO-Plus. **Left:** Rank correlation. **Right:** Energy v.s. Latent Distance.

577 **Experiment rollouts.** In Figure 9, we include rollouts from TeV to demonstrate the experimental
578 settings and provide a better intuition of how the tasks are done. For the *Water Filling*, we included a
579 case where the wrist camera was blurred.

580 **Evidence of Temporal Consistency Enhancement** Table 6 reports temporal consistency across
581 methods on the *Adjust Bottle* task. We evaluate consistency for arm pose changes, excluding the

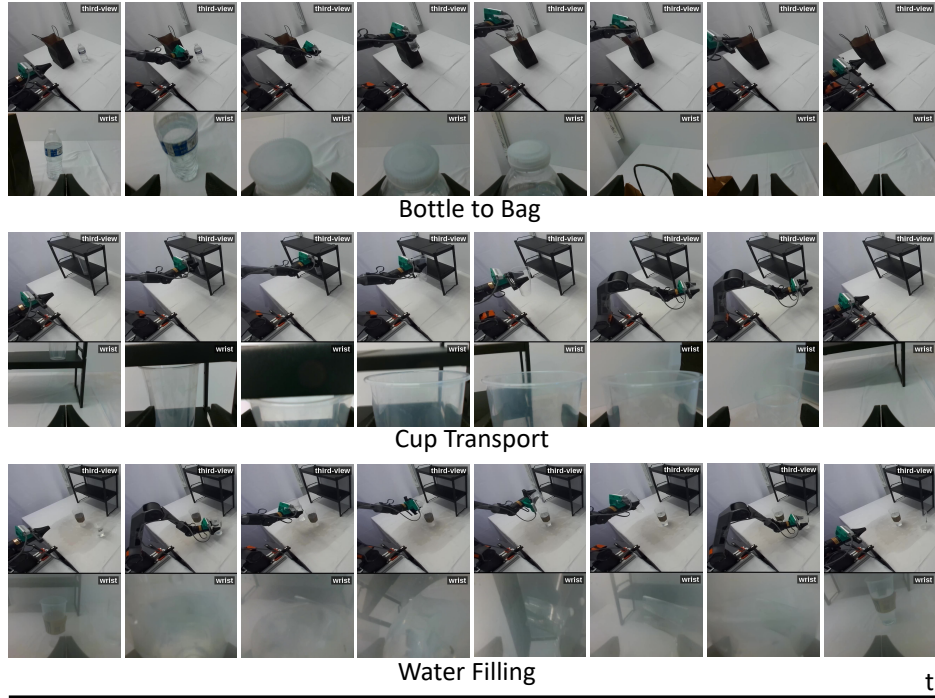


Figure 9: Three real-world tasks rollouts.

582 gripper, and for the full robot action. Following [24], we measure temporal consistency using the
 583 overlap distance between consecutive action chunks. At each decision step, the policy outputs an
 584 action chunk of length H , and the next chunk overlaps with the previous one over $H - 1$ timesteps.
 585 The overlap distance is computed as the ℓ_2 distance between the two chunks over this shared region
 586 in joint space. For each rollout, we compute the 95th percentile of overlap distances across all overlap
 587 windows and report the average across rollouts. This metric captures large plan revisions that are
 588 likely to affect the downstream controller. Lower values indicate stronger temporal consistency.

589 Results show that TeV achieves the second-best temporal consistency among the compared methods.
 590 BID performs best because it explicitly optimizes for overlap consistency, selecting the candidate
 591 with the smallest discrepancy from the previous chunk over shared timesteps. In contrast, TeV does
 592 not directly optimize this metric, yet still substantially improves temporal consistency over the base
 593 policy and other baselines. This suggests that temporally-aware verification encourages smoother
 cross-chunk behavior even without explicitly minimizing the overlap-distance objective.

Table 6: Temporal consistency comparison.

Method	Base	Random	TACO	BID	Ours
Pose	0.098	0.088	0.078	0.073	0.076
Full	0.103	0.089	0.080	0.076	0.078

594

595 E Societal Impacts

596 This work seeks to improve the reliability of flow-matching policies by adding lightweight, temporally-
 597 aware verification at test time. By encouraging temporal consistency across action chunks, the
 598 proposed method can reduce abrupt motions, unstable execution, and task failures, which may benefit
 599 various robotic manipulation tasks. However, more reliable robotic manipulation also carries potential
 600 risks. Improved deployability may accelerate automation in physical labor domains, raising concerns
 601 about workforce displacement and unequal access to robotic technologies. The method also does
 602 not provide formal safety guarantees. Real-world deployment should therefore include task-specific
 603 safety constraints, human oversight, and extensive validation across diverse conditions. We view

604 this work as a step toward more trustworthy generative robot control, rather than a replacement for
605 comprehensive safety evaluation and responsible deployment practices.

606 **F Limitations**

607 Our method is not without limitations. First, verifier-guided generation introduces additional inference
608 latency because each guided integration step requires computing energy gradients with respect to
609 intermediate action samples. This overhead can slow down closed-loop execution when multiple
610 candidates are sampled. Future work could reduce the cost using efficient gradient approximations,
611 such as simultaneous perturbation stochastic approximation (SPSA), lower-frequency guidance,
612 cached gradients, or lightweight learned guidance surrogates. Second, TeV can sometimes produce
613 more conservative motions than baseline methods. We hypothesize that the history-aware verifier
614 favors action chunks that stay close to temporally consistent trajectories, reducing abrupt or out-of-
615 distribution behavior but also lowering motion magnitude in some cases. This trade-off is beneficial
616 for stability-sensitive tasks such as liquid transfer, but may limit execution speed when faster actions
617 are safe. Future work should explore adaptive guidance strength and task-dependent objectives that
618 balance temporal consistency, stability, and speed.

619 **G LLM Usage**

620 We used LLM to assist with language polishing and presentation refinement. Specifically, it was
621 employed to improve clarity, grammar, and flow of text, as well as to adjust tone for academic writing.
622 All LLM outputs were carefully reviewed and edited by the authors.

623 **NeurIPS Paper Checklist**

624 The checklist is designed to encourage best practices for responsible machine learning research,
625 addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove
626 the checklist: **The papers not including the checklist will be desk rejected.** The checklist should
627 follow the references and follow the (optional) supplemental material. The checklist does NOT count
628 towards the page limit.

629 Please read the checklist guidelines carefully for information on how to answer these questions. For
630 each question in the checklist:

- 631 • You should answer [Yes], [No], or [N/A].
- 632 • [N/A] means either that the question is Not Applicable for that particular paper or the
633 relevant information is Not Available.
- 634 • Please provide a short (1–2 sentence) justification right after your answer (even for [N/A]).

635 **The checklist answers are an integral part of your paper submission.** They are visible to the
636 reviewers, area chairs, senior area chairs, and ethics reviewers. You will also be asked to include it
637 (after eventual revisions) with the final version of your paper, and its final version will be published
638 with the paper.

639 The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation.
640 While [Yes] is generally preferable to [No], it is perfectly acceptable to answer [No] provided a
641 proper justification is given (e.g., error bars are not reported because it would be too computationally
642 expensive” or “we were unable to find the license for the dataset we used”). In general, answering
643 [No] or [N/A] is not grounds for rejection. While the questions are phrased in a binary way, we
644 acknowledge that the true answer is often more nuanced, so please just use your best judgment and
645 write a justification to elaborate. All supporting evidence can appear either in the main paper or the
646 supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification
647 please point to the section(s) where related material for the question can be found.

648 **IMPORTANT, please:**

- 649 • **Delete this instruction block, but keep the section heading “NeurIPS Paper Checklist”.**
- 650 • **Keep the checklist subsection headings, questions/answers and guidelines below.**
- 651 • **Do not modify the questions and only use the provided macros for your answers.**

652 **1. Claims**

653 Question: Do the main claims made in the abstract and introduction accurately reflect the
654 paper’s contributions and scope?

655 Answer: [Yes]

656 Justification: The claims justify as they are.

657 Guidelines:

- 658 • The answer [N/A] means that the abstract and introduction do not include the claims
659 made in the paper.
- 660 • The abstract and/or introduction should clearly state the claims made, including the
661 contributions made in the paper and important assumptions and limitations. A [No] or
662 [N/A] answer to this question will not be perceived well by the reviewers.
- 663 • The claims made should match theoretical and experimental results, and reflect how
664 much the results can be expected to generalize to other settings.
- 665 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
666 are not attained by the paper.

667 **2. Limitations**

668 Question: Does the paper discuss the limitations of the work performed by the authors?

669 Answer: [Yes]

670 Justification: Discussed in the experiment part as well as supplementary material.

671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721

Guidelines:

- The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate “Limitations” section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren’t acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [N/A]

Justification: Not relevant.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Covered in methodology, algorithm, and experiment parts.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.

- 722
- 723
- 724
- 725
- 726
- 727
- 728
- 729
- 730
- 731
- 732
- 733
- 734
- 735
- 736
- 737
- 738
- 739
- 740
- 741
- 742
- 743
- 744
- 745
- 746
- 747
- 748
- 749
- 750
- 751
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
 - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
 - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
 - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

752 5. Open access to data and code

753 Question: Does the paper provide open access to the data and code, with sufficient instruc-
754 tions to faithfully reproduce the main experimental results, as described in supplemental
755 material?

756 Answer: [Yes]

757 Justification: Will be made publicly available.

758 Guidelines:

- 759
- 760
- 761
- 762
- 763
- 764
- 765
- 766
- 767
- 768
- 769
- 770
- 771
- 772
- 773
- 774
- 775
- The answer [N/A] means that paper does not include experiments requiring code.
 - Please see the NeurIPS code and data submission guidelines (<https://neurips.cc/public/guides/CodeSubmissionPolicy>) for more details.
 - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
 - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://neurips.cc/public/guides/CodeSubmissionPolicy>) for more details.
 - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
 - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
 - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

776 • Providing as much information as possible in supplemental material (appended to the
777 paper) is recommended, but including URLs to data and code is permitted.

778 6. Experimental setting/details

779 Question: Does the paper specify all the training and test details (e.g., data splits, hyperpa-
780 rameters, how they were chosen, type of optimizer) necessary to understand the results?

781 Answer: [Yes]

782 Justification: Detailed in experiment part.

783 Guidelines:

- 784 • The answer [N/A] means that the paper does not include experiments.
- 785 • The experimental setting should be presented in the core of the paper to a level of detail
786 that is necessary to appreciate the results and make sense of them.
- 787 • The full details can be provided either with the code, in appendix, or as supplemental
788 material.

789 7. Experiment statistical significance

790 Question: Does the paper report error bars suitably and correctly defined or other appropriate
791 information about the statistical significance of the experiments?

792 Answer: [Yes]

793 Justification: Reported in tables when relevant.

794 Guidelines:

- 795 • The answer [N/A] means that the paper does not include experiments.
- 796 • The authors should answer [Yes] if the results are accompanied by error bars, confidence
797 intervals, or statistical significance tests, at least for the experiments that support the
798 main claims of the paper.
- 799 • The factors of variability that the error bars are capturing should be clearly stated (for
800 example, train/test split, initialization, random drawing of some parameter, or overall
801 run with given experimental conditions).
- 802 • The method for calculating the error bars should be explained (closed form formula,
803 call to a library function, bootstrap, etc.)
- 804 • The assumptions made should be given (e.g., Normally distributed errors).
- 805 • It should be clear whether the error bar is the standard deviation or the standard error
806 of the mean.
- 807 • It is OK to report 1-sigma error bars, but one should state it. The authors should
808 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
809 of Normality of errors is not verified.
- 810 • For asymmetric distributions, the authors should be careful not to show in tables or
811 figures symmetric error bars that would yield results that are out of range (e.g., negative
812 error rates).
- 813 • If error bars are reported in tables or plots, the authors should explain in the text how
814 they were calculated and reference the corresponding figures or tables in the text.

815 8. Experiments compute resources

816 Question: For each experiment, does the paper provide sufficient information on the com-
817 puter resources (type of compute workers, memory, time of execution) needed to reproduce
818 the experiments?

819 Answer: [Yes]

820 Justification: Included in supplementary material.

821 Guidelines:

- 822 • The answer [N/A] means that the paper does not include experiments.
- 823 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
824 or cloud provider, including relevant memory and storage.
- 825 • The paper should provide the amount of compute required for each of the individual
826 experimental runs as well as estimate the total compute.

827 • The paper should disclose whether the full research project required more compute
828 than the experiments reported in the paper (e.g., preliminary or failed experiments that
829 didn't make it into the paper).

830 9. Code of ethics

831 Question: Does the research conducted in the paper conform, in every respect, with the
832 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

833 Answer: [Yes]

834 Justification: Confirmed to conform.

835 Guidelines:

- 836 • The answer [N/A] means that the authors have not reviewed the NeurIPS Code of
837 Ethics.
- 838 • If the authors answer [No], they should explain the special circumstances that require a
839 deviation from the Code of Ethics.
- 840 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
841 eration due to laws or regulations in their jurisdiction).

842 10. Broader impacts

843 Question: Does the paper discuss both potential positive societal impacts and negative
844 societal impacts of the work performed?

845 Answer: [Yes]

846 Justification: Included in the supplementary materials.

847 Guidelines:

- 848 • The answer [N/A] means that there is no societal impact of the work performed.
- 849 • If the authors answer [N/A] or [No], they should explain why their work has no societal
850 impact or why the paper does not address societal impact.
- 851 • Examples of negative societal impacts include potential malicious or unintended uses
852 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
853 (e.g., deployment of technologies that could make decisions that unfairly impact specific
854 groups), privacy considerations, and security considerations.
- 855 • The conference expects that many papers will be foundational research and not tied
856 to particular applications, let alone deployments. However, if there is a direct path to
857 any negative applications, the authors should point it out. For example, it is legitimate
858 to point out that an improvement in the quality of generative models could be used to
859 generate Deepfakes for disinformation. On the other hand, it is not needed to point out
860 that a generic algorithm for optimizing neural networks could enable people to train
861 models that generate Deepfakes faster.
- 862 • The authors should consider possible harms that could arise when the technology is
863 being used as intended and functioning correctly, harms that could arise when the
864 technology is being used as intended but gives incorrect results, and harms following
865 from (intentional or unintentional) misuse of the technology.
- 866 • If there are negative societal impacts, the authors could also discuss possible mitigation
867 strategies (e.g., gated release of models, providing defenses in addition to attacks,
868 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from
869 feedback over time, improving the efficiency and accessibility of ML).

870 11. Safeguards

871 Question: Does the paper describe safeguards that have been put in place for responsible
872 release of data or models that have a high risk for misuse (e.g., pre-trained language models,
873 image generators, or scraped datasets)?

874 Answer: [N/A]

875 Justification: No such risks posed.

876 Guidelines:

- 877 • The answer [N/A] means that the paper poses no such risks.

- 878
- 879
- 880
- 881
- 882
- 883
- 884
- 885
- 886
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
 - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
 - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

887 12. Licenses for existing assets

888 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
889 the paper, properly credited and are the license and terms of use explicitly mentioned and
890 properly respected?

891 Answer: [Yes]

892 Justification: Properly credited.

893 Guidelines:

- 894
- 895
- 896
- 897
- 898
- 899
- 900
- 901
- 902
- 903
- 904
- 905
- 906
- 907
- 908
- The answer [N/A] means that the paper does not use existing assets.
 - The authors should cite the original paper that produced the code package or dataset.
 - The authors should state which version of the asset is used and, if possible, include a URL.
 - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
 - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
 - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
 - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
 - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

909 13. New assets

910 Question: Are new assets introduced in the paper well documented and is the documentation
911 provided alongside the assets?

912 Answer: [Yes]

913 Justification: The data collected follows the standard LeRobot format. Data collectors are
914 all included as paper authors.

915 Guidelines:

- 916
- 917
- 918
- 919
- 920
- 921
- 922
- 923
- The answer [N/A] means that the paper does not release new assets.
 - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
 - The paper should discuss whether and how consent was obtained from people whose asset is used.
 - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

924 14. Crowdsourcing and research with human subjects

925 Question: For crowdsourcing experiments and research with human subjects, does the paper
926 include the full text of instructions given to participants and screenshots, if applicable, as
927 well as details about compensation (if any)?

928 Answer: [N/A]

929 Justification: Not relevant.

930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: Not relevant.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLM used for writing and polishing.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.