BEYOND STATE SUFFICIENCY: THE DIMENSION OF GAIN-BOUNDED PREDICTIVE REALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Once a representation fully specifies the state, can regular prediction still require additional dimensions? We show that it can for exact deterministic transitions under a prescribed bound on distance expansion. All actions must satisfy this requirement on one shared geometry. We derive verifiable dimension lower bounds, including bounds that allow small expansions. In grid experiments, every state remains identifiable after dimension reduction. Yet for the walled systems, no alternative state-preserving encoding in the reduced dimension can meet the prescribed regularity bound. Thus preserving state information alone cannot certify that a dimension is sufficient for regular prediction.

1 INTRODUCTION

Latent world models encode observations and predict the represented state's evolution under actions (Ha & Schmidhuber, 2018; Watter et al., 2015; Hafner et al., 2019; 2020). State abstraction and representation learning ask which distinctions must be retained for prediction or control (Givan et al., 2003; Ferns et al., 2004; Gelada et al., 2019; Zhang et al., 2021; Wang et al., 2022; Huang et al., 2022). Predictive-state and compact-representation approaches address related questions through predictive rank or coordinate redundancy (Littman & Sutton, 2001; Thon & Jaeger, 2015; Kobayashi & Watanuki, 2022). These perspectives help determine what information a representation must preserve.

But redundancy is relative to use. Coordinates that add no state information may be dispensable for identification; are they also dispensable for a predictor subject to regularity requirements (Asadi et al., 2018; Gelada et al., 2019; Song et al., 2024)? Complete state recovery does not itself check whether those requirements remain achievable after dimension reduction. We ask: once a representation fully specifies the state, can regular prediction still require additional Euclidean dimensions—and if so, how many?

We show that it can for exact deterministic transitions under a prescribed Lipschitz gain bound. This bound limits distance expansion under each action, and every action must satisfy it on *one shared Euclidean geometry*. Our certificates exclude every injective encoding in certain dimensions, not just a particular arrangement. We define the *predictive realization dimension* as the minimum ambient dimension admitting such a realization. This minimum ranges over all arrangements, not training outcomes. Ambient dimension can therefore serve as a *realization resource* without supplying new state information. Figure 1 shows two systems whose states admit the same scalar description but require different realization dimensions.

The standard kernel/Gram correspondence expresses realization on arbitrary state sets as minimum feasible kernel rank (Aronszajn, 1950; Dokmanić et al., 2015). A general distance–rank bound translates limited distance variation into a dimension requirement. Pair recurrence supplies such restrictions from the transitions: mutual reachability forces equal distances at gain one, while path lengths control their ratios when expansion is permitted. This yields certificates for possibly non-invertible systems and exact two-action separations between one and $n - 1$ dimensions.

A controlled re-encoding experiment preserves all state identities and transition assignments while mapping a shared simplex to grid coordinates. Non-expansiveness survives in the open grid but not the walled systems. The certificates exclude every alternative two-dimensional injective encoding of the

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.

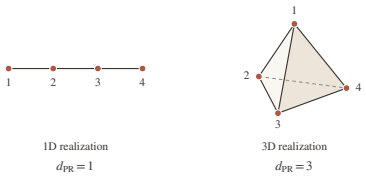


Figure 1: **State identification does not determine realization dimension.** A scalar code distinguishes the same four states in both systems. Endpoint-clipped left/right shifts admit a one-dimensional non-expansive realization (left); the cycle $(1\,2\,3\,4)$ and swap $(1\,2)$ require three dimensions (right). Section 4.2 proves the exact dimensions. The tetrahedron is shown in projection.

latter, even with a one-percent expansion allowance. Coordinate optimization supplies complementary feasible constructions.

Our contributions are:

- We formalize predictive realization dimension, separating state information from exact gain-bounded realization.
- We derive transition-based dimension certificates through shared kernel/Gram geometry, including quantitative bounds beyond gain one.
- We establish exact separations and demonstrate state-preserving dimension reductions whose regularity obstruction cannot be repaired by a different encoding in the target dimension.

2 Related Work

State information and predictive representations. Abstraction, bisimulation, and learned representations study which state distinctions matter for prediction or control (Givan et al., 2003; Ferns et al., 2004; Gelada et al., 2019; Zhang et al., 2021; Wang et al., 2022; Huang et al., 2022). Predictive-state and compact-representation approaches relate complexity to predictive rank or coordinate redundancy (Littman & Sutton, 2001; Thon & Jaeger, 2015; Kobayashi & Watanuki, 2022). DeepMDP explicitly uses reward and latent-transition prediction objectives (Gelada et al., 2019). We retain all state identities and ask whether any encoding in a given dimension can realize the fixed transitions under a common gain bound.

Compression and constrained realization. Classical realization minimizes the state order reproducing input–output behavior (Kalman, 1963); positivity can increase that order (Benvenuti, 2022). Koopman lifting, locally linear models, and augmented Neural ODEs likewise make representation choices relative to a realization class (Brunton et al., 2016; Korda & Mezić, 2018; Lusch et al., 2018; Watter et al., 2015; Dupont et al., 2019). Least Volume makes latent compression depend on decoder Lipschitz continuity (Chen & Fuge, 2024). Our additional requirement concerns the joint realizability of all action maps, rather than decoding regularity or a prescribed linear or flow model. Our state-preserving re-encoding experiment separates finite-set recoverability from this requirement; the certificates exclude alternative low-dimensional encodings, not merely the particular compression tested.

Distance geometry and rank. Euclidean distance matrices and kernels provide the standard correspondence between point configurations and Gram rank (Gower, 1966; Dokmanić et al., 2015; Aronszajn, 1950). Metric embedding and Gram completion start from prescribed distances or partial inner products (Witsenhausen, 1986; Laurent & Varvitsiotis, 2014). Few-distance bounds and perturbed-identity rank estimates are established geometric tools (Bannai et al., 1983; Petrov & Pohoata, 2021; Alon, 2009). Here the distances are variables coupled by the transitions. The unrestricted worst-case problem reduces to classical diameter-to-separation geometry and dimension reduction; we retain it as a benchmark rather than the main structural contribution (Schütte, 1963; Johnson & Lindenstrauss, 1984).

Pair dynamics and symmetric realizations. The pair graph is the classical pair-automaton construction (Saraç et al., 2021). For finite permutation actions at gain one, non-expansiveness forces distance preservation, so the generated group acts orthogonally on the centered configuration. This is a symmetric point-arrangement problem of the kind studied by Winter (2021). Our certificates also apply to non-invertible transitions: pair recurrence constrains an unknown shared geometry, and distance–rank bounds certify dimension. Path lengths yield quantitative certificates when distances need not be preserved.

Latent dimension and regularity. Related work examines latent width, representational collapse, and regular predictive models from empirical, statistical, and algorithmic perspectives (Bhalla et al., 2025; Cui et al., 2026; Lee et al., 2023; Asadi et al., 2018; Song et al., 2024; Beik-Mohammadi et al., 2024). We study the complementary existence question: after retaining state information, can any representation of the chosen dimension support the prescribed predictive regularity?

3 Predictive Realization Dimension

3.1 Setup, definition, and operational meaning

Let $\mathcal{X}$ denote a Markov state set for the controlled dynamics, and let $\mathcal{A}$ be an action set, with deterministic transition maps $\{T_a : \mathcal{X} \to \mathcal{X}\}_{a\in\mathcal{A}}$. We assume $|\mathcal{X}| \geq 2$ and $\mathcal{A} \neq \varnothing$; the state set need not be finite. We restrict attention to injective representations $\phi : \mathcal{X} \hookrightarrow \mathbb{R}^m$, so distinct states cannot be aliased. Write $S_\phi := \phi(\mathcal{X})$. The representation induces the exact target map $F_a^\phi : S_\phi \to S_\phi$ defined by $F_a^\phi(\phi(x)) := \phi(T_a(x))$. Its smallest common statewise Lipschitz gain over actions is

$$\Lambda_T(\phi) := \sup_{a\in\mathcal{A}} \sup_{x\neq y} \frac{\|\phi(T_a(x)) - \phi(T_a(y))\|_2}{\|\phi(x) - \phi(y)\|_2}. \tag{1}$$

Uniform translation and positive rescaling of the representation leave this quantity unchanged.

Consider a predictor $f : \mathbb{R}^m \times \mathcal{A} \to \mathbb{R}^m$ that must respect an action-wise Lipschitz limit $L > 0$. Exact prediction requires $f_a(\phi(x)) = \phi(T_a(x))$ for every state and action, where $f_a(\cdot) := f(\cdot, a)$. Its worst-action gain on the represented states is therefore exactly $\Lambda_T(\phi)$. Any representation supporting such an exact predictor must consequently satisfy $\Lambda_T(\phi) \leq L$. Conversely, when $\Lambda_T(\phi) \leq L$, Kirszbraun's theorem extends each exact target map F_a^ϕ from S_ϕ to $\mathbb{R}^m$ without increasing its Lipschitz constant (Kirszbraun, 1934). Thus $\Lambda_T(\phi) \leq L$ is necessary and sufficient for exact realization by otherwise unrestricted action-wise L-Lipschitz predictors. No shared neural parameterization or continuity in the action input is assumed.

This correspondence motivates measuring the minimum ambient dimension in which the predictive requirement can be met.

Definition 1 (Predictive realization dimension)**.** *For $L > 0$, define*

$$d_{\mathrm{PR}}(T; L) := \min\{m \in \mathbb{N} : \exists\, \phi : \mathcal{X} \hookrightarrow \mathbb{R}^m \text{ with } \Lambda_T(\phi) \leq L\}, \tag{2}$$

Here $\mathbb{N} = \{1, 2, \ldots\}$, and $d_{\mathrm{PR}}(T; L) = +\infty$ if no finite-dimensional realization exists.

Proposition 2 (Operational meaning)**.** *The quantity $d_{\mathrm{PR}}(T; L)$ is exactly the minimum ambient dimension admitting an injective representation and an exact predictor whose action slices are L-Lipschitz. If $m < d_{\mathrm{PR}}(T; L)$, every injective representation $\phi : \mathcal{X} \hookrightarrow \mathbb{R}^m$ has $\Lambda_T(\phi) > L$, so exact prediction necessarily requires exceeding the prescribed action-wise Lipschitz budget.*

This budget is motivated by the role of Lipschitz constants in controlling error propagation: for a fixed one-step approximation error, larger constants can yield looser multi-step prediction error bounds (Asadi et al., 2018).

3.2 One shared Euclidean geometry

Every action imposes a constraint $d_\phi(T_a(x), T_a(y)) \leq L d_\phi(x, y)$, where $d_\phi(x, y) := \|\phi(x) - \phi(y)\|_2$. Crucially, all actions use the same representation ϕ and therefore must satisfy their constraints on *one shared Euclidean geometry*. It is not enough to find a suitable arrangement for each action

separately: a single arrangement must accommodate every action simultaneously. Additional ambient coordinates enlarge the class of shared geometries available to meet these joint requirements.

To express this common geometry, associate a representation with its inner-product kernel $K_\phi(x, y) := \langle \phi(x), \phi(y) \rangle$. Each finite Gram matrix is positive semidefinite with rank at most m. A real symmetric kernel is positive semidefinite when all its finite Gram matrices are so; its rank is the supremum of their ranks. Define its induced squared distances by $D_K(x, y) := K(x, x) + K(y, y) - 2K(x, y)$. For $K = K_\phi$, this equals $\|\phi(x) - \phi(y)\|_2^2$. Thus injectivity requires $D_K(x, y) > 0$ for distinct states and every action imposes $D_K(T_a(x), T_a(y)) \le L^2 D_K(x, y)$ on the same kernel.

Conversely, the feature-map correspondence realizes a rank-r positive-semidefinite kernel as inner products in $\mathbb{R}^r$ (Aronszajn, 1950). Positive induced distances make that map injective; square roots of the transition inequalities give gain at most L. Zero-padding when $r \le m$ establishes the following equivalence.

Proposition 3 (Kernel realization). *For any state set $\mathcal{X}$, an injective realization in $\mathbb{R}^m$ with gain at most L exists if and only if there is a positive-semidefinite kernel K of rank at most m satisfying*

$$\begin{aligned} D_K(x, y) &> 0 && \text{for all } x \neq y, \\ D_K(T_a(x), T_a(y)) &\le L^2 D_K(x, y) && \text{for all } a \in \mathcal{A},\ x \neq y. \end{aligned} \tag{3}$$

Consequently, $d_{\mathrm{PR}}(T; L)$ is the minimum feasible kernel rank, with value $+\infty$ if no finite-rank kernel is feasible.

For a finite state set $\mathcal{X} = \{1, \dots, n\}$, this construction becomes a matrix factorization. Write $z_i := \phi(i)$ and $\bar{z} := n^{-1} \sum_{i=1}^n z_i$, and let $Z \in \mathbb{R}^{n \times m}$ have rows $(z_i - \bar{z})^\top$. The centered Gram matrix is $G := ZZ^\top$, whose entries are $G_{ij} = \langle z_i - \bar{z}, z_j - \bar{z} \rangle$. By construction, $G \succeq 0$, $G\mathbf{1} = 0$, and $\mathrm{rank}(G) \le m$. Centering preserves distances, so $D_{ij}(G) := G_{ii} + G_{jj} - 2G_{ij} = \|z_i - z_j\|_2^2$ (Gower, 1966; Dokmanić et al., 2015).

Finite injectivity gives a positive minimum distance $\delta := \min_{i<j} \|z_i - z_j\|_2$. Replacing Z by Z/δ and G by G/δ^2 makes every nonzero squared distance at least one without changing the rank or any gain ratio. The transition conditions then read $D_{T_a(i), T_a(j)}(G) \le L^2 D_{ij}(G)$. Conversely, any positive-semidefinite matrix of rank $r \le m$ factors as $G = ZZ^\top$ with r columns. The condition $G\mathbf{1} = 0$ centers its row vectors, $D_{ij}(G) \ge 1$ makes them distinct, and the transition inequalities give gain at most L. Zero-padding supplies the requested dimension. Thus the finite case has the following normalized form.

Proposition 4 (Finite-state Gram formulation). *For $\mathcal{X} = \{1, \dots, n\}$, an injective realization in $\mathbb{R}^m$ with gain at most L exists if and only if some $G \in \mathbb{R}^{n \times n}$ satisfies*

$$\begin{aligned} & G \succeq 0, \qquad G\mathbf{1} = 0, \qquad \mathrm{rank}(G) \le m, \\ & D_{ij}(G) \ge 1 && \text{for all } i \neq j, \\ & D_{T_a(i), T_a(j)}(G) \le L^2 D_{ij}(G) && \text{for all } a \in \mathcal{A},\ i \neq j. \end{aligned} \tag{4}$$

Thus $d_{\mathrm{PR}}(T; L)$ is the minimum feasible centered Gram rank.

The unit-separation normalization relies on finiteness; the general kernel formulation requires only positive distances between distinct states. All transition inequalities constrain one geometry, whose rank must fit within the available dimension. Our experiments use the same factorization $G = ZZ^\top$ to optimize coordinates. A regular simplex places n states at equal pairwise distances, so every action either preserves a pair's distance or sends it to zero. Hence $d_{\mathrm{PR}}(T; L) \le n - 1$ for finite n-state systems and $L \ge 1$. The next question is which distance constraints prevent lower-rank realizations.

3.3 Distance geometry and dimension

The shared-kernel formulation casts predictive realization as a rank-constrained feasibility problem on one common Euclidean geometry. What properties of its distances can force the rank to be large? An equidistant configuration gives a starting point: k distinct points with all pairwise distances equal require $k - 1$ dimensions. We ask how much of this requirement survives when the distances are allowed to vary.

Consider $k \geq 2$ distinct points $z_1, \ldots, z_k \in \mathbb{R}^m$ and their squared distances $D_{ij} := \|z_i - z_j\|_2^2$. Let μ and v be the mean and population variance of these squared distances over the $\binom{k}{2}$ pairs $i < j$. Thus v/μ^2 measures relative variation of squared distances, not variance of the state coordinates, and is unchanged by uniform rescaling.

To connect this variation to dimension, let G be the centered Gram matrix of the points. Its rank counts the independent directions in their affine span. Its trace $\operatorname{tr} G$ is the sum of its eigenvalues, while its squared Frobenius norm $\|G\|_F^2$, the sum of squared matrix entries, is the sum of their squares. Cauchy–Schwarz therefore gives $(\operatorname{tr} G)^2 \leq \operatorname{rank}(G)\|G\|_F^2$: a fixed eigenvalue sum cannot be concentrated in very few directions without a large squared sum. The distance–Gram identity gives $\operatorname{tr} G = (k-1)\mu/2$. Separating the mean squared distance from its fluctuations also gives $\|G\|_F^2 \leq (k-1)(\mu^2 + kv)/4$. Distance variation thus limits the spectral concentration that low rank would require. Combining these relations establishes the following bound; Appendix B expands the matrix identities.

Proposition 5 (Distance variation forces dimension)**.** *For $k \geq 2$ distinct points in $\mathbb{R}^m$, with μ, v, and centered Gram matrix G as above,*

$$m \geq \operatorname{rank}(G) \geq \frac{k-1}{1 + k\, v/\mu^2}. \tag{5}$$

At $v = 0$, the bound recovers $k - 1$. Small relative distance variation therefore retains a quantitative part of the equidistant-set obstruction. For the transition analysis, it will be easier to bound the extreme distances than their variance. Let $\delta > 0$ and Δ be the minimum and maximum pairwise distances of the point set. We use $R \geq 1$ for any upper bound on their ratio: $\Delta/\delta \leq R$. All squared distances then lie in $[\delta^2, \Delta^2]$. Averaging $(D_{ij} - \delta^2)(\Delta^2 - D_{ij}) \geq 0$ over pairs yields $v \leq (\Delta^2 - \mu)(\mu - \delta^2)$. The resulting scalar bound is $v/\mu^2 \leq (\Delta^2 - \delta^2)^2/(4\delta^2\Delta^2) \leq (R - R^{-1})^2/4$, as detailed in Appendix B. Substitution gives a form that requires only a distance-ratio bound.

Corollary 6 (Dimension from a bounded distance ratio)**.** *If $k \geq 2$ distinct points in $\mathbb{R}^m$ have maximum-to-minimum pairwise distance ratio at most $R \geq 1$, then*

$$m \geq \frac{k-1}{1 + \frac{k}{4}(R - R^{-1})^2}. \tag{6}$$

These are geometric statements about any finite point set. To certify realization dimension, we next need distance restrictions imposed by the transitions on every representation satisfying the gain bound.

4 Transition-Induced Dimension Certificates

We now derive distance restrictions from the transitions themselves, then return to Figure 1 for an exact separation.

4.1 From pair recurrence to dimension certificates

Each action compares a predecessor pair's distance with its successor's. An action sequence links further distances; a return path compares them in both directions. To express these relations, let $\mathcal{P}_2(\mathcal{X}) := \{\{x, y\} : x, y \in \mathcal{X},\ x \neq y\}$ be the set of unordered pairs of distinct states. Its elements p, q are state pairs, not individual states. An action sends $p = \{x, y\}$ to $a \cdot p := \{T_a(x), T_a(y)\}$ when the successors remain distinct. A merged successor pair contributes zero gain and is not a vertex of this pair graph. Write $p \rightsquigarrow_T q$ when a finite common action sequence carries p to q. Both states receive the same action at each step, although the action may change along the sequence.

For a representation with $\Lambda_T(\phi) \leq L$, write $d_\phi(p) := \|\phi(x) - \phi(y)\|_2$ for $p = \{x, y\}$. Each step satisfies $d_\phi(a \cdot p) \leq L d_\phi(p)$, so a path of r steps gives $d_\phi(q) \leq L^r d_\phi(p)$. A return path of s steps gives $d_\phi(p) \leq L^s d_\phi(q)$. The forward and return inequalities give the recurrence relation.

Lemma 7 (Quantitative pair recurrence)**.** *Suppose $\Lambda_T(\phi) \leq L$ with $L > 0$. If p reaches q in r steps, then $d_\phi(q) \leq L^r d_\phi(p)$. If q also reaches p in s steps, then $L^{-s} d_\phi(p) \leq d_\phi(q) \leq L^r d_\phi(p)$. At $L = 1$, mutually reachable pairs have equal distances.*

Call the equivalence classes of mutual reachability *dynamical pair types*; for finite systems, these are the strongly connected components (SCCs) of the directed pair graph. At gain one, each type has a single represented distance. A finite state subset $Y \subseteq \mathcal{X}$ is *single-type* when all pairs in $\mathcal{P}_2(Y)$ belong to one type. If Y has k states, all $\binom{k}{2}$ distances in $\phi(Y)$ are then equal in every gain-one realization. The zero-variance case of Proposition 5 therefore gives the certificate.

Corollary 8 (Single-type subsets)**.** *If Y is single-type and $2 \le |Y| < \infty$, every gain-one realization embeds Y as an equidistant set, so $d_{\mathrm{PR}}(T; 1) \ge |Y| - 1$.*

For $L > 1$, return paths allow some distance variation instead of forcing equality. Their lengths determine how much. Let $\ell_T(p, q)$ be the shortest directed path length from p to q, measured in primitive actions, with $\ell_T(p, p) = 0$. For a finite single-type subset Y of size k, define $h_T(Y) := \max_{p,q \in \mathcal{P}_2(Y)} \ell_T(p, q)$. This maximum is finite because Y has finitely many mutually reachable pairs. Paths are taken in the full pair graph and may leave $\mathcal{P}_2(Y)$.

In any representation with gain at most $L \ge 1$, let δ_Y and Δ_Y denote the minimum and maximum distances among the represented states of Y. A minimum-distance pair can reach a maximum-distance pair in at most $h = h_T(Y)$ steps. Lemma 7 gives $\Delta_Y/\delta_Y \le L^h$. The extremizing pairs can change with the representation, but k and h depend only on the witness and its transitions. Thus the same distance-ratio bound holds for every feasible arrangement. Applying Corollary 6 with the bound $R = L^h$ and using $(L^h - L^{-h})^2/4 = \sinh^2(h \log L)$ yields the following result.

Proposition 9 (Robust single-type certificate)**.** *Let $Y \subseteq \mathcal{X}$ be single-type, with $|Y| = k \ge 2$ finite and $h = h_T(Y)$. For every $L \ge 1$,*

$$d_{\mathrm{PR}}(T; L) \ge \left\lceil \frac{k - 1}{1 + k \sinh^2(h \log L)} \right\rceil . \tag{7}$$

The ceiling uses that feasible ambient dimensions are integers. At $L = 1$, the result recovers $k - 1$. For fixed k and $L > 1$, shorter paths give stronger bounds because less expansion can accumulate. The certificate can weaken rapidly as L grows; it quantifies a surviving obstruction rather than the exact dimension–gain profile.

For finite states and actions, these witnesses can be found directly from the transition table. For each pair-SCC, form a graph on the original states, joining two states when their pair lies in that SCC. A clique is precisely a single-type subset, so its size gives a gain-one certificate. The strongest such certificate is

$$b_1(T) := \sup_{\substack{Y \subseteq \mathcal{X},\ 2 \le |Y| < \infty \\ Y \text{ single-type}}} (|Y| - 1), \tag{8}$$

with value zero if no such subset exists. Hence $d_{\mathrm{PR}}(T; 1) \ge b_1(T)$. In a finite system, maximum-clique search computes this supremum as a maximum; the number of pair-SCC vertices alone does not determine it. Shortest-path computations for a chosen witness then give its robust bound. Exact maximum-clique search can be expensive; a supplied witness can be checked without proving that it is largest. A largest clique need not be optimal for $L > 1$, since path length also matters. Appendix G.2 gives the computational costs.

Only the witness must be finite, not the full state set. Appendix C treats subsets involving several pair types using classical few-distance bounds (Bannai et al., 1983; Petrov & Pohoata, 2021). Recurrence is one source of dimensional obstruction, not a complete characterization: the shared kernel/Gram formulation also retains one-way transition inequalities that these certificates may not exploit.

4.2 Sharp realization separations

For the exact separation in Figure 1, take $\mathcal{X} = \{1, \dots, n\}$ with $n \ge 3$ and actions $\{a, b\}$. The scalar code $s(i) = i$ distinguishes every state in both systems; only the transition maps differ.

For the low family, let $T_a^{\mathrm{low}}(i) = \min\{i + 1, n\}$ and $T_b^{\mathrm{low}}(i) = \max\{i - 1, 1\}$. These are forward and backward shifts clipped at the endpoints, not a cycle. In the placement $\phi(i) = i$, both maps are non-expansive, so one dimension suffices. Injectivity for more than one state excludes zero dimensions. For the high family, take $T_a^{\mathrm{high}} = c = (1\,2\,\dots\,n)$ and $T_b^{\mathrm{high}} = \tau = (1\,2)$, where τ swaps states 1 and 2 and fixes the others. These permutations generate all state permutations, so all

Table 1: **Certified dimensions under relaxed gain budgets.** The rightmost columns are lower bounds on $d_{\mathrm{PR}}(T; L)$. The $L = 1$ column gives the exact single-type certificate $b_1(T)$, not the exact realization dimension. Shortest paths may leave the witness pair set.

Maze	$\lvert Y^\star \rvert$	$h_T(Y^\star)$	$L = 1$	1.005	1.01	1.02
A	30	21	29	22	13	5
B	48	32	47	22	8	3
C	53	30	52	24	9	3

state pairs are mutually reachable. Corollary 8 forces dimension at least $n - 1$, attained by a regular simplex. The construction and certificate therefore meet exactly.

Corollary 10 (Two-action separation)**.** *For every $n \geq 3$, the two systems above share the same n states and two-action alphabet, but satisfy $d_{\mathrm{PR}}(T^{\mathrm{low}}; 1) = 1$ and $d_{\mathrm{PR}}(T^{\mathrm{high}}; 1) = n - 1$.*

For four states, the high family is identifiable in one coordinate, but no one- or two-dimensional injective placement realizes it non-expansively. The extra coordinates serve the transitions, not additional state information. Appendix D supplies the full construction and group-generation argument. Universal worst-case bounds remain separate geometric benchmarks in Appendix E.

5 Experiments

We distinguish state preservation, certified infeasibility, and constructive feasibility. Fully enumerated systems permit exact checks of every transition constraint without sampling uncertainty. The controlled experiments optimize coordinates rather than train a neural model; a separate trained-model analysis is summarized below.

5.1 Controlled realizations

Corollary 10 gives the analytic clipped-chain/simplex separation between dimensions 1 and $n - 1$. As complementary numerical checks, we evaluate the cyclic variant $\{c, c^{-1}\}$ and the family $\{c, \tau\}$, where $c = (1\,2\,\ldots\,n)$ and $\tau = (1\,2)$. For $n \in \{3, 4, 5, 8, 16\}$, regular polygons and simplices verify gain one up to floating-point precision. These cyclic-variant dimensions are 2 and $n - 1$ by Theorem 13; details are in Appendix G.1.

5.2 Certified dimensions for open and walled grids

The open 8×8 grid and Mazes A/B/C share 64 state identities, four actions $\{U, D, L, R\}$, and the same boundary convention. Internal walls change transitions by blocking moves, not the information specified by the row and column. For each maze, the directed graph of all $\binom{64}{2} = 2016$ state pairs and maximum-clique search within pair types give the exact single-type certificates $b_1(T_A) = 29$, $b_1(T_B) = 47$, and $b_1(T_C) = 52$. For the saved witnesses, independently checked shortest paths in the full pair graph give $h_T(Y^\star) = 21, 32, 30$. Table 1 applies Proposition 9; the relaxed-gain columns are bounds for these witnesses, not optima over all subsets.

The open grid's clipped coordinate shifts are non-expansive, giving $d_{\mathrm{PR}}(T_{\mathrm{open}}; L) \leq 2$ for $L \geq 1$. The disjoint bounds establish different realization requirements. At $L = 1.01$, A/B/C require at least 13/8/9 dimensions; at $L = 1.02$, all still exclude dimension two. The table also shows the bounds' rapid weakening. Appendix G.2 gives the construction and audits.

5.3 State-preserving re-encoding and predictive regularity

All four systems use one shared source codebook: a centered regular simplex $Z_{\mathrm{s}} \in \mathbb{R}^{64 \times 63}$. The common target codebook $Z_{\mathrm{t}} \in \mathbb{R}^{64 \times 2}$ contains centered, uniformly normalized row–column coordinates in the same state order. The 63 columns of Z_{s} span the zero-sum subspace of $\mathbb{R}^{64}$, which contains both columns of Z_{t}. Thus a matrix $W \in \mathbb{R}^{63 \times 2}$ with $Z_{\mathrm{s}} W = Z_{\mathrm{t}}$ exists exactly. The computed residual verifies this construction, not its existence. This map is injective on the 64 source states, though not on all of $\mathbb{R}^{63}$. An inverse codebook lookup from compressed coordinates recovers all

Table 2: **The same state-preserving re-encoding, different regularity outcomes.** One linear map sends a shared simplex to a shared two-dimensional grid codebook. All state identities are recovered, but only the open grid retains gain one. The gains describe these codebooks, not the optimal gains in their respective dimensions.

System	63D source gain	2D target gain	Identities recovered
Open grid	1	1	64/64
Maze A	1	2	64/64
Maze B	1	2	64/64
Maze C	1	2	64/64

64 identities, including shuffled queries. Re-encoding the source successor agrees with taking the corresponding target successor for every state and action. Thus state identification and the exact transition assignments are preserved.

Independent scalar and vectorized implementations check all 8064 ratios per codebook and system. These measurements concern the chosen map; the certificates exclude every alternative injective two-dimensional encoding of the mazes, even with one-percent expansion. Exact prediction remains possible on the codebook, but not under the prescribed regularity. The inverse is a finite-set lookup, not a trained or noise-robust decoder. Appendix G.3 gives the audit.

5.4 Constructive dimension–gain trade-offs

For each dimension m, let $L_m^\star(T) := \inf_{\phi:\mathcal{X}\hookrightarrow\mathbb{R}^m} \Lambda_T(\phi)$ denote the best gain in the infimum sense. We optimize coordinates $Z \in \mathbb{R}^{64\times m}$, parameterizing the shared Gram geometry through $G = ZZ^\top$. The candidate set combines ten random restarts per dimension, zero-padded and perturbed continuation, the grid code and its zero-padded versions, and grid-initialized optimization at $m = 2$. Every candidate is evaluated by exhaustive float64 hard gain, and its value provides a constructive upper bound on $L_m^\star(T)$. Random-restart medians and interquartile ranges exclude all deterministic and continuation candidates.

Grid-initialized optimization attains two-dimensional gains of 1.6272, 1.6857, and 1.7259 for A/B/C, improving on the raw-grid gain of two without reaching the certified excluded budgets. At $m = 24$, the retained gains are 1.0543, 1.0655, and 1.0725; a simplex gives gain one at $m = 63$. Figure 2 shows the best recorded constructions, not globally optimal frontiers. The gap to the certificates has two unresolved sources: the witness/path-length and distance-ratio bounds discard some constraints, while nonconvex optimization need not find the best geometry. We do not determine their relative contributions. In particular, a lower bound of three dimensions does not establish feasibility in three dimensions. Optimization details appear in Appendix G.4.

Gaussian and projected-simplex samples remain unoptimized geometric references, not evidence of superiority over optimized transition-independent methods (Appendix G.5). The fixed-wall study is a secondary certificate control (Appendix G.5).

Appendix G.6 examines existing LeWorldModel checkpoints on PushT. Upper quantiles of sampled same-action predictor ratios decrease with width, but sampled maxima do not decrease monotonically and state readability changes. These are exploratory measurements, not global Lipschitz bounds or a controlled demonstration of dimension necessity.

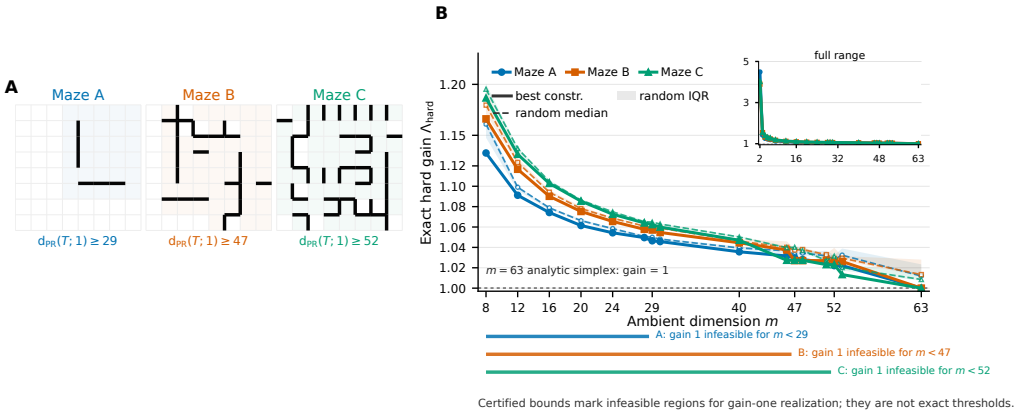


Figure 2: **Constructive dimension–gain trade-offs.** Solid curves show the best recorded gains, including grid-based candidates; dashed curves and bands summarize random restarts only. Maze labels and infeasibility annotations concern gain-one lower certificates, not exact realization thresholds. The simplex endpoint is analytic. Relaxed-gain certificates appear in Table 1.

6 Discussion and Conclusion

State information and predictive realization impose different requirements. Our kernel formulation and transition certificates show when additional Euclidean dimensions are necessary for exact gain-bounded prediction. The re-encoding experiment preserves every state identity and transition assignment, while the certificates exclude alternative encodings in the target dimension. Thus state recovery alone cannot certify that a dimension is sufficient for the required predictive geometry.

Scope and limitations. The gain budget specifies a requirement to be tested, not a recommendation that every system should be non-expansive or that smaller gain always improves control. The formal results concern exact deterministic transitions; they do not directly cover stochastic kernels, approximate prediction, or regularity across different actions. Recurrence certificates exploit only part of the full geometric constraints and need not identify the exact dimension. Finite state identity is not a claim of robust reconstruction, and the exploratory PushT measurements do not provide an architecture-specific width rule or a guarantee of downstream performance.

Beyond predictors. A latent representation is not only a description of information; it is also an input to the modules that use it. A decoder reconstructs a target, while policies and value functions use latent states for action selection and evaluation (Chen & Fuge, 2024; Hafner et al., 2020). This suggests judging coordinate redundancy relative to both the information retained and the requirements of the consuming module. Our results establish this distinction for gain-bounded predictors. Extending it to decoders or reinforcement-learning modules requires specifying their target behavior, output geometry, and regularity constraints; the present dimension bounds do not automatically apply to them. The broader perspective is therefore to design latent geometry for the operations it must support, not to infer its adequacy from information preservation alone.

AI USE STATEMENT

Generative AI tools, including ChatGPT and Codex, assisted with literature retrieval and organization, theoretical discussions, mathematical argument and proof drafting/checking, experimental design and implementation, code review, interpretation of results, and manuscript drafting and editing. The numerical audit procedures are documented in Appendix G. The author takes responsibility for the final claims, proofs, citations, code, results, and text.

REPRODUCIBILITY STATEMENT

Appendices A–F provide the detailed theoretical arguments. Appendix G specifies the grid transitions, witness checks, re-encoding audit, optimization settings, and trained-checkpoint analysis. It distinguishes exact constructions and certificates from numerical feasibility and observational evidence.

REFERENCES

Noga Alon. Perturbed identity matrices have high rank: Proof and applications. *Combinatorics, Probability and Computing*, 18(1–2):3–15, 2009. doi: 10.1017/S0963548307008917.

N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950. doi: 10.1090/S0002-9947-1950-0051437-7.

Kavosh Asadi, Dipendra Misra, and Michael Littman. Lipschitz continuity in model-based reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 264–273, 2018.

Eiichi Bannai, Etsuko Bannai, and Dennis Stanton. An upper bound for the cardinality of an s-distance subset in real euclidean space, ii. *Combinatorica*, 3(2):147–152, 1983. doi: 10.1007/BF02579288.

Hadi Beik-Mohammadi, Søren Hauberg, Georgios Arvanitidis, Nadia Figueroa, Gerhard Neumann, and Leonel Rozo. Neural contractive dynamical systems. In *The Twelfth International Conference on Learning Representations*, 2024.

Luca Benvenuti. Minimal positive realizations: A survey. *Automatica*, 143:110422, 2022. doi: 10.1016/j.automatica.2022.110422.

Brady Bhalla, Honglu Fan, Nancy Chen, and Tony Yue Yu. Higher embedding dimension creates a stronger world model for a simple sorting task, 2025.

Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLoS ONE*, 11(2):e0150171, 2016. doi: 10.1371/journal.pone.0150171.

Qiuyi Chen and Mark Fuge. Compressing latent space via least volume. In *International Conference on Learning Representations*, 2024.

Jingyi Cui, Qi Zhang, Hongwei Wen, and Yisen Wang. A generalization theory for jepa-based world models, 2026.

Ivan Dokmanić, Reza Parhizkar, Juri Ranieri, and Martin Vetterli. Euclidean distance matrices: Essential theory, algorithms and applications. *IEEE Signal Processing Magazine*, 32(6):12–30, 2015. doi: 10.1109/MSP.2015.2398954.

Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite markov decision processes. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pp. 162–169, 2004.

Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G. Bellemare. Deepmdp: Learning continuous latent space models for representation learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2170–2179, 2019.

Robert Givan, Thomas Dean, and Matthew Greig. Equivalence notions and model minimization in markov decision processes. *Artificial Intelligence*, 147(1–2):163–223, 2003. doi: 10.1016/S0004-3702(02)00376-4.

J. C. Gower. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*, 53(3–4):325–338, 1966. doi: 10.1093/biomet/53.3-4.325.

David Ha and Jürgen Schmidhuber. World models, 2018.

Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2555–2565, 2019.

Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2020.

Biwei Huang, Chaochao Lu, Liu Leqi, Jose Miguel Hernandez-Lobato, Clark Glymour, Bernhard Schölkopf, and Kun Zhang. Action-sufficient state representation learning for control with structural constraints. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9260–9279, 2022.

William B. Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. In *Conference in Modern Analysis and Probability*, volume 26 of *Contemporary Mathematics*, pp. 189–206. American Mathematical Society, 1984.

R. E. Kalman. Mathematical description of linear dynamical systems. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 1(2):152–192, 1963. doi: 10.1137/0301010.

M. Kirszbraun. Über die zusammenziehende und lipschitzsche transformationen. *Fundamenta Mathematicae*, 22(1):77–108, 1934. doi: 10.4064/fm-22-1-77-108.

Taisuke Kobayashi and Ryoma Watanuki. Sparse representation learning with modified q-vae towards minimal realization of world model, 2022.

Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93:149–160, 2018. doi: 10.1016/j.automatica.2018.03.046.

Monique Laurent and Antonios Varvitsiotis. A new graph parameter related to bounded rank positive semidefinite matrix completions. *Mathematical Programming*, 145(1–2):291–325, 2014. doi: 10.1007/s10107-013-0648-x.

Hojoon Lee, Koanho Lee, Dongyoon Hwang, Hyunho Lee, Byungkun Lee, and Jaegul Choo. On the importance of feature decorrelation for unsupervised representation learning in reinforcement learning. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 18988–19009, 2023.

Michael L. Littman and Richard S. Sutton. Predictive representations of state. In *Advances in Neural Information Processing Systems*, volume 14, 2001.

Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9:4950, 2018. doi: 10.1038/s41467-018-07210-0.

Lucas Maes, Quentin Le Lidec, Damien Scieur, Yann LeCun, and Randall Balestriero. Leworldmodel: Stable end-to-end joint-embedding predictive architecture from pixels, 2026.

Fedor Petrov and Cosmin Pohoata. A remark on sets with few distances in $\setminus^d$. *Proceedings of the American Mathematical Society*, 149(2):569–571, 2021. doi: 10.1090/proc/15231.

N. Ege Saraç, Ömer Faruk Altun, Kamil Tolga Atam, Sertaç Karahoda, Kamer Kaya, and Hüsnü Yenigün. Boosting expensive synchronizing heuristics. *Expert Systems with Applications*, 167:114203, 2021. doi: 10.1016/j.eswa.2020.114203.

Kurt Schütte. Minimale Durchmesser endlicher Punktmengen mit vorgeschriebenem Mindestabstand. *Mathematische Annalen*, 150:91–98, 1963. doi: 10.1007/BF01396584.

Yuda Song, Lili Wu, Dylan J. Foster, and Akshay Krishnamurthy. Rich-observation reinforcement learning with continuous latent dynamics. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 46181–46237. PMLR, 2024.

Michael Thon and Herbert Jaeger. Links between multiplicity automata, observable operator models and predictive state representations—a unified learning framework. *Journal of Machine Learning Research*, 16(4):103–147, 2015.

Tongzhou Wang, Simon Du, Antonio Torralba, Phillip Isola, Amy Zhang, and Yuandong Tian. Denoised mdps: Learning world models better than the world itself. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 22591–22612, 2022.

Manuel Watter, Jost Tobias Springenberg, Joschka Boedecker, and Martin A. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in Neural Information Processing Systems*, volume 28, pp. 2746–2754, 2015.

Martin Winter. Geometry and topology of symmetric point arrangements. *Linear Algebra and its Applications*, 611:1–25, 2021. doi: 10.1016/j.laa.2020.11.017.

Hans S. Witsenhausen. Minimum dimension embedding of finite metric spaces. *Journal of Combinatorial Theory, Series A*, 42(2):184–199, 1986. doi: 10.1016/0097-3165(86)90089-0.

Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. In *International Conference on Learning Representations*, 2021.

A Kernel and Gram Realization

Proof of Proposition 3. An injective representation $\phi : \mathcal{X} \hookrightarrow \mathbb{R}^m$ defines $K(x, y) = \langle \phi(x), \phi(y) \rangle$. Every finite Gram matrix of K is positive semidefinite and has rank at most m. Moreover, $D_K(x, y) = \|\phi(x) - \phi(y)\|_2^2$, so injectivity gives $D_K(x, y) > 0$ for $x \neq y$, and $\Lambda_T(\phi) \leq L$ gives the transition inequalities in Equation (3).

Conversely, a positive-semidefinite kernel has an inner-product feature representation $K(x, y) = \langle \psi(x), \psi(y) \rangle$ (Aronszajn, 1950). If K has rank $r < \infty$, the span of these feature vectors has dimension r: otherwise some $r + 1$ feature vectors would be linearly independent and their Gram matrix would have rank $r + 1$. Choosing an orthonormal basis identifies this span with $\mathbb{R}^r$. The condition $D_K(x, y) > 0$ makes ψ injective, and taking square roots of the transition inequalities gives $\Lambda_T(\psi) \leq L$. If $r \leq m$, zero-padding yields a realization in $\mathbb{R}^m$.

Thus the feasible finite dimensions are exactly those admitting a kernel of rank at most that dimension. Their minima coincide, with value $+\infty$ if neither a finite-dimensional realization nor a finite-rank feasible kernel exists. No positive lower bound on all pairwise distances is required for this general-set formulation. □

Proof of Proposition 4. Suppose $\phi : \{1, \ldots, n\} \hookrightarrow \mathbb{R}^m$ satisfies $\Lambda_T(\phi) \leq L$. Set $z_i = \phi(i)$, translate the points to have zero mean, and divide them by $\delta = \min_{i<j} \|z_i - z_j\|_2 > 0$. Translation and uniform scaling leave all gain ratios unchanged. Writing the resulting points as the rows of Z, the matrix $G = ZZ^\top$ is positive semidefinite, satisfies $G\mathbf{1} = 0$, and has rank at most m. Its squared distances satisfy $D_{ij}(G) \geq 1$ for $i \neq j$ and $D_{T_a(i),T_a(j)}(G) \leq L^2 D_{ij}(G)$.

Conversely, factor a feasible matrix as $G = ZZ^\top$ with r columns, where $r = \operatorname{rank}(G) \leq m$. The rows are centered because $\|Z^\top \mathbf{1}\|_2^2 = \mathbf{1}^\top G \mathbf{1} = 0$. They are distinct because $D_{ij}(G) \geq 1$, and the transition inequalities imply gain at most L. Padding by zeros gives the requested ambient dimension. This is the standard distance–Gram correspondence with the additional transition constraints (Gower, 1966; Dokmanić et al., 2015). □

B Distance Variation and Euclidean Dimension

Proof of Proposition 5. Let D be the squared distance matrix of the k points and define $\mu = \binom{k}{2}^{-1} \sum_{i<j} D_{ij}$ and $v = \binom{k}{2}^{-1} \sum_{i<j} (D_{ij} - \mu)^2$. Because the points are distinct, $\mu > 0$. Write $H = I - \mathbf{1}\mathbf{1}^\top / k$ and

$$D = \mu(\mathbf{1}\mathbf{1}^\top - I) + E. \tag{B.1}$$

The remainder has zero diagonal and zero sum of entries, and $\|E\|_F^2 = k(k-1)v$. Double centering therefore gives $G = -HDH/2 = \mu H/2 - HEH/2$. Since $\operatorname{tr}(HEH) = \operatorname{tr}(EH) = 0$,

$$\begin{aligned} \operatorname{tr} G &= \frac{(k-1)\mu}{2}, \\ \|G\|_F^2 &= \frac{(k-1)\mu^2 + \|HEH\|_F^2}{4} \leq \frac{(k-1)\mu^2 + k(k-1)v}{4}. \end{aligned} \tag{B.2}$$

The equality uses $\|H\|_F^2 = k - 1$ and the vanishing cross term; the inequality uses that multiplication by an orthogonal projection cannot increase the Frobenius norm.

Let $r = \operatorname{rank}(G)$ and let $\lambda_1, \ldots, \lambda_r > 0$ be its nonzero eigenvalues. Cauchy–Schwarz gives $(\sum_i \lambda_i)^2 \leq r \sum_i \lambda_i^2$, hence

$$\operatorname{rank}(G) \geq \frac{(\operatorname{tr} G)^2}{\|G\|_F^2} \geq \frac{(k-1)^2\mu^2}{(k-1)\mu^2 + k(k-1)v} = \frac{k-1}{1 + k\, v/\mu^2}. \tag{B.3}$$

Finally, G is the Gram matrix of centered points in $\mathbb{R}^m$, so $r \leq m$. □

Proof of Corollary 6. Let $a = \min_{i<j} D_{ij} > 0$ and $b = \max_{i<j} D_{ij}$. For a variable U uniformly distributed over the off-diagonal squared distances, $(U - a)(b - U) \geq 0$. Taking expectations yields $v \leq (b - \mu)(\mu - a)$, and

$$\frac{v}{\mu^2} \leq \frac{(b-\mu)(\mu-a)}{\mu^2} \leq \frac{(b-a)^2}{4ab}. \tag{B.4}$$

The second inequality follows, for example, by subtracting its left-hand side from its right-hand side; the difference equals $((a+b)\mu - 2ab)^2/(4ab\mu^2) \geq 0$. The actual distance ratio is $\sqrt{b/a} \leq R$, so $(b-a)^2/(4ab) \leq (R - R^{-1})^2/4$. Substitution into Proposition 5 proves the claim. At $R = 1$, all squared distances coincide and the same argument reduces to $m \geq k - 1$. □

C Further Certificates from Dynamical Pair Types

For a finite subset $Y \subseteq \mathcal{X}$, define $\kappa_T(Y)$ as the number of dynamical pair types touched by $\mathcal{P}_2(Y)$. This is a count of pair types, not the number $|Y|$ of original states. The single-type condition in the main text is equivalently $\kappa_T(Y) = 1$. At gain one, subsets touching several types can also be treated through classical few-distance geometry. A finite Euclidean set with at most s distinct nonzero distances has at most $\binom{m+s}{s}$ points in $\mathbb{R}^m$ (Bannai et al., 1983; Petrov & Pohoata, 2021).

Theorem 11 (Transition-computable structural lower bound)**.** *If an injective representation $\phi : \mathcal{X} \hookrightarrow \mathbb{R}^m$ satisfies $\Lambda_T(\phi) \leq 1$, then every finite $Y \subseteq \mathcal{X}$ with $|Y| \geq 2$ obeys $|Y| \leq \binom{m+\kappa_T(Y)}{\kappa_T(Y)}$. Consequently,*

$$d_{\mathrm{PR}}(T;1) \geq \sup_{\substack{Y \subseteq \mathcal{X} \\ 2 \leq |Y| < \infty}} \min\left\{ m \in \mathbb{N} : \binom{m+\kappa_T(Y)}{\kappa_T(Y)} \geq |Y| \right\}. \tag{C.1}$$

For finite $\mathcal{X}$, the outer supremum is a maximum.

Proof. Fix such a subset Y. Injectivity makes its nonzero pairwise distances positive. By Lemma 7, each pair type has one distance value at gain one; distinct types are allowed to have the same value. Thus the number s of distances realized on $\phi(Y)$ is at most $\kappa_T(Y)$. The few-distance bound yields $|Y| \leq \binom{m+s}{s} \leq \binom{m+\kappa_T(Y)}{\kappa_T(Y)}$. Every feasible dimension satisfies this for each finite Y, proving the stated bound. □

The finite-witness formulation does not require Y to be invariant under the actions. All reachability relations are inherited from the full system. Nor does the supremum over witnesses imply that the certificate is complete: one-way inequalities remain relevant to the full kernel feasibility problem.

D Sharp Separations and Generator Invariance

Proposition 12 (Generator invariance at gain one)**.** *Let $\langle T \rangle$ be the transformation semigroup generated by the primitive actions. A representation satisfies $\Lambda_T(\phi) \leq 1$ if and only if every map in $\langle T \rangle$ acts non-expansively on $\phi(\mathcal{X})$. Closing the action family under composition changes neither $d_{\mathrm{PR}}(T;1)$ nor the dynamical pair types.*

Proof. Compositions of non-expansive maps remain non-expansive, and every primitive action belongs to its generated semigroup. Adding the edge associated with a composition only shortcuts an existing pair path. If a final pair is distinct, no intermediate pair can have merged, since deterministic common actions cannot separate coincident states. Thus the reachability relation and its equivalence classes are unchanged. □

Theorem 13 (Same states, different realization dimensions)**.** *For every $n \geq 3$, the cyclic system $\mathsf{Cycle}_n : i \mapsto i + 1 \pmod n$ has $d_{\mathrm{PR}}(\mathsf{Cycle}_n;1) = 2$. Every n-state system with $\kappa_T(\mathcal{X}) = 1$ has $d_{\mathrm{PR}}(T;1) = n - 1$.*

Proof. Placing the cycle states at $z_i = (\cos(2\pi i/n), \sin(2\pi i/n))$ realizes the action as a planar rotation. This proves feasibility in two dimensions. Under the cyclic action, pair types are indexed by cyclic separation, so $\kappa_T(\mathcal{X}) = \lfloor n/2 \rfloor$. If a one-dimensional gain-one realization existed, Theorem 11 would imply $n \leq 1 + \lfloor n/2 \rfloor$, which is false for $n \geq 3$.

If instead $\kappa_T(\mathcal{X}) = 1$, Corollary 8 gives the lower bound $n - 1$. A regular simplex realizes any deterministic n-state action family with gain at most one: distinct pairs are sent either to distinct vertices at the same distance or to coincident vertices. This gives the matching upper bound. □

Proof of Corollary 10. For the low family, use $T_a^{\mathrm{low}}(i) = \min\{i+1, n\}$ and $T_b^{\mathrm{low}}(i) = \max\{i-1, 1\}$. For any $i, j \in \{1, \dots, n\}$, $|\min\{i+1, n\} - \min\{j+1, n\}| \le |i-j|$, and the corresponding inequality holds for the backward shift. Thus $\phi(i) = i$ is an injective one-dimensional realization with gain at most one. Since there is more than one state, a zero-dimensional injective realization is impossible, proving $d_{\mathrm{PR}}(T^{\mathrm{low}}; 1) = 1$.

For the high family, let $c = (1\,2\,\dots\,n)$ and $\tau = (1\,2)$. The conjugates $c^j \tau c^{-j}$ for $j = 0, \dots, n-2$ are the adjacent transpositions, which generate the symmetric group. Inverses of these finite permutations are positive powers, so every group element is a finite word in the primitive actions. Any unordered distinct state pair can therefore reach any other; no intermediate pair merges because the actions are permutations. The entire state set is single-type, and Corollary 8 gives the lower bound $n-1$. Placing the states at a regular simplex in $\mathbb{R}^{n-1}$ preserves every nonzero pairwise distance under a permutation, providing an exact gain-one realization. Hence $d_{\mathrm{PR}}(T^{\mathrm{high}}; 1) = n-1$.

Both systems use the same n state identities and two action labels, and $s(i) = i$ is a complete scalar state code in each. It is a gain-one realization only for the low family. The cyclic-system theorem above is a separate two-dimensional variant, not the low family in this corollary. □

Generator invariance here concerns gain one. At larger gain, a word of length r need only be L^r-Lipschitz, so replacing that word by a primitive action with gain budget L would change the problem.

E Universal Worst-Case Dimension–Gain Bounds

Let $\mathfrak{T}_n$ contain all deterministic n-state systems with arbitrary finite action families, where $n \ge 3$. Write $D_L(n) = \sup_{T \in \mathfrak{T}_n} d_{\mathrm{PR}}(T; L)$ and $W_m(n) = \sup_{T \in \mathfrak{T}_n} L_m^\star(T)$. These quantities vary the transition family, unlike the instance-level profile $L_m^\star(T)$.

E.1 Reduction to finite point-set geometry

For n distinct points $Z = (z_1, \dots, z_n) \subset \mathbb{R}^m$, write $\delta(Z) = \min_{i<j} \|z_i - z_j\|_2$ and $\Delta(Z) = \max_{i<j} \|z_i - z_j\|_2$. Define

$$a_m(n) := \inf_{Z \text{ distinct}} \frac{\Delta(Z)}{\delta(Z)}. \tag{E.1}$$

After setting $\delta(Z) = 1$, this is the minimum possible diameter of a finite point set with prescribed minimum separation, a classical problem in Euclidean geometry (Schütte, 1963).

Proposition 14 (Worst-case reduction)**.** *For every $m \ge 1$, $W_m(n) = a_m(n)$. For every $L \ge 1$, $D_L(n) = \min\{m \in \mathbb{N} : a_m(n) \le L\}$. Both worst cases are attained by the fixed action family containing all permutations as primitive actions.*

Proof. For any action family and any Z, every successor distance is at most $\Delta(Z)$ and every nonzero predecessor distance is at least $\delta(Z)$. Hence $\Lambda_T(Z) \le \Delta(Z)/\delta(Z)$ and $L_m^\star(T) \le a_m(n)$. For the family containing all permutations, one action sends a minimum-distance pair to a maximum-distance pair. Thus $\Lambda_T(Z) = \Delta(Z)/\delta(Z)$ for every Z, and taking infima and then the worst case proves $W_m(n) = a_m(n)$.

The infimum defining $a_m(n)$ is attained. In a minimizing sequence, translate the first point to zero and scale the minimum distance to one. A uniform diameter bound then bounds every coordinate. A convergent subsequence retains minimum distance one, and continuity of the diameter gives a minimizer. Consequently, $a_m(n) \le L$ supplies a geometry feasible for every action family, while $a_m(n) > L$ excludes the all-permutation family. Minimizing m proves the dimension statement. A simplex ensures the set of feasible dimensions is nonempty for $L \ge 1$. □

The compactness argument above is specific to the point-set ratio problem. It does not assert attainment of $L_m^\star(T)$ for an arbitrary transition family.

E.2 Worst-case dimension at fixed gain

Theorem 15 (Worst-case dimension at fixed gain)**.** *For $n \geq 3$, $D_1(n) = n - 1$. For every fixed $L > 1$, there is a universal constant $C > 0$ such that*

$$\left\lceil \frac{\log n}{\log(2L+1)} \right\rceil \leq D_L(n) \leq \min\left\{ n-1, C\left(\frac{L+1}{L-1}\right)^2 \log n \right\}. \tag{E.2}$$

In particular, $D_L(n) = \Theta_L(\log n)$ for fixed $L > 1$.

Proof. At $L = 1$, the all-permutation family forces all nonzero distances to be equal, requiring $n - 1$ dimensions. A simplex supplies the matching construction for every system.

For $L > 1$, start with a regular simplex and apply a Johnson–Lindenstrauss map preserving distances within factors $1 - \varepsilon$ and $1 + \varepsilon$ in at most $C\varepsilon^{-2} \log n$ dimensions (Johnson & Lindenstrauss, 1984). The resulting points remain distinct and have distance ratio at most $(1 + \varepsilon)/(1 - \varepsilon)$. Taking $\varepsilon = (L - 1)/(L + 1)$ and applying Proposition 14 proves the upper bound, together with the simplex construction.

For the lower bound, consider a configuration feasible for the all-permutation family, scale its minimum distance to one, and write its diameter as $\Delta \leq L$. Balls of radius $1/2$ centered at the n points have disjoint interiors. All are contained in a ball of radius $\Delta + 1/2$ centered at one of the points. Volume comparison gives $n \leq (2\Delta + 1)^m \leq (2L + 1)^m$. Rearranging and using integrality proves the lower bound. □

E.3 Worst-case gain at fixed dimension

Theorem 16 (Worst-case gain at fixed dimension)**.** *For every $m \geq 1$,*

$$\max\left\{1, \frac{n^{1/m} - 1}{2}\right\} \leq W_m(n) \leq \sqrt{m}\big(\lceil n^{1/m} \rceil - 1\big). \tag{E.3}$$

Hence $W_m(n) = \Theta_m(n^{1/m})$ for fixed m, and $W_m(n) = 1$ for $m \geq n - 1$.

Proof. For any configuration with minimum distance one and diameter Δ, the same volume argument gives $n \leq (2\Delta + 1)^m$. Also $\Delta \geq 1$. Therefore every such configuration has distance ratio at least $\max\{1, (n^{1/m} - 1)/2\}$; taking the infimum and using $W_m(n) = a_m(n)$ gives the lower bound.

For the upper bound, let $q = \lceil n^{1/m} \rceil$ and choose any n points from $\{0, \ldots, q-1\}^m$. Their minimum separation is at least one and their diameter is at most $\sqrt{m}(q - 1)$, giving the upper bound on $a_m(n)$ and hence on $W_m(n)$. For $m \geq n - 1$, a simplex has distance ratio one, which is minimal. □

These are unrestricted-action worst-case results. Their attaining family contains all permutations as primitive actions and grows with n. They do not characterize worst-case scaling under a fixed action budget. By contrast, Proposition 9 applies directly to the primitive-action paths of a given system, including the four-action mazes, without adding composed maps as new primitive actions.

F Restriction to Invariant Subsystems

Proposition 17 (Restriction principle)**.** *Let $Y \subseteq \mathcal{X}$ with $|Y| \geq 2$ satisfy $T_a(Y) \subseteq Y$ for every $a \in \mathcal{A}$. Then $d_{\mathrm{PR}}(T; L) \geq d_{\mathrm{PR}}(T|_Y; L)$.*

Proof. Every injective realization of the full system restricts to an injective realization of the invariant subsystem Y. Restriction cannot increase the required gain. Hence every ambient dimension feasible for the full system is also feasible for the restricted system, proving the inequality. □

G Additional Experimental Details

G.1 Controlled systems and maze construction

The existing controlled runs use $n \in \{3, 4, 5, 8, 16\}$ and $c = (1\,2\,\ldots\,n)$, with primitive actions $\{c, c^{-1}\}$ or $\{c, (1\,2)\}$. These check the cyclic/polygon variant with dimension two and the simplex family with dimension $n-1$, justified by Theorem 13 and Proposition 12. Their construction gains are verified by float64 enumeration. They are not numerical runs of the clipped-chain family in Corollary 10, whose one-dimensional realization is established analytically.

Each maze contains all cells $(r, c) \in \{0, \ldots, 7\}^2$, with state index $8r + c$ and actions U, D, L, R. A vertical wall (V, r, c) blocks movement between (r, c) and $(r, c+1)$ in both directions. A horizontal wall (H, r, c) blocks movement between (r, c) and $(r+1, c)$ in both directions. Blocked moves and boundary collisions leave the state unchanged. Table A1 specifies the reference systems; the open-grid control has the same boundaries and no internal walls.

Table A1: Reference wall sets. Each listed (r, c) is paired with the wall orientation in the second column.

Maze	Type	Coordinates (r, c)
A	V	$(1,3)$, $(2,3)$, $(3,3)$
A	H	$(4,4)$, $(4,5)$, $(4,6)$
B	V	$(0,1)$, $(1,1)$, $(1,2)$, $(1,4)$, $(2,1)$, $(3,1)$, $(3,5)$, $(4,1)$, $(4,5)$, $(5,5)$, $(6,5)$, $(7,4)$
B	H	$(0,1)$, $(0,2)$, $(1,3)$, $(1,4)$, $(1,5)$, $(2,3)$, $(4,5)$, $(4,7)$, $(5,1)$, $(5,2)$, $(5,3)$, $(6,5)$
C	V	$(0,1)$, $(0,2)$, $(0,3)$, $(0,4)$, $(0,5)$, $(0,6)$, $(1,1)$, $(2,0)$, $(2,5)$, $(3,0)$, $(3,6)$, $(4,1)$, $(4,2)$, $(4,3)$, $(4,5)$, $(4,6)$, $(5,0)$, $(5,6)$, $(6,3)$, $(6,4)$, $(6,5)$, $(6,6)$, $(7,1)$, $(7,6)$
C	H	$(0,0)$, $(1,1)$, $(1,4)$, $(1,5)$, $(1,7)$, $(2,5)$, $(3,1)$, $(3,4)$, $(3,5)$, $(4,5)$, $(5,1)$, $(5,3)$, $(5,4)$, $(6,1)$, $(6,3)$, $(6,5)$, $(6,6)$

G.2 Certificate computation

The directed pair graph has 2016 vertices and at most four outgoing edges per vertex. Tarjan's algorithm identifies its SCCs. For each SCC, we form the graph on original states whose edges belong to that type and solve maximum clique using Bron–Kerbosch search with pivoting, bitsets, and valid cardinality pruning. Independent Kosaraju and NetworkX checks verify the results. For each saved witness, shortest paths are computed in the full directed pair graph, not its restriction to witness pairs; a second implementation agrees on $h_T(Y^\star)$.

Table A2: Reference certificate structure. $b_1(T)$ is the exact optimum within the single-type family; $h_T(Y^\star)$ belongs to the saved witness.

Maze	Walls	Pair SCCs	Largest SCC	$\lvert Y^\star\rvert$	$h_T(Y^\star)$	$b_1(T)$
A	6	40	1476	30	21	29
B	24	48	1867	48	32	47
C	41	52	1958	53	30	52

The relaxed-gain bounds use these three saved witnesses, not an exhaustive search for optimal robust certificates. Transition tables, witnesses, and audit records are retained with the computational artifacts.

For n states and finitely many actions, the pair graph has $\binom{n}{2}$ vertices and at most $|\mathcal{A}|\binom{n}{2}$ edges. SCC decomposition is linear in this graph size. The search for the strongest single-type bound uses maximum clique in the associated graphs on original states; maximum clique is NP-hard in general (**?**). Given a candidate Y and computed SCC labels, checking all $\binom{|Y|}{2}$ labels verifies its single-type property without optimizing its size. Breadth-first search from each witness pair computes $h_T(Y)$ in the full graph. Heuristic search can therefore provide valid but possibly weaker witnesses; the quadratic pair-graph size remains a scaling limitation. We make no polynomial-time guarantee for finding the optimal certificate.

G.3 LOSSLESS STATE RE-ENCODING AUDIT

The open grid and Mazes A/B/C use identical state ordering, one centered 63-dimensional simplex Z_{s}, one centered two-dimensional grid codebook Z_{t}, and one matrix W. Because Z_{s} has full column rank and spans the zero-sum subspace, the identity is satisfied exactly by $W = (Z_{\mathrm{s}}^{\top} Z_{\mathrm{s}})^{-1} Z_{\mathrm{s}}^{\top} Z_{\mathrm{t}}$. This formula simplifies to $Z_{\mathrm{s}}^{\top} Z_{\mathrm{t}}$ only when the source columns are orthonormal. We verify $Z_{\mathrm{s}} W = Z_{\mathrm{t}}$ with maximum coordinate residual 3.89×10^{-15}. The minimum target separation is 0.3086066999. Inverse codebook matching uses only the compressed coordinate and recovers 64/64 identities before and after shuffling the queries; the source round-trip Euclidean error is zero. For each system, state, and action, re-encoding the source successor agrees with the target successor. All 1024 checks have discrepancy at most 4.05×10^{-15}. This is a finite-codebook identity and transition check, not a learned decoder generalization test.

Across two codebooks and four systems, we save all 64512 distance ratios. Scalar and vectorized implementations differ by at most 6.67×10^{-16} per ratio, and no denominator is clamped. A simplex has equal nonzero distances, so each pair ratio is zero or one. Open-grid clipped shifts are non-expansive. For a walled-grid action, each state's displacement is either zero or the same unit vector. The difference of displacements has norm at most one; an original grid distance $d \geq 1$ therefore becomes at most $d + 1 \leq 2d$. Gain two is attained under action U by predecessor pairs $(4, 4), (5, 4)$ in A, $(4, 5), (5, 5)$ in B, and $(1, 1), (2, 1)$ in C. Their raw distances change from one to two. Uniform normalization leaves these ratios unchanged. Excesses of order 10^{-15} are roundoff.

Saved witnesses are rechecked using independent SCC and shortest-path implementations, yielding $(k, h) = (30, 21), (48, 32), (53, 30)$. For $L = 1$ and $L = 101/100$, exact rational arithmetic evaluates $(k-1)/(1+\frac{k}{4}(L^h - L^{-h})^2)$ and its ceiling. No new maximum-clique search or witness optimization is performed. The source/target matrices, compression map, per-state checks, all ratios, input hashes, tolerances, software versions, and reproduction commands are retained in the audit artifacts. The inverse is defined on the finite target codebook; no linearity, Lipschitz bound, or off-codebook behavior is assumed for the decoder.

G.4 COORDINATE OPTIMIZATION AND GRID BASELINES

We optimize $Z \in \mathbb{R}^{64 \times m}$ directly and evaluate all 8064 primitive-action/state-pair ratios. Coincident successors contribute zero. The original sweep uses float64 Adam with learning rate 0.05 and a log-mean-exp surrogate of log distance ratios, with temperatures $0.10, 0.03, 0.01, 0.003$ and 2500 steps per stage. Coordinates are centered and normalized to unit global RMS after each update. The hard gain is checked every 100 steps, retaining the best evaluated coordinates.

The dimensions are 2, 3, 4, 5, 6, 8, 12, 16, 20, 24, 28, 29, 30, 40, 46, 47, 48, 51, 52, 53, and 63. Ten random restarts use seeds 0–9 at each maze/dimension. Continuation includes exact zero-padding of the retained lower-dimensional candidate and optimization after Gaussian perturbations of standard deviation 10^{-3}. The perturbation seed is $100000 + 1000j + m$, where $j = 0, 1, 2$ for A/B/C.

The candidate set also includes the original row–column coordinates, their zero-padded versions, and targeted grid-initialized optimization at $m = 2$. Raw grid gain is 2 for each walled maze; the three added optimized gains are 1.62723736512, 1.68567668223, and 1.72586806603 for A/B/C. Only the two-dimensional best values change. Enumeration verifies open-grid gain one up to floating-point precision. The $m = 63$ simplex supplies the analytic gain-one endpoint.

Evaluation and aggregation. We use the hard gain, not the surrogate, to compare candidates; minimum separation below 10^{-12} is flagged as degenerate. The reported best construction is the minimum over the retained candidate union, while random medians and interquartile ranges use random restarts only. Grid-based, continuation, and simplex candidates are excluded from those random statistics. Derived figures use the augmented candidates without rewriting historical raw results.

G.5 SECONDARY GEOMETRIC CONTROLS

At each sweep dimension, Gaussian and randomly projected-simplex baselines use 50 seeds per maze and method, without optimization. Both are centered and globally RMS-normalized and evaluated

with the same hard-gain routine. At $m = 29$, Gaussian medians are $2.0149/1.9417/1.9236$ and projected-simplex medians are $1.7670/1.7215/1.7085$ for A/B/C. No superiority over optimized transition-independent methods is established.

Fixed-wall study. With master seed zero, we draw 24 of the 112 possible internal walls uniformly and reject disconnected grids, accepting 100 systems after 166 attempts. Every accepted system has 80 self-loop state–action transitions (fraction 0.3125). Exact single-type certificates range from 44 to 59, with quartiles 50, 52, and 53. Samples 15 and 17 have equal graph diameter 16 but certificates 44 and 59. This pair was selected by maximizing the certificate gap among equal-diameter samples. The comparison shows that these coarse statistics do not determine $b_1(T)$; it does not establish different exact realization dimensions.

G.6 Exploratory analysis of trained PushT representations

We evaluate one existing LeWorldModel checkpoint per width (Maes et al., 2026), without model retraining. A common trajectory-disjoint split provides 10000/5000/5000 probe observations (train/validation/test; split seed 4101, base gain-protocol seed 4102). Whether these trajectories were unseen in world-model training is unknown.

For 50000 common observation pairs, both inputs receive an identical block of five raw actions. The sampled ratio is $r_{\text{pred}} = \|f_a(z_i) - f_a(z_j)\|_2 / \|z_i - z_j\|_2$, using the trained position-0 single-token slice; action preprocessing and positional semantics are verified. The 10^{-6} denominator exclusion removes no pairs. Separately, prediction error uses native three-token history and observed actions, reported as latent RMSE divided by target latent RMS.

Frozen-encoder ridge probes decode pusher position/velocity, block position, and orientation (sine/cosine). Validation selects regularization from $\{10^{-6}, 10^{-4}, 10^{-2}, 1, 100, 10000\}$; R^2 averages eight components. No nonlinear decoder is evaluated. NN error is standardized physical-state discrepancy for latent nearest neighbors from the validation pool. Centered covariance gives entropy effective rank $\exp(-\sum_i p_i \log p_i)$, where p_i are variance fractions.

Table A3: Existing PushT checkpoints. Gain columns describe sampled single-token predictor ratios, not global Lipschitz constants. Prediction error uses native history. Each width has one checkpoint.

m	$Q_{0.99}$	$Q_{0.999}$	Max.	Norm. error	Probe R^2	NN error	Eff. rank
8	1.2527	1.4306	1.8682	0.1014	0.2600	0.3200	7.94
16	1.2044	1.3904	2.6319	0.0669	0.3210	0.2909	15.89
32	1.1319	1.3211	3.8309	0.0676	0.2937	0.2933	31.65
64	1.0987	1.2688	3.6774	0.0691	0.5825	0.2902	62.53
192	1.0804	1.2364	3.1591	0.0694	0.7604	0.2912	118.60

Upper gain quantiles decrease, but sampled maxima do not; linear state decodability generally improves. This does not hold information fixed or establish dimension necessity, global gain reduction, or causality. Model size increases from 17.13M to 18.03M parameters, training seeds are unavailable, and checkpoint-linked success rates are unavailable.