

vla.cpp: A Unified Inference Runtime for Vision-Language-Action Models

Supplementary Appendix

Anonymous Author(s)

Affiliation

Address

email

1 About this document

2 This is the supplementary appendix to *vla.cpp: A Unified Inference Runtime for Vision-Language-*
3 *Action Models*. It collects the extended tables, ablations, and analyses that support the claims of
4 the main paper. Code is attached to the submission of this appendix. Refer to the README in the
5 attached code for downloading the model bundles. Demo videos and the reproducible benchmark
6 scaffold are available at the project website: <https://fai-modelopt-tech.github.io/vla-cpp.github.io/>.
7

8 A From Architecture to Implementation

9 A.1 The served architectures

10 Table 1 summarizes the seven architectures `vla.cpp` serves. They span four vision-encoder and six
11 language-model architectures across the seven models, yet differ in the action head along exactly the
12 two-form split above: six iterative heads, one single-pass head. BitVLA’s single-pass regression head
13 has no solver loop. The GR00T N1.5 and N1.6 vision backbones are the SigLIP2-400M encoders
14 inside the Eagle-2.5 and Eagle3-VL VLM, respectively. GR00T-N1.7 uses the native ViT of the
15 Cosmos-Reason2.

Table 1: The seven served architectures. FM = flow-matching; AltVL = AlternateVL; MLP = multi-layer perceptron; vl-self-attn = vision-language self-attention. T is the number of solver steps the iterative action head integrates per action chunk. The single-pass BitVLA head has no solver loop.

Model	Vision backbone	Language backbone	Action head	Params	T
SmolVLA	SigLIP-So400m	SmolLM2-360M	FM cross-attn expert	450M	10
Evo-1	InternViT-300M	Qwen2.5-0.5B	cross-attn DiT	770M	32
BitVLA	BitSigLIP-L	BitNet-2B	MLP regression	2.4B	-
π_0	SigLIP-So400m	Gemma-2B	FM joint-attn expert	3B	10
GR00T-N1.5	SigLIP2-400M	Qwen3-1.7B	AltVL DiT+vl-self-attn	3B	4
GR00T-N1.6	SigLIP2-400M	Qwen3-1.7B	AltVL DiT	3B	4
GR00T-N1.7	Qwen3-VL ViT	Qwen3-VL	AltVL DiT+vl-self-attn	3B	4

16 A.2 Implementation of the Prefix Extension

17 The runtime targets the two-stage structure of Section 3.1 of the main paper: a vision-language
18 backbone encodes a cached multimodal prefix, and a separate action head consumes it to produce an
19 action chunk. The prefix is exposed, masked, and cached inside the `ggml` graph.

- 20 1. **Exposing full hidden states.** A text runtime returns only logits or a pooled embedding. We
21 tap the final-layer hidden-state tensor of the language model before the output projection
22 and route the full $[\text{tokens} \times d]$ sequence to the action head as the cross-attention source.
- 23 2. **Bidirectional prefix mask.** The prefix is encoded with a bidirectional attention mask rather
24 than the causal mask used for decoding, so images, instruction, and state attend freely. We
25 construct the mask at graph-build time from the segment layout of each request.

26 3. **Cross-attention cache lifecycle.** The prefix keys and values are computed once per ob-
 27 servation and reused across all solver steps. We allocate a dedicated cross-attention cache,
 28 distinct from the language-model self-attention KV cache, that persists for the lifetime of
 29 one denoising integration and is released when the action chunk is returned.

30 The action head comes in two forms, both served behind the same prefix interface. Most families
 31 use an iterative head that cross-attends to the prefix and integrates a chunk over several solver steps:
 32 the flow-matching experts of π_0 and SmolVLA, the cross-attention flow head of Evo-1, and the
 33 diffusion-transformer head of the GR00T series. BitVLA uses the other form, a single-pass regression
 34 head that emits the chunk in one forward pass. The cache lifecycle thus amortizes prefix reuse across
 35 solver steps for iterative heads and degenerates to a single read for the single-pass head, with no
 36 change to the prefix path. This is why BitVLA recurs throughout the evaluation: its contribution is
 37 the first end-to-end ternary stack in a `ggml`-class engine and the custom tensor-core ternary GEMM
 38 of Section 4.5 of the main paper, exercising the footprint and kernel axes rather than the iterative-head
 39 axis of the other six.

40 B Per-Task Agreement Across LIBERO Suites

41 Table 2 reports success rate with 95% Wilson confidence intervals on the `libero_spatial`,
 42 `libero_goal`, and `libero_10` suites alongside `libero_object`, for the architecture-wise exper-
 43 iments. Consistent with the reference-matching statement of Section 4.2 of the main paper, the
 44 relevant quantity is the gap to the reference, not the absolute rate. Saturated suites bound run-to-run
 45 variance and make a one-episode gap meaningful. On all but one suite the runtime and reference
 46 confidence intervals overlap, so the two are statistically indistinguishable. The single exception is
 47 GR00T-N1.7 on Object task, where both policies are near ceiling (runtime 97.0%, reference 99.8%)
 48 and the intervals are narrow enough that a residual 2.8-point gap separates them. The gap is small
 49 relative to the reference’s own run-to-run spread on the harder suites and does not appear on Spatial,
 50 Goal, or Long, so we read it as a near-saturation difference of a few episodes rather than a systematic
 fidelity loss.

Table 2: Experimental results across the four LIBERO suites (SR in % with 95% Wilson intervals in brackets). `vla.cpp` is our runtime and `ref` is the reference checkpoint reproduced under consistent configuration.

Suite	BitVLA		GR00T-N1.7	
	<code>vla.cpp</code>	<code>ref</code>	<code>vla.cpp</code>	<code>ref</code>
Spatial	94.5 [90.4, 96.9]	93.6 [91.1, 95.4]	96.0 [92.3, 98.0]	97.3 [95.7, 98.4]
Object	100.0 [98.1, 100.0]	99.6 [98.6, 99.9]	97.0 [93.6, 98.6]	99.8 [99.1, 100.0]
Goal	92.5 [88.0, 95.4]	91.6 [88.8, 93.7]	97.5 [94.3, 98.9]	98.0 [96.5, 98.9]
Long	83.5 [77.7, 88.0]	85.8 [82.5, 88.6]	89.0 [83.9, 92.6]	91.7 [89.2, 93.6]

51

52 C Optimized and Cross-Stack Baselines

53 Section 4.3 of the main paper compares `vla.cpp` against the released eager-mode PyTorch ref-
 54 erence. Table 3 adds (i) additional architectures and (ii) a graph-captured PyTorch configuration
 55 (`torch.compile` with a static cache), isolating the share of the speedup attributable to the runtime
 56 rather than to eager-mode kernel dispatch. This configuration is reported only where the full for-
 57 ward path lowers into a compiled region; for the architectures whose backbone cannot be captured
 58 (discussed at the end of this section) the entry is left empty. We were unable to build an end-to-end
 59 engine for the iterative cross-attention action head with the single-vendor compiler stacks on the Orin
 60 parts. The capability matrix in Table 1 of the main paper records why each stack does not provide an
 61 apples-to-apples baseline for this inference pattern.

62 The three architectures fall into two regimes: the two compact backbones where the runtime delivers
 63 a clear compute-only speedup, and the large backbone where it matches eager mode. The forward-
 64 computation latencies are reported at the server side, excluding the small ZMQ transport overhead
 65 that the deployed runtime additionally incurs.

Table 3: Per-step latency (ms): eager PyTorch, graph-captured PyTorch (`torch.compile` with a static cache), and `v1a.cpp`. The RTX 5070 rows are the median over 100 timed steps; the GR00T-N1.6 row is the per-step latency of the on-robot ALOHA run (per-chunk inference over an eight-step chunk in Table 4 of the main paper). A dash (–) marks an architecture for which no fully graph-captured configuration is available.

Model	Device	PyTorch (eager)	PyTorch (graph-captured)	<code>v1a.cpp</code>
SmolVLA	RTX 5070	175.99	76.26	74.11
GR00T-N1.7	RTX 5070	101.67	-	101.38
GR00T-N1.6	AGX Orin	77.5	-	58.8

66 For SmolVLA the runtime delivers a $2.4\times$ compute-only speedup over eager PyTorch and edges out
 67 the graph-captured configuration (74.1 vs 76.3 ms). This confirms that most of the gain is structural
 68 rather than an artifact of eager kernel dispatch. GR00T-N1.7 is the more demanding case and matches
 69 eager mode’s result: `v1a.cpp` lands at 101.4 ms, statistically level with eager PyTorch (101.7 ms).
 70 On the embedded AGX Orin, the GR00T-N1.6 row reports the per-step latency of the on-robot
 71 ALOHA run. In this case, `v1a.cpp` reaches 58.8 ms against 77.5 ms for eager PyTorch (470 ms
 72 vs 620 ms per eight-step chunk), the gap that drives the closed-loop success difference analyzed in
 73 Table 4 of the main paper.

74 For GR00T-N1.7 the graph-captured baseline is left empty for a structural reason rather than because
 75 of insufficient tuning. Its Qwen3-VL backbone routes self-attention through an external FlashAttention
 76 CUDA extension rather than through native, traceable ATen operators. Two properties of that
 77 path defeat both capture mechanisms. First, the kernel is opaque to `torch.compile`: TorchDynamo
 78 cannot trace into the external operator and inserts a graph break at its boundary, so the backbone never
 79 lowers into a single compiled region and falls back to eager dispatch. Second, the variable-length
 80 attention path computes its cumulative-sequence-length metadata (`cu_seq_lens`) on the host and
 81 launches kernels over data-dependent, dynamic shapes with host-to-device synchronization. These
 82 operations are not permitted inside a CUDA-graph capture region, which requires fully static shapes
 83 and no host synchronization. Only the action head satisfies these constraints, as it consists of dense,
 84 static-shape attention and GEMMs over a fixed horizon and four denoiser steps. Capturing it in
 85 isolation, however, leaves the dominant backbone in eager mode and does not yield a meaningfully
 86 different end-to-end latency. Both stacks are therefore bottlenecked by the same large backbone, and
 87 the runtime’s broader advantage for this architecture remains its deployment footprint, portability
 88 across the Orin tiers, and the ternary and flow-expert paths it uniquely supports.

89 D Per-Component Roofline Breakdown

90 Table 4 decomposes the single-request forward path of π_0 , a representative large-backbone deploy-
 91 ment, into two phases: the compute-bound prefix, which combines the vision encoder and the
 92 language-model prefill, and the memory-bound action expert. Both stages of the prefix are dense
 93 many-token passes with high weight reuse, which places them in the compute-bound regime. For
 94 each phase we report the measured latency share, the operational intensity, and the regime relative to
 95 the device balance point. This makes explicit the two-phase structure discussed in Section 4.5 of the
 96 main paper.

97 We instrument three architectures that span the backbone-size range, SmolVLA, GR00T-N1.6, and π_0 ,
 98 and report their per-phase split in Table 5. The prefix phase accounts for the majority of the forward
 99 across all three models. It does so decisively for the large-backbone π_0 ($\sim 75\%$), and by a narrower
 100 margin for the compact SmolVLA and GR00T-N1.6 ($\sim 52\%$ each). For the two compact models
 101 the memory-bound action expert grows to a near-equal share ($\sim 48\%$), because its cost rises with
 102 the solver-step count and, for GR00T-N1.6, with a 32-layer cross-attention DiT. Because the device
 103 balance point (the ridge of the roofline) differs across hardware, both the per-phase latency split
 104 and, near the balance point, the regime assignment can shift between tiers. However, the operational
 105 intensity of each phase is a property of the computation itself and is device-independent. It is this
 106 quantity that places the dense prefix far above and the action expert far below the balance point on
 107 the devices we measure.

Table 4: Per-phase latency share and operational intensity (FLOPs/byte) for π_0 (PaliGemma-3B backbone + flow action expert, 10 solver steps), a representative large-backbone deployment, on the RTX 5070. The compute-bound prefix dominates the forward, while the memory-bound action expert’s share grows linearly with the solver-step count (≈ 5.4 ms/step measured).

Phase	Latency share (%)	Op. intensity (FLOPs/B)	Regime
Prefix (vision + LM)	75	$\sim 256\text{-}530$	compute-bound
Action expert	25	~ 50	memory-bound
Total forward path	100	-	prefix-dominated

108 For π_0 the prefill and the per-step denoise are separated by varying T and linear-fitting the fused
 109 inference graph ($\text{inference}(T) = \text{prefill} + T \cdot \text{per-step}$). For SmolVLA the runtime reports the two
 110 phases directly. The compute-bound prefix is the larger phase for every model. It dominates decisively
 111 for the large-backbone π_0 and for the single-pass BitVLA. For the compact SmolVLA (ten solver
 112 steps) and GR00T-N1.6, the margin is narrow, and the memory-bound action expert approaches an
 113 equal share. GR00T-N1.6 reaches this near-equal split even at only four solver steps because of its
 114 heavy 32-layer cross-attention DiT. Both the shares and the regime assignment are device-dependent
 115 near the balance point (Sec. 4.5 of the main paper); each phase’s operational intensity is not.

Table 5: Per-phase latency share (% of the forward compute) across four architectures, measured on the RTX 5070 (median over timed iterations, two camera views). T is the number of solver steps the iterative action head integrates per chunk (Table 1); the single-pass BitVLA head is marked $-$. The prefix phase combines the vision encoder and the language-model prefill.

Model	Language backbone	Size	T	Prefix	Action expert
SmolVLA	SmolLM2-360M	450M	10	52	48
GR00T-N1.6	Qwen3-1.7B	3B	4	52	48
π_0	Gemma-2B	3B	10	75	25
BitVLA	BitNet-2B	2.4B	-	99.5	0.5

116 BitVLA is the extreme of this decomposition. As mentioned in Appendix A, its single-pass action
 117 head has no solver loop. The prefix phase, formed by its one-shot 30-layer BitNet prefill together
 118 with the vision encoder, is therefore $\sim 99.5\%$ of the forward, and the action head $\sim 0.5\%$ (RTX 5070,
 119 measured). This is the cleanest instance of the prefix-compute-bound regime: nearly the entire
 120 forward is the high-arithmetic-intensity prefix that the kernel optimization targets.

121 E Reduced-Precision Ablation

122 This appendix expands the reduced-precision finding of Section 4.6 of the main paper. The vision
 123 encoder selects a row of a position-embedding table using an index computed from the patch grid,
 124 and the rounding of this index differs between 32-bit and 16-bit floating point: a value just below one
 125 rounds down in single precision but up in half precision. For SmolVLA, the model is deployed in
 126 half precision while the position index is still computed in single precision. This mismatch selects
 127 a different embedding row for almost every patch, which shifts the encoded prefix. The shift is
 128 large enough to displace the predicted action chunk by up to approximately 1.97 in the policy’s
 129 denormalized action units (Table 6). These units are the 7-DoF `libero_object` command, a delta
 130 end-effector pose plus a gripper signal recovered by the policy’s mean/std unnormalization, and they
 131 are dimensionless rather than metric. With per-degree-of-freedom action standard deviations of 0.34-
 132 0.44 (translation), 0.04-0.08 (rotation), and ≈ 1.0 (gripper), a displacement of 1.97 is a near-full-scale
 133 ($\sim 2\sigma$) error concentrated on the gripper channel. The command saturates to the opposite extreme
 134 rather than drifting, which is why the grasp fails outright. This displacement is the difference between
 135 grasping the target object and reaching past it: the `libero_object` task never completed, yet no
 136 component of the pipeline reported an error.

137 Making the index precision-aware restored numerical agreement, dropping the position-embedding
 138 relative error from 0.32 to 0.003 and recovering the task. The mechanism of reduced-precision
 139 collapse and its downstream behavioral effects are themselves documented for language and vision
 140 encoders [1, 2]. Our contribution is to quantify it in robot action space and to gate against it. To

141 separate a property of reduced precision in general from a property of the specific operation, Table 7
 142 contrasts SmolVLA with GR00T-N1.6. GR00T-N1.6’s vision encoder carries an analogous position-
 143 embedding step that is continuous rather than discrete. When its input resolution differs from the
 144 native grid, it resizes the position-embedding table by bilinear interpolation instead of selecting a
 145 row by a computed index. For each model we report `libero_object` success at the mismatched
 146 versus precision-aware configuration, together with the growth of the action-chunk displacement with
 147 solver-step count.

Table 6: Effect of the SmolVLA vision-encoder position-index precision on the encoded prefix and on `libero_object` task completion. A single rounding difference, invisible to an error-free run, moves the action chunk far enough to miss the object.

Position-index precision	pos.-embed rel. err.	action chunk max $ \Delta $	task completes
Mismatched (half model, single index)	0.32	1.97	no
Precision-aware (matched)	0.003	$< 10^{-2}$	yes

Table 7: Precision-causality ablation on `libero_object` (10 episodes/configuration, task 0). The catastrophic failure is specific to the discrete position-index selection: SmolVLA’s success collapses and its action displacement compounds with solver steps, whereas GR00T-N1.6’s continuous interpolation perturbs the position table by only $\sim 0.1\%$ and is behaviorally inert.

Model	Vision position op	SR (mismatched)	SR (aware)	action max $ \Delta $ vs. steps
SmolVLA	discrete index	20.0	90.0	0.26 \rightarrow 1.05 (grows)
GR00T-N1.6	continuous interp.	100.0	100.0	~ 0.003 (flat)

148 The comparison localizes the failure to the kind of operation, not to reduced precision as such.
 149 SmolVLA’s position index is a discrete table lookup, so a single rounding difference selects a wrong
 150 row. This is a discontinuous jump in the encoded prefix that the flow-matching solver then amplifies:
 151 the mismatched-versus-aware action displacement grows from 0.26 at one Euler step to 1.05 at eight,
 152 while task success falls from 90% to 20%. This growth is not solver noise. Across the same range
 153 of solver steps, the precision-aware policy’s own discretization error decreases toward its refined
 154 solution, so the two configurations converge to two distinct fixed points. The perturbation is therefore
 155 a systematic bias that the integrator settles into, rather than a transient. GR00T-N1.6’s bilinear table
 156 resize is by contrast a continuous operation. Evaluating it in half versus single precision moves the
 157 position table by only $\sim 0.1\%$ relative, an effect further attenuated by the bf16-resident weights. It
 158 leaves the action chunk unchanged across solver steps and does not flip a single episode (100% in both
 159 configurations). Catastrophic precision sensitivity therefore requires an operation that rounds across
 160 a discrete boundary, whereas continuous operations on the same encoded quantity degrade gracefully.
 161 The distinction is directly actionable for a runtime: `index`, `argmax`, and `bucketize` selections must
 162 be made precision-aware and gated, whereas the surrounding continuous algebra tolerates reduced
 163 precision.

164 F SimplerEnv (WidowX) Results

165 As a second simulator and embodiment, we evaluate GR00T-N1.6 on the WidowX (bridge) tasks of
 166 SimplerEnv at a 252-pixel vision resolution. GR00T-N1.6 reaches 80% on the “put carrot on plate”
 167 task and 70% on the “put spoon on towel” task. The two more complex tasks, “stack the green block
 168 on the yellow block” and “put eggplant in the yellow basket”, yield lower success rates. This pattern
 169 matches the reference policy’s own difficulty profile, consistent with the parity claim of Section 4.2
 170 of the main paper.

171 G Replan Interval and the Latency-Staleness Trade-off

172 The closed-loop argument of Section 4.7 of the main paper rests on a trade-off between inference cost
 173 and observation freshness, set by how many steps of a predicted action chunk are executed before
 174 the policy re-plans. We make this trade-off explicit by varying the replan interval S , the number of

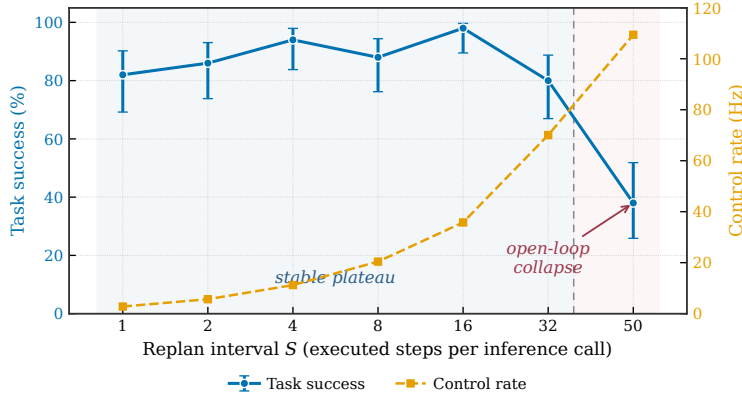


Figure 1: Replan-interval study for SmolVLA on `libero_object` (Jetson Orin Nano, 50 episodes/setting). Task success (left axis, 95% Wilson intervals) is stable across the shaded plateau and collapses once execution becomes fully open-loop at $S=50$. The amortized control rate (right axis) rises almost linearly with S .

Table 8: Replan-interval study for SmolVLA on `libero_object` (Jetson Orin Nano, 50 episodes/setting). S is the executed steps per inference call; latency is per call, the control rate is amortized over the executed steps, and success carries a 95% Wilson interval.

S	Lat. median (ms)	Lat. p95 (ms)	Control rate (Hz)	Success (%)
1	358.4	370.5	2.8	82.0 [69.2, 90.2]
2	352.7	363.4	5.7	86.0 [73.8, 93.0]
4	356.6	383.2	11.2	94.0 [83.8, 97.9]
8	391.7	429.9	20.4	88.0 [76.2, 94.4]
16	447.0	463.3	35.8	98.0 [89.5, 99.6]
32	456.7	466.8	70.1	80.0 [67.0, 88.8]
50	457.0	467.9	109.4	38.0 [25.9, 51.8]

175 executed steps per inference call, for SmolVLA on the `libero_object` suite, deployed on the Jetson
 176 Orin Nano. All other factors are held fixed: a single seed, an identical observation pipeline, fp16
 177 weights with precision-aware index gating, and 50 episodes per setting. Small S re-plans often and
 178 acts on fresh observations at high inference cost; large S executes the chunk open-loop and acts on
 179 increasingly stale observations.

180 Table 8 reports the result. The effective control rate rises almost linearly with S , from 2.8 Hz at $S=1$
 181 to 109 Hz at $S=50$, because each inference call amortizes over more executed steps while its own
 182 latency stays within a narrow band (353-457 ms). Per-call latency is roughly flat for small S and rises
 183 by about a quarter for the longest intervals, consistent with longer uninterrupted execution between
 184 cache refreshes. Task success is statistically stable across a broad plateau of intervals from $S=1$ to
 185 $S=32$, where the per-setting confidence intervals all overlap, and then collapses at $S=50$ to 38%
 186 ([25.9, 51.8]), far below every other setting. The collapse coincides with the point where the policy
 187 executes its entire chunk without a single intermediate observation, confirming that the failure is
 188 driven by observation staleness rather than by inference cost. Figure 1 plots the two axes together
 189 and shows the plateau and the open-loop collapse directly.

190 Two consequences follow for deployment. First, a wide range of replan intervals is safe, so a runtime
 191 can trade inference frequency for throughput over the plateau without sacrificing task success, which
 192 is the headroom the asynchronous execution of Section 4.7 of the main paper exploits. Second, fully
 193 open-loop execution of a long chunk is not safe, which is the same mechanism behind the lower
 194 accuracy of the long-horizon π_0 configuration in Table 2 of the main paper: a long action chunk
 195 consumed without re-planning drifts from the observation it was conditioned on.

196 **References**

- 197 [1] J. Yuan, H. Li, X. Ding, others, and Z. Liu. Give Me FP32 or Give Me Death? Challenges and
198 Solutions for Reproducible Reasoning. arXiv:2506.09501, 2025.
- 199 [2] P. Qi, Z. Liu, X. Zhou, others, and M. Lin. Defeating the Training-Inference Mismatch via FP16.
200 arXiv:2510.26788, 2025.