CONDITIONING IS NOT CONTROL: SCIENTIFIC ENDPOINT CONTROLLABILITY FOR TENSOR-CONDITIONED GENERATIVE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Property-conditioned generators are increasingly used for scientific inverse design, but a conditioning variable need not be an action that controls the requested physical response. We formalize scientific endpoint controllability for tensor-valued targets by composing a generator with an independently specified scientific evaluator and quotienting physically equivalent $\mathrm{SO}(3)$ representatives. This formulation separates three notions that are often conflated: structural recoverability, conditioning-path visibility, and local endpoint reachability; direct target tracking further requires local inverse-action calibration. On a matched clamped-piezoelectric benchmark, tensor orbits are strongly recoverable from structure (CEITNet proper-orbit cosine 0.883), yet conditioning interventions reveal a distinct actuation geometry: scalar conditioning is locally rank one, whereas full-tensor routes expose roughly twelve numerically visible directions while concentrating most response energy in about two stable-rank directions. Common-noise swaps further show that a nonzero field response can remain nearly orthogonal to the requested correction. Independently, response-state analysis shows that endpoint specification materially changes the target, while DFPT separates many-mode lattice dynamics from the pointwise Born-charge-visible response channel. These results provide a mechanism-level diagnostic criterion for locating which stage limits endpoint control despite condition responsiveness.

1 INTRODUCTION

Scientific generative models are increasingly used for inverse design, including first-principles validation and reward-guided materials discovery (Zeni et al., 2025; Liu et al., 2026; Park & Walsh, 2026). At the same time, high-order material responses are becoming predictable from structure (Yan et al., 2024; Dong et al., 2025; Jin et al., 2026). Prediction and conditioning, however, do not by themselves specify what it means to *control* a scientific response. Existing work provides complementary ingredients through generative controllability, trajectory-level control, and symmetry-aware generative modeling (Cheng et al., 2026; Ben-Hamu et al., 2024; Wang et al., 2025; Chen & Lipman, 2024; Miller et al., 2024; Xu et al., 2026). We use crystal predictors and conditional generators as mechanistic probes of a common question: which stage limits scientific target recovery or realization?

In materials design, the generated structure is generally an intermediate state rather than the quantity requested by the user. Let $\mathcal{E}^q(X)$ denote the scientific response under endpoint specification q, and let π_q identify physically equivalent tensor representatives. We therefore study the composed map

$$F^q_{\theta,z}(a) = \pi_q(\mathcal{E}^q(G_\theta(z,a)))\,. \tag{1}$$

We call control of this map *scientific endpoint controllability*. Local reachability asks whether its differential spans the quotient tangent, whereas direct tracking asks the stronger question of whether supplying a requested target as the action returns that target. The controlled object here is the scientific response rather than the generated structure, and a conditioning variable is not assumed a priori to be an inverse action.

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.

Contributions. We make three contributions. (1) We define scientific endpoint controllability for tensor-conditioned generators as controllability of the composed map $\pi_q \circ E^q \circ G_\theta$, making endpoint specification and physical symmetry part of the control object. This yields separate criteria for quotient reachability and target-as-action tracking, and makes the criterion constructive through a local inverse-action rule when endpoint gradients are available. (2) We identify conditioning-path visibility as a distinct mechanism. A differential rank bound shows that the condition representation upper-bounds the locally visible actuation directions; parameter-matched interventions then measure the resulting singular-value geometry in trained routes. (3) We connect learned actuation to endpoint physics. Paired response-state data and DFPT factorization characterize endpoint-side physical structure, while controlled endpoint systems isolate the differential conditions for reachability and tracking.

Piezoelectric tensors provide a stringent test case: the target is an $\mathrm{SO}(3)$ orbit and the endpoint combines electronic and lattice-mediated response. We distinguish response state $s \in \{\mathrm{clamped}, \mathrm{ionic}, \mathrm{relaxed}\}$, convention c, the learned field response, and the independently evaluated endpoint; application-specific structure enters through E^q, π_q, and $q = (s, c)$.

Forward prediction diagnoses recoverability from structure; conditional generation asks the stronger question of whether the learned action path can move the scientific endpoint in requested directions. Route interventions measure $S_T = D_a G_\theta$, while endpoint control additionally depends on $\bar{K}_T^q = P_{\mathrm{out}} D_X E^q S_T P_{\mathrm{in}}$ (Fig. 1).

1.1 Related work

Generative control, guidance, and inverse design. Equivariant crystal generators treat periodic coordinates and lattices jointly (Xie et al., 2022; Jiao et al., 2023); conditional and flow-based systems expose property interfaces at increasing scale, including property-guided inorganic design and downstream validation (Ye et al., 2024; Levy et al., 2025; Kelvinius et al., 2025; Miller et al., 2024; Luo et al., 2025; Zeni et al., 2025; Liu et al., 2026). Flow matching and score-based diffusion provide general trajectory and guidance mechanisms (Lipman et al., 2023; Song et al., 2021), while classifier guidance, posterior sampling, plug-and-play priors, and universal guidance steer trajectories toward auxiliary information or differentiable constraints (Dhariwal & Nichol, 2021; Ho & Salimans, 2022; Chung et al., 2023; Graikos et al., 2022; Bansal et al., 2023). D-Flow and OC-Flow differentiate through a flow or pose guidance as an optimal-control problem (Ben-Hamu et al., 2024; Wang et al., 2025); GenCtrl formalizes reachable and controllable sets for generic generative systems (Cheng et al., 2026). We study the missing scientific composition: whether an endpoint-specific tensor response after a downstream evaluator is locally reachable, and whether an injected condition is itself an inverse action.

Symmetry-aware generation and tensor prediction. Riemannian and general-geometry flow models place generated states or their dynamics on geometric spaces, while quotient-space diffusion removes symmetry redundancy in a generated state or distribution (Chen & Lipman, 2024; Miller et al., 2024; Xu et al., 2026). The same symmetry principles constrain tensor-valued predictors: foundational equivariant architectures formalize how features and output tensors transform (Cohen & Welling, 2016; Kondor & Trivedi, 2018; Thomas et al., 2018), and modern E(3)-equivariant networks make these constraints explicit at scale, including higher-degree representations (Bronstein et al., 2021; Geiger & Smidt, 2022; Batzner et al., 2022; Liao et al., 2024). Tensor Field Networks and E(n)-equivariant graph networks provide the geometric machinery for tensor prediction (Thomas et al., 2018; Satorras et al., 2021); GMTNet, EATGNN, CEITNet, and MatTen establish that high-order material responses can be predicted from structure (Yan et al., 2024; Dong et al., 2025; Jin et al., 2026; Wen et al., 2024), with recent extensions improving canonicalized and strain-conditioned tensor recovery (Hua et al., 2026; Yu et al., 2026). Strain-conditioned response models further expose the difference between a static label and its physical pathway (Yu et al., 2026). These approaches address symmetry or forward recoverability; our quotient acts on the controlled scientific response target and therefore enters the controllability criterion.

Identifiability, selection, and intervention. Multiple Choice Learning motivates diverse hypotheses for downstream selection (Lee et al., 2016), but a target-independent selector cannot choose a branch using an unobserved response. Crystal tensors are therefore compared through representatives,

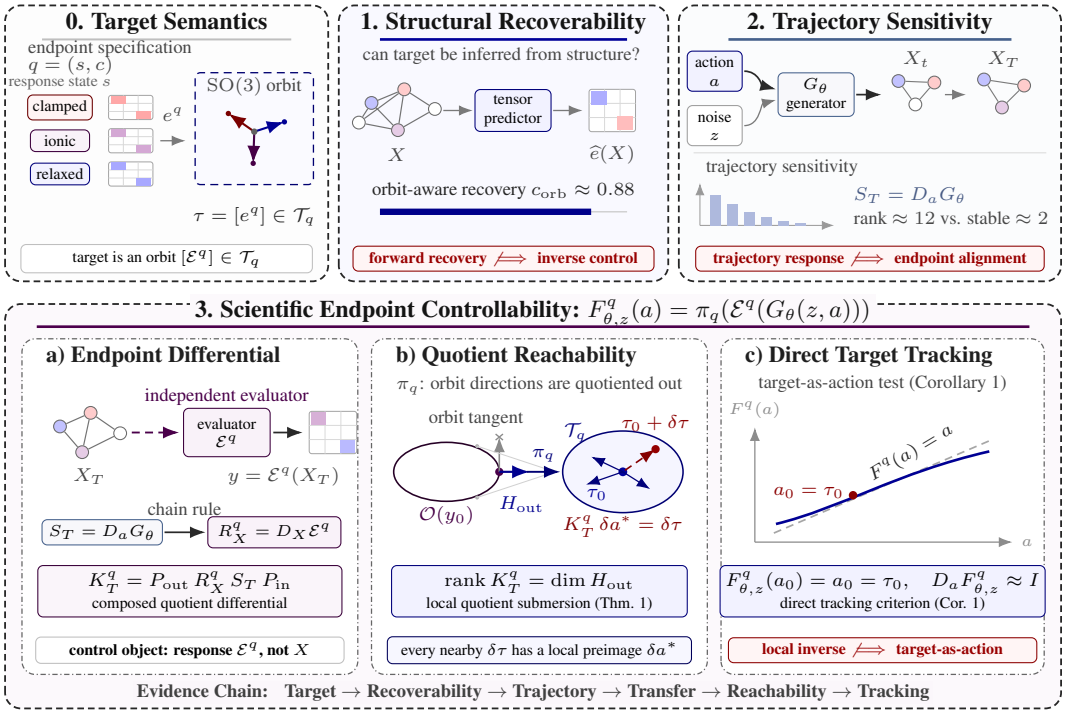


Figure 1: **Scientific endpoint controllability requires a chain of distinct control criteria.** A scientific tensor target is defined by its endpoint specification $q = (s, c)$ and symmetry quotient. Structural recoverability tests whether the target information is inferable from the generated structure, while the trained action path measures local trajectory visibility through $S_T = D_a G_\theta$. Physical endpoint control additionally requires transfer through an independently defined scientific evaluator E^q, yielding the quotient endpoint differential $K^q_T = P_{\text{out}} D_X E^q S_T P_{\text{in}}$. Full row rank guarantees local quotient reachability; direct target tracking further requires basepoint agreement and a near-identity action-to-endpoint map.

axial classes, or proper-rotation orbits; canonicalization can simplify a selected frame but does not guarantee preservation of the deployment target (Hua et al., 2026). Causal representation learning and inverse-materials design likewise emphasize the distinction between association, intervention, and verified property realization (Saengkyongam et al., 2024; Morioka & Hyvarinen, 2024; Schölkopf et al., 2021; Park et al., 2024). First-principles calculations remain an independent endpoint (Choudhary et al., 2020b). JARVIS provides the source infrastructure used here (Choudhary et al., 2020a).

2 Method

2.1 Scientific target, endpoint specification, and symmetry

Scientific endpoint control object. A crystal state is $\mathbf{x} = (a, f, L)$, with species $a = (a_i)_{i=1}^N$, fractional coordinates $f \in [0, 1)^{N\times 3}$, and a lattice matrix L with positive volume. Cartesian positions are $r = fL$. Physical identity is unchanged by periodic translation, site permutation, integer basis changes, and global proper rotations. With P acting on the ordered site list, the joint quotient relation is

$$(a, f, L) \sim (Pa,\ PfB^{-1} + \mathbf{1}t^\top,\ BLR^\top), \qquad B \in \mathrm{GL}(3, \mathbb{Z}),\ R \in \mathrm{SO}(3),\ t \in \mathbb{T}^3. \tag{2}$$

The regular-stratum construction follows standard smooth group-action and quotient-manifold geometry: orbit directions are removed, while a horizontal complement represents the quotient tangent locally (Lee, 2013; Absil et al., 2008); the same quotient view of generation underlies recent geometric and quotient-space flow models (Chen & Lipman, 2024; Xu et al., 2026). The piezoelectric

stress tensor has $e_{ijk} = e_{ikj}$ and is represented in an orthonormal Kelvin form $E_K \in \mathbb{R}^{3\times 6}$. The response space is $V_e = V_{1a} \oplus V_{1b} \oplus V_2 \oplus V_3 \equiv 2V_1 + V_2 + V_3$ (18 dimensions). Its scalar magnitude coordinate is $e_{\max} = \sigma_{\max}(E_K)$, whereas the orientation-free tensor target is the proper-rotation orbit $[e] = \{\rho(R)e : R \in \mathrm{SO}(3)\}$. A lab-frame task would retain a representative instead of quotienting it. Conditional pairs transform jointly as $(X, e) \mapsto (RX, \rho(R)e)$. If G_X and H_e are the proper stabilizers of the crystal and tensor, the residual relative frame is identified only up to

$$[R] \in G_X \backslash \mathrm{SO}(3)/H_e. \tag{3}$$

For normalized tensor representatives, the orientation-free comparison used by the benchmark is

$$c_{\mathrm{orb}}(u, v) = \sup_{R\in \mathrm{SO}(3)} \langle u, \rho(R)v\rangle, \tag{4}$$

with the reported implementation using a canonicalized finite proper-rotation grid. Signed cosine retains representative-level information and axial cosine identifies antipodal representatives; both are therefore auxiliary metrics, while the proper-orbit cosine is the primary metric for the orientation-free target.

2.2 Local quotient reachability and direct tracking

Let $G_\theta(z, a)$ be a fixed-noise conditional generator and let $\mathcal{E}^q(X)$ be the scientific endpoint under endpoint specification q. The controlled map is the composed endpoint map $F^q_{\theta,z}$ in Eq. (1), where π_q maps a response representative to its quotient class. Let $\mathcal{Y}_q$ denote the local representative space for the response under endpoint specification q, and let $\mathcal{T}_q = T_{[y_0]}(\mathcal{Y}_q/\sim_q)$ be the quotient target tangent at a regular point. The projector P_{out} is the orthogonal projection onto the horizontal complement of the output orbit tangent. P_{in} is the analogous projection for the action chart; all route experiments use a fixed local chart without input quotienting, so $P_{\mathrm{in}} = I$ there. Let a_0 be an action, $X_0 = G_\theta(z, a_0)$, and $y_0 = \mathcal{E}^q(X_0)$ in regular symmetry strata. Define

$$S_T = D_a G_\theta(z, a_0), \qquad \bar{K}^q_T = P_{\mathrm{out}} D_X \mathcal{E}^q(X_0) S_T P_{\mathrm{in}}. \tag{5}$$

Theorem 1 (Local controllability on a regular tensor-orbit quotient)**.** *Assume that the fixed-noise action map and $\mathcal{E}^q$ are continuously differentiable and equivariant, and that the input and output actions have locally constant orbit type. Then $F^q_{\theta,z}$ descends locally to the regular quotient strata with differential $\bar{K}^q_T$. If $\mathrm{rank}(\bar{K}^q_T) = \dim H_{\mathrm{out}}$, the induced endpoint map is a submersion and every sufficiently small quotient-tangent perturbation has a local preimage. If the rank is smaller, a nonzero cokernel covector gives an infinitesimally unreachable quotient perturbation.*

For direct target-as-action analysis, we use matched action and target charts with the same local coordinates and identify the basepoint coordinates as $a_0 = \tau_0$.

Corollary 1 (Direct target tracking requires more than reachability)**.** *In matched local charts, nonsingularity of $\bar{K}^q_T$ gives a local inverse, but direct tracking additionally requires basepoint agreement $F^q_{\theta,z}(a_0) = a_0 = \tau_0$. If $\|D_a F^q_{\theta,z} - I\|_{\mathrm{op}} \le \varepsilon < 1$ on a convex neighborhood, then*

$$\|F^q_{\theta,z}(a) - a\| \le \|F^q_{\theta,z}(a_0) - a_0\| + \varepsilon\|a - a_0\|. \tag{6}$$

Thus full rank permits a locally computed inverse action, whereas feeding the requested target itself is guaranteed only under basepoint and near-identity tracking conditions.

Local control synthesis. When $\bar{K}^q_T$ has full row rank, the minimum-norm first-order action for a quotient-tangent displacement $\delta\tau$ is $\delta a^* = (\bar{K}^q_T)^\top \big(\bar{K}^q_T(\bar{K}^q_T)^\top\big)^{-1}\delta\tau$. For numerical regularization we use the damped form

$$\delta a_\lambda = (\bar{K}^q_T)^\top \big(\bar{K}^q_T(\bar{K}^q_T)^\top + \lambda I\big)^{-1}\delta\tau, \qquad \lambda > 0, \tag{7}$$

which trades exactness of the first-order preimage for stability. Thus the same endpoint differential used for diagnosis also defines a local controller once an independent evaluator Jacobian is available. Direct target-as-action conditioning is the special calibrated case in which this inverse is locally represented by the identity action chart.

The argument is local and slice-based. On a regular orbit-type stratum, the slice theorem supplies local slices transverse to the input and output group orbits, whose tangent spaces are the horizontal complements selected by P_{in} and P_{out}. Equivariance of G_θ and $\mathcal{E}^q$ makes the composed map descend to these slices, and the chain rule there gives exactly $\bar{K}_T^q$ as the differential of the descended map. Full row rank is then the hypothesis of the submersion theorem, yielding a local preimage for every sufficiently small quotient-tangent perturbation; if the rank is deficient, finite-dimensional duality supplies a nonzero cokernel covector that annihilates every first-order quotient response.

An empirical learned Jacobian constitutes endpoint evidence only when $D_X\mathcal{E}^q$ is available from an independent evaluator. The complete prediction and selector bounds, together with the symmetry-conditioned irrep rank bound, are given in Supplementary Sections B.3 and B.4.

This rank condition is the local reachability statement for the composed input–output map studied in recent generative-control formalisms (Cheng et al., 2026; Ben-Hamu et al., 2024; Wang et al., 2025); here the output is the endpoint-specific quotient response rather than the generated state itself.

2.3 Conditioning-path visibility

Let the condition enter the field through a collection of maps $\Psi(e) = (\psi_1(e), \dots, \psi_L(e))$ and write $v_\theta(X, t; e) = \mathcal{V}_\theta(X, t, \Psi(e))$. The chain rule gives the central visibility relation

$$D_e v_\theta = D_\Psi \mathcal{V}_\theta \, D_e \Psi, \qquad \operatorname{rank}(D_e v_\theta) \le \operatorname{rank}(D_e \Psi). \tag{8}$$

Equation (8) makes the representation bottleneck explicit: the local action space accessible through conditioning cannot have rank greater than the differential of the condition representation. A scalar summary therefore exposes at most one conditioning-visible direction at any active point, irrespective of backbone depth or width; a rotation-invariant summary additionally removes orbit-tangent directions. Full-tensor conditioning removes this representation-level rank bottleneck, but does not imply an isotropic response spectrum, endpoint reachability, or target alignment. The full proof and the symmetry-conditioned irrep bound are given in Supplementary Section B.4.

2.4 Evaluation protocol and metrics

For a local Jacobian with singular values $s_1 \ge s_2 \ge \cdots$, we report

$$r_{1\%} = \#\{i : s_i \ge 0.01 s_1\}, \qquad r_{\mathrm{stable}} = \frac{\sum_i s_i^2}{s_1^2}. \tag{9}$$

The first quantity counts numerically visible directions at the predefined 1% threshold; stable rank measures how broadly response energy is distributed and therefore need not track numerical rank. A record whose Jacobian vanishes identically has $r_{1\%} = r_{\mathrm{stable}} = 0$ by convention and enters population averages at zero. In the common-noise intervention, structure and noise context are held fixed while only the active condition is changed; the induced field response and its overlap with the requested target direction are then measured separately.

2.5 Benchmarks and evaluation panels

The primary case study uses the TensorOrbit–JARVIS source population of 3,889 crystals from 2,726 formula families, with a fixed 3165/363/361 train/development/test split, a maximum of 20 atoms, and zero family leakage. The formal matched-orbit comparison uses an independent 1,728-record panel with a fixed $1536/192/192$ train/development/test split and three model seeds. A common offline evaluator recomputes all material-level predictions from Cartesian prediction/target arrays. A separate 5,000-record curated source supplies complete clamped/ionic/relaxed (C/I/R) targets for response-semantics analysis and is never pooled with the primary prediction panel. Finite-grid proper-orbit cosine is the primary prediction metric; signed and axial cosine are auxiliary metrics. We use Q_{25} for the empirical 25th percentile. For the common-noise intervention, each Q_{25} lower bound is the 5th percentile of 2,000 record-level bootstrap replicates, so the reported interval is a percentile 95% lower bound. Paired model differences use material-level bootstrap resampling with 95% percentile intervals. Exact denominators and dataset definitions are given in Supplementary Sections A and A.2.

Table 1: Matched full-tensor prediction benchmark on the fixed clamped-piezoelectric panel. Values are mean±sample standard deviation over three model seeds, except the deterministic all-zero baseline. The finite-grid proper-orbit score is the primary orientation-free measure; axial and signed scores diagnose antipodal and representative information. Component errors are strongly influenced by the mass of near-zero responses; directional metrics are therefore evaluated on nonzero tensors and report the recoverable orientation information separately. Component MAE and RMSE are reported in $\mathrm{C\,m^{-2}}$.

method	finite-grid proper-orbit cosine ↑	axial cosine ↑	signed cosine ↑	component MAE ↓	RMSE ↓
All-zero baseline	—	—	—	0.0947	0.4644
GMTNet	0.6066 ± 0.0099	0.2512 ± 0.0295	0.0870 ± 0.0289	0.1032 ± 0.0066	0.4661 ± 0.0039
EATGNN	0.8031 ± 0.0160	0.7348 ± 0.0322	0.4665 ± 0.0905	0.0787 ± 0.0025	0.4439 ± 0.0061
CEITNet	0.8833 ± 0.0063	0.8476 ± 0.0089	0.6147 ± 0.0026	0.0656 ± 0.0021	0.4431 ± 0.0097

3 Results

3.1 Target semantics

Response state is part of the endpoint specification. Across the fixed 5,000-record clamped/ionic/relaxed (C/I/R) population, changing from clamped to relaxed response activates 1,928 previously zero tensors, reverses at least one representative component in 2,007 records, and changes the dominant irrep sector in 1,680 records. Among the 3,071 nonzero pairs, median signed and finite-grid proper-orbit cosines are 0.448 and 0.599, while the median tensor norm rises from 0.0361 to 0.1011. The response state therefore changes activity, representative direction, irreducible content, and quotient-level similarity.

The response convention is scientifically consequential: polarization and piezoelectric response in periodic solids depend on the chosen branch and proper-response definition (King-Smith & Vanderbilt, 1993; Resta, 1994; Vanderbilt, 2000; Wu et al., 2005).

3.2 Structural recoverability

The matched panel asks whether tensor information is present in the visible structure before any claim about generative actuation. The finite-grid proper-orbit score is primary because the design target is orientation-free; axial and signed scores indicate antipodal and representative information.

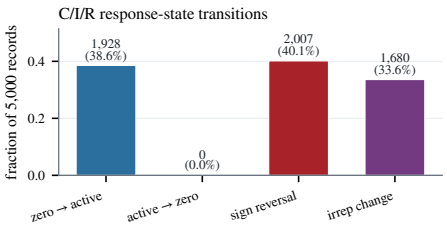


Figure 2: C/I/R transitions in the 5,000-record curated source.

The matched comparison shows that directional recovery varies substantially across evaluated model families under a common data and evaluation protocol. While this panel reflects standardized training recipes rather than isolated architectural ablations, the higher-performing models demonstrate that the target orbit is substantially recoverable from the visible structure. A paired material bootstrap over the common test IDs gives orbit improvements of $+0.1965$ [0.1656, 0.2251] for EATGNN over GMTNet and $+0.0802$ [0.0564, 0.1059] for CEITNet over EATGNN (95% percentile intervals for this matched piezoelectric panel).

3.3 Conditioned trajectories and actuation geometry

To isolate the capacity of the conditioning interface from backbone variation, we benchmark three parameter-matched routing architectures (scalar/invariant, Flat18-FiLM, and OperatorSpan, each with approximately 2.25k route parameters) on an identical structural base, probing each route on 32 test materials under four fixed noise conditions. Jacobian–vector products are validated numerically against finite-difference directional derivatives (details in Supplementary Section D). Numerical rank uses the 1% threshold in Eq. 9.

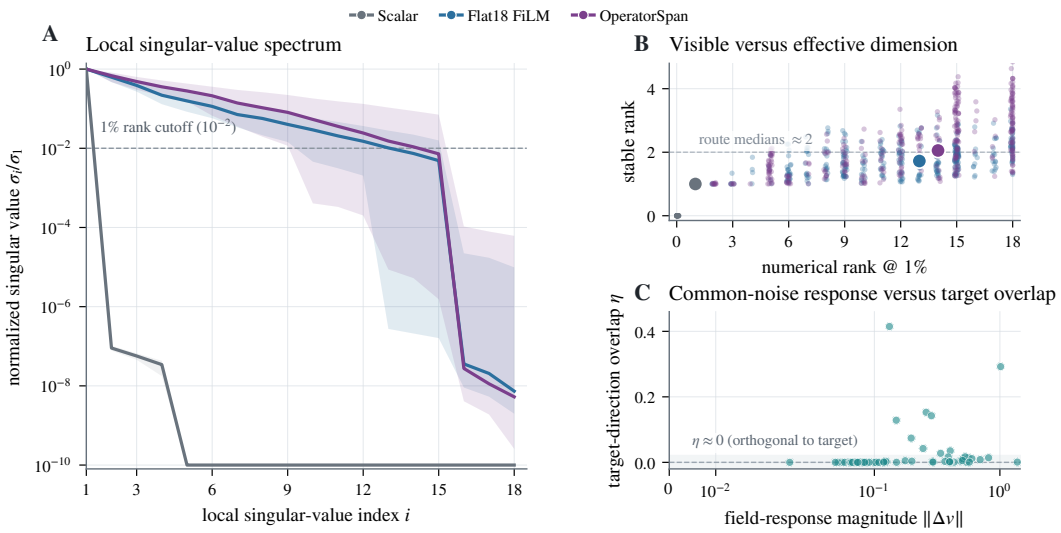


Figure 3: Trained conditioning routes expose different local trajectory geometries. (A) Median normalized singular-value spectra of the trained trajectory Jacobian $S_T = D_a G_\theta$ over material–noise records, with interquartile bands. (B) Each point is one material–noise trajectory record; large markers show route medians. (C) A separate common-noise residual-control analysis shows conditioning-induced field-response magnitude against the target-direction overlap η.

Conditioning visibility. Under a common structural backbone and approximately matched route capacity, scalar conditioning is rank one at every active point. Full-tensor Flat18-FiLM and OperatorSpan increase the numerical trajectory rank to roughly twelve, but their stable ranks remain near two. Full-tensor routing therefore opens many locally visible directions while concentrating most response energy in a much smaller dominant subspace (Fig. 3A–B). Across the three training runs, scalar $e_{\max}$ conditioning has numerical rank and stable rank exactly one wherever the local Jacobian is nonzero; records with exactly zero local response (36 of 384, concentrated in three of the 32 test materials) enter the mean at zero and reduce the across-record means to 0.90625. Full-tensor Flat18-FiLM reaches mean numerical ranks of 11.750–12.484 across runs, with stable ranks of 1.702–1.829. Raising the input dimensionality of the condition is therefore necessary to lift the rank bottleneck of the conditioning path, but insufficient to yield isotropic trajectory control. The distribution-level comparison is shown in Fig. 3 and Supplementary Fig. S2; the denoising losses of these routes are given as reproducibility context in Supplementary Section G.

In paired common-noise swaps, changing only the tensor condition induces a field response measured against the matched-condition baseline. Here field effect is the sum of graph RMS changes in the coordinate score and element logits and the Frobenius change in the symmetric strain-rate output. Target margin is $L_{\text{geom}}(\text{wrong condition}) - L_{\text{geom}}(\text{matched condition})$, where L_{geom} is the mean of the coordinate and six-component strain geometric losses with unit channel scales.

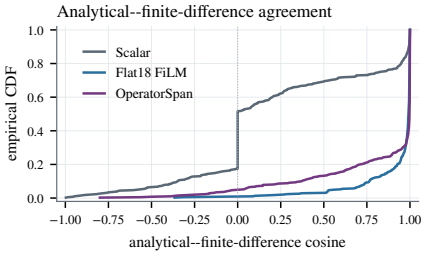


Figure 4: Analytical–finite-difference agreement in the route evaluation.

Target alignment. Common-noise swaps change only the active tensor condition while preserving the structural and noise context. The induced field response is nonzero ($Q_{25} = 0.0359$), whereas the target-margin lower quartile is essentially zero ($Q_{25} = 0.0000$); the direct projection statistic η likewise concentrates near zero (Fig. 3C). Conditioning sensitivity and target-directed actuation are therefore empirically distinct. The structural prior, noise, and zero-response fraction (0.688) are held fixed.

For Fig. 3, we additionally report the target-projection statistic

$$\eta = \frac{\|P_{\mathrm{Im}(A_\theta)} v_{\mathrm{target}}\|^2}{\|v_{\mathrm{target}}\|^2 + 10^{-12}}, \tag{10}$$

where A_θ is the residual Jacobian with respect to the active tensor condition representation and $P_{\mathrm{Im}(A_\theta)} = A_\theta A_\theta^\dagger$ is the orthogonal projector onto its image. The term 10^{-12} is numerical stabilization only. This statistic is distinct from the residual-level target-margin statistic above.

The paired training data identify an observational conditional, whereas the post-training common-noise swaps intervene directly on the conditioning input. Symmetry leaves relative-frame choices through Eq. (3). We therefore interpret the resulting quantities as trajectory-geometry measurements and evaluate physical endpoint transfer separately.

All three OperatorSpan evaluations produce complete finite trajectories on the same 32 test materials with four fixed noise conditions per material (384 trajectories each; Supplementary Fig. S2). Across the three per-seed evaluations, the mean numerical rank at the 1% threshold ranges from 12.242 to 12.391 and the mean stable rank from 2.203 to 2.271; pooling the records gives a median numerical rank of 14.0, median stable rank of 2.0549, and median finite-difference cosine of 0.9970, with $282/384$ rows exceeding a cosine of 0.9. The common-noise analysis evaluates target alignment; endpoint-side physics is examined separately below.

3.4 Endpoint transfer and controlled evidence

Mechanistic decomposition. Equations (5) and (8) expose two distinct rank bottlenecks. The condition representation first restricts the image of S_T; the scientific evaluator then transfers that image through $D_X E^q$ before output-orbit directions are removed by P_{out}. Increasing condition dimensionality can therefore enlarge the trajectory manifold without increasing endpoint rank or target alignment. The route analysis characterizes S_T, whereas the DFPT analysis resolves pointwise physical structure within E^q. The control-relevant transfer between them is the structural differential $D_X E^q$, whose composition with S_T yields the endpoint controllability differential $\bar{K}_T^q$. This decomposition identifies distinct loci at which controllability can fail: at the conditioning interface, along the generated trajectory, or during physical endpoint transfer.

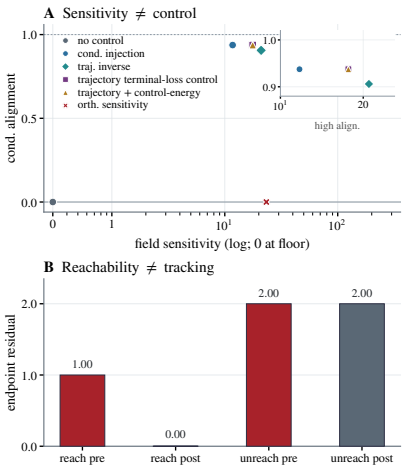


Figure 5: **Controlled endpoint analysis.** Sensitivity, reachability, and direct tracking are evaluated separately. The trajectory terminal-loss and control-energy variants are analytic controls inspired by the optimization principles of D-Flow and OC-Flow, rather than implementations of those methods.

Physical endpoint structure. For the relaxed-ion response, the standard DFPT factorization (Gonze & Lee, 1997; Baroni et al., 2001; Wu et al., 2005) gives $e^{\mathrm{ionic}} = (e_0/\Omega) A^\top B$. At each structure, only the projection of the internal-strain response onto the Born-charge-visible mode subspace $\mathrm{Col}(A)$ contributes, with $\mathrm{rank}\, A \leq 3$. Across 3,371 stable records, the factorization closes to numerical precision, and the projected response rank $r_{\mathrm{resp}} = \mathrm{rank}(P_A B)$ takes values $0/1/2/3$ with counts $97/72/481/2{,}721$ (Fig. 6), even though many more optical modes can participate. Structural endpoint sensitivity is governed separately by $D_X E^q$, which differentiates the structure dependence of A, B, and Ω (Supplementary Sec. E.1). Source-matched finite-displacement/Berry checks on GaN and ZnO agree with the factorized response within 1.5% and validate the local response channel; step-size and significance diagnostics are reported in Supplementary Sections E.2–E.3.

Controlled endpoint geometry. The analytic endpoint separates sensitivity from control directly. For the periodic polar tensor $E(X) = \sum_{i<j} w_{ij} r_{ij}^{\otimes 3}$ with exact automatic differentiation, an orthogonal sensitivity direction produces field effect 23.34 but essentially zero target alignment, whereas

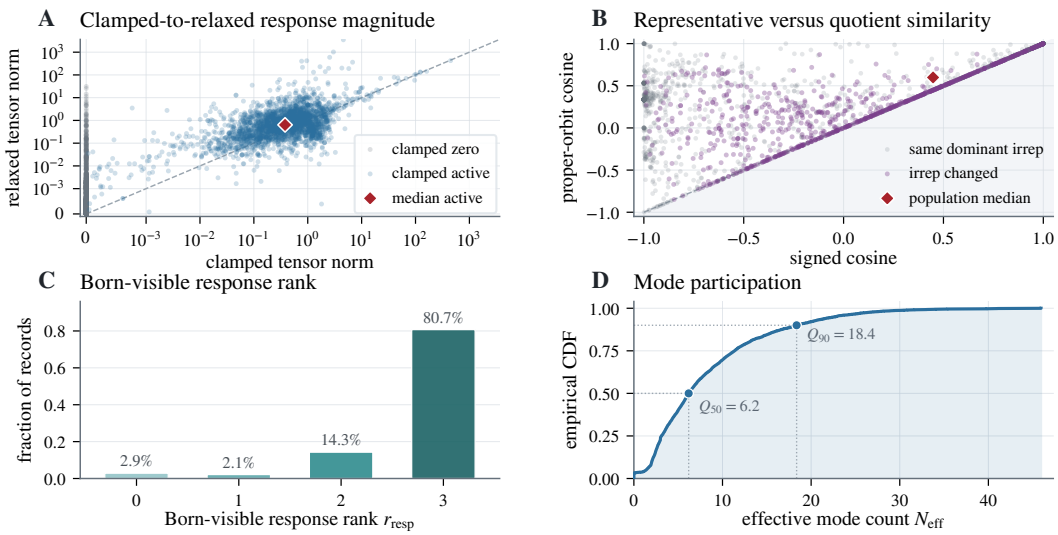


Figure 6: Endpoint response semantics and DFPT channel structure. (A) Paired clamped and relaxed tensor norms for the 5,000-record curated source. (B) Per-record signed versus proper-orbit cosine. (C) Fraction of the 3,371 stable DFPT records in each Born-visible response-rank stratum. (D) Empirical CDF of the effective-mode-count field, with one value per stable material record in the same DFPT population.

a damped local inverse reaches alignment 0.999 with endpoint error 0.0032. On 18 clean JARVIS geometries, response-aligned reachable targets achieve cosine 1.000 and relative error 1.9×10^{-5}, while orthogonal targets remain unreachable. These experiments instantiate the distinction in Sec. 2.2: reachability depends on the image of the endpoint Jacobian, whereas direct tracking requires an aligned inverse action. The controlled systems also instantiate the synthesis rule of Eq. (7) and show that the criterion is operational, rather than purely diagnostic.

4 Discussion

Viewed diagnostically, forward prediction and conditional generation expose different bottlenecks: target semantics and structural recoverability constrain inference; conditioning visibility and target alignment constrain learned actuation; endpoint-side physics enters through the structural differential $D_X E^q$. Piezoelectricity provides a stringent test case because its target is symmetry-quotiented and its relaxed response combines electronic and lattice-mediated channels. The same composition $\pi_q \circ E^q \circ G_\theta$ applies to other differentiable scientific endpoints with known physical equivalents, including elasticity and dielectric tensors, molecular observables, and discretized PDE responses. The clamped/ionic/relaxed decomposition and Born-visible factorization are specific to piezoelectricity.

The criterion remains local. For any differentiable tracked path $\tau(t) = F^q_{\theta,z}(a(t))$, $\dot\tau(t) = \bar K^q_T(a(t))\dot a(t)$, so every smoothly trackable target direction must lie in the image of the endpoint differential at the operating point. Rank deficiency therefore certifies missing instantaneous control directions: any target path whose tangent leaves the image of the endpoint differential cannot be tracked smoothly through that operating point. It does not exclude a nonlocal route that detours through other configurations, exploits nonlinear effects, or crosses symmetry strata. Global controllability additionally requires path connectivity and analysis across stratified transitions. Accordingly, a scientific control claim should specify the endpoint semantics and symmetry, the action-to-trajectory response, the locally reachable endpoint directions, and whether the supplied target is a calibrated inverse action; where possible, these criteria should be checked against an independent physical endpoint.

AI USE STATEMENT

We used generative AI tools for language editing, manuscript organization, and assistance with code and scripts for data processing and experimental workflows. The scientific questions, experimental protocols, analyses, interpretations, and final claims were determined and verified by the authors. All AI-assisted text and code were reviewed by the authors, who take full responsibility for the manuscript and accompanying artifacts.

REFERENCES

P.-A. Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. ISBN 9780691132983. URL https://press.princeton.edu/books/hardcover/9780691132983/optimization-algorithms-on-matrix-manifolds.

Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 843–852, 2023. doi: 10.1109/CVPRW59228.2023.00091. URL https://openaccess.thecvf.com/content/CVPR2023W/GCV/html/Bansal_Universal_Guidance_for_Diffusion_Models_CVPRW_2023_paper.html.

Stefano Baroni, Stefano de Gironcoli, Andrea Dal Corso, and Paolo Giannozzi. Phonons and related crystal properties from density-functional perturbation theory. *Reviews of Modern Physics*, 73(2):515–562, 2001. doi: 10.1103/RevModPhys.73.515. URL https://doi.org/10.1103/RevModPhys.73.515.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13:2453, 2022. doi: 10.1038/s41467-022-29939-5. URL https://doi.org/10.1038/s41467-022-29939-5.

Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-flow: Differentiating through flows for controlled generation. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 3462–3483. PMLR, 2024. URL https://proceedings.mlr.press/v235/ben-hamu24a.html.

Rabi N. Bhattacharya and Vic Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds. *The Annals of Statistics*, 31(1):1–29, 2003.

Glen E. Bredon. *Introduction to Compact Transformation Groups*, volume 46 of *Pure and Applied Mathematics*. Academic Press, 1972. ISBN 9780080873596.

Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021. doi: 10.48550/arXiv.2104.13478. URL https://arxiv.org/abs/2104.13478.

Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations*, 2024.

Emily Cheng, Carmen Amo Alonso, Federico Danieli, Arno Blaas, Luca Zappella, Pau Rodriguez, and Xavier Suau. GenCtrl: A formal controllability toolkit for generative models. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=HJTFgDYoLO.

Kamal Choudhary, Kevin F Garrity, Andrew CE Reid, Brian DeCost, Adam J Biacchi, Angela R Hight Walker, Zachary Trautt, Jason Hattrick-Simpers, A Gilad Kusne, Andrea Centrone, et al. The joint automated repository for various integrated simulations (jarvis) for data-driven materials design. *npj computational materials*, 6(1):173, 2020a. doi: 10.1038/s41524-020-00440-1. URL https://doi.org/10.1038/s41524-020-00440-1.

Kamal Choudhary, Kevin F. Garrity, Vinit Sharma, Adam J. Biacchi, Angela R. Hight Walker, and Francesca Tavazza. High-throughput density functional perturbation theory and machine learning predictions of infrared, piezoelectric and dielectric responses. *npj Computational Materials*, 6:64, 2020b. doi: 10.1038/s41524-020-0337-2.

Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=OnD9zGAGT0k.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2990–2999. PMLR, 2016. URL https://proceedings.mlr.press/v48/cohenc16.html.

Prafulla Dhariwal and Alexander Q. Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794, 2021.

Luqi Dong, Xuanlin Zhang, Ziduo Yang, Lei Shen, and Yunhao Lu. Accurate piezoelectric tensor prediction with equivariant attention tensor graph neural network. *npj Computational Materials*, 11(1), 2025. doi: 10.1038/s41524-025-01546-0.

Mario Geiger and Tess Smidt. e3nn: Euclidean neural networks. *arXiv preprint arXiv:2207.09453*, 2022. doi: 10.48550/arXiv.2207.09453. URL https://arxiv.org/abs/2207.09453.

Xavier Gonze and Changyol Lee. Dynamical matrices, born effective charges, dielectric permittivity tensors, and interatomic force constants from density-functional perturbation theory. *Physical Review B*, 55(16):10355–10368, 1997. doi: 10.1103/PhysRevB.55.10355. URL https://doi.org/10.1103/PhysRevB.55.10355.

Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors. In *Advances in Neural Information Processing Systems*, volume 35, pp. 14715–14728, 2022. doi: 10.52202/068431-1070. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/5e6cec2a9520708381fe520246018e8b-Abstract.html.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. doi: 10.48550/arXiv.2207.12598. URL https://arxiv.org/abs/2207.12598.

Haowei Hua, Jingwen Yang, Wanyu Lin, and Pan Zhou. Revisiting the canonicalization for fast and accurate crystal tensor property prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 417–425, 2026. doi: 10.1609/aaai.v40i1.37004. URL https://ojs.aaai.org/index.php/AAAI/article/view/37004.

Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. *Advances in Neural Information Processing Systems*, 36:17464–17497, 2023.

Dian Jin, Yancheng Yuan, and Xiaoming Tao. Efficient equivariant high-order crystal tensor prediction via cartesian local-environment many-body coupling. In *International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2602.04323. ICML 2026 Poster.

Filip Ekstrom Kelvinius, Oskar B. Andersson, Abhijith S. Parackal, Dong Qian, Rickard Armiento, and Fredrik Lindsten. Wyckoffdiff: A generative diffusion model for crystal symmetry. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267, 2025.

R. D. King-Smith and David Vanderbilt. Theory of polarization of crystalline solids. *Physical Review B*, 47(3): 1651–1654, 1993. doi: 10.1103/PhysRevB.47.1651. URL https://doi.org/10.1103/PhysRevB.47.1651.

Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2747–2755. PMLR, 2018. URL https://proceedings.mlr.press/v80/kondor18a.html.

John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 2013. doi: 10.1007/978-1-4419-9982-5. URL https://doi.org/10.1007/978-1-4419-9982-5.

Stefan Lee, Senthil Purushwalkam Shiva Prakash, Michael Cogswell, Viresh Ranjan, David Crandall, and Dhruv Batra. Stochastic multiple choice learning for training diverse deep ensembles. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper/2016/hash/20d135f0f28185b84a4cf7aa51f29500-Abstract.html.

Daniel Levy, Siba Smarak Panigrahi, Sekou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. Symmcd: Symmetry-preserving crystal generation with diffusion models. In *International Conference on Learning Representations*, 2025.

Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/ab12e8f3443c1a789f595b18d8c597b4-Abstract-Conference.html.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Zicheng Liu, Ben Fei, and Di Huang. PRO-MOF: Policy optimization with universal atomistic models for controllable MOF generation. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=BIzrFlp0hv.

Xiaoshan Luo, Zhenyu Wang, Qingchang Wang, Xuechen Shao, Jian Lv, Lei Wang, Yanchao Wang, and Yanming Ma. Crystalflow: A flow-based generative model for crystalline materials. *Nature Communications*, 16(1):9267, 2025. doi: 10.1038/s41467-025-64364-4.

Kanti V. Mardia and Peter E. Jupp. *Directional Statistics*. Wiley, 2 edition, 2000.

Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M. Wood. Flowmm: Generating materials with riemannian flow matching. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 35664–35686. PMLR, 2024. URL https://proceedings.mlr.press/v235/miller24a.html.

Hiroshi Morioka and Aapo Hyvarinen. Causal Representation Learning Made Identifiable by Grouping of Observational Variables. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 36249–36293. PMLR, 2024. URL https://proceedings.mlr.press/v235/morioka24a.html.

Hyunsoo Park and Aron Walsh. Guiding generative models to uncover diverse and novel crystals via reinforcement learning. *Nature Machine Intelligence*, 2026. doi: 10.1038/s42256-026-01262-4. URL

https://www.nature.com/articles/s42256-026-01262-4.

Hyunsoo Park, Zhenzhu Li, and Aron Walsh. Has generative artificial intelligence solved inverse materials design? *Matter*, 7(7):2355–2367, 2024. doi: 10.1016/j.matt.2024.05.017.

Raffaele Resta. Macroscopic polarization in crystalline dielectrics: The geometric phase approach. *Reviews of Modern Physics*, 66(3):899–915, 1994. doi: 10.1103/RevModPhys.66.899. URL https://doi.org/10.1103/RevModPhys.66.899.

Sorawit (James) Saengkyongam, Elan Rosenfeld, Pradeep K. Ravikumar, Niklas Pfister, and Jonas Peters. Identifying Representations for Intervention Extrapolation. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/aa9508b2577e0d9c824da9e2e27fe9d6-Abstract-Conference.html.

Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n)-equivariant graph neural networks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 9323–9332. PMLR, 2021.

Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. doi: 10.1109/JPROC.2021.3058954. URL https://doi.org/10.1109/JPROC.2021.3058954.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

David Vanderbilt. Berry-phase theory of proper piezoelectric response. *Journal of Physics and Chemistry of Solids*, 61(2):147–151, 2000. doi: 10.1016/S0022-3697(99)00273-5. URL https://doi.org/10.1016/S0022-3697(99)00273-5.

Luran Wang, Chaoran Cheng, Yizhen Liao, Yanru Qu, and Ge Liu. Training free guided flow matching with optimal control. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=61ss5RA1MM.

Mingjian Wen, Matthew K. Horton, Jason M. Munro, Patrick Huck, and Kristin A. Persson. An equivariant graph neural network for the elasticity tensors of all seven crystal systems. *Digital Discovery*, 3(5):869–882, 2024. doi: 10.1039/D3DD00233K. URL https://doi.org/10.1039/D3DD00233K.

Xifan Wu, David Vanderbilt, and D. R. Hamann. Systematic treatment of displacements, strains, and electric fields in density-functional perturbation theory. *Physical Review B*, 72:035105, 2005. doi: 10.1103/PhysRevB.72.035105.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022.

Yixian Xu, Yusong Wang, Shengjie Luo, Kaiyuan Gao, Tianyu He, Di He, and Chang Liu. Quotient-space diffusion models. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2604.21809. Oral presentation.

Keqiang Yan, Alexandra Saxton, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. A space group symmetry informed network for O(3)-equivariant crystal tensor prediction. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Cai-Yuan Ye, Hong-Ming Weng, and Quan-Sheng Wu. Con-cdvae: A method for the conditional generation of crystal structures. *Computational Materials Today*, 1:100003, 2024.

Chunlin Yu, Chao Liang, Junhao Liang, Yuchong Kang, Jian Yuan, and Huashan Li. Learning piezoelectric tensors through strain-conditioned polarization clusters. *Nano Letters*, 26(19):6474–6481, 2026. doi: 10.1021/acs.nanolett.6c01295.

Claudio Zeni, Robert Pinsler, Daniel Zuegner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbe, Shoko Ueda, Roberto Sordillo, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Chunlei Yang, Wenjie Li, Ryota Tomioka, and Tian Xie. A generative model for inorganic materials design. *Nature*, 639:624–632, 2025. doi: 10.1038/s41586-025-08628-5.

A SUPPLEMENTARY SETUP AND NOTATION

The supplement follows the mechanism decomposition of the main paper. Sections A–B formalize target semantics, quotient geometry, and identifiability; Sections C–F diagnose structural recovery, conditioning geometry, endpoint physics, and controlled reachability in turn. Section G collects implementation and reproducibility details that are not required for the scientific argument. All reported prediction metrics are recomputed by the common offline evaluator from saved Cartesian prediction and target arrays.

A.1 TENSOR REPRESENTATION AND RESPONSE CONVENTIONS

The piezoelectric response is represented in the declared engineering order, expanded to Cartesian components, and then mapped to the Kelvin/irrep basis used by the equivariant model. The response space is written $V_e = V_{1a} \oplus V_{1b} \oplus V_2 \oplus V_3$, of dimension 18. The two V_1 copies have different parity labels in the implementation; writing them as $2V_1^-$ is a bookkeeping convention and does not change the total dimension.

For a nonzero response e, let $U = e/\|e\|_F$. We distinguish three comparisons, following standard directional statistics for signed and axial observations (Mardia & Jupp, 2000; Bhattacharya & Patrangenaru, 2003):

$$c_{\text{signed}}(u, v) = \langle u, v \rangle, \tag{11}$$

$$c_{\text{axial}}(u, v) = |\langle u, v \rangle|, \tag{12}$$

$$c_{\text{orb}}(u, v) = \sup_{R \in \text{SO}(3)} \langle u, \rho(R) v \rangle. \tag{13}$$

The finite-grid evaluator used for the reported prediction results replaces the last supremum by a canonicalized proper-rotation grid. Orbit cosine is the primary metric because it identifies physically equivalent representatives; signed and axial cosine remain useful auxiliary metrics of representative-level information.

The endpoint specification is indexed by $q = (s, c)$. The response state $s \in \{\text{clamped}, \text{ionic}, \text{relaxed}\}$ selects the C/I/R decomposition,

$$e^{\text{relaxed}} = e^{\text{clamped}} + e^{\text{ionic}}. \tag{14}$$

The response convention c specifies the piezoelectric response definition (proper versus improper) and the declared component ordering. Response state and response convention are separate choices; we do not infer one from the other. For a structure X, $\mathcal{E}^q(X) = \mathcal{E}^{(s,c)}(X)$ denotes the corresponding scientific endpoint.

A.2 DATASETS AND SPLITS

The primary TensorOrbit–JARVIS source population contains 3,889 crystals from 2,726 formula families, with a fixed 3165/363/361 train/development/test split, maximum 20 atoms, and zero family leakage. The formal matched-orbit comparison panel is an independent 1,728-record panel with a fixed 1536/192/192 train/development/test split and three model seeds. The common evaluator keeps material identifiers and returns per-record Cartesian predictions, targets, norms, and signed/axial/orbit scores. The 5,000-record curated C/I/R source is a separate response-state panel: it contains 4,000/499/499 TensorOrbit train/validation/test records, two excluded records, and 3,542 formula families. It is never pooled with either prediction panel.

A.3 EVALUATION METRICS

Component MAE/RMSE are computed in Cartesian coordinates. Directional metrics are computed after normalizing each nonzero tensor. Orbit scores use the same fixed grid for every model and split; no model sees the target while the grid score is computed. Norm correlation is Pearson correlation of per-record Frobenius norms. For a local Jacobian with singular values $s_1 \geq \cdots \geq s_d$, numerical rank at 1% is the number of values above $0.01 s_1$, and stable rank is $\sum_i s_i^2 / s_1^2$. These are different notions of visible and effective dimension. A record whose Jacobian vanishes identically has both quantities equal to zero by convention and enters population averages at zero.

B PROOFS

B.1 QUOTIENT GEOMETRY PRELIMINARIES

Let $G = \text{SO}(3)$ act on response representatives through ρ. The orbit of e is $[e] = G \cdot e$, and the quotient tangent at a regular point is represented by a horizontal complement to $T_e(G \cdot e)$. If a crystal has stabilizer G_X, a deterministic canonical frame may be discontinuous at symmetry-enhanced strata. The orbit metric used conceptually is

$$d_{\text{orb}}(e_1, e_2) = \min_{R \in G} \|e_1 - \rho(R) e_2\|_F. \tag{15}$$

The finite-grid score serves as a computationally tractable proxy for this quotient comparison over G.

B.2 LOCAL QUOTIENT CONTROLLABILITY

Let $a \in \mathbb{R}^p$ be the action, $G_\theta(z, a)$ the generated state, and $F^q_{\theta,z}(a) = \pi_q(\mathcal{E}^q(G_\theta(z, a)))$ the composed endpoint map. Here the quotient target tangent is $\mathcal{T}_q = T_{[y_0]}(\mathcal{Y}_q/\sim_q)$, with π_q mapping a representative to its quotient class. P_{out} projects onto the horizontal output tangent and P_{in} is the analogous action-chart projection; it is the identity in the fixed local charts used by the route evaluations. At a regular point define

$$S_T = D_a G_\theta(z, a_0), \qquad \bar{K}^q_T = P_{\mathrm{out}} D_X \mathcal{E}^q(X_0) S_T P_{\mathrm{in}}. \tag{16}$$

where $X_0 = G_\theta(z, a_0)$, and P_{in} and P_{out} project onto the horizontal input and output tangents. This is the differential of the descended endpoint map in local quotient charts; π_q is not differentiated with respect to the structure variable.

Proof of Theorem 1. If $\bar{K}^q_T$ is surjective onto the quotient tangent at $F^q_{\theta,z}(a_0)$, then the endpoint image contains a neighborhood of that quotient point. If its rank is deficient, every nonzero covector in the cokernel defines a quotient-tangent perturbation that is unreachable to first order.

Proof. On a regular orbit-type stratum, the slice theorem supplies local slices transverse to the input and output group orbits (Bredon, 1972). Their tangent spaces are the horizontal spaces selected by P_{in} and P_{out}. Equivariance of G_θ and $\mathcal{E}^q$ makes their composition descend to these local quotient charts. Applying the chain rule on the slices gives $P_{\mathrm{out}} D_X \mathcal{E}^q D_a G_\theta P_{\mathrm{in}} = \bar{K}^q_T$. Full row rank is exactly the submersion condition, so the submersion theorem gives a local preimage for every sufficiently small quotient-tangent perturbation. If the rank is deficient, finite-dimensional duality supplies a nonzero output covector in the cokernel that annihilates every first-order quotient response. The statement is confined to the regular-stratum chart. □

Proof of Corollary 1. Local quotient reachability does not imply direct tracking. In matched action and target charts with the same local coordinates, identifying the basepoint coordinates as $a_0 = \tau_0$, let $h(a) = F^q_{\theta,z}(a) - a$. Along the segment from a_0 to a, the fundamental theorem of calculus gives

$$h(a) = h(a_0) + \int_0^1 \big(D_a F^q_{\theta,z}(a_0 + s(a - a_0)) - I\big)(a - a_0)\,\mathrm{d}s. \tag{17}$$

Taking norms and using the uniform operator-norm bound yields $\|h(a)\| \le \|h(a_0)\| + \varepsilon\|a - a_0\|$, which is the main-text bound. Thus basepoint agreement and a near-identity differential certify local direct tracking; surjectivity of $\bar{K}^q_T$ alone only certifies the existence of some nearby inverse action.

B.3 PREDICTION, SELECTION, AND AXIAL INFORMATION

Let $Z = \phi(X)$ contain the structural information supplied to a predictor, and let $U = e/\|e\|_F$ on the nonzero population.

Define the signed and orbit information quantities in parallel:

$$\mathcal{I}_{\mathrm{dir}}(\phi) \equiv \sup_{\|g(Z)\|=1} \mathbb{E}\langle g(Z), U\rangle, \qquad \mathcal{I}_{\mathrm{orb}}(\phi) \equiv \sup_g \mathbb{E}\, c_{\mathrm{orb}}(g(Z), U). \tag{18}$$

Proposition 2 (Bayes-optimal signed directional prediction)**.** *If regular conditional distributions exist, then*

$$\sup_{\|g(Z)\|=1} \mathbb{E}\langle g(Z), U\rangle = \mathbb{E}\,\|\mathbb{E}[U \mid Z]\| = \mathcal{I}_{\mathrm{dir}}(\phi). \tag{19}$$

The optimum is the normalized conditional mean wherever it is nonzero.

Proof. Condition on Z and apply Cauchy–Schwarz to $\langle g(Z), \mathbb{E}[U \mid Z]\rangle$. Integrating gives the upper bound, and choosing the normalized conditional mean attains it almost surely. □

The corresponding conditional expression for $\mathcal{I}_{\mathrm{orb}}$ is

$$\mathcal{I}_{\mathrm{orb}}(\phi) = \mathbb{E}\left[\sup_{[v]} \mathbb{E}[c_{\mathrm{orb}}(v, U) \mid Z]\right]. \tag{20}$$

For a finite candidate set $g_1(Z), \ldots, g_K(Z)$, a structure-only selector has value

$$C_{\mathrm{sel}} = \sup_q \mathbb{E} \sum_k q(k \mid Z)\langle g_k(Z), U\rangle, \qquad C_{\mathrm{oracle}} = \mathbb{E} \max_k \langle g_k(Z), U\rangle. \tag{21}$$

Proposition 3 (Coverage does not imply target-independent selection)**.** *If the selector cannot observe U, then $C_{\mathrm{sel}} \le \mathcal{I}_{\mathrm{dir}}(\phi)$, and the analogous orbit-valued selector is bounded by $\mathcal{I}_{\mathrm{orb}}(\phi)$. The target-dependent oracle is a coverage metric, not a deployable selector.*

Proof. Condition on Z. The selector's convex combination is bounded by the conditional supremum over all unit predictions; then apply the definition of the corresponding Bayes information quantity and integrate. □

Theorem 4 (Axial directional information)**.** *For a hidden sign branch $U = BU_0(Z)$, with $\Pr(B = +1 \mid Z) = p(Z)$, signed identifiability is $\mathbb{E}|2p(Z) - 1|$. If the two representatives lie in one proper rotation orbit, axial/orbit information can remain maximal while signed information vanishes at $p = 1/2$.*

Proof. The conditional mean is $(2p(Z) - 1)U_0(Z)$, giving the signed result by Proposition 2. A proper-orbit comparison identifies the two representatives, so its quotient score is one in the idealized pair. □

B.4 Conditioning-path visibility and irrep rank

For condition maps $\Psi(e) = (\psi_1(e), \dots, \psi_L(e))$ and $v_\theta(X, t; e) = \mathcal{V}_\theta(X, t, \Psi(e))$,

Theorem 5 (Group-orbit visibility)**.**

$$D_e v_\theta = D_\Psi \mathcal{V}_\theta \, D_e \Psi, \qquad \operatorname{rank}(D_e v_\theta) \le \operatorname{rank}(D_e \Psi). \tag{22}$$

If every ψ_l is rotation-invariant, then $T_e(G \cdot e) \subseteq \ker(D_e v_\theta)$. More generally, every subspace discarded by $D_e\Psi$ is invisible to the conditioning path.

Proof. The first statement is the chain rule. Differentiating $\psi_l(\rho(\exp(s\xi))e) = \psi_l(e)$ at zero gives a zero directional derivative along every orbit tangent, which lies in the kernel of the product. □

Theorem 6 (Symmetry-conditioned irrep rank bound)**.** *Let $V_e = \bigoplus_\lambda m_\lambda V_\lambda$ and $T_X\mathcal{M} = \bigoplus_\lambda n_\lambda V_\lambda$ be decompositions into pairwise non-isomorphic real irreducible G_X-representations. If $A = D_e v_\theta$ is G_X-equivariant, then*

$$\operatorname{rank}_{\mathbb{R}}(A) \le \sum_\lambda d_\lambda \min(m_\lambda, n_\lambda), \qquad d_\lambda = \dim_{\mathbb{R}} V_\lambda. \tag{23}$$

where multiplicities are counted over the appropriate real Schur division algebra $\mathbb{K}_\lambda \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$. There are no cross-irrep blocks; the bound is governed by representation multiplicities rather than imposing an intrinsic low-rank constraint via equivariance.

Proof. Equivariance makes A an intertwiner. Schur's lemma removes blocks between inequivalent isotypic components. On the λ-component, the remaining intertwiner is an $n_\lambda \times m_\lambda$ matrix over the commutant division algebra $\mathbb{K}_\lambda$, so its algebraic rank is at most $\min(m_\lambda, n_\lambda)$. Each nonzero multiplicity copy contributes at most d_λ real dimensions to the image. Summing the dimensions of the surviving image blocks gives Eq. 23. □

B.5 Compact counterexamples

Sensitivity does not imply target alignment: for a nonzero target tangent u, choose a nonzero orthogonal response w. A condition intervention with $\Delta v = w$ has positive field sensitivity but zero first-order target projection. Likewise, a scalar condition summary $\psi(e)$ gives $D_e v = D_\psi v \, D_e \psi$ and therefore has rank at most the local dimension of the summary; the rank-one conclusion applies only to a single scalar.

For a compact latent branch group, replacing a hidden sign by a group element replaces the scalar conditional moment by the corresponding orbit-averaged moment. This yields the same separation: representative prediction can be ambiguous even when the quotient target is identifiable.

Implication. These constructions isolate two independent failure modes: a conditioning path may be sensitive without aligning with the requested target, and a low-dimensional condition representation can impose a rank bottleneck before any endpoint physics is considered.

C Additional target-semantics and structural-recovery results

C.1 C/I/R response-state distributions

Figure S1 uses all 5,000 curated records. The additive reconstruction residual for $e^{\mathrm{relaxed}} = e^{\mathrm{clamped}} + e^{\mathrm{ionic}}$ is zero to numerical precision. The main text reports the transition counts; this figure retains the complementary norm, representative/orbit-score, and irrep distributions. These source statistics characterize response-state heterogeneity at the population level.

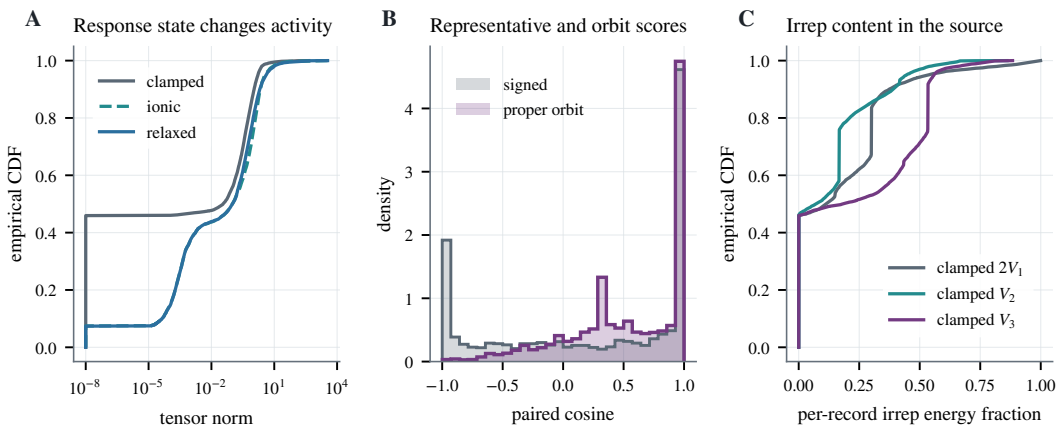


Figure S1: Response-state semantics in the 5,000-record curated source. The three response states differ in magnitude, representative direction, and irrep composition; the distributions are source-level statistics rather than generator predictions.

C.2 Matched frame comparison

Frame canonicalization is often used to simplify orbit comparison; the question here is whether it preserves the deployment-relevant response.

Table S1: Matched raw-frame versus target-free canonicalized GMTNet. Values are the three-seed means on the common test panel; the final column is canonicalized minus raw.

variant	signed	axial	orbit	Δ orbit
raw frame	0.129	0.509	0.679	—
canonicalized frame	0.108	0.499	0.660	−0.019
difference	−0.021	−0.010	−0.019	—

The matched raw-frame and target-free canonicalized GMTNet comparison uses the same split, objective, seeds, checkpoint rule, and offline evaluator. Only the frame treatment changes. Canonicalization changes signed/axial/orbit cosine by $-0.021/-0.010/-0.019$, respectively, and therefore does not support a stable improvement claim. The weighted-frame analysis starts from 1,888 curated nonzero rows; 594 have at least one atom-token-compatible candidate and 1,294 have none. Applying a joint site/lattice RMS threshold of 0.1 Å yields 144 aligned pairs and 450 compatible residuals. The residual full-tensor cosine is -0.004; 84 aligned pairs admit both parity classes, and seven directed pairs cross the tensor-cosine sign. These statistics characterize the local geometric correspondence within this compatible subset.

C.3 Objective-level prediction diagnostics

For a squared unit-direction loss with target magnitude a and prediction magnitude $\hat{a}$, the direction-gradient magnitude is $2a\hat{a}\sqrt{1-\cos^2\theta}$. Thus magnitude can suppress directional gradients. An independent relaxed-target decomposed-objective run (500 predictions, including 469 active test records) reported test signed and axial cosines of (0.0316) and (0.3812). Together with the matched direction-only control below, this experiment isolates objective-level mechanisms that can degrade directional recovery without implying an information limit.

A matched direction-only objective control on the same formal panel (192 test records per run, three runs; direction-matched loss against the identical Huber baseline) shifts the test signed cosine by $+0.072$, the axial cosine by $+0.046$, and the proper-orbit cosine by $+0.011$, with component MAE moving from 0.103 to 0.175. This control complements the matched prediction comparison in Section 3.2. The two objective-control runs are complementary: the relaxed decomposed objective isolates magnitude–direction coupling, whereas the matched direction-only control isolates the directional channel under identical training conditions.

C.4 Orbit evaluator convergence

The finite-grid evaluator was compared at coarse, standard, and high precision resolutions using the same predictions, test records, and rotation seed. The standard setting uses a grid level of 128, 16 restarts, 96

refinement steps, and tolerance 10^{-8}. The high-precision minus standard changes are $+0.00003$, $+0.00007$, and $+0.00014$ for GMTNet, EATGNN, and CEITNet, respectively; the maximum change is below 1.5×10^{-4}, and the ranking is unchanged (CEITNet $>$ EATGNN $>$ GMTNet). All models use the same post-hoc search budget.

D CONDITIONING MECHANISM AND TRAJECTORY GEOMETRY

D.1 TRAINED ROUTE GEOMETRY

The trained-route analysis asks whether increasing condition dimensionality expands the locally visible action space or merely redistributes response energy. Scalar conditioning exposes one active direction, whereas full-tensor routes open a much broader numerical span while retaining a strongly anisotropic effective spectrum. Figure S2 separates these two effects through numerical rank and stable rank.

Figure S2 is generated from 384 scalar records, 384 Flat18-FiLM records, and 384 OperatorSpan records from three complete per-seed evaluations. Each seed contains 128 material/noise trajectories under the fixed 32-material, four-noise protocol. The route comparison is a post-training local-geometry measurement of the conditioning path. Jacobian–vector products are computed by automatic differentiation and cross-validated against a fixed-step finite-difference directional derivative; trajectories encountering numerical stiffness in these products are evaluated without synthetic interpolation.

The trained intervention isolates the trajectory-visibility layer. Across three scalar seeds, rank at 1% and stable rank are both 0.90625: every record with a nonzero Jacobian has both quantities exactly equal to one, while the 36 records with exactly zero local response (three of the 32 test materials across all four noise conditions and three runs) enter the mean at zero. Flat18-FiLM has rank-at-1% values of approximately 11.75–12.48 and stable ranks of approximately 1.70–1.83. Thus full-tensor conditioning opens many more locally visible directions while the response remains strongly anisotropic. OperatorSpan has seed-level mean stable ranks of 2.252, 2.271, and 2.203, with a pooled median of 2.0549 across the three per-seed evaluations. Thus the route exposes a broad numerical span while retaining a strongly anisotropic effective spectrum. This characterizes the trajectory-geometry layer; target-alignment and endpoint-transfer criteria are examined separately in the controlled-endpoint analysis.

The parameter-matched random-initialization/synthetic operator mechanism study contains the following seven routes: scalar/invariant, flat18 FiLM, irrep FiLM, tensor product, cross-attention, GeoAda-style, and OperatorSpan. Each route uses the same interface and approximately 2.25k route parameters. The reported teacher columns are cosine recovery on synthetic operators of ranks 1, 2, 4, and 8; they measure mechanism capacity in isolation from real-data endpoints.

D.2 SYNTHETIC JACOBIAN BOTTLENECK

Figure S3 is a controlled synthetic mechanism study of the tensor encoder, fusion, and final field map, isolating the rank loss that is otherwise entangled with real-data supervision. The spectra exhibit the intended rank loss across the composed path, reported through singular spectra and irrep coverage of each stage.

The synthetic path uses an 18-dimensional input, a 4-dimensional intermediate representation, and a 1-dimensional output map. The encoder and fusion maps are fixed linear maps with prescribed singular values on their retained directions and exact null spaces on the discarded directions; the final map keeps one singular direction. Input coordinates are assigned to the four response irreps before the maps are composed, so the plotted coverage records which blocks survive each stage. The nonzero singular values are normalized to one, making numerical rank and stable rank coincide in this construction. Exact zeros are drawn at a plotting floor for visibility only and are not treated as nonzero singular values.

D.3 BR-RATC: COVERAGE AND SELECTION

This experiment separates candidate coverage from executable selection. Increasing the number of branches can place a high-quality hypothesis in the candidate set, but a deployable selector that observes only structure need not identify that branch.

For each nonzero response, branch directions are normalized in the 18-dimensional unit sphere. The $K \in \{1, 2, 4\}$ branch models use the same fixed structural features and branch-label exports; mixture and MAP scores use the branch probabilities, whereas the learned selector is trained on TRAIN branch labels, selected on DEV, and evaluated once on TEST. It never receives the target tensor. The target-dependent oracle chooses the highest-cosine branch only as a coverage metric, and the target-conditioned positive control is reported separately for the same reason. Thus the selector and oracle answer different questions on the same evaluation records.

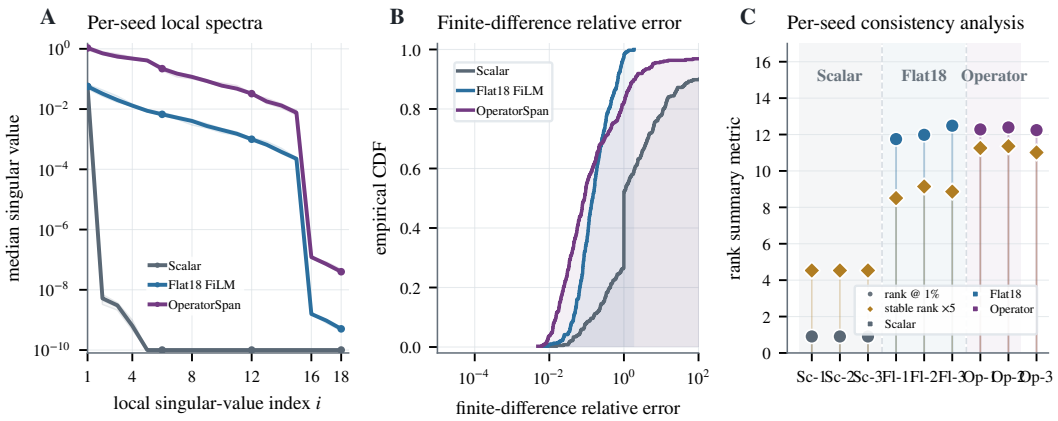


Figure S2: **Numerical derivative validation and multi-run route stability analysis.** (A) Median singular-value spectra across independent model seeds for scalar, Flat18-FiLM, and OperatorSpan routes. (B) Empirical CDF of the finite-difference relative error against analytical Jacobian-vector products. (C) Per-seed means of numerical rank (at 1%) and scaled stable rank. OperatorSpan summarizes three complete per-seed evaluations (384 trajectories in total); short labels denote route initializations (Sc: Scalar, Fl: Flat18 FiLM, Op: OperatorSpan). The directional-cosine validation is promoted to the main text.

Table S2: Parameter-matched random-initialization and synthetic operator mechanism study. All routes use the same interface and approximately 2.25k parameters.

route	parameters	rank	stable rank	$r_{1\%}$	teacher cosine (1/2/4/8)
scalar/invariant	2250	1	1.0000	1	0.0253/0.2819/0.0830/0.3477
flat18 FiLM	2257	15	2.4651	13	0.6871/0.8990/0.8856/0.9207
irrep FiLM	2249	4	1.4971	4	0.4013/0.5650/0.3504/0.4248
tensor product	2262	15	2.7652	15	0.9034/0.7885/0.8647/0.9245
cross-attention	2257	3	1.0702	3	0.3916/0.6137/0.3091/0.3206
GeoAda-style	2257	15	2.3869	14	0.8797/0.9102/0.9216/0.9117
OperatorSpan	2259	15	4.9371	15	0.9216/0.8571/0.9459/0.9237

For normalized branch directions $\mu_k(Z)$, the structure-only selector must choose without observing the target. The target-dependent oracle is reported only to quantify candidate coverage. On the evaluation panel, the mean signed cosines for $K = 1, 2, 4$ are respectively $0.1750/0.1788/0.0651$ (mixture), $0.1750/0.1665/0.0174$ (MAP), $0.1750/0.2001/0.0725$ (learned structure-only selector), and $0.1750/0.6459/0.7652$ (target-dependent oracle). Figure S4 provides the comparison visually. The $K = 4$ learned selector has mean resultant length 0.336 and branch entropy 1.189. A target-conditioned positive control reaches 0.790, but it uses the supplied target and is therefore not a structure-only selector. On the quotient proper-$\mathrm{SO}(3)$ evaluator, the $K = 4$ target-free selectors remain well below the target-dependent ceiling: mean proper-orbit cosines are 0.7499 (learned linear), 0.7456 (learned MLP), and 0.8946 (target-dependent oracle; 191 test materials, three runs). These quotient metrics complement the signed representative-level metrics above: they measure orbit-level selection, whereas the signed metrics measure representative-level selection and the oracle quantifies candidate coverage. Figure S4 shows the per-record coverage/selection geometry; no target is supplied to the deployable selector.

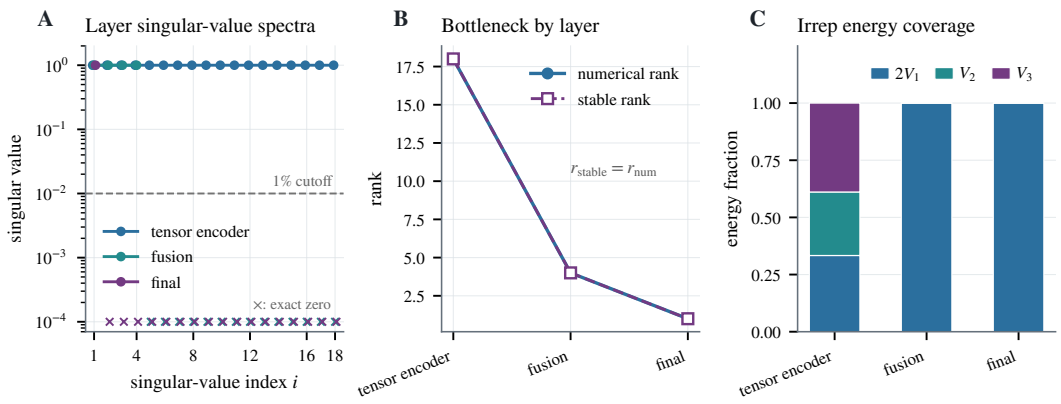


Figure S3: **Synthetic Jacobian-path analysis.** (A) Singular-value spectra across the tensor-encoder, fusion, and final-field stages; exact null directions are shown at the plotting floor. (B) The corresponding numerical and stable ranks collapse from 18 to 4 to 1. (C) Irrep-energy coverage of the surviving response subspace. The controlled construction demonstrates how a conditioning bottleneck arises before endpoint transfer in the absence of empirical tangent supervision.

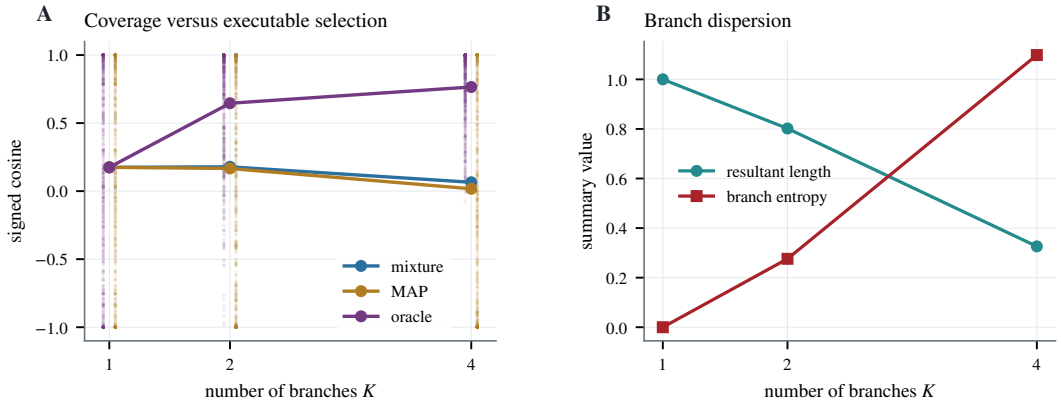


Figure S4: BR-RATC candidate coverage versus target-independent selection. Candidate sets can contain a high-cosine branch even when a selector that sees only structure cannot identify it.

E Endpoint-side physics

E.1 DFPT factorization and Born-visible mode geometry

For the relaxed-ion response state, the lattice-mediated contribution is

$$e^{\mathrm{ionic}} = \frac{e_0}{\Omega}(Z^*)^\top \Phi_o^\dagger \Lambda. \tag{24}$$

After acoustic projection, let $A = \Phi_o^{-1/2} Z^*$ and $B = \Phi_o^{-1/2}\Lambda$, so that $e^{\mathrm{ionic}} = \frac{e_0}{\Omega} A^\top B$. Let P_A denote the orthogonal projector onto $\mathrm{Col}(A)$. Since $A^\top(I - P_A) = 0$, $A^\top B = A^\top P_A B$. Thus, pointwise in structure, only the component of the internal-strain response lying in the Born-charge-visible mode subspace $\mathrm{Col}(A)$ contributes to the lattice-mediated response. We define its visible mode rank as $r_{\mathrm{Born}}(X) = \mathrm{rank}\, A(X) \le 3$. This rank is independent of the number of participating optical modes. The population counts reported in the main text enumerate the Born-visible response rank $r_{\mathrm{resp}}(X) = \mathrm{rank}(P_A B) \le r_{\mathrm{Born}}(X)$, the rank of the internal-strain response after projection onto the Born-visible mode subspace.

Structural differentiation acts on the factorization itself. Writing $c(X) = e_0/\Omega(X)$, $D_X e^{\mathrm{ionic}}[\delta X] = D_X c[\delta X]\, A^\top B + c\left(D_X A[\delta X]^\top B + A^\top D_X B[\delta X]\right)$. Accordingly, the structural endpoint differential contains variation of the Born-charge map, force constants, internal-strain coupling, and volume in addition to the pointwise visible-mode geometry.

On 3,371 dynamically stable, well-conditioned records, reconstruction identity error has median 1.04×10^{-15}, p90 9.55×10^{-13}, p99 4.90×10^{-12}, and maximum 4.74×10^{-11}. The stable-core Born-visible response-rank counts are 97/72/481/2,721 for ranks 0/1/2/3. Figure S5 reports the effective mode count, optical rank, and conditioning/error relationship.

E.2 GaN/ZnO local endpoint checks

The fixed GaN and ZnO systems provide independent local checks of the factorized response on source-matched structures and fixed branch coordinates. Berry derivatives agree with the factorized Born prediction within approximately 1.5% for GaN and 0.1% for ZnO. Centered force and energy residuals exhibit the expected $O(h^2)$ and $O(h^3)$ scaling under step halving, supporting the pointwise response calculation used in Sec. E.1.

E.3 Response-active versus design-active

This check asks whether the selected e_{33} projection of the validated GaN branch is statistically design-active at the measurement resolution. The estimated clamped derivative is -2.857 C m^{-2} Å^{-1} with propagated noise $\sigma = 0.960$; it is below both five-sigma significance thresholds. The branch is therefore response-active in the measured channel, while the selected e_{33} projection is not statistically design-active at this resolution.

F Controlled analytic and semi-synthetic systems

The analytic endpoint system separates three notions by construction: a field can be sensitive but orthogonal to the target, a target can be reachable in an explicit endpoint map, and direct tracking additionally requires aligned inverse action. On the 18-structure semi-synthetic panel, the reachable and unreachable subsets are balanced by construction. The response-aligned controller has reachable mean signed/axial cosine 1.000/1.000 and relative endpoint error 1.85×10^{-5}; the sensitivity-only controller has signed cosine 0.818 but relative endpoint error 4.847. The contrast between the two controllers quantifies, on an exactly tractable endpoint map, the gap between field sensitivity and target tracking.

F.1 Hidden branch and endpoint construction

In the hidden-branch calibration, the visible feature is a normalized 18-dimensional base direction $U_0(Z)$, while an independent sign $B \in \{-1, +1\}$ is sampled with $\Pr(B = +1) = p$. Each of 4,096 training and 4,096 disjoint test structures is repeated 32 times. A ridge estimator is fit to the averaged training targets and evaluated on all test repetitions for $p \in \{1.0, 0.9, 0.75, 0.6, 0.5\}$. The population optimum is $|2p-1|$; the largest observed absolute deviation is 7.94×10^{-4}. The disjoint split tests branch ambiguity rather than memorization of structure identifiers.

For the periodic endpoint, fractional pair displacements use the minimum-image rule $\Delta f_{ij} \leftarrow \Delta f_{ij} - \mathrm{round}(\Delta f_{ij})$ for every unordered pair $i < j$, followed by $r_{ij} = \Delta f_{ij} L$. The pair weight is $w_{ij} = s_i s_j$, with species weights $(1.0, 1.7, 2.3, 3.1)$; there is no cutoff or additional normalization. The endpoint is the 27-component Cartesian tensor

$$E(X) = \sum_{i<j} w_{ij} r_{ij}^{\otimes 3}. \tag{25}$$

The fixed lattice is $\mathrm{diag}(4.3, 4.8, 5.2)$, and the four fractional positions are sampled independently in $[0.15, 0.80]^3$. Jacobians are computed in float64 with automatic differentiation using a Jacobian-by-reverse mode implementation.

The 64 target cases use a structural perturbation $\Delta X^\star$ of norm 0.015 in the span of the first three right-singular directions of the local endpoint Jacobian, and define the corresponding endpoint target from $E(X_0 + \Delta X^\star)$. The additional 18-geometry panel contains one response-aligned target and one orthogonal target per clean JARVIS geometry; the former lies in the local Jacobian image and the latter in its orthogonal complement. This construction makes reachable and unreachable cases explicit without using a learned tangent label.

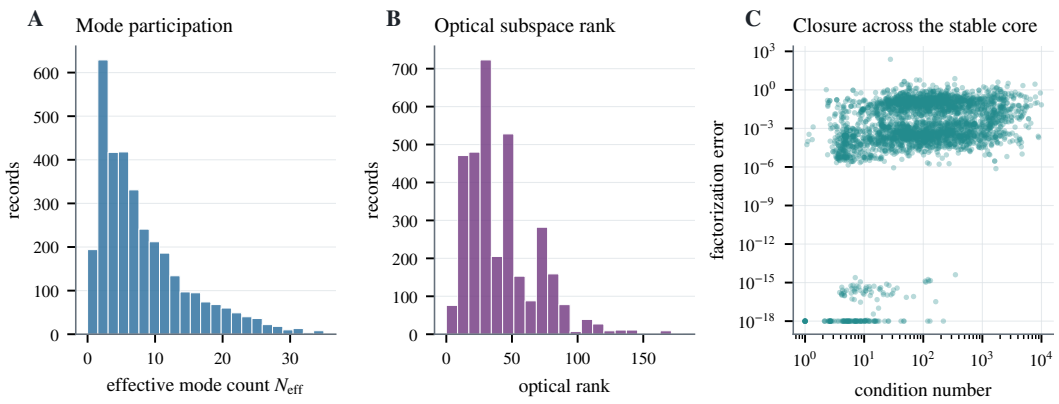


Figure S5: Population-level DFPT factorization analysis on 3,371 stable, well-conditioned records. At each fixed structure, the Born-charge-visible mode subspace has dimension at most three, set by the three Cartesian polarization components; the panels contrast many-mode participation with this low-dimensional pointwise observation channel, while effective mode count and numerical conditioning are separate from the pointwise rank.

F.2 Controller rules and metrics

Controllers. The no-control endpoint is $E(X_0)$. Condition injection evaluates $E(X_0 + 0.5\,\Delta X_\star)$, where $\Delta X_\star$ is the target perturbation. The sensitivity-only control uses a nonzero coordinate direction orthogonal to the target residual. The local inverse uses the damped pseudoinverse

$$\Delta X = J^\top (JJ^\top + 10^{-4} I)^{-1} (y^\star - E(X_0)), \tag{26}$$

where $J = D_X E(X_0)$. D-Flow-style control minimizes terminal squared error through the differentiable eight-step Euler trajectory with zero control energy weight. OC-Flow-style control uses the same terminal objective plus an energy penalty of 10^{-3}. These are analytic analogues of the cited principles, not full implementations of either external method. On the same 64 targets, the trajectory-level inverse-response baseline, the D-Flow-style control, and the OC-Flow-style control attain endpoint-relative errors of 0.1329, 0.0329, and 0.0329, respectively; for reference, no control gives error 1.000.

Optimization details. For the trajectory controls, eight zero-initialized control increments are optimized for 12 LBFGS iterations with learning rate 1, strong-Wolfe line search, gradient tolerance 10^{-9}, and change tolerance 10^{-12}.

For an endpoint $y = E(X)$, target $y^\star$, and initial endpoint y_0, we report field sensitivity $\|y - y_0\|_2$, target alignment $\langle y - y_0, y^\star - y_0\rangle / (\|y - y_0\|\|y^\star - y_0\|)$, and endpoint-relative error $\|y - y^\star\|_2 / \|y^\star - y_0\|_2$. A controlled trial is counted as a successful target hit when alignment is at least 0.99. Reachable and unreachable targets are defined by whether the constructed target perturbation lies in the local Jacobian image or its orthogonal complement.

The structural tangent and reachability-vs-margin views remain here as complementary analyses; the residual pre/post control comparison is shown in the main controlled-endpoint figure.

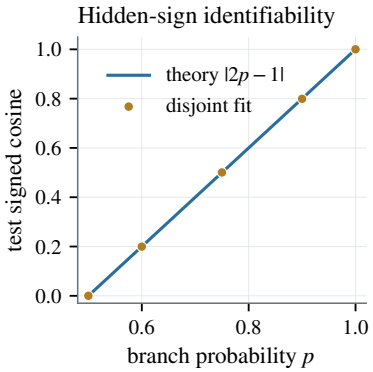


Figure S6: Controlled analytic benchmarks generated from the predefined 64-target periodic-endpoint system. The hidden-sign calibration follows the predicted $|2p-1|$ dependence. Field sensitivity/target alignment and the local inverse comparison are promoted to the main controlled-endpoint figure.

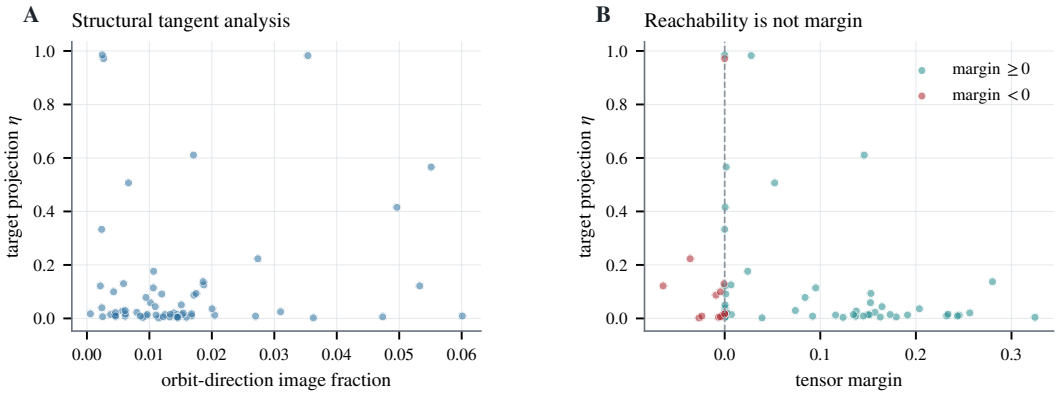


Figure S7: Complementary controlled endpoint geometry. (A) Structural tangent image fraction versus target projection. (B) Tensor margin versus target projection. The pre/post residual comparison is shown in the main text.

G Implementation and reproducibility details

The structural state uses categorical species, periodic fractional coordinates, and a positive lattice chart. The response interface keeps response-present, present-zero, nonzero magnitude, and normalized tensor direction distinct. All predictions are saved with material ID, split, target endpoint specification $q = (s, c)$, Cartesian target, Cartesian prediction, model identity, and evaluator version before metrics are recomputed offline. Model comparison uses the same checkpoint selection rule, split, and finite proper-$\mathrm{SO}(3)$ evaluator.

The route evaluation uses a fixed structural base and changes only the conditioning path. Scalar, Flat18-FiLM, and OperatorSpan route records are analyzed from saved trained checkpoints. The seven random-initialization/synthetic routes are scalar/invariant, flat18 FiLM, irrep FiLM, tensor product, cross-attention, GeoAda-style, and OperatorSpan; all use approximately 2.25k route parameters and the same interface. The synthetic Jacobian and operator-teacher tests measure capacity and numerical behavior on synthetic operators, complementing the real-data trajectory analysis. The common-noise field effect is computed as graph RMS coordinate-score change plus Frobenius strain-rate change plus graph RMS element-logit change. Its target margin is the wrong-condition minus matched-condition geometric loss, with equal coordinate and strain channel weights. The route archive provides the common-noise fields at the distribution level, so the target-margin comparison is reported at the distribution level.

The response-source composition is summarized as follows:

The generator panel is observational, so the endpoint-side physics relevant to transfer is characterized through two qualified channels: the fixed-material GaN/ZnO calculations probe the local response channel, while the DFPT atlas characterizes the Born-visible mode geometry. Extending target-alignment training to real data

Table S3: Response-source definitions used in this study.

source block	interpretation	evidence
public non-ionic DFPT	clamped-ion response	source/value agreement
curated ionic block	internal-relaxation contribution	additive reconstruction
relaxed-ion target	clamped plus ionic response	explicit response-state definition

would require an independently sourced coordinate/strain tangent for this panel; the present controlled systems establish the geometric criteria against which such a tangent can be evaluated in future work.