ON THE DEMAND-SUPPLY GEOMETRY OF REPRESENTATION COLLAPSE

Anonymous authors
Paper under double-blind review

ABSTRACT

Preventing representation collapse remains a central challenge in Joint-Embedding Predictive Architectures (JEPAs). Endpoint-only constraints cannot fully prevent or repair collapse when intermediate representations lose sample distinctions. We introduce **SADL** (Supply And Demand Loss), an adaptive one-sided bridge that matches endpoint demand to the available capacity of the representation path. **SADL** compares demand from the output objective and realized endpoint geometry with the weaker of encoder excitation and cross-endpoint transmission capacity, applying a correction only when demand exceeds supply. It works with different output regularization functions without auxiliary networks, learned gates, or a separate schedule. In single-task and jointly trained four-task LeWorldModel (LeWM) and INTACT experiments, the collapse rate falls from 50% to 0% in a ten-seed collapse test. After only 1 epoch of training with **SADL**, encoder effective rank increases by 3.4–6.4$\times$ and latent prediction error decreases by 20–303$\times$ relative to the corresponding endpoint-only baselines. Under multiple regularization functions, INTACT also achieves over 95% mean search-free control success. Across four image datasets, each with multiple paired seeds, **SADL** delivers consistent seed-wise improvements in accuracy and stability at representative checkpoints, with its largest gains emerging early in training. These findings indicate that demand–supply matching along the representation path is a general principle for faster and more stable representation learning across control and image classification.

1 INTRODUCTION

Joint-Embedding Predictive Architectures (JEPAs) learn from images and video by predicting latent targets instead of reconstructing every pixel (Assran et al., 2023; Bardes et al., 2024). Predictability alone, however, does not ensure a useful representation: suppressing distinctions between samples can also make targets easier to predict. Existing anti-collapse objectives typically specify statistics or geometry at a projected endpoint Z (Zbontar et al., 2021; Bardes et al., 2022a; Balestriero & LeCun, 2025). Yet an apparently regular endpoint does not, on its own, certify that its encoder representation H retains sample variation, or that such variation survives the projection from H to Z. In a world model, losing these distinctions can turn different states or action consequences into nearly indistinguishable control inputs.

The projection interface makes this gap concrete. Endpoint gradients reach H through the Jacobian of the projector; they can shape upstream features, as prior analysis shows (Mialon et al., 2022), but their action depends on the directions that the projector transmits. In our audited endpoint-only INTACT–VISReg run, Official success falls from 55.67% at epoch one to 24.75% at epoch five, while the effective rank of H falls from 13.74 to 6.64 (Table 2). This example does not imply that every endpoint objective collapses. It shows why output-side regularization needs a further condition: the representation path must be able to supply and transmit the variation demanded at its endpoint.

This leads us to ask not which endpoint density is universally preferable, but *whether the endpoint's present demand matches the path's attainable supply*. The answer changes over training and across endpoint objectives. Upstream supply has two distinct limitations: the encoder may fail to excite sample-dependent directions, or the projector may fail to transmit them. We therefore in-

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

troduce **SADL** (**S**upply **A**nd **D**emand **L**oss), a general supply–demand bridge that can **Shape Any Demand-aligned Latent** representation. **SADL** reads demand from both the endpoint objective and the structure already attained at Z, measures supply through the weaker of encoder excitation and cross-endpoint transmission, and corrects H only while supply falls short. OT-based endpoint losses shape the distribution of Z; **SADL** instead repairs insufficient excitation and transmission along $H \to Z$, without optimizing a transport plan. The selected endpoint objective remains unchanged; **SADL** adds no auxiliary networks, learnable gates, or separate regularization schedule.

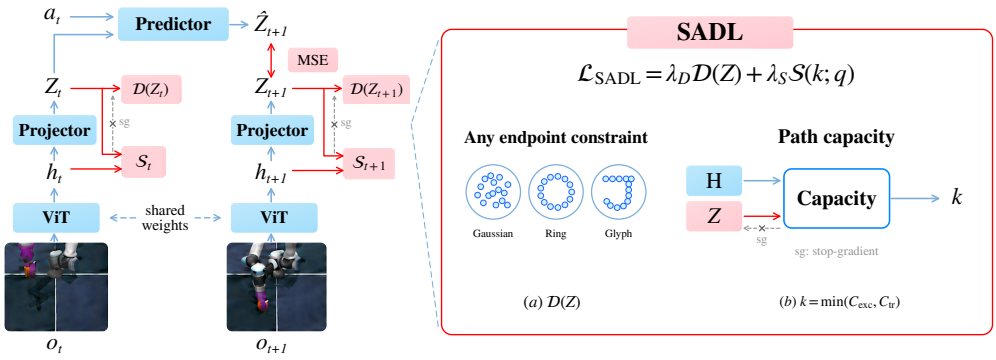

Figure 1: Overview of **SADL**. The endpoint constraint $\mathcal{D}(Z)$ and realized endpoint geometry $C_{\text{exc}}(\text{sg}(Z))$ jointly determine demand q; encoder excitation and cross-endpoint transmission determine path capacity k. Their mismatch is penalized by the one-sided supply loss $\mathcal{S}$ alongside the task objective. The figure displays the scalar dependence on q and k; Section 3.4 gives its gradient-routing arguments.

We evaluate **SADL** on LeWorldModel (LeWM) and INTACT with multiple endpoint constraints in single-task and jointly trained four-task control, and examine its generalization to image classification. Across more than ten tested endpoint target distributions and geometries, one epoch of joint four-task INTACT training yields at least 92% mean search-free success under the reported evaluation protocols (Appendix Figure 6). By epoch three, the INTACT curves reach approximately 96% mean Official Direct success (Figure 3). Across image self-supervised learning datasets, **SADL** also serves as a near-zero-overhead plug-in that accelerates convergence while yielding more stable and higher-quality representations.

Our contributions are:

- We explicitly separate endpoint statistical regularity, encoder excitation, and cross-endpoint transmission. Across endpoint choices, our analysis identifies a mismatch between endpoint demand and upstream path capacity as a mechanism of collapse; endpoint loss form alone does not diagnose it.
- We introduce **SADL**, a low-overhead, adaptive, one-sided correction loss function that attaches to any endpoint objective satisfying the stated conditions and targets an observed supply deficit rather than imposing a fixed upstream distribution.
- Across more than ten tested endpoint objectives and several representation-learning settings, we study early control, convergence, representation probes, and image classification. Across the evaluated tasks, the results support path-capacity matching as a transferable design principle and show faster convergence and more stable representations.

2 Related Work

Predictive representations and control world models. Latent-dynamics models learn compact states for planning or policy improvement (Hafner et al., 2019; 2023; Hansen et al., 2024); iCEM and self-predictive representations refine planning and predictive training (Pinneri et al., 2021; Schwarzer et al., 2021). I-JEPA and V-JEPA predict in learned image/video feature spaces (Assran et al., 2023; Bardes et al., 2024), with later video models extending toward planning and dense

features (Assran et al., 2025; Mur-Labadia et al., 2026). Other JEPA-style world models explore compact dynamics, object-level latent masking, and value-guided planning (Zhou et al., 2024; Nam et al., 2026; Destrade et al., 2026). LeWM and INTACT provide the planning and search-free control interfaces studied here (Maes et al., 2026; Sun et al., 2026); PhyLatent distinguishes statistical non-collapse from dynamics-relevant information (Zeng et al., 2026). **SADL** focuses on the encoder–projector path capacity supporting a chosen endpoint objective.

Self-supervised objectives and collapse. Contrastive image learning uses pairwise or assignment relations (Chen et al., 2020; He et al., 2020; Caron et al., 2020), whereas BYOL, SimSiam, and DINO learn without explicit negative pairs through architectural or optimization asymmetries (Grill et al., 2020; Chen & He, 2021; Caron et al., 2021). Masked prediction and self-distillation provide other routes (He et al., 2022; Zhou et al., 2022; Assran et al., 2022; Baevski et al., 2022; Oquab et al., 2023). These methods differ in their prediction targets and collapse-avoidance mechanisms; **SADL** instead studies how upstream supply is transmitted to meet endpoint demand along the representation path.

Endpoint geometry and statistical regularization. Alignment and uniformity, whitening, Barlow Twins, VICReg, and VICRegL constrain embedding geometry through distributional or moment statistics (Wang & Isola, 2020; Ermolov et al., 2021; Zbontar et al., 2021; Bardes et al., 2022a;b; Balestriero & LeCun, 2025; Zhao et al., 2026). VISReg's shape term uses sliced W_2^2 matching, while QQWorld's quantile objective approximates one-dimensional W_2^2 (Wu et al., 2026; Yu et al., 2026). In the balanced setting, Earth Mover's Distance corresponds to W_1 for normalized discrete measures with a metric ground cost (Peyré & Cuturi, 2019; Rubner et al., 2000). The transport-based objectives above constrain endpoint geometry; **SADL** retains the selected loss but compares its demand with upstream excitation and cross-endpoint association, without optimizing a transport coupling.

Projection interfaces and intermediate regularization. DirectCLR and Guillotine Regularization show that projection and readout choices affect dimensional collapse and downstream utility (Jing et al., 2022; Bordes et al., 2023). Under specific projector conditions, endpoint variance–covariance control can induce pre-projection independence; VCReg and Seq-VCR also constrain intermediate representations (Mialon et al., 2022; Zhu et al., 2023; Arefin et al., 2024). Analyses of non-contrastive Siamese learning examine architectural and optimization routes around collapse (Zhang et al., 2022; Li et al., 2022), while the contrastive–non-contrastive duality concerns relations between loss families under stated assumptions (Garrido et al., 2023). Unlike fixed intermediate constraints, **SADL** conditions its correction on demand from the selected endpoint objective.

3 METHOD

For N aligned samples, an encoder f_θ and projector g_ϕ produce $H = f_\theta(X) \in \mathbb{R}^{N\times d_h}$ and $Z = g_\phi(H) \in \mathbb{R}^{N\times d_z}$. We retain the task loss $\mathcal{L}_{\text{task}}$ and an endpoint constraint $\mathcal{D}(Z) \geq 0$, where lower values indicate greater endpoint satisfaction and $0 < \mathcal{D}(0) < \infty$. For a transport-based $\mathcal{D}$, **SADL** calibrates q_R from its detached scalar value; the bridge solves no transport problem (Peyré & Cuturi, 2019; Wu et al., 2026; Yu et al., 2026). **SADL** reads a detached endpoint demand q, measures the available path supply k, and corrects H only when $k < q$. Figure 2 summarizes the demand and supply branches and their gradient routes.

3.1 ENDPOINT CONSTRAINTS AND REPRESENTATION-PATH CAPACITY

For the batch-level projector $Z = \mathcal{G}_\phi(H)$, backpropagation gives $\nabla_H \mathcal{D}(Z) = J_{\mathcal{G}_\phi}(H)^\top \nabla_Z \mathcal{D}(Z)$. Thus, for any perturbation ΔH,

$$|\langle \nabla_H \mathcal{D}(Z), \Delta H\rangle_F| \leq \|\nabla_Z \mathcal{D}(Z)\|_F \left\|J_{\mathcal{G}_\phi}(H)[\Delta H]\right\|_F . \tag{1}$$

The corrective signal along an upstream direction is limited by the projector's local response and vanishes in its Jacobian null space. This motivates a direct constraint on H, coordinated with the structure formed at Z. We characterize representation-path capacity through encoder excitation, which measures sample-dependent variation, and cross-endpoint transmission, which measures its

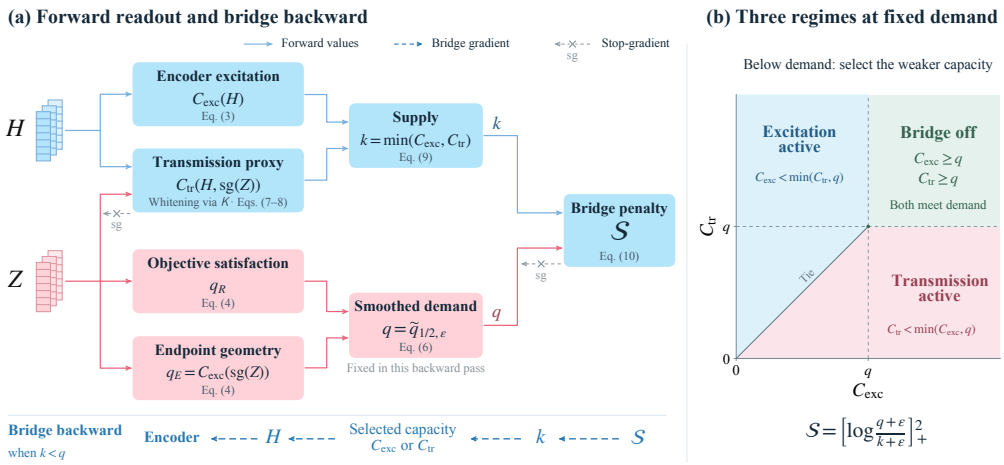

Figure 2: Demand–supply mechanism. Detached endpoint readouts determine demand q, while the weaker of encoder excitation and cross-endpoint transmission determines supply k. The $\tilde{q}_{1/2,\epsilon}$ diagram depicts the smoothed symmetric demand used in training, as defined in equation 6. The bridge acts only for $k < q$ and updates the encoder-side representation; ordinary task and endpoint gradients remain in place.

linear association with Z. This association score operationalizes path support but does not estimate the projector Jacobian.

3.2 Normalized Geometry and Endpoint Demand

For $U \in \mathbb{R}^{N \times d_U}$ and $\epsilon > 0$, define

$$\mu_U = \frac{1}{N} U^\top \mathbf{1}, \quad \bar{U} = U - \mathbf{1}\mu_U^\top,$$
$$\Sigma_U = \frac{1}{N} \bar{U}^\top \bar{U}, \quad e_U = \max\left\{ \frac{\|U\|_F^2}{N d_U}, \epsilon \right\}. \tag{2}$$

The same normalized excitation score is read at Z to calibrate demand and at H to measure encoder-side supply:

$$C_{\text{exc}}(U) = \frac{\log\det(I + \gamma \Sigma_U / e_U)}{d_U \log(1+\gamma)}, \qquad \gamma > 0. \tag{3}$$

This score lies in $[0, 1]$, is orthogonally invariant, and is scale invariant whenever the energy floor remains inactive. Normalization by uncentered energy makes it sensitive to common-mean energy and spectral concentration rather than uniform feature amplification.

The endpoint supplies an objective-based reference and a realized-geometry reference. With $\mathcal{D}_0 = \mathcal{D}(0)$, define

$$q_R := \text{clip}\left(1 - \frac{\mathcal{D}(Z)}{\mathcal{D}_0}, 0, 1\right), \qquad q_E = C_{\text{exc}}(\text{sg}(Z)). \tag{4}$$

Here q_R increases as the endpoint objective is satisfied, while q_E measures current endpoint excitation. The calibration is invariant to common positive scaling of $\mathcal{D}(Z)$ and $\mathcal{D}_0$, but depends on the prescribed loss zero point. For positive references, the ideal log-space interpolation is

$$q_\alpha^\star = q_R^{1-\alpha} q_E^\alpha, \qquad \alpha \in [0, 1]. \tag{5}$$

The implementation evaluates its numerically smoothed counterpart

$$\tilde{q}_{\alpha,\epsilon} := \exp\left((1-\alpha)\log(q_R + \epsilon) + \alpha \log(q_E + \epsilon)\right). \tag{6}$$

We use the symmetric default $q = \tilde{q}_{1/2,\epsilon}$; $q_0 = q_R$ and $q_1 = q_E$ denote separately defined boundary methods. The ideal demand extends continuously to zero when either positive-weight reference is zero, while implementation retains an ϵ-scale numerical floor. Adaptation follows the evolving endpoint statistics without a learned coefficient.

3.3 Excitation, Transmission, and Path Supply

Let $\Sigma_{HZ} = N^{-1}\bar{H}^\top \bar{Z}$. Using a regularized form of canonical correlation (Hotelling, 1936), define

$$\begin{aligned} A_H &= \Sigma_H + \rho\tau_H I, \qquad A_Z = \Sigma_Z + \rho\tau_Z I, \\ \tau_U &= \max\left\{\frac{\operatorname{tr}(\Sigma_U)}{d_U}, \epsilon\right\}, \qquad K = A_H^{-1/2}\Sigma_{HZ}A_Z^{-1/2}, \end{aligned} \tag{7}$$

where $\rho > 0$. The endpoint entering this score is $\mathrm{sg}(Z)$, whereas the relative ridge scale on the H side remains differentiable. This preserves the backward, as well as forward, invariance of the channel score to global rescaling of H whenever the floor is inactive. We measure transmission by

$$C_{\mathrm{tr}}(H, \mathrm{sg}(Z)) = \frac{\log\det(I + \gamma K^\top K)}{r\log(1+\gamma)}, \qquad r = \min(d_h, d_z). \tag{8}$$

Normalization by all potential interface dimensions distinguishes broad association from strong correlation in only a few directions. This second-order association proxy does not capture general nonlinear dependence; at zero cross-covariance, both its value and first-order gradient vanish.

The path supply is the limiting capacity:

$$k = \min\{C_{\mathrm{exc}}(H), C_{\mathrm{tr}}(H, \mathrm{sg}(Z))\}. \tag{9}$$

Hence $k \geq q$ requires both scores to reach q. In this two-stage path-capacity model, the weaker of excitation and transmission is the active cut.

3.4 One-Sided Demand–Supply Matching

We penalize insufficient path supply through the one-sided supply-shortfall correction

$$\mathcal{S}\big(H, \mathrm{sg}(Z); \mathrm{sg}(q)\big) = \left[\log\frac{q+\epsilon}{k+\epsilon}\right]_+^2, \qquad [x]_+ = \max(x, 0). \tag{10}$$

This supply-side constraint covers both excitation and cross-endpoint transmission. In log-capacity coordinates, let $E_D = \log(q+\epsilon)$ and $E_j = \log(k_j + \epsilon)$ for $j \in \{\mathrm{exc}, \mathrm{tr}\}$. The same correction has the energy representation

$$\mathcal{S} = \big[E_D - \min(E_{\mathrm{exc}}, E_{\mathrm{tr}})\big]_+^2 = \max_{j\in\{\mathrm{exc},\mathrm{tr}\}} \big[E_D - E_j\big]_+^2. \tag{11}$$

Geometrically, this is the squared ℓ_∞ distance to the joint support region in smoothed log-capacity coordinates. The correction is zero by construction whenever supply satisfies demand.

The complete objective separates the task loss from the two weighted parts of **SADL**:

$$\begin{aligned} \mathcal{L}_{\mathrm{total}} &= \mathcal{L}_{\mathrm{task}} + \mathcal{L}_{\mathrm{SADL}}, \\ \mathcal{L}_{\mathrm{SADL}} &= \lambda_D \mathcal{D}(Z) + \lambda_S \mathcal{S}\big(H, \mathrm{sg}(Z); \mathrm{sg}(q)\big). \end{aligned} \tag{12}$$

Here λ_D weights the chosen endpoint constraint, λ_S weights the supply-shortfall loss, and α remains solely the demand-interpolation coefficient. Setting $\lambda_S = 0$ removes the supply correction and recovers the selected endpoint-only objective. Detached demand and endpoint references direct the correction toward upstream supply; the task and endpoint terms retain their normal gradients. The ridge scale derived from H is not detached. The ϵ-smoothed demand is used in the correction, and the ideal construction is recovered as $\epsilon \to 0$ away from the zero boundary. Because the correction remains active whenever a nonzero demand exceeds the available path supply, it can also repair an already degraded checkpoint, rather than only prevent collapse during initial training. We evaluate this post-collapse setting in Appendix D.6.

3.5 Gradient Action of Conditional Correction

Let $G_k = \nabla_H^{\mathrm{AD}} k$ denote the supply gradient under the stated stop-gradient conventions. Where the selected branch is differentiable,

$$\begin{aligned} \nabla_H^{\mathrm{AD}}\big(\lambda_S \mathcal{S}\big) &= -a(H, Z) G_k, \\ a(H, Z) &= \frac{2\lambda_S}{k+\epsilon}\left[\log\frac{q+\epsilon}{k+\epsilon}\right]_+. \end{aligned} \tag{13}$$

The deficit sets the correction strength a, and the active bottleneck selects its direction: excitation or cross-endpoint association. Once $k \geq q$, $a = 0$. The encoder gradient is $-aJ_{f_\theta}^\top G_k$, bypassing the projector Jacobian while retaining dependence on the encoder Jacobian and capacity gradient. The correction has no direct gradient on separate projector parameters.

To characterize the excitation branch, treat H as a directly optimized variable and define $r_i = \lambda_i(\Sigma_H)/e_H$.

Proposition 1 (Spectral equalization under excitation limitation)**.** *Assume $\lambda_S > 0$, an inactive energy floor, and $0 < C_{\text{exc}}(H) < \min\{q, C_{\text{tr}}(H, \text{sg}(Z))\}$ throughout the interval considered. The correction-only feature flow $\dot{H} = a\nabla_H C_{\text{exc}}$, with $a > 0$, preserves total feature energy, makes the common-mean norm non-increasing, and satisfies*

$$\frac{\mathrm{d}}{\mathrm{d}t} \log \frac{r_i}{r_j} = \frac{4a\gamma}{Nd_h e_H \log(1+\gamma)} \left(\frac{1}{1+\gamma r_i} - \frac{1}{1+\gamma r_j} \right) \leq 0 \tag{14}$$

for positive relative eigenvalues $r_i \geq r_j > 0$.

Thus, weaker positive directions receive larger relative growth rates, making the positive-spectrum condition number non-increasing while preserving total energy. Exact zero modes remain zero. Under an active excitation bottleneck, **SADL** preferentially strengthens weak but nonzero spectral directions, counteracting incipient collapse before sample variation is completely lost; Appendix D.6 shows that the same deficit-triggered correction can also reopen several near-collapsed checkpoints. This result describes the correction-only feature flow; joint training also incorporates task and endpoint updates.

Sequence features are aligned in batch–time order. Centering, covariance, whitening, and log-determinant calculations use FP32 within each task's local minibatch before shared-parameter gradients are averaged across tasks. Excitation energy and covariance-relative whitening scales retain their gradients, while $\mathcal{D}_0$ is cached from a shape-matched zero input. The architecture and endpoint computation paths remain unchanged.

4 Experiments

Our experiments test whether **SADL** prevents path collapse, accelerates convergence, improves representation quality, and generalizes across endpoint objectives and tasks. We report E1 control and fixed probes, then examine mechanism, endpoint geometries, and image classification; extended benchmarks and per-task results appear in the appendix.

4.1 Experimental setup

We evaluate INTACT and LeWorldModel (LeWM) (Maes et al., 2026; Sun et al., 2026) on PushT (Chi et al., 2023), Cube, Reacher, and TwoRoom, both separately and under joint four-task training, with three endpoint objectives: VISReg, SIGReg, and QQWorld. In the primary comparison, joint models share a ViT-Tiny encoder, projector, and BatchNorm; the BN configuration used for the independent Flash collapse diagnostic is specified in Table 2. Unless noted otherwise, **SADL** uses $\alpha = 0.5$ and $\lambda_S = 0.01$, and control models use BF16. INTACT uses search-free Direct control, whereas LeWM uses Pure-CEM. All control scores below are four-task macro success on the Official benchmark. The primary performance comparison uses three training seeds per configuration and evaluates 300 episodes per task at each checkpoint; we report the mean and sample standard deviation across training seeds. The independent collapse diagnostic uses ten training seeds, with its protocol specified in the table caption.

Table 2(b) holds the encoder and predictor backends, endpoint objective, and coefficient fixed between baseline and +**SADL** (SIGReg 0.03, VISReg 1.5, QQWorld 3.5). Table 1 retains the original single-task protocol; its backend details, additional benchmarks, and multi-epoch results are provided in Appendix E.

4.2 Convergence speed and representation quality

We first test whether path correction makes control-relevant representations available early in training. Each of PushT, Cube, Reacher, and TwoRoom is trained separately, and Official success is

Table 1: Single-task E1 Official macro success rate (SR, %; mean ± SD over three seeds) and matched-stage fixed probes. SR_0/SR_A denote baseline/**+SADL**; r is effective rank, c is mean off-diagonal sample cosine, and subscripts h/z indicate representation stage. Probe entries show baseline → **+SADL**.

Model	Endpoint	SR_0	SR_A	r_h	c_h	r_z	c_z	L_{pred}
LeWM	SIGReg	45.25 ± 6.50	**72.19 ± 0.97**	15.08 → **51.90**	.963 → **.013**	22.97 → **46.32**	.421 → **.007**	3.247 → **.021**
	VISReg	55.67 ± 5.46	**76.31 ± 1.23**	13.30 → **45.60**	.943 → **.013**	21.36 → **34.59**	.214 → **.007**	.586 → **.011**
	QQWorld	51.19 ± 6.98	**75.28 ± 1.73**	15.33 → **51.49**	.956 → **.013**	26.25 → **47.54**	.287 → **.007**	.869 → **.018**
INTACT	SIGReg	58.64 ± 8.14	**95.33 ± 0.44**	15.28 → **51.96**	.976 → **.011**	25.86 → **47.95**	.363 → **.007**	3.160 → **.021**
	VISReg	71.97 ± 12.79	**95.75 ± 0.08**	13.74 → **46.07**	.963 → **.014**	23.77 → **37.39**	.264 → **.007**	.620 → **.012**
	QQWorld	61.56 ± 8.61	**95.47 ± 0.21**	15.16 → **51.26**	.970 → **.013**	27.35 → **49.88**	.350 → **.006**	2.279 → **.019**

macro-averaged across the four tasks. After one epoch, across all six pairs, **SADL** adds 20.6–36.7 points of SR, raises h/z effective rank by 3.35–3.44×/1.57–2.02×, and lowers prediction loss by 48–155× (Table 1).

Joint four-task training tests one shared encoder and projector across all tasks. Table 2 pairs six low-diversity Flash E1 runs with same-seed **+SADL** runs and reports the three-seed comparison. Across the six Flash-matched model–objective pairs in Table 2(b), **SADL** raises E1 Official SR by 23.9–67.4 percentage points: LeWM reaches 63.0–71.6% and INTACT 95.4–95.6%. Effective rank at h/z rises 3.4–6.4/1.6–5.5 times, mean sample cosine falls 57–71/41–93 times, and prediction loss falls 20–303 times. **Both single-task and joint four-task training yield more useful control policies and more discriminative representations after only one epoch.** Per-task results are in Appendix E.

Table 2: Joint four-task E1 Official control and fixed-probe comparison. (a) Six matched single-seed low-diversity Flash examples. (b) Three-seed Flash-matched comparison. SR/SR_{SADL} are without/with **SADL**; arrows show the corresponding probe changes. In (b), SR is mean ± SD and probes are three-seed means.

Model	Endpoint	SR	SR_{SADL}	r_h	c_h	r_z	c_z	L_{pred}
(a) Collapse comparison								
LeWM	SIGReg s42/E1	31.3	**61.8**	17.4 → **54.0**	.926 → **.012**	23.5 → **50.4**	.601 → **.007**	.912 → **.017**
	SIGReg s3072/E1	46.7	**62.9**	12.3 → **53.8**	.960 → **.014**	24.9 → **51.1**	.103 → **.008**	.301 → **.018**
	QQWorld s3072/E1	23.9	**64.0**	14.1 → **55.2**	.992 → **.014**	13.0 → **53.5**	.864 → **.009**	3.752 → **.018**
INTACT	SIGReg s0/E1	26.9	**95.8**	5.4 → **53.5**	.995 → **.013**	9.4 → **53.3**	.795 → **.007**	4.430 → **.017**
	SIGReg s3072/E1	22.5	**95.6**	8.1 → **52.8**	.991 → **.015**	7.2 → **53.4**	.885 → **.009**	10.041 → **.018**
	QQWorld s3072/E1	24.2	**96.0**	4.6 → **53.7**	1.000 → **.015**	7.9 → **54.9**	.833 → **.009**	4.495 → **.019**
(b) Normal comparison								
LeWM	SIGReg	37.2 ± 8.3	**63.0 ± 1.3**	15.7 → **54.0**	.937 → **.013**	24.4 → **50.9**	.419 → **.008**	.651 → **.018**
	VISReg	47.7 ± 5.1	**71.6 ± 3.0**	13.7 → **48.0**	.853 → **.014**	22.7 → **37.7**	.376 → **.009**	.267 → **.014**
	QQWorld	32.9 ± 8.3	**68.4 ± 4.2**	15.8 → **54.5**	.944 → **.013**	21.7 → **52.7**	.619 → **.008**	1.643 → **.017**
INTACT	SIGReg	28.2 ± 6.4	**95.6 ± .2**	8.3 → **53.1**	.994 → **.014**	9.7 → **53.1**	.744 → **.008**	5.300 → **.017**
	VISReg	54.4 ± 2.9	**95.4 ± .6**	13.9 → **48.2**	.912 → **.016**	21.8 → **39.8**	.492 → **.010**	.526 → **.016**
	QQWorld	38.6 ± 12.8	**95.6 ± .3**	12.9 → **53.6**	.952 → **.014**	18.2 → **54.4**	.661 → **.008**	2.044 → **.019**

4.3 Collapse prevention and mechanistic ablations

Endpoint regularization alone does not ensure a healthy representation path. Across six low-diversity E1 Flash checkpoints spanning LeWM, INTACT, SIGReg, and QQWorld, matched **SADL** training restores control and representation diversity (Table 2(a)). In a ten-seed Flash+SIGReg audit, relative path-degradation events fall from 5/10 to 0/10 with full **SADL**. Formal definitions and broader VISReg evidence are in Appendix D.

In four-task LeWM, excitation drives most of the observed gain, so excitation-only approaches Full **SADL** in Table 3(b); transmission remains necessary when the cross-endpoint channel is the limiting capacity, as the controlled optimization and forward replay in Appendix Figure 5 demonstrate.

Table 3: Path collapse and mechanistic ablations. (a) Collapse incidence for four-task LeWM at E1: fraction of observed events among ten seeds; C0 serves only as the matched reference, and ranks and losses are at E1. (b) Independent SIGReg four-task E2 Official SR (%): mean ± SD over three seeds.

(a) Matched collapse events and path metrics

Cell	Incidence	S_{path}	r_h/r_z	L_{pred}
C0 MATH+SIG	reference	.031	5.9/27.4	.7767
C1 Flash	100%	.005	1.3/1.1	27.0896
C2 C1+SIG	50%	.032	6.2/21.5	1.8586
C3 C2+Exc	0%	.370	113.2/71.0	.0174
C4 +**SADL**	0%	.373	113.1/71.6	.0173

C4 adds **SADL** to C2. C1–C4 are judged against seed-matched C0; C0 has no self-incidence. 0% means no event was observed in these ten seeds; C3 is excitation-only.

(b) SIGReg internal ablation: three-seed Official E2

Cell	LeWM CEM	INTACT Direct
Excitation-only	68.86 ± .38	96.25 ± .30
Transmission-only	46.17 ± 22.13	48.67 ± 39.49
Full **SADL**	**69.06 ± 1.16**	**96.56 ± .17**

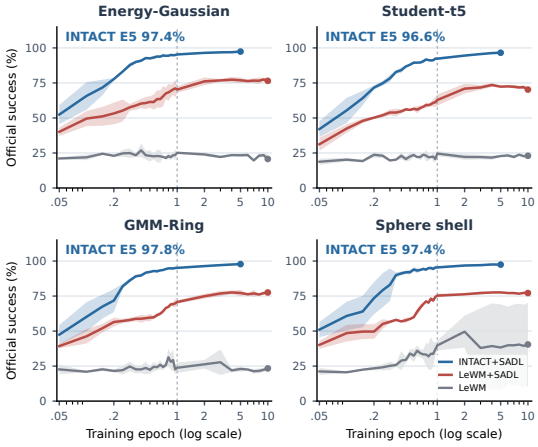

Figure 3: Four-task Official success across four endpoint objectives: INTACT+**SADL** (blue), LeWM+**SADL** (red), and endpoint-only LeWM (gray). Endpoint-only training is collapse-prone and unstable across seeds; with **SADL**, diverse endpoint geometries yield stable, high-quality representations.

4.4 Generality across endpoint objectives

We test **SADL** with eight endpoint objectives rather than assuming a single Gaussian target. After one epoch, all eight give INTACT over 92% search-free Direct success (seven reach 95.03%–95.81%), while LeWM+**SADL** Pure-CEM reaches 62.81%–75.33% (Table 4). Against matched endpoint-only LeWM at 23.81%–47.69%, **SADL** gains 23.89–46.94 points (Appendix Table 33). Student-t_5, GMM-Ring, Sphere shell, and VIC moments further cover non-Gaussian or density-unspecified constraints. **SADL therefore adapts across endpoint objectives without implying that all geometries are equivalent.**

Table 4: Eight-endpoint E1 Official joint control and LeJEPA-related diagnostics. Entries are three-seed means (Macro: mean ± SD); LeWM uses Pure-CEM 300×30, INTACT Direct, and $\lambda_S = .01$. $\widehat{J}_{\text{proj}}$ and b_{1024} are projected Fisher and finite-sample kNN-bias proxies; full definitions, uncertainty, and evaluation protocol appear in Appendix G.

Endpoint	λ_D	PushT	Cube	Reacher	TwoRoom	Macro	$\widehat{J}_{\text{proj}}$	b_{1024}
LeWM – Gaussian targets								
SIGReg	0.0114	56.00	67.11	82.33	80.78	71.56 ± 1.27	1.289	.1772
VISReg	1.5000	47.33	68.44	81.00	89.56	71.58 ± 4.67	1.184	.1717
QQWorld	1.2729	60.89	68.00	80.33	88.22	74.36 ± 2.90	1.179	.1699
Energy-Gaussian	2.3383	49.44	61.56	79.44	90.56	70.25 ± 2.35	1.161	.1712
LeWM – non-Gaussian or density-free targets								
VIC moments	1.9130	35.56	62.00	79.67	92.89	67.53 ± 6.09	1.308	.2228
Sphere	1.1318	73.11	59.56	81.89	86.78	75.33 ± 1.36	1.189	.1660
Student-t_5	1.2104	30.11	69.56	70.44	81.11	62.81 ± 3.99	1.591	.2034
GMM-Ring	2.3345	51.00	63.33	79.78	88.89	70.75 ± 1.23	1.232	.1756
INTACT – Gaussian targets								
SIGReg	0.0117	85.44	100.00	96.89	100.00	95.58 ± 0.00	1.287	.1762
VISReg	1.5000	83.00	99.89	97.44	100.00	95.08 ± 0.17	1.250	.1716
QQWorld	1.2746	85.89	100.00	97.44	99.89	95.81 ± 0.13	1.213	.1745
Energy-Gaussian	2.2074	86.33	100.00	95.11	99.89	95.33 ± 0.63	1.259	.1683

Table 4 continued

Endpoint	λ_D	PushT	Cube	Reacher	TwoRoom	Macro	$\widehat{J}_{\mathrm{proj}}$	b_{1024}
INTACT – non-Gaussian or density-free targets								
VIC moments	1.7339	83.78	100.00	96.33	100.00	95.03 ± 0.99	1.330	.2287
Sphere	1.0725	86.33	100.00	95.89	99.56	95.44 ± 0.82	1.251	.1590
Student-t_5	1.0662	76.67	99.78	93.00	99.89	92.33 ± 1.04	1.660	.2052
GMM-Ring	2.1996	84.89	100.00	95.67	99.89	95.11 ± 0.35	1.220	.1661

4.5 Image representation learning

Beyond control, **SADL** improves ImageNet-LT (Liu et al., 2019), Galaxy10 (Leung & Bovy, 2024), CIFAR-10, and Imagenette under matched budgets and endpoint objectives (Figure 4; Appendix Figure 7, Table 34). It adds no trainable parameters and shows no systematic step-time overhead in a separate control-model audit (Appendix I). Early paired gains span 1.56–8.91 points; under a LeJEPA-style ViT-S/8 SIGReg recipe, all seven Imagenette seeds improve at E5, E20, and E100 (+8.91 points at E20).

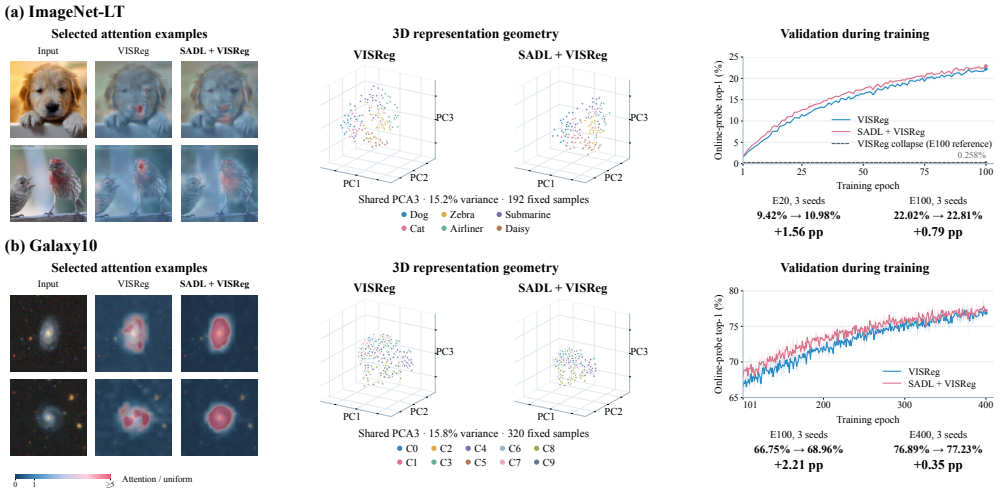

Figure 4: ImageNet-LT (top) and Galaxy10 (bottom): representative attention maps, shared three-dimensional feature geometry, and paired online-validation curves. The baseline remains vulnerable to path degradation; **SADL** converges faster and more stably.

5 Conclusion

We introduce **SADL**, an adaptive one-sided bridge that matches endpoint demand to encoder excitation and cross-endpoint transmission capacity. Across JEPA-style world models and endpoint regularization functions, it accelerates learning and yields more stable, useful representations. Ablations distinguish excitation-driven expansion from correction of transmission bottlenecks; image experiments show benefits beyond control. Together, these results support demand–supply matching as a low-overhead principle for coordinating endpoint objectives with the representation path that must realize them.

Future work. The present transmission score captures a linear, regularized association between hidden and endpoint representations. A natural next step is to generalize path capacity to nonlinear, multimodal, and temporally structured transmission operators. We also aim to develop online diagnostics that anticipate emerging collapse modes and trigger state-adaptive repair, extending **SADL** from early prevention and checkpoint rescue toward prediction and automatic recovery for a broader class of representation failures.

REPRODUCIBILITY STATEMENT

Section 3 and Appendices A–B specify the **SADL** objective, its stop-gradient paths, numerical stabilization, and gradient computation. Appendix C documents model configurations, training and evaluation protocols, success criteria, and aggregation across seeds. Appendices E–H report the complete control comparisons, mechanism ablations, endpoint objectives, and image-classification protocols; the corresponding tables state evaluation budgets and sample counts.

ETHICS STATEMENT

We conduct this work in accordance with the ICLR Code of Ethics. The reported experiments use established image datasets and simulated-control benchmarks and do not involve human participants. Real-world deployment of learned control policies requires task-specific safety evaluation.

AI USE STATEMENT

Generative AI tools were used to aid and polish manuscript writing, refine LaTeX formatting and page layout, and provide suggestions on coding style and implementation. The authors reviewed and verified all resulting text and code and take full responsibility for the final submission.

REFERENCES

Md Rifat Arefin, Gopeshh Subbaraj, Nicolas Gontier, Yann LeCun, Irina Rish, Ravid Shwartz-Ziv, and Christopher Pal. Seq-VCR: Preventing collapse in intermediate transformer representations for enhanced reasoning, 2024. URL https://arxiv.org/abs/2411.02344.

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15619–15629, 2023. URL https://openaccess.thecvf.com/content/CVPR2023/html/Assran_Self-Supervised_Learning_From_Images_With_a_Joint-Embedding_Predictive_Architecture_CVPR_2023_paper.html.

Mahmoud Assran et al. Masked siamese networks for label-efficient learning, 2022. URL https://arxiv.org/abs/2204.07141.

Mido Assran et al. V-JEPA 2: Self-supervised video models enable understanding, prediction and planning, 2025. URL https://arxiv.org/abs/2506.09985.

Alexei Baevski et al. data2vec: A general framework for self-supervised learning in speech, vision and language, 2022. URL https://arxiv.org/abs/2202.03555.

Randall Balestriero and Yann LeCun. LeJEPA: Provable and scalable self-supervised learning without the heuristics. *arXiv preprint arXiv:2511.08544*, 2025. URL https://arxiv.org/abs/2511.08544.

Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*, 2022a. URL https://openreview.net/forum?id=xm6YD62D1Ub.

Adrien Bardes, Jean Ponce, and Yann LeCun. VICRegL: Self-supervised learning of local visual features, 2022b. URL https://arxiv.org/abs/2210.01571.

Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=QaCCuDfBk2.

Florian Bordes, Randall Balestriero, Quentin Garrido, Adrien Bardes, and Pascal Vincent. Guillotine regularization: Why removing layers is needed to improve generalization in self-supervised learning. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=ZgXfXSz51n.

Mathilde Caron et al. Unsupervised learning of visual features by contrasting cluster assignments, 2020. URL https://arxiv.org/abs/2006.09882.

Mathilde Caron et al. Emerging properties in self-supervised vision transformers, 2021. URL https://arxiv.org/abs/2104.14294.

Ting Chen et al. A simple framework for contrastive learning of visual representations, 2020. URL https://arxiv.org/abs/2002.05709.

Xinlei Chen and Kaiming He. Exploring simple siamese representation learning, 2021. URL https://arxiv.org/abs/2011.10566.

Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin C. M. Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.026. URL https://roboticsproceedings.org/rss19/p026.html.

Matthieu Destrade et al. Value-guided action planning with JEPA world models, 2026. URL https://arxiv.org/abs/2601.00844.

Aleksandr Ermolov et al. Whitening for self-supervised representation learning, 2021. URL https://arxiv.org/abs/2007.06346.

Quentin Garrido, Yubei Chen, Adrien Bardes, Laurent Najman, and Yann LeCun. On the duality between contrastive and non-contrastive self-supervised learning. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2206.02574.

Jean-Bastien Grill et al. Bootstrap your own latent: A new approach to self-supervised learning, 2020. URL https://arxiv.org/abs/2006.07733.

Danijar Hafner et al. Learning latent dynamics for planning from pixels, 2019. URL https://arxiv.org/abs/1811.04551.

Danijar Hafner et al. Mastering diverse domains through world models, 2023. URL https://arxiv.org/abs/2301.04104.

Nicklas Hansen et al. TD-MPC2: Scalable, robust world models for continuous control. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/cf73d57b6dcda32b293df7c2d5341f49-Abstract-Conference.html.

Kaiming He et al. Momentum contrast for unsupervised visual representation learning, 2020. URL https://arxiv.org/abs/1911.05722.

Kaiming He et al. Masked autoencoders are scalable vision learners, 2022. URL https://arxiv.org/abs/2111.06377.

Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3–4):321–377, 1936. doi: 10.1093/biomet/28.3-4.321. URL https://doi.org/10.1093/biomet/28.3-4.321.

Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=YevsQ05DEN7.

W. Henry Leung and Jo Bovy. Galaxy10 decals, 2024. URL https://zenodo.org/records/10845026.

Alexander C. Li, Alexei A. Efros, and Deepak Pathak. Understanding collapse in non-contrastive siamese representation learning, 2022. URL https://arxiv.org/abs/2209.15007.

Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. URL https://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Large-Scale_Long-Tailed_Recognition_in_an_Open_World_CVPR_2019_paper.html.

Lucas Maes, Quentin Le Lidec, Damien Scieur, Yann LeCun, and Randall Balestriero. LeWorld-Model: Stable end-to-end joint-embedding predictive architecture from pixels. *arXiv preprint arXiv:2603.19312*, 2026. URL https://arxiv.org/abs/2603.19312.

Grégoire Mialon, Randall Balestriero, and Yann LeCun. Variance covariance regularization enforces pairwise independence in self-supervised representations. *arXiv preprint arXiv:2209.14905*, 2022. URL https://arxiv.org/abs/2209.14905.

Lorenzo Mur-Labadia et al. V-JEPA 2.1: Unlocking dense features in video self-supervised learning, 2026. URL https://arxiv.org/abs/2603.14482.

Heejeong Nam et al. Causal-JEPA: Learning world models through object-level latent masking, 2026. URL https://arxiv.org/abs/2602.11389.

Maxime Oquab et al. DINOv2: Learning robust visual features without supervision, 2023. URL https://arxiv.org/abs/2304.07193.

Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 2019. doi: 10.1561/2200000073.

Cristina Pinneri et al. Sample-efficient cross-entropy method for real-time planning. In *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pp. 1049–1065. PMLR, 2021. URL https://proceedings.mlr.press/v155/pinneri21a.html.

Olivier Roy and Martin Vetterli. The effective rank: A measure of effective dimensionality. In *15th European Signal Processing Conference (EUSIPCO)*, pp. 606–610, 2007. URL https://eurasip.org/Proceedings/Eusipco/Eusipco2007/Papers/a5p-h05.pdf.

Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. The Earth Mover's Distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000. doi: 10.1023/A:1026543900054.

Max Schwarzer et al. Data-efficient reinforcement learning with self-predictive representations, 2021. URL https://arxiv.org/abs/2007.05929.

Junhan Sun, Hao Zhao, and Guofeng Zhang. INTACT: Isomorphic intent-to-action learning for search-free world models. *arXiv preprint arXiv:2607.26056*, 2026. URL https://arxiv.org/abs/2607.26056.

Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere, 2020. URL https://arxiv.org/abs/2005.10242.

Haiyu Wu, Randall Balestriero, and Morgan Levine. VISReg: Variance-invariance-sketching regularization for JEPA training. *arXiv preprint arXiv:2606.02572*, 2026. URL https://arxiv.org/abs/2606.02572.

Zhoushun Yu, Xiaoyu Hu, and Xiangyu Xu. QQWorld: Quantile-quantile matching for world model regularization. *arXiv preprint arXiv:2607.28415*, 2026. URL https://arxiv.org/abs/2607.28415.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow Twins: Self-supervised learning via redundancy reduction. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12310–12320. PMLR, 2021. URL https://proceedings.mlr.press/v139/zbontar21a.html.

Xi Zeng, Haojie Ren, and Ziying Song. PhyLatent: Learning dynamics-relevant representations for JEPA world models, 2026. URL https://arxiv.org/abs/2608.05720.

Chaoning Zhang et al. How does SimSiam avoid collapse without negative samples? a unified understanding with self-supervised contrastive learning, 2022. URL https://arxiv.org/abs/2203.16262.

Kai Zhao, Dongliang Nie, Yuchen Lin, Zhehan Luo, Yixiao Gu, Deng-Ping Fan, and Dan Zeng. Sub-JEPA: Subspace gaussian regularization for stable end-to-end world models. *arXiv preprint arXiv:2605.09241*, 2026. URL https://arxiv.org/abs/2605.09241.

Gaoyue Zhou et al. DINO-WM: World models on pre-trained visual features enable zero-shot planning, 2024. URL https://arxiv.org/abs/2411.04983.

Jinghao Zhou et al. iBOT: Image BERT pre-training with online tokenizer, 2022. URL https://arxiv.org/abs/2111.07832.

Jiachen Zhu, Katrina Evtimova, Yubei Chen, Ravid Shwartz-Ziv, and Yann LeCun. Variance-covariance regularization improves representation learning. *arXiv preprint arXiv:2306.13292*, 2023. URL https://arxiv.org/abs/2306.13292.

APPENDIX

This appendix gives the theoretical derivations, numerical implementation, training and evaluation protocols, collapse diagnosis, complete control results, mechanism ablations, alternative endpoint geometries, image experiments, and computational costs. Unless stated otherwise, control comparisons use three training seeds; each task is evaluated with three evaluation seeds and 100 episodes per seed. Other sample counts are stated with their respective results.

A THEORETICAL DERIVATIONS AND IMPLEMENTATION DETAILS

This appendix provides the proofs, gradient expressions, and numerical details for Section 3.

Table 5: Notation and experimental labels used throughout the appendix.

Symbol	Meaning
N, d_h, d_z	Aligned sample count and encoder/projector dimensions.
$\mathbf{1}, P$	All-ones vector $\mathbf{1} \in \mathbb{R}^N$ and sample-centering matrix $P = I_N - \mathbf{1}\mathbf{1}^\top/N$.
H, Z; h_i, z_i	Aligned representation matrices and their sample rows.
$\hat{Z}$	Endpoint representation predicted by the forward model.
$\bar{U}, \Sigma_U, \Sigma_{HZ}, \mu_U, e_U, v_U, \gamma$	Centered features, sample covariance, cross-covariance, mean, normalized energy, centered-energy fraction, and spectral sensitivity for U.
$\mathcal{D}, \mathcal{D}_0$	Endpoint loss $\mathcal{D}(Z)$ and its zero-representation reference $\mathcal{D}_0 = \mathcal{D}(0)$.
$\mathcal{L}_{\text{task}}, \mathcal{L}_{\text{SADL}}, \mathcal{S}$	Task loss, combined SADL regularizer, and one-sided supply-shortfall loss. The latter takes $(H, \text{sg}(Z); \text{sg}(q))$ and also uses q_{img} in image experiments.
$q_R, q_E, q^\star_\alpha, \widetilde{q}_{\alpha,\epsilon}, q, q_{\text{img}}$	Calibrated, realized, ideal interpolated, numerically smoothed, default, and image-capped demands.
$C_{\text{exc}}, C_{\text{tr}}$	Encoder excitation and cross-endpoint transmission capacities.
$\tau_U, \rho, A_H, A_Z, K, r$	Relative ridge scale and coefficient, regularized endpoint covariances, whitened cross-covariance, and interface dimension $r = \min(d_h, d_z)$.
k	Available path capacity $k = \min\{C_{\text{exc}}(H), C_{\text{tr}}(H, \text{sg}(Z))\}$.
sg, ϵ	Stop-gradient and numerical floor/smoothing parameter.
$\alpha, \lambda_D, \lambda_S$	Log-space demand interpolation, endpoint weight, and bridge weight.
$r_{\text{eff}}, \eta_1, \bar{\kappa}$	Effective covariance rank, leading-eigenvalue share, and mean absolute standardized excess kurtosis.
$\widehat{J}_{\text{proj}}, J(p), n_{\text{nbr}}$	Smoothed projected Fisher proxy, full Fisher functional, and kNN neighbor count; the first two are not interchangeable.
$p, \mathcal{F}, \mathcal{P}_T$	Target endpoint density, downstream task family, and endpoints attainable within training budget T.
E1/E5/E10, Official	Epoch-1/5/10 checkpoints and the control-evaluation protocol defined below.
Endpoint-only, Full SADL	Endpoint regularization alone versus the same endpoint plus the complete SADL bridge. Excitation-only and Transmission-only retain one supply branch each.
MATH, Flash	Explicit versus fused memory-efficient encoder-attention backends, not distinct regularization objectives.

A.1 NOTATION AND BASIC SETUP

Let $H \in \mathbb{R}^{N \times d_h}$ and $Z \in \mathbb{R}^{N \times d_z}$ contain encoder and projected representations of the same samples in the same row order. Define the centering matrix

$$P = I_N - \frac{1}{N}\mathbf{1}\mathbf{1}^\top.$$

For any representation matrix $U \in \mathbb{R}^{N \times d_U}$, write

$$\mu_U = \frac{1}{N}U^\top\mathbf{1}, \qquad \Sigma_U = \frac{1}{N}U^\top PU,$$

$$e_U^{\text{raw}} = \frac{\|U\|_F^2}{Nd_U}, \qquad e_U = \max\{e_U^{\text{raw}}, \epsilon\},$$

where $\epsilon > 0$ is a numerical floor. The total energy decomposes as

$$d_U e_U^{\text{raw}} = \text{tr}(\Sigma_U) + \|\mu_U\|_2^2.$$

Consequently, the centered-energy fraction

$$v_U = \frac{\text{tr}(\Sigma_U)}{d_U e_U}$$

satisfies $0 \leq v_U \leq 1$. When the energy floor is inactive,

$$v_U = 1 - \frac{\|\mu_U\|_2^2}{d_U e_U}.$$

All results below concern finite-batch empirical statistics. They require neither Gaussian representations nor independence between samples within the batch.

A.2 Properties of Excitation Capacity

Recall that

$$C_{\text{exc}}(U) = \frac{\log\det\left(I + \gamma\Sigma_U/e_U\right)}{d_U \log(1+\gamma)}, \qquad \gamma > 0.$$

Let $r_i = \lambda_i(\Sigma_U)/e_U$. Then

$$C_{\text{exc}}(U) = \frac{\sum_{i=1}^{d_U} \log(1+\gamma r_i)}{d_U \log(1+\gamma)}, \qquad \frac{1}{d_U}\sum_{i=1}^{d_U} r_i = v_U.$$

Range and equality conditions. Each logarithmic term is nonnegative because $r_i \geq 0$. By the strict concavity of $\log(1+\gamma r)$,

$$\begin{aligned}\frac{1}{d_U}\sum_{i=1}^{d_U}\log(1+\gamma r_i) &\leq \log\left(1 + \frac{\gamma}{d_U}\sum_{i=1}^{d_U} r_i\right)\\ &= \log(1+\gamma v_U)\\ &\leq \log(1+\gamma).\end{aligned}$$

Therefore,

$$0 \leq C_{\text{exc}}(U) \leq 1.$$

The lower bound is attained if and only if $\Sigma_U = 0$. Equality at the upper bound requires every relative eigenvalue to equal one, equivalently

$$\mu_U = 0, \qquad \Sigma_U = e_U I, \qquad e_U^{\text{raw}} = e_U.$$

Since $\text{rank}(\Sigma_U) \leq N - 1$, attaining the upper bound also requires $N - 1 \geq d_U$.

Coordinate and scale properties. For any orthogonal matrix Q, the transformation $U \mapsto UQ$ preserves total energy and orthogonally conjugates the covariance. Hence

$$C_{\text{exc}}(UQ) = C_{\text{exc}}(U).$$

Now consider $U \mapsto cUQ$, where $c \neq 0$. If the energy floor is inactive before and after the transformation, then

$$\Sigma_{cUQ} = c^2 Q^\top \Sigma_U Q, \qquad e_{cUQ} = c^2 e_U.$$

The normalized covariance changes only by orthogonal conjugation, which gives

$$C_{\text{exc}}(cUQ) = C_{\text{exc}}(U).$$

The normalization retains sensitivity to the common feature mean. At fixed centered covariance, increasing the mean energy increases e_U and decreases the relative covariance spectrum.

Rank–energy bound. Let $k_U = \operatorname{rank}(\Sigma_U) > 0$. Only k_U relative eigenvalues are nonzero, and their sum is $d_U v_U$. Applying Jensen's inequality to these nonzero terms gives

$$\sum_{i=1}^{k_U} \log(1+\gamma r_i) \le k_U \log\left(1 + \frac{\gamma d_U v_U}{k_U}\right).$$

Thus,

$$C_{\mathrm{exc}}(U) \le \frac{k_U}{d_U} \frac{\log\left(1+\gamma d_U v_U / k_U\right)}{\log(1+\gamma)}. \tag{15}$$

Equality holds if and only if all positive covariance eigenvalues are equal. If $k_U = 0$, the excitation capacity is zero.

This bound captures both the number of available representation directions and the fraction of total energy devoted to sample-dependent variation. At fixed rank and centered energy, a uniform nonzero spectrum maximizes excitation capacity.

A.3 Transmission Capacity and the Path-Supply Bound

Define

$$\Sigma_{HZ} = \frac{1}{N} H^\top P Z,$$

$$\tau_H = \max\left\{\frac{\operatorname{tr}(\Sigma_H)}{d_h}, \epsilon\right\}, \qquad \tau_Z = \max\left\{\frac{\operatorname{tr}(\Sigma_Z)}{d_z}, \epsilon\right\},$$

and

$$A_H = \Sigma_H + \rho\tau_H I, \qquad A_Z = \Sigma_Z + \rho\tau_Z I, \qquad \rho > 0.$$

The regularized whitened cross-covariance is

$$K = A_H^{-1/2} \Sigma_{HZ} A_Z^{-1/2}.$$

Let $r = \min(d_h, d_z)$ and let $\sigma_1, \ldots, \sigma_r$ denote the singular values of K. The transmission capacity can be written as

$$C_{\mathrm{tr}}(H, \mathrm{sg}(Z)) = \frac{\sum_{i=1}^{r} \log(1+\gamma\sigma_i^2)}{r\log(1+\gamma)}.$$

Singular-value bound. The joint empirical covariance is positive semidefinite:

$$\begin{pmatrix} \Sigma_H & \Sigma_{HZ} \\ \Sigma_{HZ}^\top & \Sigma_Z \end{pmatrix} = \frac{1}{N} \begin{pmatrix} H^\top \\ Z^\top \end{pmatrix} P \begin{pmatrix} H & Z \end{pmatrix} \succeq 0.$$

Adding the two positive ridge terms yields

$$\begin{pmatrix} A_H & \Sigma_{HZ} \\ \Sigma_{HZ}^\top & A_Z \end{pmatrix} \succ 0.$$

Its Schur complement satisfies

$$A_Z - \Sigma_{HZ}^\top A_H^{-1} \Sigma_{HZ} \succ 0.$$

Congruence by $A_Z^{-1/2}$ then gives

$$I - K^\top K \succ 0.$$

Therefore, $0 \le \sigma_i < 1$ for every i.

Let $k_t = \operatorname{rank}(\Sigma_{HZ})$. Both whitening matrices are invertible, so K has exactly k_t nonzero singular values. It follows that

$$0 \le C_{\mathrm{tr}}(H, \mathrm{sg}(Z)) \le \frac{k_t}{r} \le 1. \tag{16}$$

For $\rho > 0$ and $k_t > 0$, the middle inequality is strict. Transmission capacity is zero if and only if the cross-covariance is zero.

The denominator counts all potential interface modes. Strong transmission within a small active subspace therefore retains a capacity deficit associated with the missing directions.

Finite ridge and invariance. Transmission capacity is invariant to translations and independent orthogonal transformations at the two endpoints. When the corresponding scale floors are inactive, it is also invariant to independent nonzero global rescalings.

A finite ridge contracts the whitened correlation strengths. For example, if $H = Z$, the covariance is isotropic, and the scale floor is inactive, then

$$K = \frac{1}{1+\rho} I$$

and

$$C_{\mathrm{tr}} = \frac{\log\left(1+\gamma/(1+\rho)^2\right)}{\log(1+\gamma)} < 1.$$

The normalization thus refers to ideal transmission in the vanishing-ridge limit.

The ridge is relative to the current covariance scale and remains on the automatic-differentiation path. In particular, under $H \mapsto aH$ with an inactive floor, $\Sigma_H \mapsto a^2\Sigma_H$, $\tau_H \mapsto a^2\tau_H$, and $\Sigma_{HZ} \mapsto a\Sigma_{HZ}$, so the channel score and its radial directional derivative are invariant. Detaching τ_H would leave the forward value unchanged at the current point but would break this backward invariance, creating a spurious radial gradient whenever transmission is the active bottleneck.

The score measures second-order cross-endpoint transmission. Zero cross-covariance does not exclude nonlinear dependence between the representations.

Supply-bound derivation. By definition,

$$k(H, \mathrm{sg}(Z)) = \min\left\{C_{\mathrm{exc}}(H), C_{\mathrm{tr}}(H, \mathrm{sg}(Z))\right\}.$$

Combining Equation 15 with Equation 16 gives

$$k(H, \mathrm{sg}(Z)) \le \min\left\{\frac{k_h}{d_h}\frac{\log\left(1+\gamma d_h v_h/k_h\right)}{\log(1+\gamma)}, \frac{k_t}{r}\right\}.$$

If $k_h = 0$, then $C_{\mathrm{exc}}(H) = 0$ and hence $k(H, \mathrm{sg}(Z)) = 0$.

The minimum corresponds to simultaneous support by both capacity scores:

$$k \ge q \quad \Longleftrightarrow \quad C_{\mathrm{exc}}(H) \ge q \text{ and } C_{\mathrm{tr}}(H, \mathrm{sg}(Z)) \ge q.$$

This is the joint support criterion used in the construction. □

A.4 Demand Calibration and Log-Space Interpolation

Assume that $\mathcal{D}(Z) \ge 0$ and that the zero-representation reference is finite and positive:

$$\mathcal{D}_0 = \mathcal{D}(0) > 0.$$

Reference calibration. Consider an affine calibration of endpoint energy satisfying

$$\mathcal{D}(Z) = 0 \mapsto q_R = 1, \qquad \mathcal{D}(Z) = \mathcal{D}_0 \mapsto q_R = 0.$$

These conditions uniquely determine the affine map $1 - \mathcal{D}(Z)/\mathcal{D}_0$. Restricting its range to the unit interval gives

$$q_R = \mathrm{clip}\left(1 - \frac{\mathcal{D}(Z)}{\mathcal{D}_0}, 0, 1\right).$$

For any constant $c > 0$,

$$1 - \frac{c\mathcal{D}(Z)}{c\mathcal{D}_0} = 1 - \frac{\mathcal{D}(Z)}{\mathcal{D}_0}.$$

The calibrated demand is therefore invariant to the overall units of the endpoint energy.

The second demand readout is

$$q_E = C_{\mathrm{exc}}(\mathrm{sg}(Z)).$$

It uses the relative covariance structure already formed at the endpoint to calibrate external demand. Excitation and transmission have different spectral normalizations, so their values need not obey a fixed ordering.

Log-space barycenter. For $q_R, q_E > 0$, consider

$$F_\alpha(q) = (1-\alpha)(\log q - \log q_R)^2 + \alpha(\log q - \log q_E)^2, \qquad \alpha \in [0,1].$$

Let $y = \log q$, $a = \log q_R$, and $b = \log q_E$. Completing the square gives

$$F_\alpha(q) = (y - [(1-\alpha)a + \alpha b])^2 + \alpha(1-\alpha)(a-b)^2.$$

The second term is independent of q. The unique minimizer therefore satisfies

$$\log q_\alpha^\star = (1-\alpha)\log q_R + \alpha \log q_E,$$

or equivalently,

$$q_\alpha^\star = q_R^{1-\alpha} q_E^\alpha.$$

Under the metric $\mathrm{d}\ell^2 = (\mathrm{d}q/q)^2$ on the positive half-line, this interpolation is a geodesic between the two demand readouts. It also satisfies

$$\min(q_R, q_E) \le q_\alpha^\star \le \max(q_R, q_E).$$

The external and realized-geometry demands are recovered at $\alpha = 0$ and $\alpha = 1$, respectively.

Minimax interpretation of the symmetric choice. For any $y \in \mathbb{R}$,

$$|a-b| \le |a-y| + |y-b| \le 2\max\{|a-y|, |y-b|\}.$$

The largest deviation is thus at least $|a-b|/2$, with equality only at $y = (a+b)/2$. Consequently,

$$\sqrt{q_R q_E} = \operatorname*{argmin}_{q>0} \max\left\{\left|\log\frac{q}{q_R}\right|, \left|\log\frac{q}{q_E}\right|\right\}.$$

The choice $\alpha = \frac{1}{2}$ is therefore the symmetric center of the two readouts in log space. Adaptation follows from changes in the endpoint energy and geometry, without learning the interpolation weight.

Zero-demand boundaries. For fixed $0 < \alpha < 1$, the geometric demand has a continuous extension to the nonnegative quadrant:

$$q_\alpha^\star = 0 \qquad \text{whenever} \qquad q_R q_E = 0.$$

The two boundary methods are defined separately by $q_0 = q_R$ and $q_1 = q_E$.

For numerical evaluation at the boundary, the implementation uses Equation 6. Thus $\widetilde{q}_{\alpha,\epsilon} > 0$ for $\epsilon > 0$, while $\widetilde{q}_{\alpha,\epsilon} \to q_\alpha^\star$ as $\epsilon \to 0$ wherever the ideal demand is positive. This distinguishes the geometric object from its finite-precision realization. For the symmetric default, $\epsilon \le q \le 1+\epsilon$: $q_R = q_E = 1$ gives $q = 1+\epsilon$, $q_R = 0$ and $q_E = c > 0$ give $q = \sqrt{\epsilon(c+\epsilon)} = O(\sqrt{\epsilon})$, and $q_R = q_E = 0$ gives $q = \epsilon$. At this last corner, $k = 0$ yields $\mathcal{S} = (\log 2)^2$; an exactly constant H still has zero first-order capacity gradient, as shown below.

A.5 Geometric Derivation of the One-Sided Bridge

Fix a demand $q \ge 0$ and write

$$c_1 = C_{\mathrm{exc}}(H), \qquad c_2 = C_{\mathrm{tr}}(H, \mathrm{sg}(Z)), \qquad k = \min(c_1, c_2).$$

Introduce log-supply coordinates

$$u_i = \log(c_i + \epsilon), \qquad b_q = \log(q + \epsilon).$$

The two support conditions define the region

$$\mathcal{C}_q = [b_q, \infty)^2.$$

For any $v \in \mathcal{C}_q$,

$$\|u - v\|_\infty \ge \max_{i=1,2} [b_q - u_i]_+.$$

The choice $v_i = \max(u_i, b_q)$ attains this lower bound. Therefore,

$$\begin{aligned}\operatorname{dist}_\infty^2(u, \mathcal{C}_q) &= \max_{i=1,2}[b_q - u_i]_+^2 \\ &= [b_q - \min(u_1, u_2)]_+^2 \\ &= \left[\log \frac{q+\epsilon}{k+\epsilon}\right]_+^2.\end{aligned}$$

This proves the distance representation in Equation 10.

The bridge equals the largest squared log-violation of the two supply constraints and satisfies

$$\mathcal{S} = 0 \quad \Longleftrightarrow \quad c_1 \geq q \text{ and } c_2 \geq q.$$

It exerts no additional geometric pressure once both support conditions hold. Otherwise, its correction is determined by the current bottleneck.

The distance is defined in capacity coordinates. The network parameterization determines how these coordinates change under parameter updates.

A.6 Gradient Paths and Boundary Behavior

Response to insufficient supply. For fixed demand, let

$$\Delta(k) = \left[\log \frac{q+\epsilon}{k+\epsilon}\right]_+.$$

Then

$$\frac{\partial \mathcal{S}}{\partial k} = -\frac{2\Delta(k)}{k+\epsilon}.$$

Within the active region $k < q$,

$$\frac{\partial^2 \mathcal{S}}{\partial k^2} = \frac{2\left(1 + \log \frac{q+\epsilon}{k+\epsilon}\right)}{(k+\epsilon)^2} > 0.$$

At fixed demand, the magnitude of the negative scalar derivative therefore increases as supply decreases. The first derivative is zero for $k \geq q$ and is continuous at $k = q$.

When the two capacities differ,

$$\nabla_H \mathcal{S} = -\frac{2\Delta(k)}{k+\epsilon} \begin{cases} \nabla_H C_{\text{exc}}, & C_{\text{exc}} < C_{\text{tr}}, \\ \nabla_H C_{\text{tr}}, & C_{\text{tr}} < C_{\text{exc}}. \end{cases}$$

At a tie, a generalized gradient is selected from the convex hull of the active branch gradients.

Detached references and endpoint roles. The bridge uses a stopped endpoint representation and demand:

$$\mathcal{S}\big(H, \operatorname{sg}(Z); \operatorname{sg}(q)\big) = \left[\log \frac{q+\epsilon}{k+\epsilon}\right]_+^2.$$

The endpoint tensor and demand are fixed references during the current backward pass. The covariance-relative ridge scale derived from H remains differentiable; the Z side receives no gradient because the endpoint tensor has already been stopped.

With separate encoder and projector parameters θ and ϕ, the bridge's automatic-differentiation gradients satisfy

$$\nabla_\phi^{\text{AD}} \mathcal{S} = 0, \qquad \nabla_\theta^{\text{AD}} \mathcal{S} = J_{f_\theta}^\top \nabla_H \mathcal{S}.$$

The same loss decomposition as Equation 12 remains valid in automatic differentiation:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \mathcal{L}_{\text{SADL}}, \qquad \mathcal{L}_{\text{SADL}} = \lambda_D \mathcal{D}(Z) + \lambda_S \mathcal{S}\big(H, \operatorname{sg}(Z); \operatorname{sg}(q)\big),$$

The task objective and endpoint regularizer retain their normal gradients. Stopping the references directs the bridge toward current upstream supply, rather than directly reducing demand to eliminate the mismatch.

Feature gradient of excitation capacity. Assume that the energy floor is inactive, with $e_H^{\mathrm{raw}} > \epsilon$. Define

$$M = I + \gamma \Sigma_H / e_H, \qquad a_H = \frac{1}{d_h} \operatorname{tr}\left(\frac{\Sigma_H}{e_H} M^{-1} \right).$$

Differentiating excitation capacity gives

$$\mathrm{d}C_{\mathrm{exc}} = \frac{\gamma}{d_h \log(1+\gamma)} \left[\frac{\operatorname{tr}(M^{-1} \mathrm{d}\Sigma_H)}{e_H} - \frac{\operatorname{tr}(M^{-1}\Sigma_H)}{e_H^2} \, \mathrm{d}e_H \right].$$

Substituting the covariance and energy differentials yields

$$\nabla_H C_{\mathrm{exc}} = \frac{2\gamma}{N d_h e_H \log(1+\gamma)} \left(P H M^{-1} - a_H H \right). \tag{17}$$

Along a relative covariance eigenmode r_i, the spectral response of M^{-1} is

$$\frac{1}{1+\gamma r_i}.$$

Weaker directions therefore receive larger spectral response weights. The second term compensates for changes along the global scaling direction, giving

$$\langle H, \nabla_H C_{\mathrm{exc}} \rangle_F = 0.$$

This identity describes the local scale property of excitation capacity itself in representation space.

Zero-demand limit and exact degeneracy. The ideal and implemented demands are, respectively,

$$q^\star = q_R^{1-\alpha} q_E^{\alpha}, \qquad q = \exp\left((1-\alpha) \log(q_R + \epsilon) + \alpha \log(q_E + \epsilon) \right),$$

and the bridge uses

$$\mathcal{S} = [\log(q+\epsilon) - \log(k+\epsilon)]_+^2.$$

If either reference is zero, then $q^\star = 0$ for an interior interpolation weight, but the implemented q remains at an ϵ-controlled positive scale. Exact bridge deactivation is therefore determined by $k \geq q$ in the smoothed implementation, not by substituting the ideal boundary value into the code.

An exactly constant encoder representation has a separate boundary behavior. If $H = \mathbf{1} h_0^\top$, then

$$PH = 0, \qquad \Sigma_H = 0, \qquad \Sigma_{HZ} = 0.$$

Both capacity scores have zero first-order gradients with respect to H at this point. The scalar supply response and the resulting feature gradient must therefore be distinguished. The original task objective and endpoint regularizer remain in the complete objective, while the bridge responds to the nonzero demand and structure formed during training.

B NUMERICAL IMPLEMENTATION

Sample alignment and statistical precision. Sequence representations are flattened in the same batch–time order before computing cross-covariance. If the encoder output has shape $B \times S \times d_h$ and the projected output has shape $S \times B \times d_z$, the first two projected dimensions are exchanged before both tensors are flattened into $N = BS$ samples.

Centering, covariance estimation, whitening, and log-determinant operations within the bridge are performed in FP32 with autocast disabled. Endpoint regularizers retain their respective computation paths.

Cholesky whitening. The implementation uses Cholesky factorizations and triangular solves instead of explicitly forming inverse square roots. Let

$$A = L_H L_H^\top, \qquad B = L_Z L_Z^\top,$$

and compute

$$\widetilde{K} = L_H^{-1} \Sigma_{HZ} L_Z^{-\top}.$$

The matrices $\widetilde{K}$ and K differ only by orthogonal transformations on the left and right, and therefore have the same singular values.

Using

$$\det\left(I + \gamma \widetilde{K}^\top \widetilde{K}\right) = \det\left(I + \gamma \widetilde{K} \widetilde{K}^\top\right),$$

the log-determinant can be computed from the smaller Gram matrix, whose dimension is $r = \min(d_h, d_z)$.

The ridge scales are computed from differentiable mean covariance energies:

$$\tau_U = \max\left\{\frac{\mathrm{tr}(\Sigma_U)}{d_U}, \epsilon\right\}.$$

Both this relative ridge scale and the total energy in excitation capacity retain their gradients, subject to their numerical floors.

Gradient-semantic audit. The release implementation keeps the relative ridge scale differentiable. A historical control runner detached this scalar, which preserves the forward score but can create a radial gradient when transmission is the active bottleneck. Across 25,396 audited capacity records, excitation was the bottleneck in every case. A matched E1 replay further produced identical values for all 204 common diagnostics through step 5,650 of 5,695, excluding wall time. Thus the audited headline configurations are unchanged, while the released implementation follows the scale-invariant gradient semantics stated above.

Reference energy and smoothing. The reference energy $\mathcal{D}_0$ is evaluated on an all-zero tensor with the same shape as the endpoint representation and cached for the current configuration. The evaluation does not update model parameters. Its random-number state is saved and restored so that reference calibration does not alter the subsequent training random sequence.

The energy floors and logarithmic smoothing serve different numerical roles. The former define stable capacity normalizations, while the latter keep the comparison finite at zero supply. Equation 6 also places ϵ inside the implemented interpolation. This is a numerical regularization of the ideal demand, not a change to the ideal log-geodesic definition.

Configuration in the robotic control experiments. Table 6 lists the bridge settings used in the robotic control experiments. All three endpoint regularizers share the same SADL parameters.

Algorithm 1 Single-batch computation of SADL

Require: Aligned representations H, Z; task loss $\mathcal{L}_{\text{task}}$; endpoint regularizer $\mathcal{D}$
Require: Reference $\mathcal{D}_0$; weights λ_D, λ_S; smoothing ϵ
Ensure: Total objective $\mathcal{L}_{\text{total}}$
1: Evaluate $\mathcal{D}(Z)$ with its endpoint gradient retained
2: $q_R \leftarrow \text{clip}(1 - \mathcal{D}(Z)/\mathcal{D}_0, 0, 1)$
3: $q_E \leftarrow C_{\text{exc}}(\text{sg}(Z))$
4: $q \leftarrow \text{sg}(\exp\{[\log(q_R + \epsilon) + \log(q_E + \epsilon)]/2\})$
5: $k \leftarrow \min\{C_{\text{exc}}(H), C_{\text{tr}}(H, \text{sg}(Z))\}$
6: $\mathcal{S} \leftarrow [\log(q + \epsilon) - \log(k + \epsilon)]_+^2$
7: $\mathcal{L}_{\text{SADL}} \leftarrow \lambda_D \mathcal{D}(Z) + \lambda_S \mathcal{S}$
8: **return** $\mathcal{L}_{\text{task}} + \mathcal{L}_{\text{SADL}}$
Capacity statistics and matrix factorizations use FP32 with autocast disabled.

Table 6: SADL configuration in the robotic control experiments. Endpoint weights multiply the corresponding raw regularizer implementations.

Parameter	Value
Demand interpolation weight α	0.5
Capacity parameter γ	9
Ridge coefficient ρ	10^{-3}
Numerical floor and log smoothing ϵ	10^{-6}
Bridge weight λ_S	0.01
SIGReg endpoint weight	0.03
VISReg endpoint weight	1.5
QQ endpoint weight	3.5

C Experimental Setup and Statistical Protocols

C.1 World-Model Training

Single-task experiments train one model per task. Joint experiments share the encoder and projector while retaining task-specific prediction and control modules. BatchNorm running buffers and shared-parameter gradients are averaged across task workers.

Capacity statistics are computed from each task's local batch before gradient aggregation. The bridge therefore acts on within-task statistics, rather than a covariance pooled across tasks.

Table 7: Core world-model training configuration.

Configuration	Setting
Tasks	PushT, Cube, Reacher, TwoRoom
Training seeds	0, 42, 3072
Encoder	ViT-Tiny, trained from scratch
Encoder representation dimension	192
Input image size	224×224
Patch size	14×14
Projector	MLP with BatchNorm; hidden dimension 2,048
Projected dimension	192
Predictor	Six transformer blocks
Training clip length	Eight frames
Conditioning history	Three steps
Numerical precision	BF16 mixed precision
Batch size per task	256 clips
Joint batch size	1,024 clips across four workers
Optimizer	AdamW
Learning rate	5×10^{-4}
Weight decay	10^{-3}
Endpoint-only encoder backend	MATH
SADL and structural-ablation encoder backend	Flash
Predictor backend	MATH

INTACT retains its forward-prediction, inverse-action, and goal-displacement objectives with weights 1.0, 0.1, and 0.05, respectively. LeWM uses forward prediction without the INTACT inverse-action and goal-action objectives. The shared bridge parameters are given in Table 6.

Joint endpoint-only runs extend to ten epochs. Full SADL and excitation-only runs extend to five epochs. The transmission-only, bilateral, one-sided-linear, and bridge-only comparisons use two-epoch runs. Single-task comparisons use the fixed first-epoch checkpoint.

C.2 Control Interfaces and Success Criteria

Control execution. INTACT uses Direct control with $H = 5/RH = 5$. Each replanning step executes five distinct predicted actions. The initial context contains the actual actions preceding the evaluation start, with zero padding only before a true episode boundary. Subsequent context entries contain executed actions.

LeWM uses Pure-CEM with 300 candidate sequences per iteration, 30 iterations, and 30 retained candidates. Both horizons are five, and no actor warm start is used.

CLEAR-LeWM. Moderate and Strict evaluation use a goal offset of 25 primitive steps and an evaluation budget of 50 primitive actions. Sampling is episode-balanced over the `all` split, and initially successful pairs are excluded. These evaluations are not held-out-episode tests. TwoRoom additionally uses cross-room pairs, clean source windows, and swept-collision handling.

Table 8: Task-level CLEAR-LeWM success predicates. "Hold" gives the number of consecutive primitive steps satisfying the predicate. Thresholds use the respective environment coordinate systems.

Task	Moderate	Strict
PushT	Joint agent-and-block position error < 20 native pixel units; block angle error $< 20^\circ$; hold one step.	Block position error < 10; block angle error $< 10^\circ$; hold three steps.
Cube	Position error ≤ 0.04 m; no orientation constraint; hold one step.	Position error ≤ 0.03 m and cube-symmetry-aware orientation error $\leq 15^\circ$; hold three steps.
Reacher	Maximum joint error < 0.05 rad; periodic shoulder and unwrapped elbow error; hold one step.	End-effector distance ≤ 0.01 m; hold two steps.
TwoRoom	Goal distance < 16 native units; hold one step.	Goal distance < 8, valid collision-free route, doorway crossing, and correct goal room; hold one step.

Official. Official retains the upstream evaluator's row-uniform sampling, initially successful pairs, and first-hit success semantics. Its base thresholds are 20 position units and 20° for PushT, 0.04 m without an orientation condition for Cube, 0.05 rad using raw joint errors for Reacher, and 16 native units for TwoRoom. The configured goal offset and evaluation budget are 25 and 50, respectively. The evaluator retains its historical action-loop compatibility convention. Official and CLEAR results are reported separately.

C.3 Aggregation and Pairing

For each training seed and task, evaluation uses seeds 0, 1, and 42, with 100 episodes per evaluation seed. Let $s_{k,t,e}$ denote the corresponding success percentage. We compute

$$s_{k,t} = \frac{1}{3}\sum_{e=1}^{3} s_{k,t,e}, \qquad m_k = \frac{1}{4}\sum_{t=1}^{4} s_{k,t},$$

and report

$$\bar{m} = \frac{1}{3}\sum_{k=1}^{3} m_k, \qquad \mathrm{SD}(m) = \sqrt{\frac{1}{2}\sum_{k=1}^{3}(m_k - \bar{m})^2}.$$

Task-level standard deviations use the same training-seed aggregation. Evaluation episodes are not treated as independent training replicates.

For paired comparisons, differences are computed before aggregation:

$$\Delta_k = m_k^A - m_k^B.$$

We report the mean and sample standard deviation of Δ_k. Rounding is applied only after aggregation.

A complete four-task result contains all 36 task–training-seed–evaluation-seed records. Checkpoint identity, training epoch, control interface, and success criterion are checked before aggregation.

Table 9: Paired 100-step component swaps from identical initialization, input order, random state, and optimizer. Seeds: 3072, 3073, 3074. The first divergence is a numerical trajectory event, not itself a collapse event.

Visual encoder	Latent predictor	Bitwise match	First divergence	Max. parameter gap at step 100
MATH	Flash	$3/3$	Not observed	0
Flash	MATH	$0/3$	Step 1, $3/3$	.0556–.0671

D Collapse Phenomenology and Path Diagnosis

We proceed from the observed failure through controlled component and constraint-placement tests to an operational definition and a path-capacity explanation. Fixed probes identify where variation is lost; a paired ten-seed study estimates relative path-degradation events.

D.1 Observed failure: endpoint constraints do not certify the path

In one jointly trained four-task INTACT–VISReg run, Official Direct success decreases from 55.67% at E1 to 24.75% at E5. At the same checkpoints, $r_{\text{eff}}(H)$ decreases from 13.74 to 6.64, while the mean inter-sample cosine c_H increases from .910 to .997. The control failure coincides with increasingly indistinguishable encoder outputs; endpoint regularization loss alone does not establish representation utility.

The layer-wise audit fixes 256 trajectories per task and eight frames per trajectory. At matched observation indices we probe successive visual-encoder layers, $H = h_{\text{cls}}$, projected Z, and predicted $\hat{Z}$, centering within each task before comparison. Across 37 runs, the cross-run correlation between effective ranks of H and Z is .993. In all 13 observed severe low-rank runs, the earliest abnormal boundary is already in the visual encoder; we find no case in which H remains healthy while Z or the predictor alone first becomes low-rank.

D.2 Controlled localization of the first numerical divergence

Component swap. We clone each initialization, input sequence, random state, and optimizer, changing only the attention backend of the visual encoder or latent predictor. Keeping predictor Flash while switching the encoder to MATH gives bitwise agreement for all three seeds over 100 steps; the reverse swap diverges on step one in all three. In one seed, the first detected difference is an encoder backward-gradient perturbation of 1.49×10^{-8}, which is amplified over subsequent parameter updates. This locates the first measurable numerical perturbation; long-run training separately establishes the representation outcome.

Visual-encoder layer position. With predictor MATH and endpoint SIGReg weight .03, we vary which encoder layers use MATH. In independent run A, protecting the final six layers preserves substantially more variation in $H/Z/\hat{Z}$ than protecting the first six. In run B, MATH in layers 10–12 approaches the all-MATH reference, whereas changing only layers 7–9 does not recover it. Late attention is thus a sensitive location under this architecture and protocol.

Constraint location. Taking the MATH encoder as a healthy reference, we compare endpoint-only SIGReg, a fixed variance/spectral barrier on H alone, and their direct sum with Flash encoder and MATH predictor. The H-only barrier raises $r_{\text{eff}}(H)$ to 27.90, but $Z/\hat{Z}$ reach only $9.67/2.96$; SIGReg(Z) alone gives $4.65/7.65/7.74$. Directly stacking the two constraints does not automatically repair the entire path. Neither *expanding H alone* nor *matching the marginal distribution of Z alone* guarantees a usable readout.

D.3 Numerical trigger versus representational bottleneck

The component swap identifies the first perturbation in visual-encoder Flash-SDPA backward. Shared four-task optimization, BF16, and a relatively large learning rate can amplify this small

Table 10: Encoder-layer backend ablation. Runs A and B are independently trained four-task INTACT models with seed 3072, SIGReg(Z) weight .03, MATH predictor, and the same fixed PushT probe.

Run	Encoder MATH layers	$r_{\text{eff}}(H)$	$r_{\text{eff}}(Z)$	$r_{\text{eff}}(\hat{Z})$
A	All 12	21.13	50.18	49.14
A	None (all Flash)	4.65	7.65	7.74
A	1–6	10.58	20.72	23.00
A	7–12	20.62	51.22	51.32
B	All 12	26.84	66.63	64.78
B	None (all Flash)	13.80	14.86	4.73
B	7–9	7.53	12.08	2.64
B	10–12	27.09	64.61	64.39

Table 11: Fixed-threshold H barrier and endpoint SIGReg placement. Four-task INTACT, seed 3072, E2 fixed PushT probe; predictor MATH throughout. The H barrier penalizes mean standard deviation below .10 or leading eigenvalue share above .35.

Encoder	Constraint	$r_{\text{eff}}(H)$	$r_{\text{eff}}(Z)$	$r_{\text{eff}}(\hat{Z})$
MATH	SIGReg(Z) .03	21.13	50.18	49.14
Flash	SIGReg(Z) .03	4.65	7.65	7.74
Flash	H barrier .01, no Z constraint	27.90	9.67	2.96
Flash	H barrier .01 + SIGReg(Z) .03	2.12	6.14	3.33

difference into distinct non-convex training trajectories. Predictor Flash alone remains bitwise reproducible. Conversely, enabling deterministic algorithms while retaining Flash still yields a low-rank phase in a single-seed run (mean four-task E1 rank 7.36, versus 25.42 with MATH). Short-run divergence and long-run rank loss together connect a numerical trigger to representation degradation.

Once the encoder has erased distinctions between samples, a deterministic downstream map cannot restore them from identical inputs. Moreover, an endpoint gradient reaches H through the projector Jacobian,

$$\nabla_H \mathcal{D}(Z) = J_P(H)^\top \nabla_Z \mathcal{D}(Z).$$

An endpoint objective may therefore produce a gradient at Z without sufficiently exciting or transmitting useful differences in H. Conversely, expanding H alone does not ensure a useful endpoint or predictive readout.

D.4 Operational definition of collapse and degradation

For a fixed probe $U \in \mathbb{R}^{N \times d_U}$, define the centered sample covariance $\Sigma_U = (PU)^\top (PU)/(N-1)$. An exactly constant representation has $\text{rank}(\Sigma_U) = 0$. Training more often produces a concentrated spectrum, nearly vanishing scale, or deteriorated predictive readout, rather than exact rank zero. We therefore report effective rank, mean inter-sample cosine, scale, and prediction error. The 13 severe low-rank runs above were screened by effective rank below half the healthy same-epoch reference.

For the fixed-probe path dimension, let

$$S_{\text{path}} = \min\left\{ \frac{r_{\text{eff}}(H)}{\min(N-1, d_h)}, \frac{r_{\text{eff}}(Z)}{\min(N-1, d_z)} \right\}.$$

This is an external diagnostic, distinct from the training-time supply k. In the formal paired ten-seed audit, each candidate is compared with a same-seed, same-epoch MATH reference. We bootstrap trajectories 10,000 times and record a *relative path-degradation event* only when the upper endpoint of the one-sided 95% confidence interval of $\log(S^{\text{cand}}_{\text{path}}/S^{\text{ref}}_{\text{path}})$ is below zero. This event is distinct from exact rank-zero collapse and task success loss.

Table 12: Paired ten-training-seed E1 four-task LeWM audit. A degradation event is defined relative to the same-seed MATH reference by the fixed-probe bootstrap criterion above; it is not an exact rank-zero event. Success is Official Pure-CEM macro SR (%).

Training condition	Relative events	S_{path}	Prediction loss	Official SR
MATH + SIGReg (reference)	–	.031	.7767	32.02
Flash, no endpoint constraint	10/10	.005	27.0896	24.11
Flash + SIGReg	5/10	.032	1.8586	31.92
Flash + SIGReg + excitation-only SADL	0/10	.370	.0174	66.67
Flash + SIGReg + full SADL	0/10	.373	.0173	64.97

D.5 Matching endpoint demand to path supply

The central remedy is not to force latents onto a fixed manifold, but to ensure that sample excitation in the hidden states and transmission across the projection interface can realize the endpoint's requested representation demand. The endpoint function $\mathcal{D}(Z)$ contributes demand q. The path must both preserve differences in H (excitation capacity C_{exc}) and transmit them to Z (transmission capacity C_{tr}), giving bottleneck supply $k = \min(C_{\text{exc}}, C_{\text{tr}})$. If $q > k$, the request remains unsupported; if $k \geq q$, an auxiliary bridge should not continue forcing H to expand. Excess supply is not itself noise: task readout determines which directions are relevant.

SADL implements this principle with a bridge active only at a deficit,

$$\mathcal{S} = \left[\log(q + \epsilon) - \log(k + \epsilon)\right]_+^2.$$

Unlike an endpoint-only objective, it addresses the upstream path; unlike a fixed H-expansion term, it switches off when supply suffices. In an independent paired ten-seed four-task LeWM audit, Flash+SIGReg has $5/10$ relative degradation events, versus $0/10$ after adding full SADL. S_{path} rises from .032 to .373, prediction loss falls from 1.8586 to .0173, and Official Pure-CEM macro success rises from 31.92% to 64.97%.

Excitation alone also reaches $0/10$ in this training regime, consistent with an excitation-limited path; a separate intervention that makes transmission limiting is given in Section F.3. The paired audit shows that SADL reduces relative path degradation under Flash from $5/10$ to $0/10$, retaining the faster and lower-memory fused attention implementation rather than reverting to MATH (FlashAttention). The bridge adds no model parameters; per-step overhead is reported in Table 35. **SADL also accelerates representation formation and improves E1 path dimension, predictive quality, and control success.**

D.6 SADL Repairs Near-Collapsed Checkpoints

We next ask whether SADL can repair an already degraded path, rather than only prevent degradation from the beginning of training. Starting from fixed near-collapsed checkpoints, we resume training with either the original endpoint objective or the same objective augmented by SADL. Each reported control value averages four tasks and evaluation seeds $\{0, 1, 42\}$, with 100 episodes per task–seed cell. LeWM uses Official Pure-CEM (300×30, $H = R_H = 5$), whereas INTACT uses Official Direct control ($H = R_H = 5$). Within each row, the checkpoint, continuation protocol, training host, and evaluation seeds are matched; rows from different training hosts are not pooled.

The native bridge rescues three of four LeWM checkpoints only weakly or strongly, but fails on seed 42. For that same LeWM state, moderate increases to .03 or .10 make the macro improvement positive, whereas larger or intermediate doses can again fail. INTACT is similarly checkpoint-dependent: seed 0 is not robustly rescued, while seeds 3072 and 42 improve substantially. For seed 42, the Pro6000D matched group is strongest at the native weight and every larger tested dose weakens recovery. The H200 replication confirms a non-monotone response—.15 rebounds above native, while .10 and .20 largely erase the gain—so absolute scores are not combined across hosts.

SADL therefore has genuine post-collapse repair capability, but recovery depends on the checkpoint state and bridge strength; increasing λ_S is not monotonically beneficial. This is more informative than a binary "can rescue" claim: a deficit-triggered correction can reopen a dam-

Table 13: Post-collapse rescue and bridge-strength sensitivity. Panel (a) describes the two seed-42 source checkpoints before continuation using task-averaged fixed probes; r_H, r_Z are effective ranks, c_H, c_Z are mean off-diagonal cosines, and $\mathcal{L}_{\text{pred}}$ is latent prediction loss. Both states retain nonzero numerical rank and are therefore termed *near-collapsed*, not exact rank-zero. Panel (b) reports matched Official macro success rates. Native denotes $\lambda_S = .01$. Exploratory doses are reported to expose non-monotone rescue rather than to select a best result.

(a) Pre-rescue checkpoint diagnostics

Family	Seed / state	r_H	c_H	r_Z	c_Z	$\mathcal{L}_{\text{pred}}$	Diagnosis
LeWM	42 / C2 E2	15.01	.972	33.65	.128	.346	low-rank, highly aligned H
INTACT	42 / E1	18.14	.945	23.62	.599	1.126	low-rank path and poor readout

(b) Matched continuation and dose response

Family	Train seed / start	Continuation group	λ_S	Control SR	SADL SR	Δ (pp)	Outcome
LeWM	21 / E2→E4	matched native	.01	25.250	46.917	+21.667	clear rescue
LeWM	7 / E2→E4	matched native	.01	61.000	72.167	+11.167	clear rescue
LeWM	3072 / E3→E5	matched native	.01	21.500	23.083	+1.583	small positive
LeWM	42 / E2→E4	matched native	.01	24.167	22.500	-1.667	not rescued
LeWM	42 / E2→E4	same source	.03	24.167	28.333	+4.167	rescued
LeWM	42 / E2→E4	post-hoc dose	.05	24.167	23.167	-1.000	not rescued
LeWM	42 / E2→E4	post-hoc dose	.07	24.167	25.667	+1.500	marginal
LeWM	42 / E2→E4	same source	.10	24.167	29.083	+4.917	rescued
LeWM	42 / E2→E4	same source	.15	24.167	21.917	-2.250	not rescued
LeWM	42 / E2→E4	same source	.20	24.167	24.167	+0.000	neutral
LeWM	42 / E2→E4	hardware replicate	.30	24.167	22.250	-1.917	not rescued
LeWM	42 / E2→E4	hardware replicate	1.00	24.167	20.250	-3.917	not rescued
INTACT	0 / E1→E2	matched native	.01	28.917	26.500	-2.417	not rescued
INTACT	3072 / E1→E2	matched native	.01	24.417	53.917	+29.500	clear rescue
INTACT	42 / E1→E2	Pro6000D native	.01	18.917	68.333	+49.417	clear rescue
INTACT	42 / E1→E2	Pro6000D matched	.03	18.917	39.500	+20.583	weakened
INTACT	42 / E1→E2	Pro6000D matched	.10	18.917	29.750	+10.833	weakened
INTACT	42 / E1→E2	Pro6000D matched	.15	18.917	27.000	+8.083	weakened
INTACT	42 / E1→E2	Pro6000D matched	.20	18.917	27.667	+8.750	weakened
INTACT	42 / E1→E2	H200 replicate	.01	23.083	69.083	+46.000	clear rescue
INTACT	42 / E1→E2	H200 replicate	.03	23.083	55.833	+32.750	weakened
INTACT	42 / E1→E2	H200 replicate	.10	23.083	28.917	+5.834	weakened
INTACT	42 / E1→E2	H200 replicate	.15	23.083	74.583	+51.500	rebound
INTACT	42 / E1→E2	H200 replicate	.20	23.083	27.583	+4.500	weakened

aged representation path, yet the intervention must remain commensurate with the path's attainable capacity.

E Additional Control Results

We separate Official from CLEAR-LeWM v0.8 so that success rates are compared only within the same evaluation protocol. The tables report single-task and jointly trained four-task results for LeWM and INTACT under SIGReg, VISReg, and QQWorld. Each task value is the mean ± sample SD across three training seeds after averaging the complete three-evaluation-seed triplet; the macro is first averaged over the four tasks within each training seed and then summarized across training seeds. E1 and E5 are equal-budget comparisons, while endpoint-only E10 is a longer-budget reference. The additional tables here retain the original backend protocol: endpoint-only uses a MATH encoder, +SADL uses a Flash encoder, and both use MATH predictors. The joint E1 Flash-matched comparison appears in Table 2(b).

E.1 Official

Table 14: Single-task LeWM Official success (%). All rows are evaluated after one epoch. LeWM uses Pure-CEM with $H = 5/RH = 5$. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. Values are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	67.00 ± 13.02	61.78 ± 1.64	26.44 ± 12.70	25.78 ± 8.15	45.25 ± 6.50
SIGReg	+SADL	E1	65.22 ± 3.86	58.89 ± 1.84	79.44 ± 1.71	85.22 ± 4.82	$\mathbf{72.19 \pm 0.97}$
VISReg	Endpoint-only	E1	74.11 ± 4.62	63.22 ± 6.45	32.56 ± 2.71	52.78 ± 23.40	55.67 ± 5.46
VISReg	+SADL	E1	76.22 ± 2.17	58.22 ± 0.51	80.89 ± 2.04	89.89 ± 2.55	$\mathbf{76.31 \pm 1.23}$
QQWorld	Endpoint-only	E1	72.22 ± 5.32	64.56 ± 4.86	34.33 ± 16.50	33.67 ± 5.36	51.19 ± 6.98
QQWorld	+SADL	E1	71.78 ± 2.17	61.00 ± 4.63	81.44 ± 1.17	86.89 ± 1.26	$\mathbf{75.28 \pm 1.73}$

Table 15: Single-task INTACT Official success (%). All rows are evaluated after one epoch. INTACT uses search-free Direct control. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. Values are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	80.89 ± 4.35	97.67 ± 2.03	28.78 ± 22.52	27.22 ± 8.98	58.64 ± 8.14
SIGReg	+SADL	E1	85.67 ± 1.45	100.00 ± 0.00	96.89 ± 2.14	98.78 ± 1.07	$\mathbf{95.33 \pm 0.44}$
VISReg	Endpoint-only	E1	84.22 ± 2.50	97.67 ± 1.15	51.89 ± 40.50	54.11 ± 14.99	71.97 ± 12.79
VISReg	+SADL	E1	86.89 ± 0.19	100.00 ± 0.00	97.00 ± 0.58	99.11 ± 0.19	$\mathbf{95.75 \pm 0.08}$
QQWorld	Endpoint-only	E1	82.67 ± 1.20	97.22 ± 3.10	37.33 ± 30.55	29.00 ± 6.17	61.56 ± 8.61
QQWorld	+SADL	E1	87.11 ± 1.39	100.00 ± 0.00	95.22 ± 1.58	99.56 ± 0.19	$\mathbf{95.47 \pm 0.21}$

Table 16: Joint four-task LeWM Official success (%). E1 and E5 compare equal training budgets; endpoint-only E10 is a longer-budget reference. LeWM uses Pure-CEM with $H = 5/RH = 5$. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. Values are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	13.00 ± 1.86	53.11 ± 2.22	26.22 ± 4.11	35.67 ± 7.42	32.00 ± 2.52
SIGReg	+SADL	E1	28.78 ± 5.32	70.22 ± 3.02	79.89 ± 0.77	73.00 ± 1.45	$\mathbf{62.97 \pm 1.25}$
SIGReg	Endpoint-only	E5	75.67 ± 3.71	67.89 ± 2.22	82.00 ± 1.53	35.11 ± 14.42	65.17 ± 2.21
SIGReg	+SADL	E5	67.00 ± 1.76	67.56 ± 2.50	83.33 ± 1.67	64.67 ± 1.76	$\mathbf{70.64 \pm 0.77}$
SIGReg	Endpoint-only	E10	78.78 ± 3.02	67.44 ± 1.95	83.11 ± 1.71	44.78 ± 8.19	68.53 ± 0.75
VISReg	Endpoint-only	E1	53.00 ± 4.84	53.67 ± 0.33	31.56 ± 8.69	55.44 ± 12.62	48.42 ± 3.76
VISReg	+SADL	E1	53.56 ± 11.64	66.44 ± 2.41	78.44 ± 0.84	88.11 ± 1.07	$\mathbf{71.64 \pm 2.96}$
VISReg	Endpoint-only	E5	83.78 ± 1.39	60.00 ± 1.20	83.56 ± 3.10	61.44 ± 7.20	72.19 ± 1.89
VISReg	+SADL	E5	81.44 ± 2.01	66.11 ± 2.55	84.56 ± 2.27	82.22 ± 1.64	$\mathbf{78.58 \pm 0.98}$
VISReg	Endpoint-only	E10	86.33 ± 2.08	58.78 ± 1.26	85.89 ± 0.51	65.00 ± 4.33	74.00 ± 1.08
QQWorld	Endpoint-only	E1	16.67 ± 5.13	55.33 ± 4.10	28.44 ± 3.79	41.78 ± 5.85	35.56 ± 3.97
QQWorld	+SADL	E1	45.67 ± 11.15	67.44 ± 2.87	78.00 ± 3.53	82.56 ± 0.69	$\mathbf{68.42 \pm 4.19}$
QQWorld	Endpoint-only	E5	80.89 ± 3.86	67.11 ± 4.50	85.00 ± 3.53	52.22 ± 12.29	71.31 ± 4.18
QQWorld	+SADL	E5	76.89 ± 2.46	68.00 ± 0.33	83.67 ± 0.88	71.89 ± 2.36	$\mathbf{75.11 \pm 1.34}$
QQWorld	Endpoint-only	E10	86.11 ± 1.17	65.78 ± 2.50	82.78 ± 1.07	56.00 ± 4.51	72.67 ± 0.65

Table 17: Joint four-task INTACT Official success (%). E1 and E5 compare equal training budgets; endpoint-only E10 is a longer-budget reference. INTACT uses search-free Direct control. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. Values are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	37.67 ± 9.29	71.67 ± 6.51	13.89 ± 4.48	55.00 ± 8.99	44.56 ± 2.80
SIGReg	+SADL	E1	85.22 ± 0.77	100.00 ± 0.00	97.11 ± 0.38	100.00 ± 0.00	$\mathbf{95.58 \pm 0.17}$
SIGReg	Endpoint-only	E5	85.78 ± 3.83	99.89 ± 0.19	95.33 ± 0.58	66.00 ± 16.29	86.75 ± 3.20
SIGReg	+SADL	E5	90.78 ± 1.68	100.00 ± 0.00	98.22 ± 0.51	99.78 ± 0.19	$\mathbf{97.19 \pm 0.51}$
SIGReg	Endpoint-only	E10	89.44 ± 3.79	100.00 ± 0.00	96.00 ± 0.58	80.89 ± 11.19	91.58 ± 2.01
VISReg	Endpoint-only	E1	54.67 ± 3.28	81.33 ± 9.85	34.78 ± 11.65	64.11 ± 6.35	58.72 ± 4.79
VISReg	+SADL	E1	83.78 ± 2.04	100.00 ± 0.00	97.78 ± 1.17	100.00 ± 0.00	$\mathbf{95.39 \pm 0.65}$
VISReg	Endpoint-only	E5	89.67 ± 1.76	99.33 ± 0.33	97.67 ± 0.88	76.44 ± 1.95	90.78 ± 0.38
VISReg	+SADL	E5	90.89 ± 1.50	100.00 ± 0.00	99.33 ± 0.33	99.89 ± 0.19	$\mathbf{97.53 \pm 0.25}$
VISReg	Endpoint-only	E10	91.00 ± 0.58	100.00 ± 0.00	97.11 ± 0.19	88.56 ± 5.98	94.17 ± 1.45
QQWorld	Endpoint-only	E1	36.89 ± 11.65	74.22 ± 2.50	21.00 ± 6.77	50.78 ± 9.50	45.72 ± 4.33
QQWorld	+SADL	E1	86.00 ± 0.88	100.00 ± 0.00	96.56 ± 0.51	100.00 ± 0.00	$\mathbf{95.64 \pm 0.32}$
QQWorld	Endpoint-only	E5	89.11 ± 1.68	99.89 ± 0.19	97.44 ± 0.84	82.56 ± 14.20	92.25 ± 3.20
QQWorld	+SADL	E5	91.67 ± 0.88	100.00 ± 0.00	98.56 ± 0.19	100.00 ± 0.00	$\mathbf{97.56 \pm 0.17}$
QQWorld	Endpoint-only	E10	89.56 ± 2.22	100.00 ± 0.00	97.22 ± 1.54	87.11 ± 16.28	93.47 ± 4.69

E.2 CLEAR-LEWM V0.8

E.2.1 SINGLE-TASK MODERATE AND STRICT

Table 18: Single-task LeWM CLEAR-LeWM v0.8 success (%) at E1 under Moderate and Strict criteria. LeWM uses Pure-CEM with $H = 5/R_H = 5$; entries are mean ± sample SD over three training seeds.

Criterion	Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
Moderate	SIGReg	Endpoint-only	E1	70.22 ± 15.32	41.33 ± 3.84	14.56 ± 7.90	10.44 ± 5.80	34.14 ± 6.55
	SIGReg	+SADL	E1	67.89 ± 4.22	38.78 ± 0.51	76.89 ± 6.36	81.67 ± 14.62	$\mathbf{66.31 \pm 2.38}$
	VISReg	Endpoint-only	E1	78.67 ± 3.84	43.33 ± 7.26	18.78 ± 2.87	45.67 ± 32.30	46.61 ± 7.94
	VISReg	+SADL	E1	77.11 ± 4.19	37.67 ± 1.20	77.78 ± 1.54	83.00 ± 5.20	$\mathbf{68.89 \pm 0.61}$
	QQWorld	Endpoint-only	E1	74.33 ± 3.76	42.33 ± 3.93	27.44 ± 16.88	18.33 ± 13.17	40.61 ± 8.59
	QQWorld	+SADL	E1	71.44 ± 3.24	42.33 ± 7.23	75.78 ± 0.69	79.67 ± 3.51	$\mathbf{67.31 \pm 2.56}$
Strict	SIGReg	Endpoint-only	E1	50.67 ± 12.91	18.44 ± 2.01	18.00 ± 7.88	3.89 ± 2.36	22.75 ± 5.40
	SIGReg	+SADL	E1	51.22 ± 0.19	19.00 ± 0.58	82.33 ± 1.15	73.11 ± 9.90	$\mathbf{56.42 \pm 2.43}$
	VISReg	Endpoint-only	E1	55.11 ± 3.02	22.22 ± 5.01	26.22 ± 1.90	30.56 ± 26.34	33.53 ± 6.44
	VISReg	+SADL	E1	57.78 ± 2.71	17.67 ± 2.73	80.11 ± 2.87	73.78 ± 2.91	$\mathbf{57.33 \pm 1.01}$
	QQWorld	Endpoint-only	E1	56.22 ± 6.05	19.33 ± 6.08	36.78 ± 20.35	11.22 ± 12.28	30.89 ± 9.64
	QQWorld	+SADL	E1	58.44 ± 1.02	22.67 ± 4.36	78.78 ± 3.89	71.22 ± 3.66	$\mathbf{57.78 \pm 3.08}$

Table 19: Single-task INTACT CLEAR-LeWM v0.8 success (%) at E1 under Moderate and Strict criteria with search-free Direct control; entries are mean ± sample SD over three training seeds.

Criterion	Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
Moderate	SIGReg	Endpoint-only	E1	80.89 ± 6.95	95.00 ± 5.03	25.11 ± 20.97	10.78 ± 4.67	52.94 ± 7.19
	SIGReg	+SADL	E1	86.67 ± 2.91	100.00 ± 0.00	97.00 ± 0.88	97.22 ± 2.36	$\mathbf{95.22 \pm 1.02}$
	VISReg	Endpoint-only	E1	86.56 ± 3.67	96.22 ± 2.83	47.00 ± 39.40	37.00 ± 41.01	66.69 ± 19.72
	VISReg	+SADL	E1	87.56 ± 2.01	100.00 ± 0.00	97.22 ± 1.02	98.00 ± 1.20	$\mathbf{95.69 \pm 0.97}$
	QQWorld	Endpoint-only	E1	85.00 ± 2.33	96.67 ± 4.36	35.78 ± 31.62	9.11 ± 5.32	56.64 ± 6.16
	QQWorld	+SADL	E1	89.00 ± 0.58	100.00 ± 0.00	97.11 ± 1.50	98.89 ± 1.35	$\mathbf{96.25 \pm 0.55}$
Strict	SIGReg	Endpoint-only	E1	70.67 ± 3.93	80.78 ± 6.38	40.22 ± 29.92	3.00 ± 1.73	48.67 ± 7.91
	SIGReg	+SADL	E1	76.00 ± 2.67	86.33 ± 0.67	98.67 ± 0.88	92.89 ± 1.95	$\mathbf{88.47 \pm 0.84}$
	VISReg	Endpoint-only	E1	73.56 ± 0.51	80.11 ± 5.32	58.67 ± 38.23	23.67 ± 33.50	59.00 ± 16.07
	VISReg	+SADL	E1	78.00 ± 2.73	86.56 ± 0.19	98.22 ± 0.84	93.33 ± 2.89	$\mathbf{89.03 \pm 1.25}$
	QQWorld	Endpoint-only	E1	74.22 ± 1.50	82.33 ± 6.66	51.78 ± 33.01	3.67 ± 4.93	53.00 ± 6.07
	QQWorld	+SADL	E1	79.89 ± 1.26	86.67 ± 1.76	98.11 ± 0.19	92.44 ± 3.91	$\mathbf{89.28 \pm 1.64}$

E.2.2 JOINT MODERATE

Table 20: Joint four-task LeWM Moderate success (%). E1 and E5 compare equal training budgets; endpoint-only E10 provides a longer-budget reference. LeWM uses Pure-CEM with $H = 5/RH = 5$. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. All entries are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	12.89 ± 3.34	29.89 ± 1.35	14.67 ± 5.78	20.33 ± 9.26	19.44 ± 3.59
SIGReg	+SADL	E1	69.00 ± 2.96	49.00 ± 2.00	81.44 ± 2.01	58.11 ± 1.58	$\mathbf{64.39 \pm 1.44}$
SIGReg	Endpoint-only	E5	79.00 ± 3.48	50.44 ± 4.50	61.56 ± 16.78	22.33 ± 18.05	53.33 ± 3.63
SIGReg	+SADL	E5	68.33 ± 2.73	49.00 ± 2.00	67.67 ± 13.86	58.11 ± 1.58	$\mathbf{60.78 \pm 3.63}$
SIGReg	Endpoint-only	E10	76.67 ± 3.76	50.78 ± 1.35	70.44 ± 18.03	28.56 ± 11.03	56.61 ± 6.49
VISReg	Endpoint-only	E1	52.67 ± 8.17	32.56 ± 3.10	21.11 ± 6.30	32.22 ± 15.87	34.64 ± 3.67
VISReg	+SADL	E1	62.56 ± 21.91	45.89 ± 7.34	79.56 ± 3.29	80.00 ± 8.97	$\mathbf{67.00 \pm 3.67}$
VISReg	Endpoint-only	E5	84.78 ± 0.51	39.67 ± 3.84	79.89 ± 0.69	40.56 ± 10.21	61.22 ± 2.94
VISReg	+SADL	E5	83.22 ± 2.27	47.67 ± 3.48	82.56 ± 1.35	68.00 ± 4.06	$\mathbf{70.36 \pm 0.53}$
VISReg	Endpoint-only	E10	86.33 ± 2.19	39.33 ± 1.53	70.56 ± 20.69	47.44 ± 6.16	60.92 ± 4.79
QQWorld	Endpoint-only	E1	14.56 ± 4.86	31.89 ± 6.55	19.89 ± 3.69	26.22 ± 10.87	23.14 ± 5.32
QQWorld	+SADL	E1	71.44 ± 9.31	50.78 ± 1.35	79.44 ± 3.53	67.67 ± 14.47	$\mathbf{67.33 \pm 1.04}$
QQWorld	Endpoint-only	E5	80.67 ± 5.49	49.67 ± 3.18	68.78 ± 21.17	35.11 ± 14.26	58.56 ± 10.60
QQWorld	+SADL	E5	76.78 ± 3.67	52.11 ± 1.17	80.33 ± 0.88	58.56 ± 1.68	$\mathbf{66.94 \pm 0.97}$
QQWorld	Endpoint-only	E10	87.33 ± 1.67	49.00 ± 5.03	68.67 ± 13.89	38.44 ± 7.12	60.86 ± 3.57

Table 21: Joint four-task INTACT Moderate success (%) with search-free Direct control. E1 and E5 compare equal training budgets; endpoint-only E10 provides a longer-budget reference. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. All entries are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	39.11 ± 9.14	54.11 ± 12.04	9.89 ± 4.17	41.33 ± 8.21	36.11 ± 5.28
SIGReg	+SADL	E1	86.22 ± 0.38	100.00 ± 0.00	97.00 ± 0.58	99.67 ± 0.00	$\mathbf{95.72 \pm 0.05}$
SIGReg	Endpoint-only	E5	87.56 ± 5.59	99.67 ± 0.58	97.00 ± 1.45	58.89 ± 20.42	85.78 ± 3.78
SIGReg	+SADL	E5	92.89 ± 0.38	100.00 ± 0.00	99.44 ± 0.69	99.78 ± 0.19	$\mathbf{98.03 \pm 0.13}$
SIGReg	Endpoint-only	E10	91.00 ± 3.18	100.00 ± 0.00	96.89 ± 0.69	73.89 ± 15.08	90.44 ± 3.14
VISReg	Endpoint-only	E1	55.78 ± 3.83	69.78 ± 13.62	31.11 ± 9.67	55.00 ± 11.32	52.92 ± 4.08
VISReg	+SADL	E1	86.78 ± 1.84	100.00 ± 0.00	97.56 ± 0.84	99.67 ± 0.00	$\mathbf{96.00 \pm 0.46}$
VISReg	Endpoint-only	E5	91.89 ± 1.84	98.56 ± 0.51	97.67 ± 0.58	69.22 ± 3.91	89.33 ± 0.92
VISReg	+SADL	E5	93.22 ± 0.38	100.00 ± 0.00	98.89 ± 0.77	99.78 ± 0.38	$\mathbf{97.97 \pm 0.10}$
VISReg	Endpoint-only	E10	93.78 ± 1.64	100.00 ± 0.00	97.78 ± 0.84	89.11 ± 5.30	95.17 ± 1.10
QQWorld	Endpoint-only	E1	37.78 ± 10.67	60.78 ± 3.15	16.56 ± 7.07	30.89 ± 15.12	36.50 ± 4.59
QQWorld	+SADL	E1	88.78 ± 0.84	100.00 ± 0.00	97.44 ± 0.38	99.44 ± 0.19	$\mathbf{96.42 \pm 0.22}$
QQWorld	Endpoint-only	E5	90.11 ± 1.26	99.89 ± 0.19	97.67 ± 0.58	77.11 ± 22.04	91.19 ± 5.22
QQWorld	+SADL	E5	92.00 ± 1.15	100.00 ± 0.00	99.00 ± 0.58	99.67 ± 0.33	$\mathbf{97.67 \pm 0.08}$
QQWorld	Endpoint-only	E10	92.56 ± 1.17	100.00 ± 0.00	97.33 ± 0.58	85.22 ± 22.13	93.78 ± 5.38

E.2.3 JOINT STRICT

Table 22: Joint four-task LeWM CLEAR-LeWM v0.8 Strict success (%). E1 and E5 compare equal training budgets; endpoint-only E10 is a longer-budget reference. LeWM uses Pure-CEM with $H = 5/RH = 5$. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. Values are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	16.00 ± 4.93	9.56 ± 0.19	21.11 ± 3.56	10.33 ± 7.09	14.25 ± 3.44
SIGReg	+SADL	E1	61.78 ± 1.90	24.56 ± 2.52	82.67 ± 1.76	40.00 ± 1.45	$\mathbf{52.25 \pm 0.30}$
SIGReg	Endpoint-only	E5	63.00 ± 3.61	26.67 ± 2.96	64.22 ± 19.33	11.78 ± 11.21	41.42 ± 4.50
SIGReg	+SADL	E5	62.22 ± 2.14	24.56 ± 2.52	72.44 ± 7.71	40.00 ± 1.45	$\mathbf{49.81 \pm 2.14}$
SIGReg	Endpoint-only	E10	62.44 ± 7.32	26.33 ± 3.61	69.00 ± 24.29	17.22 ± 5.74	43.75 ± 6.94
VISReg	Endpoint-only	E1	35.56 ± 5.60	12.22 ± 2.59	30.00 ± 11.62	18.00 ± 6.66	23.94 ± 4.15
VISReg	+SADL	E1	54.11 ± 15.35	25.22 ± 6.36	81.67 ± 2.73	64.56 ± 15.91	$\mathbf{56.39 \pm 1.55}$
VISReg	Endpoint-only	E5	73.78 ± 4.91	20.78 ± 0.77	82.22 ± 0.19	23.11 ± 7.63	49.97 ± 1.63
VISReg	+SADL	E5	71.78 ± 1.17	28.44 ± 1.35	83.22 ± 2.22	46.33 ± 0.33	$\mathbf{57.44 \pm 1.04}$
VISReg	Endpoint-only	E10	76.11 ± 2.69	22.11 ± 0.84	68.89 ± 26.69	24.11 ± 1.50	47.81 ± 6.06
QQWorld	Endpoint-only	E1	21.11 ± 6.19	13.00 ± 3.18	26.33 ± 4.73	14.56 ± 8.44	18.75 ± 5.28
QQWorld	+SADL	E1	61.78 ± 11.48	30.33 ± 3.06	84.11 ± 2.46	47.89 ± 17.44	$\mathbf{56.03 \pm 1.74}$
QQWorld	Endpoint-only	E5	66.22 ± 6.27	29.22 ± 3.29	79.56 ± 6.58	19.56 ± 9.22	48.64 ± 5.98
QQWorld	+SADL	E5	69.00 ± 2.00	29.44 ± 1.68	85.11 ± 0.84	37.78 ± 0.84	$\mathbf{55.33 \pm 0.71}$
QQWorld	Endpoint-only	E10	72.22 ± 4.29	28.33 ± 2.40	74.11 ± 9.71	21.89 ± 3.95	49.14 ± 3.74

Table 23: Joint four-task INTACT CLEAR-LeWM v0.8 Strict success (%). E1 and E5 compare equal training budgets; endpoint-only E10 is a longer-budget reference. INTACT uses search-free Direct control. Endpoint-only uses a MATH encoder; +SADL uses a Flash encoder; both predictors use MATH. Values are mean ± sample SD over three training seeds.

Endpoint	Method	Epoch	PushT	Cube	Reacher	TwoRoom	Macro
SIGReg	Endpoint-only	E1	36.00 ± 7.23	43.33 ± 8.95	20.67 ± 7.54	24.33 ± 10.79	31.08 ± 2.04
SIGReg	+SADL	E1	77.56 ± 0.51	88.89 ± 1.17	98.00 ± 0.88	97.33 ± 0.67	$\mathbf{90.44 \pm 0.55}$
SIGReg	Endpoint-only	E5	74.78 ± 4.02	89.44 ± 1.07	98.22 ± 0.51	37.33 ± 20.80	74.94 ± 3.81
SIGReg	+SADL	E5	84.33 ± 1.20	96.67 ± 5.20	98.56 ± 0.38	97.44 ± 0.96	$\mathbf{94.25 \pm 1.42}$
SIGReg	Endpoint-only	E10	81.00 ± 2.19	88.67 ± 0.58	98.44 ± 0.19	59.33 ± 19.78	81.86 ± 4.32
VISReg	Endpoint-only	E1	42.56 ± 3.53	55.33 ± 12.55	54.78 ± 14.28	32.11 ± 14.05	46.19 ± 3.44
VISReg	+SADL	E1	74.89 ± 2.01	87.67 ± 1.20	97.89 ± 0.77	95.22 ± 1.02	$\mathbf{88.92 \pm 0.22}$
VISReg	Endpoint-only	E5	78.11 ± 1.64	86.56 ± 1.17	99.44 ± 0.38	52.56 ± 7.07	79.17 ± 1.83
VISReg	+SADL	E5	83.56 ± 0.38	89.56 ± 0.19	98.44 ± 0.69	96.22 ± 1.17	$\mathbf{91.94 \pm 0.29}$
VISReg	Endpoint-only	E10	81.78 ± 2.59	88.56 ± 0.51	98.56 ± 1.17	74.78 ± 6.06	85.92 ± 2.08
QQWorld	Endpoint-only	E1	40.56 ± 5.42	47.11 ± 6.11	33.89 ± 12.85	11.00 ± 4.10	33.14 ± 1.88
QQWorld	+SADL	E1	76.67 ± 2.03	88.22 ± 1.26	97.89 ± 0.69	96.11 ± 1.17	$\mathbf{89.72 \pm 0.96}$
QQWorld	Endpoint-only	E5	79.11 ± 2.91	89.56 ± 0.69	99.00 ± 0.33	59.22 ± 28.55	81.72 ± 7.39
QQWorld	+SADL	E5	83.56 ± 2.69	96.89 ± 5.10	98.44 ± 0.51	96.22 ± 0.19	$\mathbf{93.78 \pm 1.25}$
QQWorld	Endpoint-only	E10	81.33 ± 1.73	92.00 ± 0.58	98.56 ± 0.19	73.89 ± 27.35	86.44 ± 7.00

F Structural Ablations and Mechanism Diagnostics

F.1 Supply branches and bridge potentials

All structural ablations use Flash, SIGReg weight .03, bridge weight $\lambda_S = .01$, and $\alpha = .5$. Excitation-only and Transmission-only replace k by C_{exc} or C_{tr}, respectively. Potential variants compare $[u]_+^2$, u^2, and $[u]_+$ for $u = \log(q + \epsilon) - \log(k + \epsilon)$. Bridge-only still evaluates detached demand but removes the endpoint regularizer's direct training gradient.

Table 24: E2 supply-branch and endpoint-gradient ablations under SIGReg. Entries are four-task CLEAR-LeWM v0.8 macro success, mean ± sample SD (%) over three training seeds, with three evaluation seeds per task. All use Flash.

Configuration	INTACT Mod.	INTACT Strict	LeWM Mod.	LeWM Strict
Excitation-only	97.14 ± 0.17	91.03 ± 0.61	62.78 ± 2.00	50.94 ± 3.26
Transmission-only	42.94 ± 45.45	38.33 ± 41.64	35.28 ± 27.30	26.97 ± 20.63
Bridge-only	47.00 ± 5.59	35.33 ± 1.86	11.72 ± 1.80	7.81 ± 0.49
Full SADL	$\mathbf{97.14 \pm 0.17}$	$\mathbf{92.06 \pm 1.04}$	$\mathbf{65.50 \pm 0.85}$	$\mathbf{53.03 \pm 1.42}$

The excitation-only branch matches Full SADL's INTACT E2 Moderate macro success (97.14%) and trails it by 2.72 percentage points in LeWM (62.78% versus 65.50%). Transmission-only is markedly weaker and more seed-variable in both models. Thus excitation is the dominant observed bottleneck in these four-task control runs; this does not make transmission capacity dispensable. Bridge-only falls to 47.00% in INTACT and 11.72% in LeWM on Moderate, showing that the path correction complements rather than replaces the direct endpoint constraint. In the independently constructed transmission-deficit setting of Table 28(B), Full SADL raises C_{tr} from .180 to .746 and held-out R^2 from $-.018$ to .535, whereas excitation-only leaves both essentially unchanged.

Table 25: Complete A2 bridge-potential ablation under SIGReg on four-task CLEAR-LeWM v0.8. Within each model and epoch, the canonical one-sided square $[u]_+^2$, bilateral square u^2, and one-sided linear $[u]_+$ use matched training and evaluation protocols. Entries are macro success mean ± sample SD (%) over three training seeds, with three evaluation seeds per task. LeWM uses Pure-CEM; INTACT uses Direct control.

Potential	Epoch	Moderate	Strict
LeWM			
Canonical $[u]_+^2$	E1	$\mathbf{64.39 \pm 1.44}$	$\mathbf{52.25 \pm 0.30}$
Bilateral u^2	E1	59.58 ± 4.38	52.11 ± 2.66
Linear $[u]_+$	E1	47.17 ± 5.55	39.92 ± 6.45
Canonical $[u]_+^2$	E2	$\mathbf{65.50 \pm 0.85}$	$\mathbf{53.03 \pm 1.42}$
Bilateral u^2	E2	61.00 ± 4.86	47.44 ± 6.22
Linear $[u]_+$	E2	53.53 ± 4.19	42.03 ± 6.65
INTACT			
Canonical $[u]_+^2$	E1	95.72 ± 0.05	$\mathbf{90.44 \pm 0.55}$
Bilateral u^2	E1	95.61 ± 0.47	89.97 ± 1.80
Linear $[u]_+$	E1	$\mathbf{96.06 \pm 0.34}$	89.44 ± 0.59
Canonical $[u]_+^2$	E2	$\mathbf{97.14 \pm 0.17}$	$\mathbf{92.06 \pm 1.04}$
Bilateral u^2	E2	93.97 ± 5.76	88.03 ± 7.90
Linear $[u]_+$	E2	$\mathbf{97.14 \pm 0.43}$	90.86 ± 0.67

Table 25 isolates potential shape. The canonical and bilateral squares have identical values and gradients for $u > 0$; only the bilateral potential continues to penalize the already-sufficient regime $u < 0$, while the canonical bridge is zero there by construction. At equal E1/E2 budgets, canonical exceeds bilateral in LeWM Moderate macro success by $4.81/4.50$ percentage points; the E1 gain is

positive for all three paired training seeds. INTACT is nearly tied at E1 (95.72% versus 95.61%), but canonical leads by 3.17 points at E2 (97.14% versus 93.97%). The large bilateral E2 seed SD in INTACT means this last mean difference is a robustness signal, not a uniform gain across seeds. The one-sided linear form remains competitive in INTACT but loses $17.22/11.97$ points to canonical in LeWM Moderate success at E1/E2. Together, the matched A2 results support the practical value of a one-sided gate and the squared potential across the two control models.

In the fixed LeWM and INTACT probes, $C_{\mathrm{exc}} < C_{\mathrm{tr}}$, so Full SADL and excitation-only select the same instantaneous bottleneck: both bridge gradients repair C_{exc} at the same state. This equivalence concerns the observed state and bridge term, not necessarily the complete parameter trajectories of independently trained models. Persistent excitation growth explains the early expansion and convergence improvement. Transmission-only has high seed-level variance due to bimodal seed failures, not episode noise. Its E2 Moderate macro results by training seed are

$$\begin{aligned} &\text{INTACT:} \quad 95.42\%,\ 17.33\%,\ 16.08\%,\\ &\text{LeWM:} \quad 65.67\%,\ 12.83\%,\ 27.33\%. \end{aligned}$$

Excitation is the active bottleneck in this regime; an independent transmission-limited test follows in Section F.3.

F.2 Demand, supply, and gradients at fixed checkpoints

The audit contains 168 observations: two model families, seven configurations, three training seeds, and four tasks. Within each model–seed–task group, configurations use the same deterministic 64-sample probe; statistics are computed task-locally. Define deficit, active rate, and transmission-active rate by

$$\Delta = \left[\log \frac{q+\epsilon}{k+\epsilon}\right]_+, \quad p_{\mathrm{act}} = \Pr(k < q), \quad p_{\mathrm{tr,act}} = \Pr(C_{\mathrm{tr}} < C_{\mathrm{exc}},\ k < q).$$

The fixed-checkpoint diagnostics use E10/E5/E2; endpoint-only bridge gradients are computed post hoc as counterfactuals. Tables 26 and 27 report the observed states, whereas Table 24 reports the equal-budget E2 control ablation. All observed checkpoints have $C_{\mathrm{exc}} < C_{\mathrm{tr}}$; all 24 Full SADL snapshots have an active bridge.

Table 26: Demand, supply, and control success at fixed checkpoints. E is the audited epoch; capacity entries average 12 task–seed cells. $p_{\mathrm{act}} = \Pr(q > k)$. SR is independently evaluated four-task CLEAR-LeWM v0.8 Moderate macro success (%; mean ± sample SD over three training seeds, with three evaluation seeds per task) at the listed epoch. The rows have different training budgets and are not an equal-budget SR comparison. † denotes a counterfactual bridge diagnostic on endpoint-only checkpoints; its SR is the actual endpoint-only result. Exc./Trans./Bilateral/Linear denote excitation-only, transmission-only, bilateral squared, and one-sided linear variants.

Setting	E	q_R	q_E	q	C_{exc}	C_{tr}	p_{act} (%)	SR (%)
LeWM								
Endpoint†	10	.8049	.6541	.7174	.3177	.6948	100.0	56.61 ± 6.49
Bridge-only	2	.0000	.0781	.0000	.0439	.1858	0.0	11.72 ± 1.80
Trans.	2	.5416	.3150	.3886	.2578	.5062	41.7	35.28 ± 27.30
Bilateral	2	.9496	.5307	.7098	.5882	.6944	100.0	61.00 ± 4.86
Linear	2	.9479	.5439	.7179	.6640	.7287	100.0	53.53 ± 4.19
Exc.	5	.9517	.6127	.7632	.6426	.7510	100.0	63.72 ± 0.92
Full SADL	5	.9521	.6133	.7637	.6430	.7509	100.0	60.78 ± 3.63
INTACT								
Endpoint†	10	.7381	.6366	.6789	.2622	.6918	100.0	90.44 ± 3.14
Bridge-only	2	.0000	.0791	.0000	.1196	.3486	0.0	47.00 ± 5.59
Trans.	2	.2947	.2171	.2225	.2687	.5673	25.0	42.94 ± 45.45
Bilateral	2	.9343	.5392	.7096	.5697	.7246	100.0	93.97 ± 5.76
Linear	2	.9497	.5591	.7286	.6710	.7539	100.0	97.14 ± 0.43
Exc.	5	.9524	.6411	.7812	.6384	.7964	100.0	97.50 ± 0.22
Full SADL	5	.9529	.6362	.7784	.6331	.7918	100.0	98.03 ± 0.13

Table 27: Bridge deficit and encoder-parameter gradients on the same fixed probe. Epoch identities and independently evaluated CLEAR-LeWM v0.8 Moderate SR (%; mean ± sample SD over three training seeds) match Table 26; different rows have different training budgets. G_T, G_B are task and bridge gradient norms; G_B/G_T averages per-cell ratios. Bridge-only has zero bridge gradient, so its cosine is undefined (n/a). † denotes a counterfactual bridge diagnostic.

Configuration	Δ	G_T	G_B	G_B/G_T (%)	$\cos(G_T, G_B)$	SR (%)
LeWM						
Endpoint†	.8180	.25432	.00756	3.843	-.0251	56.61 ± 6.49
Bridge-only	.0000	.00092	.00000	.000	n/a	11.72 ± 1.80
Trans.	.0233	15.74698	.00023	.017	.0987	35.28 ± 27.30
Bilateral	.1886	.16440	.00196	1.763	.0815	61.00 ± 4.86
Linear	.0790	.14240	.00408	3.697	.1885	53.53 ± 4.19
Exc.	.1740	.05307	.00080	1.852	.0588	63.72 ± 0.92
Full SADL	.1739	.05342	.00081	1.946	.0592	60.78 ± 3.63
INTACT						
Endpoint†	.9730	.69619	.01579	2.393	.1254	90.44 ± 3.14
Bridge-only	.0000	1.10121	.00000	.000	n/a	47.00 ± 5.59
Trans.	.0053	314.61248	.00003	.003	.0966	42.94 ± 45.45
Bilateral	.2223	.97499	.00455	.948	.2988	93.97 ± 5.76
Linear	.0829	.24236	.00472	2.123	.1361	97.14 ± 0.43
Exc.	.2032	.12964	.00110	.916	.0648	97.50 ± 0.22
Full SADL	.2083	.15412	.00119	.887	.0416	98.03 ± 0.13

The average bridge-to-task encoder-gradient norm ratio is $1.946 \pm 1.205\%$ for LeWM and $0.887 \pm 0.310\%$ for INTACT. These are pre-update parameter gradients on the same probe, $G_T = \|\nabla_\theta \mathcal{L}_{\text{task}}\|_2$ and $G_B = \|\nabla_\theta(\lambda_S \mathcal{S})\|_2$; the INTACT task term also contains action negative log-likelihood. For Full SADL, task–bridge cosine is 0.059 ± 0.173 in LeWM and 0.042 ± 0.131 in INTACT; bridge–endpoint cosine is 0.338 ± 0.197 and 0.320 ± 0.164, respectively, with SD across 12 task–seed cells. A small mid-path correction thus changes path geometry without overwhelming the task gradient. All Bridge-only snapshots have $q_R = 0$ and zero ideal demand; smoothed demand stays below available supply despite nonzero endpoint excitation, so the bridge gradient is zero.

F.3 Independent roles of the two supply capacities

Table 28(A) compares two initial capacity deficits. In separate paired ten-seed, 200-step tests, panels B/C construct $C_{\text{tr}} < q < C_{\text{exc}}$ and $C_{\text{exc}} < q < C_{\text{tr}}$, respectively. Each runs No SADL, Excitation-only, Transmission-only, and Full SADL; the table displays the ineffective single branch and Full, since the unshown paired arms coincide seed by seed in the corresponding controlled bottleneck. The three setups respectively optimize H, encoder weights, and a shared bias with fixed informative weights. B measures held-out readout; C holds the information channel fixed and reports H RMS. Full uses $k = \min(C_{\text{exc}}, C_{\text{tr}})$; each single-branch variant uses one capacity alone.

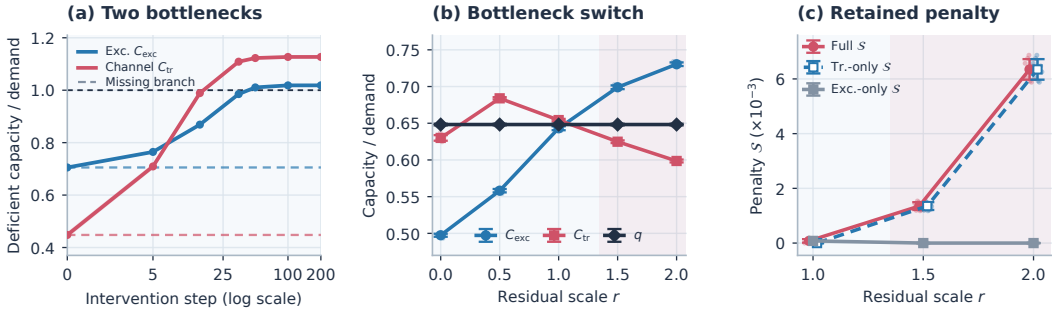

Figure 5: Independent capacity deficits and the transmission branch. (a) Matched-initialization controlled optimization of H: blue denotes excitation deficit, red transmission deficit; solid curves are Full SADL, dashed curves lack the needed branch. Curves show mean ± seed SD over 20 seeds; the horizontal dashed line marks capacity equal to demand. (b–c) Independent forward replay of an ImageNet-LT E100 checkpoint: endpoint and image mean are fixed while view-residual multiplier r varies. Pale dots show the same eight batches; filled dots and error bars show mean ± batch SD. The hollow blue squares in (c) are same-state Transmission-only counterfactual penalties, shifted horizontally only to separate overlapping error bars. At $r \geq 1.5$, $C_{\mathrm{tr}} < q < C_{\mathrm{exc}}$: Full and Transmission-only penalties coincide, whereas Excitation-only is zero. Panels (b–c) are forward-replay mechanism readouts, not retrained downstream success.

Table 28: Paired mechanism tests of two capacity deficits with fixed demand $q = .70$. (A) Direct optimization of H with 20 seeds. (B) Training encoder weights under a transmission deficit with held-out readout. (C) Training only a shared bias, with the informative weight channel fixed, under an excitation deficit. B/C use ten paired seeds each and display the ineffective single branch and Full SADL. Entries are initial → final mean ± seed SD.

(A) Controlled capacity optimization, 20 seeds

Initial deficit	Bridge target	C_{exc}, 0→200	C_{tr}, 0→200	$\mathcal{S}$, 0→200
Excitation	Transmission-only	.494 → .494 ± .002	.999 → .999 ± .000	.000 → .000
Excitation	Full SADL	.494 → **.713 ± .001**	.999 → .975 ± .001	.122 → .000
Transmission	Excitation-only	.976 → .976 ± .001	.314 → .314 ± .009	.000 → .000
Transmission	Full SADL	.976 → .978 ± .001	.314 → **.789 ± .003**	.646 → .000

(B) Transmission deficit, independent held-out readout, 10 seeds

Bridge target	C_{exc}, initial→final	C_{tr}, initial→final	Held-out R^2, initial→final
Excitation-only	.779 → .779 ± .003	.180 → .180 ± .005	−0.018 → −0.018±.002
Full SADL	.779 → .861 ± .003	.180 → **.746 ± .011**	−0.018 → .535 ± .018

(C) Excitation deficit, fixed weights/shared-bias training, 10 seeds

Bridge target	C_{exc}, initial→final	C_{tr}, initial→final	H RMS, initial→final
Transmission-only	.442 → .442 ± .002	.999 → .999 ± .000	2.234 → 2.234 ± .004
Full SADL	.442 → **.710 ± .002**	.999 → .999 ± .000	2.234 → 1.459 ± .003

Figure 5 juxtaposes the controlled deficits with a forward intervention on a real visual checkpoint. The former tests whether each capacity can independently become limiting; the latter tests whether the bridge still responds when transmission falls below demand. The same-state calculation shows why raising excitation alone cannot guarantee that the bridge has switched off: once transmission falls below demand, C_{tr} still sets the minimum capacity k.

F.4 DEMAND INTERPOLATION COEFFICIENT α

$\alpha = 0$ emphasizes external demand and is generally more aggressive; $\alpha = 1$ emphasizes realized endpoint demand and is more conservative. Table 29 gives realized demand and outcome for the same ImageNet-LT seed.

Table 29: Single-seed development ablation of demand interpolation α on ImageNet-LT (seed 42). q is the realized smoothed E20 demand, additionally capped by attainable supply in this image branch; online accuracy is at E100. $\alpha = 0$ favors the more aggressive external demand, whereas $\alpha = 1$ favors the conservative realized endpoint demand. In this setting, the highest observed result is at $\alpha = .5$.

Configuration	E20 demand q	E100 Top-1 (%)	vs. baseline (pp)
VISReg baseline	n/a	22.034	0.000
SADL $\alpha = 0$	.772	22.904	+0.870
SADL $\alpha = .25$	.772	22.348	+0.314
SADL $\alpha = .5$	.667	**23.026**	**+0.992**
SADL $\alpha = .75$	.554	22.144	+0.110
SADL $\alpha = 1$	.460	22.064	+0.030

F.5 DEMAND–SUPPLY GEOMETRY AND DUAL-ENDPOINT CONSTRAINTS

We compare a simple H-side barrier, direct dual-endpoint regularization, and SADL. For centered H covariance Σ_H, the barrier penalizes low average feature standard deviation or excessive spectral concentration:

$$\mathcal{R}_{\text{bar}}(H) = [.10 - \overline{\sigma}(H)]_+^2 + [\lambda_{\max}(\Sigma_H)/\operatorname{tr}(\Sigma_H) - .35]_+^2. \tag{18}$$

The tested barrier configuration retains SIGReg(Z) at weight .03; its H weight is .03 in E1 and .003 thereafter. It is therefore an H-side regularizer added to an endpoint objective, not an H-only training objective. The dual recipe instead applies the same distributional penalty $\mathcal{R}$ at H and Z, together with the tested scale floors F_γ:

$$\mathcal{L}_{\text{dual}} = \mathcal{L}_{\text{task}} + \lambda_R[\mathcal{R}(H) + \mathcal{R}(Z)] + 0.02[F_{0.5}(H) + F_{1.0}(Z)]. \tag{19}$$

Here $\mathcal{R}$ is SIGReg, VISReg, or QQWorld; F_γ is the fixed scale-floor term with threshold γ. The three tested λ_R values are .02, 1.5, and 3.5, respectively. This is a *dual-plus-floor* control, not a pure two-endpoint penalty. SADL instead retains the selected endpoint loss $\lambda_D\mathcal{D}(Z)$ and applies the one-sided supply loss $\lambda_S\mathcal{S}(H, \text{sg}(Z); \text{sg}(q))$ only for a measured demand–supply deficit (Section 3).

Table 30 reports four-task Official control success for three training seeds per recipe. LeWM uses Pure-CEM and INTACT uses Direct; each task score is averaged over three evaluation seeds. The late checkpoint is E2 for dual LeWM and E5 for SADL LeWM and both INTACT recipes; LeWM late scores therefore are not equal-budget comparisons. The barrier, dual-plus-floor, and SADL recipes were trained in separate experimental matrices, so the table compares observed recipes rather than estimating a paired causal effect of replacing one loss term.

Table 30: Simple H-side barrier, direct dual-endpoint regularization with scale floors, and SADL in jointly trained four-task control. Official success rates are percentages; macro entries are mean $\pm$ sample SD over three training seeds. Superscripts in the late column give the checkpoint epoch.

Model	Endpoint	Recipe	E1 PushT	E1 Cube	E1 Reacher	E1 TwoRoom	E1 Macro	Late Macro
LeWM	SIGReg	Dual+floor	7.56	61.00	80.67	74.67	$55.97 \pm .17$	$58.06 \pm .89^{\mathrm{E2}}$
		SADL	28.78	70.22	79.89	73.00	$\mathbf{62.97 \pm 1.25}$	$70.64 \pm .77^{\mathrm{E5}}$
	VISReg	Dual+floor	8.89	65.33	77.67	81.44	58.33 ± 1.09	$59.56 \pm .35^{\mathrm{E2}}$
		SADL	53.56	66.44	78.44	88.11	$\mathbf{71.64 \pm 2.96}$	$78.58 \pm .98^{\mathrm{E5}}$
	QQWorld	Dual+floor	7.22	64.44	78.89	71.67	55.56 ± 1.65	$56.31 \pm 2.12^{\mathrm{E2}}$
		SADL	45.67	67.44	78.00	82.56	$\mathbf{68.42 \pm 4.19}$	$75.11 \pm 1.34^{\mathrm{E5}}$
INTACT	SIGReg	H barrier	26.56	70.78	11.22	48.44	39.25 ± 12.74	$48.00 \pm 41.68^{\mathrm{E5}}$
INTACT	SIGReg	Dual+floor	42.11	99.11	94.78	98.00	83.50 ± 1.99	$95.81 \pm .43^{\mathrm{E5}}$
		SADL	85.22	100.00	97.11	100.00	$\mathbf{95.58 \pm .17}$	$97.19 \pm .51^{\mathrm{E5}}$
	VISReg	Dual+floor	61.11	99.44	96.67	97.56	88.69 ± 3.58	$96.67 \pm .36^{\mathrm{E5}}$
		SADL	83.78	100.00	97.78	100.00	$\mathbf{95.39 \pm .65}$	$97.53 \pm .25^{\mathrm{E5}}$
	QQWorld	Dual+floor	50.00	99.67	95.89	96.45	85.50 ± 1.38	$95.28 \pm .46^{\mathrm{E5}}$
		SADL	86.00	100.00	96.56	100.00	$\mathbf{95.64 \pm .32}$	$97.56 \pm .17^{\mathrm{E5}}$

The simple barrier is not uniformly ineffective: its three E5 training-seed macros are 28.67, 19.50, and 95.83%, but the aggregate remains $48.00 \pm 41.68\%$. At E2 its macro is $62.00 \pm 28.10\%$, compared with $96.56 \pm .17\%$ for three-seed SIGReg+SADL in the separate E2 mechanism audit (Table 3(b)). These are independent runs, not a paired loss-only ablation. Dual-endpoint regularization is also a viable control: at E5 it attains 95.28–96.67% INTACT Direct success. However, SADL is higher for all six model–endpoint pairs at E1, with the largest early gap on INTACT–SIGReg (83.50% versus 95.58%). The gain is most visible on PushT. At the matched INTACT E5 checkpoint, SADL remains higher by .86–2.28 percentage points; the margin has narrowed, but is not restricted to E1. These observations are consistent with adapting a path correction to measured demand and available supply rather than enforcing a fixed geometry at both ends, but the separately trained recipes do not isolate that mechanism on their own. Fixed-probe effective-rank, cosine, and prediction-loss evidence for SADL is in Table 2; the dual-plus-floor runs do not have an audited matched fixed-probe readout.

G Endpoint Distributions and Representation Geometry

G.1 Eight endpoint objectives and designable two-dimensional geometry

Let $Z \in \mathbb{R}^{T\times B\times d}$ and let v_j be random unit projection directions. Define $Q_\nu(Z)$ by sorting each projected batch against the quantiles of a reference distribution ν:

$$Q_\nu(Z) = \frac{1}{TBK}\sum_{t=1}^{T}\sum_{j=1}^{K}\sum_{i=1}^{B}\left((Z_t v_j)_{(i)} - F_\nu^{-1}\left(\frac{i-1/2}{B}\right)\right)^2.$$

For independent reference samples $y, y' \sim \nu$, the empirical energy distance is

$$E_\nu(Z) = \frac{1}{\sqrt{d}}\left[2\mathbb{E}\|z-y\|_2 - \mathbb{E}\|z-z'\|_2 - \mathbb{E}\|y-y'\|_2\right]_+.$$

At most 96 samples per temporal position are used to compute the energy distance.

Four Gaussian objectives. Write $\nu_G = \mathcal{N}(0, I)$, let $\widehat{\varphi}_Z$ be the empirical characteristic function, and let μ_j, σ_j denote coordinate-wise means and standard deviations. Set $\widetilde{Z} = (Z-\mu)/\operatorname{sg}(\sigma)$. The four objectives are

$$\mathcal{D}_{\mathrm{SIG}}(Z) \propto \left\langle |\widehat{\varphi}_Z(\omega) - e^{-\|\omega\|^2/2}|^2\right\rangle_\omega, \qquad \mathcal{D}_{\mathrm{QQ}}(Z) = Q_{\nu_G}(Z),$$
$$\mathcal{D}_{\mathrm{VIS}}(Z) = \left\langle(\sigma_j-1)^2 + \mu_j^2\right\rangle_j + Q_{\nu_G}(\widetilde{Z}), \qquad \mathcal{D}_{\mathrm{EG}}(Z) = E_{\nu_G}(Z).$$

All share a Gaussian reference, but respectively constrain the characteristic function, separated scale and shape, raw sorted projections, and sample distances; their gradients and finite-sample behavior differ.

Four non-Gaussian or partially specified objectives. Let C be the sample covariance of the TB endpoint samples and $u = z/\|z\|_2$. VIC moments specifies no full probability density:

$$\mathcal{D}_{\mathrm{VIC}}(Z) = \left\langle\left[1 - \sqrt{\mathrm{Var}(Z_j) + 10^{-4}}\right]_+\right\rangle_j + \left\langle C_{ij}^2\right\rangle_{i\neq j} + \left\langle\mu_j^2\right\rangle_j.$$

Sphere, Student-t_5, and GMM-Ring use, respectively,

$$\mathcal{D}_{\mathrm{Sphere}}(Z) = \left\langle\left(\|z\|_2/\sqrt{d} - 1\right)^2\right\rangle_z + \left\langle e^{-2\|u_i-u_j\|_2^2}\right\rangle_{t,i\neq j},$$
$$\mathcal{D}_{t_5}(Z) = Q_{\nu_{t_5}}(Z), \qquad F^{-1}_{\nu_{t_5}}(p) = \sqrt{3/5}\,F^{-1}_{t_5}(p),$$
$$\mathcal{D}_{\mathrm{Ring}}(Z) = E_{\nu_{\mathrm{Ring}}}(Z), \qquad \nu_{\mathrm{Ring}} = W_\#\left[\frac{1}{16}\sum_{\ell=0}^{15}\mathcal{N}\big(2.5m_\ell,\, 0.35^2 I\big)\right],$$
$$m_\ell = \cos\frac{2\pi\ell}{16}u_1 + \sin\frac{2\pi\ell}{16}u_2.$$

Here $W_\#$ analytically whitens the two-dimensional ring subspace and u_1, u_2 form an orthonormal basis. The two Sphere terms are added directly; Student-t_5 uses 128 projection directions. At $Z = 0$, Sphere sets $u = 0$ through its positive norm floor, and VISReg sets $\widetilde{Z} = 0$ through its positive standard-deviation floor; the cached $D_0 = \mathcal{D}(0)$ is evaluated with the same conventions.

G.2 Learned endpoint geometry and designable glyphs

The eight-objective control results are reported in Table 4 in the main text. For SIGReg and QQ-World, the first-step encoder-gradient calibration was performed separately for LeWM and INTACT, with VISReg anchored at $\lambda_D = 1.5$; the five additional objectives were likewise calibrated at the first step of seed 0. Each model-specific endpoint weight was then held fixed during training.

Figure 6 displays the corresponding E1 representations on held-out clips. Its first two rows show learned endpoint geometries under eight objectives; the third row is a separate demonstration of

designable two-dimensional glyph endpoints. The learned representations have no visually obvious collapse in these visualizations, although neither a low-dimensional projection nor control SR alone establishes a zero incidence of strict collapse.

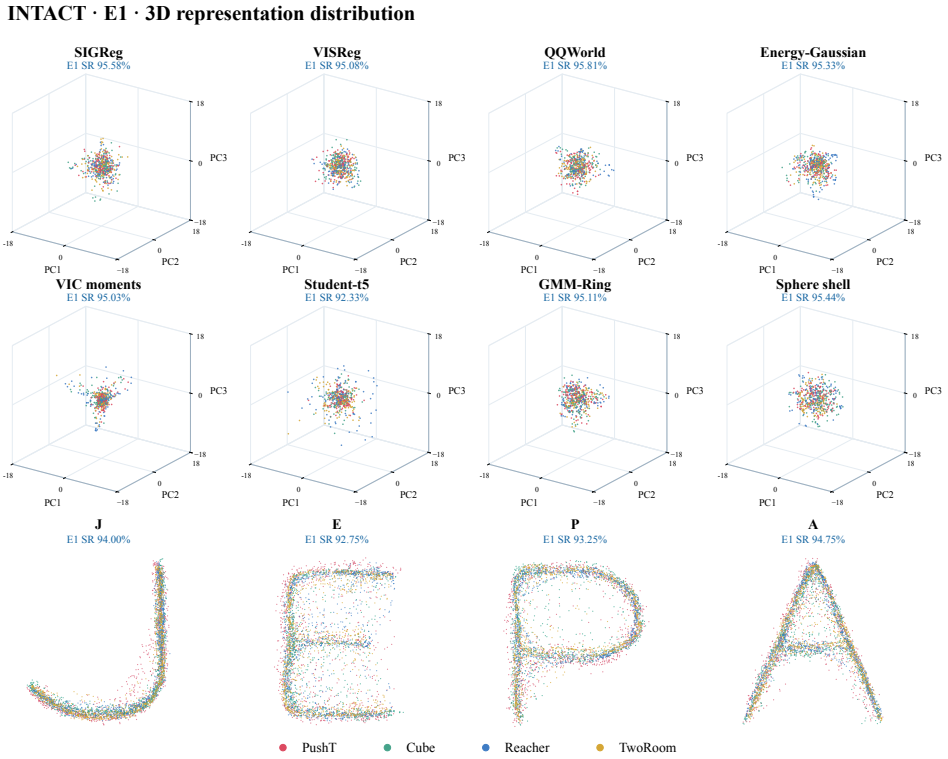

Figure 6: Real E1 representations with four-task color coding; all particles are vector objects. The first two rows show independent 3-D PCA projections of INTACT+SADL z on a fixed probe; blue SR values are three-seed Official Direct means. The third row shows independently trained raw 2-D glyph representations; blue SR values are single-seed Direct macro success, evaluated with 100 episodes per task, and are therefore not directly comparable to those above. The image illustrates learned local geometry, not the original high-dimensional density.

G.3 Eight-endpoint latent-geometry diagnostic

The diagnostic uses E1 INTACT checkpoints from three training seeds (0, 42, 3072); Official success is the corresponding three-seed aggregate. For each task, 128 segments with eight positions yield 4,096 pooled observations. A reference subset of 1,024 fits whitening, and the remaining 3,072 observations are analyzed. With nonzero covariance spectrum, we adapt the entropy-based effective rank of Roy & Vetterli (2007) to covariance eigenvalues:

$$r_{\text{eff}}(U) = \exp\left(-\sum_i p_i \log p_i\right), \qquad p_i = \frac{\lambda_i(\Sigma_U)}{\operatorname{tr}(\Sigma_U)},$$

where $0 \log 0 = 0$; raw effective rank is computed before whitening. We also report leading-eigenvalue share and average absolute excess kurtosis across random standardized projections:

$$\eta_1 = \frac{\lambda_{\max}(\Sigma_Z)}{\operatorname{tr}(\Sigma_Z)}, \qquad \bar{\kappa} = \langle |\operatorname{kurt}(Zv)| \rangle_v.$$

Each checkpoint fits its own whitening transform, while observation indices, 32 random projection directions, and the split are shared. Four-task observations are pooled within each seed before

reporting mean and sample standard deviation across seeds. Training uses within-task covariance, whereas pooled covariance obeys

$$\mathrm{Cov}(Z) = \mathbb{E}_T[\mathrm{Cov}(Z \mid T)] + \mathrm{Cov}_T(\mathbb{E}[Z \mid T]).$$

Function readout uses 32 smooth random Fourier functions $f_j(x) = \sin(a_j \omega_j^\top x + b_j)$ in each checkpoint's whitened coordinates. Directions, frequencies, and phases are shared, while x changes with the representation; these are isomorphic coordinate-defined function classes, not a shared physical-state target. For 128 fixed queries, $N = 128$ or 1,024 training samples, and eight independent resamples, kNN with $n_{\mathrm{nbr}} = 15$ and fixed-bandwidth RBF predictions are averaged before integrated squared bias is computed. Table 31 gathers the eight-objective E1 geometry, projected Fisher proxy, kNN bias, and control outcomes.

G.4 Asymptotic endpoint optimality and finite-training attainability

LeJEPA analyzes linear and nonlinear probe risk under fixed covariance scale, smooth full-dimensional densities, and unknown downstream tasks. Its Fisher-information argument selects an *isotropic Gaussian* endpoint, not an arbitrary anisotropic Gaussian. SADL does not dispute this endpoint choice; it asks whether the encoder–projector path can produce and transmit enough sample variation to attain it within a training budget T. Conceptually, the evaluation levels can be separated as

$$R_T(p; \mathcal{F}) \;\approx\; R_{\text{asymptotic probe}}(p) + \Delta_{\text{path/optimization},T}(p) + \Delta_{\text{task}}(p; \mathcal{F}), \qquad p \in \mathcal{P}_T.$$

Here $\mathcal{P}_T$ is the set of endpoints attainable in T steps and $\mathcal{F}$ is the task family. The expression distinguishes asymptotic probing, finite-training path effects, and task fit, with interactions assigned to the final term. Consequently, asymptotic, task-unknown Gaussian optimality does not automatically imply system-level optimality for a restricted path, finite budget, and specific control tasks.

We first calibrate synthetic distributions at matched total covariance, observation indices, smooth function family, and kNN-15 readout, then examine learned E1 endpoints. For synthetic Student-t_5, the full-dimensional Fisher ratio has the closed-form benchmark

$$\frac{J(t_\nu)}{J(\mathcal{N}(0, I_d))} = \frac{\nu(\nu + d)}{(\nu - 2)(\nu + d + 2)} = 1.650 \quad (d = 192,\ \nu = 5).$$

Synthetic full Fisher ratios are 1.000 for Gaussian, 1.650 for Student-t_5, and 2.283 for a two-scale full-dimensional Gaussian mixture, consistent with LeJEPA's theorem. Sphere lies on a singular lower-dimensional manifold outside its full-density assumptions.

The eight-objective E1 three-seed experiment reveals a distinct finite-time perspective (Table 31): the mean projected Fisher proxy is $1.220 \pm .083$ for GMM-Ring versus $1.287 \pm .053$ for SIGReg; finite-sample b_{1024} is $.1590 \pm .0025$ for Sphere versus $.1762 \pm .0049$ for SIGReg, and is lower in all three paired seeds. Together with Table 4, this shows that SADL's anti-collapse behavior and high control success do not depend on a Gaussian endpoint. LeJEPA's asymptotic Fisher conclusion remains applicable under its stated conditions. All geometry entries are E1 measurements; E5 appears only for control success.

Table 31: Three-training-seed E1 geometry, finite-sample readout, and Official Direct control across eight endpoints. All models are jointly trained four-task INTACT+SADL; each seed pools 4,096 positions from the same fixed four-task observations. Geometry and bias are mean ± sample seed SD; E1/E5 SR also aggregates three training seeds. The two panels use identical endpoint order. $10^3\mathrm{SW}^2$ denotes squared sliced-Wasserstein distance to a Gaussian reference, multiplied by 10^3. Geometry and readout are measured at E1; the E5 column contains only control SR from E5 checkpoints.

A. Path and endpoint geometry: mean ± seed SD

Endpoint	$r_{\mathrm{eff}}(h)$	$r_{\mathrm{eff}}(z)$	η_1 (%)	$\bar{\kappa}$	$\widehat{J}_{\mathrm{proj}}$
SIGReg	118.85 ± 1.19	58.68 ± 1.77	3.45 ± 0.22	$.573 \pm .103$	$1.287 \pm .053$
VISReg	122.11 ± 1.09	60.44 ± 1.26	3.34 ± 0.06	$.473 \pm .074$	$1.250 \pm .077$
QQWorld	120.25 ± 0.90	58.82 ± 0.76	3.47 ± 0.11	$.604 \pm .063$	$1.213 \pm .004$
Energy-Gaussian	112.86 ± 1.69	48.10 ± 0.56	4.26 ± 0.08	$.677 \pm .098$	$1.259 \pm .077$
VIC moments	106.17 ± 1.53	66.42 ± 0.35	3.01 ± 0.22	$1.035 \pm .158$	$1.330 \pm .066$
Sphere	102.61 ± 1.29	42.30 ± 0.60	4.96 ± 0.15	$.920 \pm .062$	$1.251 \pm .081$
Student-t_5	89.91 ± 3.03	35.11 ± 1.27	6.38 ± 0.47	$2.840 \pm .319$	$1.660 \pm .135$
GMM-Ring	112.46 ± 0.34	48.29 ± 0.57	4.20 ± 0.10	$.657 \pm .066$	$1.220 \pm .083$

B. Distribution distance, function readout, and control: mean ± seed SD

Endpoint	$10^3\mathrm{SW}^2$	b_{128}	b_{1024}	E1 SR (%)	E5 SR (%)
SIGReg	5.19 ± 0.97	$.2558 \pm .0056$	$.1762 \pm .0049$	$95.58 \pm .00$	$97.19 \pm .51$
VISReg	4.02 ± 0.46	$.2505 \pm .0049$	$.1716 \pm .0044$	$95.08 \pm .17$	$97.53 \pm .25$
QQWorld	5.37 ± 0.98	$.2552 \pm .0045$	$.1745 \pm .0035$	$95.81 \pm .13$	$97.56 \pm .17$
Energy-Gaussian	5.88 ± 0.58	$.2483 \pm .0020$	$.1683 \pm .0012$	$95.33 \pm .63$	$97.39 \pm .42$
VIC moments	9.43 ± 1.58	$.2836 \pm .0182$	$.2287 \pm .0167$	$95.03 \pm .99$	$96.97 \pm .19$
Sphere	8.35 ± 0.77	$.2399 \pm .0085$	$.1590 \pm .0025$	$95.44 \pm .82$	$97.44 \pm .10$
Student-t_5	29.76 ± 3.33	$.2804 \pm .0052$	$.2052 \pm .0043$	92.33 ± 1.04	$96.56 \pm .55$
GMM-Ring	5.57 ± 0.32	$.2403 \pm .0061$	$.1661 \pm .0037$	$95.11 \pm .35$	$97.81 \pm .38$

Table 32: LeWM+SADL E1 path and endpoint diagnostics underlying the upper half of Table 4. Each training seed pools 4,096 positions from the same fixed four-task observations as the INTACT probe in Table 31. Entries are mean ± sample SD across three training seeds. $\widehat{J}_{\mathrm{proj}}$ is a projected Fisher proxy, not the full-dimensional asymptotic functional; b_{1024} is finite-sample kNN integrated squared bias after per-checkpoint whitening.

Endpoint	$r_{\mathrm{eff}}(h)$	$r_{\mathrm{eff}}(z)$	$\widehat{J}_{\mathrm{proj}}$	b_{1024}
SIGReg	116.71 ± 2.84	48.84 ± 2.23	$1.289 \pm .049$	$.1772 \pm .0075$
VISReg	123.62 ± 2.15	55.24 ± 2.65	$1.184 \pm .020$	$.1717 \pm .0027$
QQWorld	114.11 ± 9.39	48.18 ± 3.52	$1.179 \pm .031$	$.1699 \pm .0020$
Energy-Gaussian	103.69 ± 10.34	40.40 ± 2.31	$1.161 \pm .037$	$.1712 \pm .0019$
VIC moments	105.05 ± 2.75	64.58 ± 2.00	$1.308 \pm .117$	$.2228 \pm .0157$
Sphere	95.48 ± 2.14	$33.57 \pm .93$	$1.189 \pm .025$	$.1660 \pm .0022$
Student-t_5	84.26 ± 1.06	30.48 ± 1.29	$1.591 \pm .093$	$.2034 \pm .0184$
GMM-Ring	106.73 ± 2.29	39.83 ± 1.86	$1.232 \pm .063$	$.1756 \pm .0042$

Table 33: LeWM E1 endpoint-only versus +SADL under the same Official Pure-CEM protocol. Both configurations use a Flash encoder, MATH predictor, shared four-task training, and the same endpoint coefficient for each objective. Pure-CEM uses 300 samples, 30 iterations, $H = 5$, and $RH = 5$ without actor warm-start. Task columns show endpoint-only success rates (%); macro columns are mean $\pm$ sample SD across three training seeds, each averaged over three evaluation seeds of 100 episodes per task. Δ is the difference of the two macro means in percentage points.

Endpoint	PushT$_0$	Cube$_0$	Reacher$_0$	TwoRoom$_0$	Macro$_0$	Macro$_{\mathrm{SADL}}$	Δ pp
SIGReg	3.44	49.89	20.00	36.44	27.44 ± 3.44	71.56 ± 1.27	$+44.12$
VISReg	53.00	55.11	27.33	55.33	47.69 ± 5.14	71.58 ± 4.67	$+23.89$
QQWorld	19.22	52.56	18.78	40.33	32.72 ± 11.23	74.36 ± 2.90	$+41.64$
Energy-Gaussian	3.78	45.11	14.78	36.89	25.14 ± 0.67	70.25 ± 2.35	$+45.11$
VIC moments	5.67	47.44	15.56	29.56	24.56 ± 4.21	67.53 ± 6.09	$+42.97$
Sphere	34.33	54.33	23.22	47.11	39.75 ± 3.54	75.33 ± 1.36	$+35.58$
Student-t_5	3.11	45.78	13.44	36.00	24.58 ± 1.78	62.81 ± 3.99	$+38.23$
GMM-Ring	3.00	43.44	12.78	36.00	23.81 ± 3.71	70.75 ± 1.23	$+46.94$

H IMAGE CLASSIFICATION EVALUATION

We test whether the same path-capacity correction helps image representations across long-tailed natural images (ImageNet-LT), galaxies (Galaxy10), small natural images (CIFAR-10), and the Imagenette subset of ImageNet. Each comparison pairs the endpoint-only baseline with the same model plus SADL, holding the data, architecture, optimization, and evaluation protocol fixed. SADL estimates excitation and realized demand from the mean representation of each image's augmented views (the *image-mean* variant) in all four datasets. Table 34 reports fixed-epoch classification accuracy: an online head for ImageNet-LT, Galaxy10, and Imagenette, and a held-out ridge readout for CIFAR-10. Each mean and standard deviation is over independent training seeds; the last column counts seeds improved by SADL.

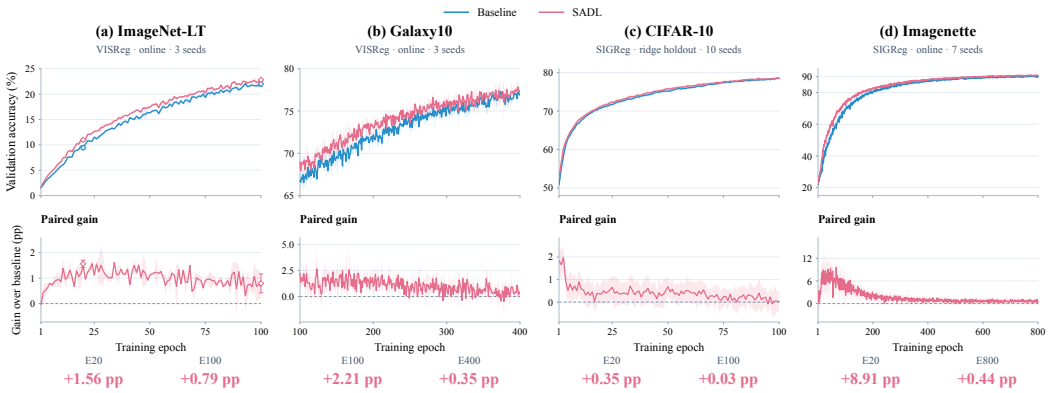

Figure 7: Training trajectories and paired accuracy gains across four image-classification datasets. Top: classification accuracy; bottom: SADL minus the paired baseline, in percentage points. Lines and bands show means and standard deviations over paired training seeds (3, 3, 10, and 7, respectively). ImageNet-LT, Galaxy10, and Imagenette use online heads; CIFAR-10 uses a held-out ridge readout.

Table 34: Image-classification accuracy (%). Baseline and SADL use matched training seeds and the same endpoint objective. Δ is the mean paired gain in percentage points; "wins" counts positive seed-wise differences.

Dataset (n)	Endpoint	Epoch	Baseline	+SADL	Δ	Wins
ImageNet-LT (3)	VISReg	100	22.02 ± 0.41	$\mathbf{22.81 \pm 0.20}$	$+0.794$	3/3
Galaxy10 (3)	VISReg	400	76.89 ± 0.45	$\mathbf{77.23 \pm 0.53}$	$+0.348$	3/3
CIFAR-10 (10)	SIGReg	1	50.81 ± 0.73	$\mathbf{52.64 \pm 0.80}$	$+1.830$	10/10
		10	66.87 ± 0.40	$\mathbf{67.76 \pm 0.67}$	$+0.894$	9/10
		50	75.16 ± 0.36	$\mathbf{75.84 \pm 0.36}$	$+0.680$	9/10
		100	78.50 ± 0.30	78.52 ± 0.32	$+0.026$	4/10
Imagenette (7)	SIGReg	5	24.23 ± 1.49	$\mathbf{27.68 \pm 1.43}$	$+3.450$	7/7
		20	35.60 ± 4.48	$\mathbf{44.51 \pm 3.13}$	$+8.906$	7/7
		100	69.34 ± 1.99	$\mathbf{73.15 \pm 1.89}$	$+3.807$	7/7
		800	90.25 ± 0.17	$\mathbf{90.69 \pm 0.31}$	$+0.444$	6/7

ImageNet-LT and Galaxy10. On ImageNet-LT, ViT-S/8 with VISReg is trained for 100 epochs using a global batch of 256, four 128×128 global crops, and 4,096 projection directions. At E100, SADL improves the matched three-seed online accuracy by 0.794 points. Galaxy10 uses the same backbone and endpoint family; its three-seed gain is 2.207 points at E100 and remains positive at E400 (Table 34). The fixed E400 comparison uses each arm's own E100 model and optimizer state.

CIFAR-10 and Imagenette. CIFAR-10 pairs a ResNet-18 SIGReg baseline with image-mean SADL and the same held-out ridge readout over ten seeds. The E1 gain is 1.830 points, and nine of ten seeds still improve at E50; by E100 the mean accuracies are nearly equal. The Imagenette baseline follows the LeJEPA ViT-S/8 SIGReg recipe (Balestriero & LeCun, 2025), with four global views and a global batch of 256. Seven paired seeds show the largest gain at E20 ($+8.906$ points), and the gain remains positive at E800 ($+0.444$ points). These results show the clearest and most consistent advantage in the early training regime, with improvements also present at the fixed final checkpoint on ImageNet-LT, Galaxy10, and Imagenette.

Image-specific demand. With n image means and encoder width d_h, the covariance rank is at most $r = \min(n-1, d_h)$. We cap demand at

$$q_{\max} = \frac{r \log(1 + \gamma d_h / r)}{d_h \log(1+\gamma)}, \qquad q_{\text{img}} = \min\{q_{\max}, \widetilde{q}_\alpha\}.$$

As in the general demand interpolation, the smoothed image demand is

$$\widetilde{q}_\alpha = \begin{cases} q_R, & \alpha = 0, \\ (q_R + \epsilon)^{1-\alpha}(q_E + \epsilon)^\alpha, & 0 < \alpha < 1, \\ q_E, & \alpha = 1, \end{cases} \qquad \mathcal{S}^{\text{img}} = \left[\log(q_{\text{img}} + \epsilon) - \log(k + \epsilon)\right]_+^2,$$

where $\epsilon = 10^{-6}$ and the demand is detached. For ImageNet-LT, $n = 256$, $d_h = 384$, and $\gamma = 9$, giving $q_{\max} \simeq 0.7723$. In the matched five-point interpolation sweep, $\alpha = 0.5$ gives the highest E100 online accuracy.

Which capacity limits the image path? Unlike the measured LeWM and INTACT fixed probes, which are excitation-limited ($C_{\text{exc}} < C_{\text{tr}}$), the separate ImageNet-LT E100 view-residual replay can become transmission-limited ($C_{\text{tr}} < q < C_{\text{exc}}$; Figure 5). This bottleneck switch motivates retaining both capacities in k; the classification comparisons above use the image-mean protocol.

I Computational Cost and Inference Workload

I.1 Short-Run SDPA Hardware Matrix

We benchmark the incremental training cost of SADL across MATH and Flash attention in the encoder and predictor. Here M and F denote MATH and Flash SDPA, respectively; B denotes endpoint-only LeWM and S denotes the same model with full SADL. Both arms use the same SIGReg endpoint weight 0.03, SADL bridge weight 0.01 where applicable, and $\alpha = 0.5$. The fixed protocol uses BF16, gradient checkpointing, batch size 256 per task, 10 data-loader workers, and 40 optimizer steps. PushT runs on one GPU; four-task training uses four GPUs, one task per rank. The reported step time is $(t_{30} - t_{10})/20$, including data loading, optimization, and synchronization; peak device-wide memory is sampled by `nvidia-smi` every 500 ms. Each non-dagger cell averages three independent repetitions.

Table 35: Training runtime and peak memory on RTX 4090, H200, and H20. Each B/S pair gives endpoint-only/SADL means; $\Delta_t = 100(t_S/t_B - 1)$, where positive means measured SADL time is higher. Times are ms/step and memory is GiB. M/M for H200 and H20 uses earlier same-protocol measurements; H20 M/M† pools four hosts and is not host-matched to the other H20 rows. Three-repeat SDs and telemetry are archived.

		Four tasks, four GPUs			*PushT, one GPU*		
Hardware	Enc./pred.	Step B/S	Δ_t (%)	Peak B/S	Step B/S	Δ_t (%)	Peak B/S
RTX 4090	M/M	2528/2442	-3.4	23.38/23.28	2352/2326	-1.1	22.37/22.38
	M/F	2428/2422	-0.3	23.51/23.51	2283/2300	$+0.7$	22.24/22.24
	F/M	1996/1856	-7.0	17.40/17.42	1694/1669	-1.5	16.10/16.10
	F/F	1861/1865	$+0.2$	17.21/17.18	1631/1671	$+2.5$	15.85/15.87
H200	M/M	1195/1131	-5.3	25.11/25.26	1148/1112	-3.1	22.48/22.87
	M/F	1065/1018	-4.4	25.18/24.94	1070/1023	-4.4	22.79/22.55
	F/M	654/569	-13.0	19.97/19.46	622/637	$+2.4$	17.12/17.14
	F/F	579/516	-11.0	19.18/19.13	577/616	$+6.6$	16.87/16.90
H20	M/M†	1355/1354	-0.1	24.24/24.18	1398/1353	-3.3	22.26/22.26
	M/F	1323/1335	$+0.9$	24.77/24.49	1398/1323	-5.4	22.40/22.38
	F/M	1077/921	-14.4	19.28/19.03	1133/918	-19.0	16.92/16.93
	F/F	1078/871	-19.2	19.09/18.76	1063/1053	-0.9	16.69/16.70

Across the 24 within-backend B/S comparisons in Table 35, the median measured SADL step-time change is -3.2%; six comparisons are positive, with a maximum of $+6.6\%$. These short runs therefore show no systematic added per-step latency from the bridge, rather than evidence that the bridge itself accelerates an optimizer step. Flash in the encoder accounts for most of the hardware benefit: in four-task endpoint-only runs, M/M to F/M reduces step time from 2528 to 1996 ms and peak memory from 23.38 to 17.40 GiB on RTX 4090, and from 1195 to 654 ms and 25.11 to 19.97 GiB on H200. Enabling Flash in both modules further reaches 1861 ms/17.21 GiB and 579 ms/19.18 GiB, respectively. The corresponding H20 trend is consistent but uses the historical M/M reference from different hosts. The RTX 4090 M/M third repeat completed all 40 steps and wrote its measurements, despite a post-training distributed shutdown error documented in the archive.

I.2 Search Workload of the Control Interfaces

INTACT's Direct interface performs no candidate-action search. The Official records for the alternative endpoint experiments contain zero candidate-action evaluations and zero planning-cost calls. LeWM Pure-CEM uses 300 candidates per iteration and 30 iterations, corresponding to 9,000 candidate evaluations per planning invocation.

SADL adds a training-only objective and retains these existing control interfaces. It introduces no additional inference network or candidate-search procedure.