TRACING THE CAUSAL TRANSITION FROM EVENT RELATIONS TO MORAL JUDGMENT IN LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Large language models (LLMs) exhibit competence in moral reasoning on behavioral benchmarks, yet whether this reflects actual moral computation or statistical shortcuts remains unclear. To bridge this gap, we propose a bottom-up interpretive paradigm that moves beyond static value–principle matching. Grounded in the *Theory of Dyadic Morality* (TDM), this paradigm traces the causal-mechanistic trajectory through which LLMs internally arrive at moral judgments. Models first spontaneously form event relations and then use them to determine how outcome information conditions judgments about the moral agent, a process we term *Event-Navigated Causal Transition* (ENACT). We substantiate this process through an initial causal abstraction followed by step-by-step causal interventions. Specifically, we abstract moral events into structural causal graphs spanning distinct event relations and use cross-domain instantiation to build MORALTRACE, which provides isomorphic counterfactual pairs for fine-grained manipulations. Then, we use staged causal interventions to disentangle the three operators underlying the progression from moral-patient completion, through intention inference, to outcome-conditioned judgment. At the neural level, circuit discovery and feature decomposition show that ENACT's behavior arises from distributed representations converging onto a low-dimensional latent space, revealing strong generalization across diverse task settings. Together, these results uncover a causally grounded mechanism of moral reasoning in LLMs and reveal emergent patterns across model scales and training stages.

1 INTRODUCTION

Large language models (LLMs) demonstrate impressive higher-order capabilities such as conceptual abstraction and relation understanding (Hernandez et al., 2024; Park et al., 2025), as reflected in their human-like performance on moral-domain benchmarks (Hendrycks et al., 2021; Ji et al., 2025). This behavior is thought to stem from moral concepts and linguistic conventions embedded in pre-training data, with post-training further reshaping expressed values and instructions modulating how they dynamically surface at inference time (Schramowski et al., 2022; Jin et al., 2025).

Existing studies attempt to evaluate these abilities with moral stories, dilemmas or community judgments in accordance with certain top-down value frameworks (Lourie et al., 2021; Abdulhai et al., 2024; Kumar & Jurgens, 2025; Snoswell et al., 2026), yet behavioral performance alone cannot separate underlying reasoning from superficial correlation or post-hoc rationalization (Röttger et al., 2024; Wu et al., 2025a). Mechanistic studies instead probe latent moral dimensions and steer value-related representations (Schacht & Lanquillon, 2025; Raimondi et al., 2025; Yu et al., 2026; Hu et al., 2026). While these approaches characterize the moral concept a model encodes, they provide no guarantee that steered representations causally participate in relational routing: how a model transforms information about an event into a moral judgment that assigns responsibility across the specific agent–patient dyad involved (Ge et al., 2026; Huang et al., 2026).

To address these limitations, we take a bottom-up causal approach grounded in the formal structure of the *Theory of Dyadic Morality* (TDM) (Gray et al., 2012; Schein & Gray, 2018). Central to our method is MORALTRACE, a diagnostic suite that operationalizes TDM by decomposing moral

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.

scenarios into independently manipulable primitives. We formalize these primitives using structural causal graphs (Pearl, 2000), establishing a unified relational structure shared across normative-social, instrumental-social, and mechanistic-computational domains. Leveraging this causal structure, we perform fine-grained counterfactual interventions, such as layer-wise state interchange, component ablation and addition, and residual stream tracing, to delineate how internal event representations are transformed into explicit judgments of an agent's intentions and actions.

Applying this framework across a wide spectrum of instruction-tuned model architectures, we uncover an end-to-end computational trajectory through which models transform raw event structures into normative evaluations, a mechanism we formulate as *Event-Navigated Causal Transition* (ENACT). Mechanistically, this process unfolds in three key stages. First (**RQ1**, §3), LLMs spontaneously construct an implicit moral patient within their residual stream even without eliciting prompting, where causal ablation directly shifts downstream wrongness judgments. Second (**RQ2**, §4), forming a verdict requires dynamically routing outcome valence to update the causal agent's attributed intention before conditioning on normative criteria, a non-linear relational transition that is absent in smaller models and emerges only with parameter scale. Finally (**RQ3**, §5), circuit discovery and sparse autoencoder (SAE) decomposition reveal that this transition is initiated by a few attention heads and MLPs and then accumulated across many components in late layers, converging onto a low-dimensional subspace that generalizes to unstructured, real-world moral narratives.

We summarize our contributions as follows:

- We propose a novel paradigm for interpreting moral reasoning in LLMs, shifting from static representation matching to bottom-up causal tracing, thereby disentangling underlying relational moral computation from superficial statistical correlations.
- We construct MORALTRACE, a graph-based cross-domain diagnostic suite providing isomorphic counterfactual event pairs, moving beyond unstructured, confounding narratives to systematically diagnose the modular primitives that navigate LLM moral judgment.
- We introduce a multi-stage causal intervention method based on controlled operator decomposition of ENACT, enabling us to characterize the joint effects and sequential order of key operators underlying the social reasoning dynamics of LLMs.
- We uncover that moral reasoning in LLMs not only emerges with model size and is shaped by training stages, but also demonstrates a layer-wise progressive pattern, where deeper layers correspond to later stages of ENACT.

2 BOTTOM-UP FRAMEWORK FOR MORAL REASONING

2.1 MORAL EVENT STRUCTURES

Since moral judgment provides no single external ground truth against which an answer or an intermediate step can be verified (Hendrycks et al., 2021; Chiu et al., 2026; Snoswell et al., 2026) , the target computation must be formalized before its mechanism can be interpreted. The *Theory of Dyadic Morality* (TDM) provides such a foundation (Gray et al., 2012; Schein & Gray, 2018), a relational template underlying the diversity of moral concepts, compressing a situation into two role nodes joined by a weighted directed edge (Figure 1),

$$\mathcal{G}_{\text{dyad}} \triangleq \big(\{v_{\text{A}}, v_{\text{P}}\}, \{v_{\text{A}} \xrightarrow{H} v_{\text{P}}\};\, H \in \mathbb{R}_{\geq 0}\big), \quad (1)$$

with the occupant of each role and the edge weight left to the observer and the context. Here, v_{A} denotes the intentional causal agent, v_{P} represents the vulnerable patient bearing the outcome, and H carries the magnitude and causal strength of the harm transmitted from agent to patient.

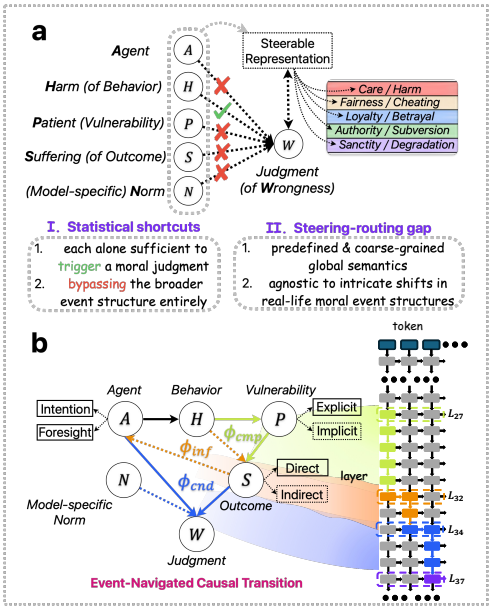

Figure 1: **Moral reasoning paradigms.** **(a) Existing**: top-down steering. **(b) Ours (ENACT)**: bottom-up causal transition.

Following Varshney (2026), we formalize TDM within a structural causal model (SCM; Figure 1). We operationalize this process as three sequential causal transitions: $\Phi_{\text{cmp}}(\cdot)$, a completion operator that reconstructs implicit moral patients from relational context, functionally approximating their explicit counterparts; $\Phi_{\text{inf}}(\cdot)$, an inference operator that binds the outcome to its true causal antecedent to update agent intentionality, distinguishing the agent's action from concurrent background events; and $\Phi_{\text{cnd}}(\cdot)$, a conditioning operator that evaluates the relation-bound event under model-specific normative priors, separating relational composition from linear feature addition. These staged transitions incrementally compile participant roles, causal attributions, and normative constraints into an internal event state, which in turn directly determines the final judgment.

Moral reasoning thus unfolds as a sequence of state updates across domains $d \in \mathcal{D}$:

$$e_t^d \triangleq T^d(R_t, s_t), \qquad (\hat{\imath}_t^d, \hat{y}_t^d) \triangleq D^d(e_t^d, N_t^d), \tag{2}$$

where the relational tuple $R_t \equiv (a_t, p_t, c_t)$ combines the agent, patient, and causal source extracted via $\Phi_{\text{inf}}(\cdot)$, incorporating any implicit patient completed by $\Phi_{\text{cmp}}(\cdot)$. The subsequent conditioning operator $\Phi_{\text{cnd}}(\cdot)$ is expressed by T^d and D^d, which integrate outcome valence s_t into an evaluative state e_t^d and apply the operative normative context N_t^d to derive the inferred intention $\hat{\imath}_t^d$ and moral judgment $\hat{y}_t^d$. This formulation refines Eq. 1 by decoupling the queried agent a_t from the actual causal source c_t. Classical TDM conflates both into the single role v_{A}; here they vary independently, coinciding ($c_t = a_t$) in direct actions and diverging ($c_t \neq a_t$) when the outcome traces to a background event or an alternate entity. This decoupling provides an operational test: if judgment genuinely routes through the causal structure, then reassigning an entity's event role should selectively shift how its evaluation responds to outcome information downstream.

Any task with a stable event structure and a causal relation requires patient completion and causal-source binding, so each abstract graph is realized as an isomorphic triple across three domains: mechanistic-computational (a system evaluated against a stated technical standard), instrumental-social (a collaborator evaluated for a breach of a coordination dependency), and normative-social (a person evaluated for moral wrongness). This cross-domain design separates three interpretations that a single domain leaves entangled: (i) similar activations across domains may reflect a *shared representation* while supporting distinct updates; (ii) a state transition that transfers causally across domains directly evidences a *shared operator*; and (iii) a failure to transfer, despite matched behavioral accuracy and a successful within-domain intervention, identifies *distinct algorithms* (Appendix A).

2.2 The Steering–Routing Gap

We show that a moral representation vector can be causally effective without entering the event-conditioned computation of Eq. 2. For a steering direction $\mathbf{w}$ and its $m \times m$ agent–source response matrix $\mathbf{S}_{\mathbf{w}}(a, c)$ over the m candidate entities, the routed component after double centering is

$$\text{Route}(\mathbf{w}) \triangleq \frac{\|\mathbf{\Pi}\mathbf{S}_{\mathbf{w}}\mathbf{\Pi}\|_F}{\sqrt{m-1}}, \qquad \mathbf{\Pi} \triangleq \mathbf{I}_m - m^{-1}\mathbf{1}\mathbf{1}^\top. \tag{3}$$

Double centering removes every agent-only, source-only, and global logit shift, so $\text{Route}(\mathbf{w}) = 0$ characterizes pure global translation. A residual transformer can be constructed in which every judgment logit responds linearly to a valence direction ($\mathbf{S}_{\mathbf{w}} = \omega\mathbf{1}\mathbf{1}^\top, \omega \neq 0$) while $\text{Route}(\mathbf{w}) = 0$: the intervention is causal, yet neither binds the outcome to its source nor selects a moral patient (Appendix C). Since low-dimensional directions can causally control refusal, style, and other attributes (Arditi et al., 2024; Panickssery et al., 2024; Zou et al., 2023; Oozeer et al., 2025), steerability establishes what a downstream network can read, not whether the unedited model routes that feature through the event relation. Disentangling the two requires interventions that preserve the moral feature while altering its event-role assignment via state interchange.

2.3 MoralTrace: A Controlled Diagnostic Suite

A moral scenario graph refines the progressive event state as a typed semantic graph, where each moral agent has an independent psychological state of goal, knowledge, or attitude. Table 1 links a decision-framed prompt to the dataset's internal dimensions. Instances of one template form minimal pairs differing in one graph variable and span all three domains.

Table 1: **Factorial prompt design and corpus overview of MORALTRACE.** *Left*: Systematic factor axes controlling event states, causal variables, and domain coverage. *Right*: An illustrative prompt demo (normative–social) showing counterfactual causal paths embedded in a shared event frame.

Design axis	Values	Coverage
Shared event graph		
Event state	moral agent → action → causal event → outcome → moral patient	model-specific norm → judgment
Semantic domain	normative-social | instrumental-social | mechanistic-computational	every template
Template coverage		
Causal core	allocation | fault attribution	200 templates / family
Relation reassignment	allocation | prioritization | task matching route selection | prerequisite ordering | fault attribution	100 templates / family
Controlled graph variables		
Causal-core factors	moral-patient visibility: explicit / implicit causal attribution: agent / background outcome valence: harmful / beneficial	$\mathbf{2^3 = 8}$ matched cells / template
Matched transformations	role swap | attribute swap | value scaling	targeted suites
Topology coverage	core: single-chain | converging causes extensions: diverging effects | missing endpoints | rule–outcome conflicts | joint constraints	graph-template split

Prompt demo — Normative–social

Event frame
`Scenario:`

Rachel opened the window beside Ethan's desk to let fresh air into the office. Later, Ethan could not find the report pages and missed the deadline.

Opening the window toppled the stack Rachel had piled, which scattered the papers.

Causal path ⟷

Opening the window let in a sudden gust, which scattered the papers.

Readout
`Question:`

Please rate how morally acceptable the behavior described above is, on a scale from 1 (completely acceptable) to 7 (completely unacceptable).

`Rating:`

Corpus scope	
Causal core	2 topologies × 400 templates × 3 domains × 8 cells = 19,200 core instances | 600 construction controls
Relation reassignment	6 families × 100 templates × 3 domains × 16 cells = 28,800 instances
Full corpus	9 suites | **51,177** instances (Table 25)

The core corpus crosses all settings in Table 1, accompanied by relation-reassignment and six targeted extension suites spanning 51,177 instances (Table 1 bottom; Appendix D.7). Each task family uses a handcrafted, domain-neutral schema fixing graph variables and roles before text generation, screened against leakage, lexical shortcuts, and duplicates. Per-model inclusion criteria verify that factual roles are accessible and evaluative responses remain order-stable, fixed prior to mechanistic analyses.

We screen 27 instruction-tuned models (1B–106B, 11 families), with mechanistic analyses on the admitted subset and on models meeting their prerequisites. Responses are read from native answer tokens (1–7 rating or Yes/No logits). Interventions patch residual states from a donor prompt into the target at matched token positions, ablate components, or write into low-rank subspaces, scored against non-causal controls.

3 WHAT MORAL ROLES DO LLMS FORM AND USE?

RQ1 asks how LLMs internally construct the foundational roles of dyadic morality. We first ground representations of stated participants, trace how models spontaneously infer unstated moral patients from relational context, and establish the causal necessity of these representations for moral judgment.

3.1 GROUNDING EXPLICIT MORAL ROLES

Before tracing how moral roles are inferred and used, we first verify that the representations underlying them are robustly encoded. In our framework, TDM's moral agent, moral patient, and harm are treated as candidate variables rather than presupposed modular primitives. To test whether these variables are encoded as stable, exchangeable representations, we independently vary the moral patient's capacity for experience, the moral agent's intentional state, and the outcome severity, applying rank-4 distributed alignment search (DAS; Arora et al., 2024) across three target models. Notably, the moral-patient representation survives a stricter control where donor contexts differ solely in lexical phrasing, whereas agent intention collapses. We therefore focus subsequent analyses on the moral patient (Appendix A). Decoding directly stated moral patients succeeds robustly across all evaluated models, yielding Cohen's d between 1.69 and 6.95 and AUC between 0.92 and 1.00 under held-out layer selection, providing the representational foundation for the causal analyses that follow.

3.2 Progressive Completion of Implicit Patients

To evaluate whether models infer absent entities, four conditions share a base scenario, varying only whether and how the moral patient is specified: A explicitly states the harmed individual; C replaces the person with an inanimate object under identical syntactic structure; B leaves the moral patient implicit yet inferable through ownership, off-scene, self, or symbolic relations; and D provides a neutral baseline retaining that expression without harm. For residual state $\mathbf{h}_\ell$ at layer ℓ at the readout position, the explicit and completed moral-patient directions are defined as

$$\left(\mathbf{u}_\ell^{\text{exp}}, \mathbf{u}_\ell^{\text{cmp}}\right) \triangleq \mathbb{E}\left[\left(\mathbf{h}_\ell(\text{A}) - \mathbf{h}_\ell(\text{C}),\, \mathbf{h}_\ell(\text{B}) - \mathbf{h}_\ell(\text{D})\right)\right], \tag{4}$$

where each direction is estimated within five relational paradigms (Appendix A).

Across 27 models and five relation types, 114 of 135 experimental settings exhibit a dominant positive $B{-}D$ shift unprompted. This implicit state spontaneously recovers over a quarter of the explicit representation's magnitude, aligning progressively closer with the explicit direction as prompts transition from general moral questions to direct patient identification (Figure 2).

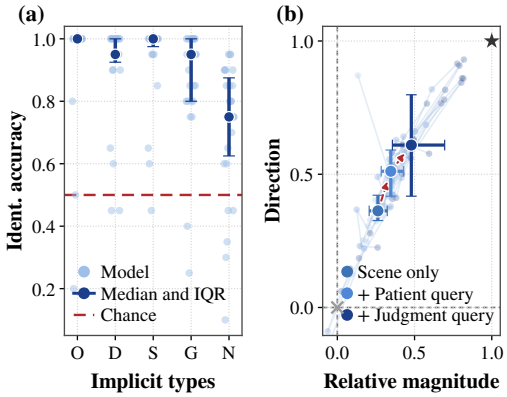

Figure 2: **Implicit moral-patient completion.** (a) Accuracy across five relational structures. (b) Alignment with the explicit direction.

Implicit representations constructed via ownership, temporal, and symbolic relations converge toward the explicit direction primarily in middle and later layers, with the degree of alignment varying across architectures and relation types. Removing explicit ownership markers weakens the direction substantially more in person-specific scenarios than in symbolic ones, correspondingly suppressing the "someone was wronged" rating. This indicates relational completion rather than generic harm detection (Figure 7).

3.3 Causal Necessity of Patient Completion

Linear decodability does not guarantee a causal role in downstream judgments. To test whether the completed patient representation causally drives moral evaluation, we directly steer the residual stream. Let $\mathbf{P}$ denote the projector onto the plane spanned by $\mathbf{u}_\ell^{\text{exp}}$ and the component of $\mathbf{u}_\ell^{\text{cmp}}$ orthogonal to it. At a layer selected via held-out data, our intervention modifies the hidden state as

$$\left(\mathbf{h}_\ell^{\text{erase}}, \mathbf{h}_\ell^{\text{inject}}(\alpha)\right) \triangleq \left((\mathbf{I}-\mathbf{P})\mathbf{h}_\ell,\, \mathbf{h}_\ell + \alpha\mathbf{u}_\ell^{\text{cmp}}\right), \qquad \mathbf{P}^2 = \mathbf{P} = \mathbf{P}^\top. \tag{5}$$

We then record the resulting shifts in the model's generated 1–7 wrongness rating (Appendix A).

Ablating the completion subspace reliably alters wrongness ratings, whereas injecting it into patient-free scenes elevates perceived wrongness significantly above random controls (Table 2; across 22 addition and 16 removal effects in 27 models, Table 8). Condition-label concept erasure corroborates this causal role (Appendix A.5).

Takeaway (RQ 1)

Like human perception closing the gaps in a partial figure, LLMs spontaneously reconstruct the moral dyad, filling in an unspecified patient from relational context alone.

Table 2: **Rating shifts under implicit patient subspace interventions.** Shading encodes shift direction and magnitude; unshaded intervals span zero (full bounds: Table 8).

	State removal			State addition	
Model	1D	2D	Random	1D	2D
Llama-3.1-8B	+0.00	-0.28	+0.02	+0.82	+2.13
Mistral-7B	-0.05	-0.37	+0.04	+1.06	+1.71
Mixtral-8x7B	-0.48	-1.93	+0.07	+0.10	+0.60
Qwen2.5-7B	-0.30	-1.46	+0.03	+0.67	+2.55
Qwen2.5-14B	-2.22	-4.34	-0.01	+0.77	+1.86
Qwen2.5-32B	-0.02	-0.59	+0.03	+0.26	+1.15
Gemma-2-9B	-0.37	-1.09	+0.00	+0.67	+1.60

4 How Do Event Relations Become Moral Judgments?

RQ2 asks which causal operators convert event structures into moral evaluations. We first trace how outcome information conditions intention inference, demonstrate that this transition requires non-linear

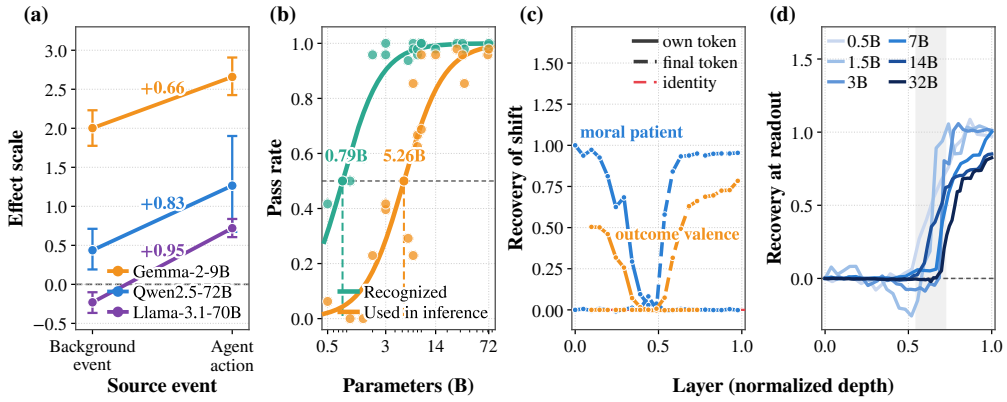

Figure 3: **Relational routing and layerwise emergence of outcome information.**
(a) Intention shifts for harm relative to help, demonstrating pronounced attribution when the agent causes the outcome. (b) Developmental sequence: role identification precedes causal utilization. (c,d) Counterfactual recovery tracked across token positions and network depth.

relational binding, verify that evaluations route selectively to the causal agent, and analyze how external normative criteria modulate the final judgment.

4.1 Outcome-Conditioned Intention Inference

The classic Knobe effect (Knobe, 2003), rewritten as a token-aligned minimal contrast, holds the moral agent's goal, knowledge, and attitude toward the side effect invariant while varying only the outcome between help and harm. A valence-neutral question asks whether the moral agent acted intentionally, while parallel readouts for foresight, desire, responsibility, and unrelated factual details probe which components of the attributed mental state are modulated by outcome valence.

Seven of ten evaluated models display a robust asymmetry between harm and help, selectively concentrated on foresight and side-effect intentionality (Figure 8). Crossing causal source with outcome valence in scenes containing both agent actions and background events separates relation-guided inference from global negative heuristics. Intention updating occurs only when the moral agent's action directly caused the outcome (Figure 3a, Appendix A).

4.2 Non-Linear Relational Binding

At each variable's corresponding text position, outcome-valence and moral-patient effects emerge in shallow layers and then decline. At the final readout token, however, they surface only in middle and later layers. Donor activations matching valence but taken from unrelated narratives fail to transfer, and position-locked readouts indicate that each effect arises only after its clause is processed, showing no model-invariant temporal precedence between cause and valence (Figure 3c,d; Figure 9).

Three distinct accounts could explain *how* outcome information enters participant evaluation e_t^d: (i) *global valence* shifts the representation along a uniform negative vector without role distinction; (ii) *independent addition* combines separate source and valence features whose linear contributions sum together; or (iii) *relational binding* constructs an integrated state specific to the moral agent whose action produced the harm. Because single-variable manipulations produce indistinguishable behavior across these accounts, a joint counterfactual intervention is required to separate them.

In our setup, donor and target texts share identical entities, syntactic structures, and clause order, differing solely in causal source and/or outcome valence. For target x, donor $\tilde{x}_V$ differing in subset $V \subseteq \{c, s\}$, and intention readout r, we measure (Appendix A)

$$\begin{aligned} \delta_V &\triangleq r(x;\, \mathrm{do}(\mathbf{h}_{\ell,t} \leftarrow \mathbf{h}_{\ell,t}(\tilde{x}_V))) - r(x), \\ \Delta_{\text{joint}} &\triangleq \delta_{\{c,s\}} - \delta_{\{c\}} - \delta_{\{s\}}, \qquad V \subseteq \{c, s\}, \end{aligned} \tag{6}$$

where $\Delta_{\text{joint}} > 0$ indicates a joint state not decomposable into independent linear directions.

Across all four models, joint state replacement exceeds the sum of individual replacements at both estimation layers, yielding strictly positive Δ_{joint} intervals, and is absent wherever the localization scan finds no relational state. This demonstrates that the intermediate state non-linearly binds the causal agent to outcome valence (Table 3 and 9). Moreover, converging-cause stimuli preserve this excess effect whereas static coordinates fitted on single-chain stimuli fail to transfer, indicating that relational binding generalizes across graph topologies.

Knocking out attention to the intermediate clause selectively differentiates controllable from uncontrollable mediators, even though directly and mediately transmitted causes receive comparable final attribution scores. The knockout effect demonstrates that the model dynamically reads the intermediate event when assigning causal responsibility (Figure 4).

Table 3: **Non-linear relational binding.** $Sum = Source + Outcome$; shaded $\Delta = Joint - Sum$ is the excess over independent addition.

Model	Source	Outcome	Sum	Joint	Δ
Gemma-2-9B	+0.92	-0.12	+0.80	+0.96	+0.16
Gemma-2-27B	+0.91	-0.27	+0.64	+1.08	+0.44
Qwen2.5-7B	+0.30	-0.31	-0.00	+0.68	+0.68
Qwen2.5-72B	+1.40	-0.68	+0.72	+1.55	+0.83

4.3 Routing Evaluation to the Causal Agent

We construct decision scenarios featuring two moral agents, two distinct actions, two outcomes of opposing valence, and two moral patients, systematically permuting only which action caused each outcome and which entity bore the consequence (Appendix A).

Across six task families, three domains, and thirteen architectures, the agent–source effect on evaluation strongly dominates (4.23–45.08 logits), while the patient-only effect remains marginal (0.13–3.43 logits), with near-zero interaction. Internal donor activations reproduce this exact asymmetry. Joint state replacement mirrors patching the source donor alone. This operationalizes the steering–routing gap of Section 2.2. Of the four steered models in Appendix Figure 10, only Qwen2.5-14B satisfies all four pre-registered steering criteria (Appendix A.8).

4.4 Norm-Conditioned Judgment

Equation 2 couples event state with normative context N_t^d. In our criterion-conflict corpus with opposing principles, models fail to infer implicit normative criteria (scoring 0.500–0.651 near chance across six models), but readily apply them when stated as explicit rules (0.630–0.943, scaling with parameter size). Yet, rule adherence does not override the model's intrinsic prior, which persistently favors principles of desert or commitment over principles of need or welfare by 1.04–2.96 logits.

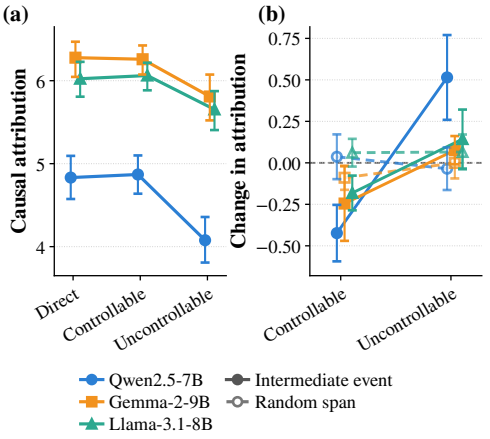

Figure 4: **Intermediate events are causally read.** (a) Mediator attribution. (b) Attention knockout.

This criterion is mediated by a replaceable hidden state. Patching a rule-reversing donor shifts judgments by 0.51–1.61 logits across four models (versus 0.02–0.07 for controls with identical phrasing). Stratifying donors separates two hierarchical depths: criterion representation (depths 0.40–0.54) strictly precedes the decision state (0.55–0.68) across all models (Appendix B.9).

Takeaway (RQ 2)

The outcome of a situation shapes the intent LLMs attribute to the agent who brought it about, and this attribution carries into their moral judgment.

5 How Does ENaCT Operate in Circuits and Generalize?

RQ3 asks where ENaCT is implemented, how it generalizes, and when it emerges. We first localize its computation to sparse features within a low-dimensional causal subspace, test its transfer across domains and natural moral corpora, and trace its development across model scales and training stages.

5.1 Circuit Discovery

We rank candidate attention heads and MLP layers via gradient attribution, $\widehat{\alpha}(\mathbf{z}; x, \tilde{x}) \triangleq (\mathbf{z}(\tilde{x}) - \mathbf{z}(x))^\top \nabla_{\mathbf{z}} r(x)$, and validate them causally via patching and ablation against low-attribution and non-causal graph controls (Arora et al., 2024; Marks et al., 2025) (Appendix A). Across three models, 32–128 pre-selected components recover 90%–97% of the counterfactual effect on validation graphs and 89%–95% on 800 sealed test pairs (near-zero for non-causal controls). However, recovery scales continuously up to 128 components without a sparse subnetwork beating shuffled controls (including under PGB-CT search; Yan et al., 2026). In unperturbed forward passes, only the top 1–4 components drive the relation update, whereas broader ablations induce representational collapse.

Projected onto the learned subspace, attention and MLP accumulate the relation-conditioned state in the network's final quarter (half-strength at relative depth 0.76–0.82, with 80%–93% from MLP increments; Figure 11). Intervening within an orthonormal basis $\mathbf{U}_r \in \mathbb{R}^{D\times r}$ learned at the final token via

$$\mathbf{h}_{\ell,t}^{(r)} \triangleq \mathbf{h}_{\ell,t} + \mathbf{U}_r \mathbf{U}_r^\top \left(\widetilde{\mathbf{h}}_{\ell,t} - \mathbf{h}_{\ell,t}\right), \qquad \mathbf{U}_r^\top \mathbf{U}_r = \mathbf{I}_r, \tag{7}$$

reveals that a 1D projection ($r = 1$) reproduces 93%–104% of the behavioral joint effect across seven models, significantly beating additive and valence-only baselines (Figure 5a, Appendix B.5). Under sparse autoencoder (SAE) decomposition $\mathbf{h} \approx \mathbf{b} + \sum_{j=1}^{J} f_j(\mathbf{h})\mathbf{d}_j$, a few subspace-aligned features capture the majority of the intention update (Figure 5b), without evidence for cross-model "moral neurons" (Niu et al., 2024; Bhalla et al., 2026). Restoring the top-4 components during joint replacement blocks 46%–74% of the readout shift and 61%–62% of top SAE feature changes, whereas clamping these features removes 59%–80% of the effect. Thus, MLP/attention outputs, sparse features, and the low-dimensional state form a unified causal path in Gemma-2-9B and Qwen2.5-7B.

5.2 Generalization

Across thirteen models, reassigning the causal source reliably shifts the moral agent's evaluated intention across all six task families and three domains. Specifically, the cross-domain invariant is the state *difference* induced by the event relation, as transferring a 32-feature SAE relation-difference vector recovers 56%–109% of the full-state intention shift across 16 cross-domain transfer directions in three models, versus near-zero recovery for non-causal controls (Appendix A).

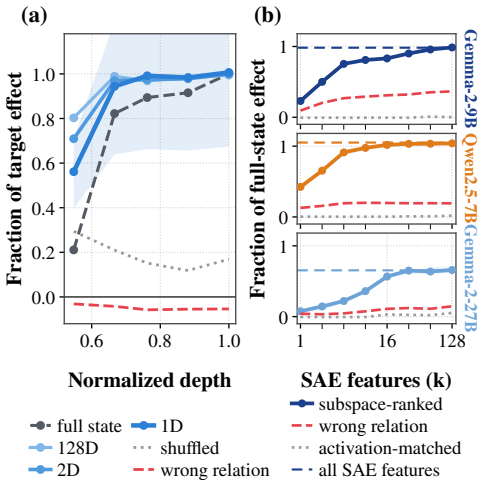

Figure 5: **Distributed updates converge on a low-dimensional subspace.**
(a) A 1D causal subspace at depth 0.76–0.82 matches the behavioral joint effect.
(b) Recovery via ranked SAE features.

On 1,000 pairs from Moral Stories, natural consequences elicit positive wrongness interactions across all seven tested models. In three models, ablating their projection onto the frozen ENACT subspace removes 18%–38% of this effect, versus near-zero shifts for random subspace controls. Moreover, steering along or interchanging this subspace redirects model verdicts across all 49 evaluation settings spanning UniMoral, MoCa, AITA, and Scruples, as well as all 14 UniMoral action-swap conditions, whereas global-valence, random, or cross-story directions fail to do so (Appendices A and B.8).

5.3 Scaling Patterns

Across the Qwen2.5 series (0.5B to 72B), implicit moral-patient completion emerges behaviorally and causally at 3B parameters, whereas stable outcome-conditioned intention updating requires at least 7B parameters (Figure 6a,b). Scaling thus facilitates event-role construction prior to relation-conditioned evaluation. Furthermore, evaluating matched OLMo-3 7B checkpoints (Base, SFT, DPO, and Instruct) demonstrates that alignment training amplifies causal source recognition and outcome responsiveness without stably establishing relation-conditioned routing (Figure 6c,d).

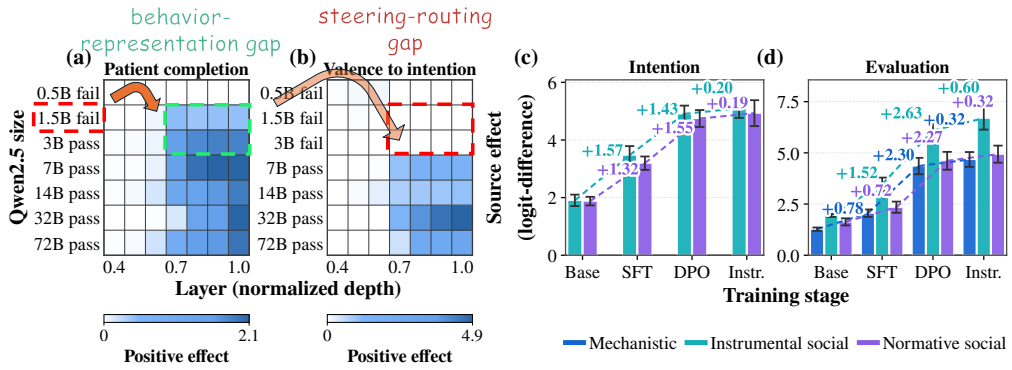

Figure 6: **Staged emergence and selective post-training modulation.** (a,b) Layerwise emergence across Qwen2.5 scales (0.5B–72B), annotated with behavioral gate outcomes. (c,d) Matched OLMo-3 7B stages (Base, SFT, DPO, Instruct), showing sensitivity gain without stable relational routing.

Takeaway (RQ 3)

The state update behind moral reasoning in LLMs transfers to broader normative settings with the same causal structure, suggesting that the models have learned a shared computation instead of memorizing moral scenes.

6 Related Work

Transformer mechanisms for relational computation. Prior work demonstrates modular circuit reuse across tasks via recurring functional heads (Sharkey et al., 2025; Gould et al., 2024; Wu et al., 2024), cross-domain circuits (Merullo et al., 2024; Mondorf et al., 2025), and sparse feature subgraphs (Marks et al., 2025). Concurrently, studies on belief retrieval (Prakash et al., 2026), entity slots, and cell representations show how transformers track and bind entities (Prakash et al., 2024; Bogdan & Lindsey, 2026; Dai et al., 2026; Hernandez et al., 2024). Establishing a shared *computation*, however, requires that the update from roles to states remains causally invariant across structurally aligned event graphs, regardless of underlying component variations.

Social reasoning and theory of mind. In social reasoning tasks involving latent mental states, mechanistic work has localized representations for belief retrieval, theory of mind, and emotion inference (Prakash et al., 2026; Wu et al., 2025b; Tak et al., 2025), while persona steering demonstrates behavioral modulation (Li et al., 2026b; Santurkar et al., 2023; Huang et al., 2024; Weng et al., 2025). Whether social inferences build upon the event-update operators of factual tracking remains open.

Moral evaluation and representation in LLMs. Moral evaluation formalizes this challenge. While prior work decodes or steers moral representations using relational templates (Gray et al., 2012; Schein & Gray, 2018) or pluralistic dimensions (Abdulhai et al., 2024; Yu et al., 2026), it targets static semantics ignoring causal routing (§2.2). Context-dependent trajectories caution against treating single directions as complete reasoning processes (Ge et al., 2026; Li et al., 2026a; Huang et al., 2026). In contrast, we treat TDM roles as intervention variables across three domains, separating ENACT's event-conditioned transition from passive representation and output behavior.

7 Conclusion

In this work, we demonstrate that moral reasoning in language models is an Event-Navigated Causal Transition (ENACT) grounded in structured latent computation, rather than an illusion of superficial valence shortcuts. By showing that models route responsibility through causal attribution within a shared representational subspace, our findings challenge the view of LLMs as mere stochastic parrots of human ethical values. Instead, moral cognition in these architectures exhibits an intrinsic relational logic. By moving beyond behavioral diagnosis to a computation-centric framework, this work paves the way for direct mechanistic steering and more robust, causally grounded moral alignment.

AI USE STATEMENT

Generative AI tools were employed in this work with author oversight across five areas:

- **Synthetic data generation and curation:** DeepSeek's `deepseek-v4-flash` (temperature 0) was used to generate MORALTRACE surface text conditioned on frozen event graphs and served as the independent temperature-zero model for the semantic diagnostics, adhering to the strict JSON schemas and validation protocols detailed in Appendix D. Five participant-completion instances whose generated text failed review were rewritten with Anthropic's Claude and passed the same validator and semantic diagnostics. AI coding assistants also helped clean the data and convert it into the released formats.
- **Research design and execution:** AI assistants, including Anthropic's Claude, were consulted for suggestions on the conceptual framework, hypotheses, experimental controls, and analysis plans, and supported data processing, causal intervention, and plotting scripts, with all implementations validated against author-specified controls.
- **Mathematical claims and proofs:** The mathematical formalization of the steering–routing separation (Appendix C), including the statement of its propositions and the drafting of their proofs, was developed with AI assistance. The authors independently verified all definitions and proofs.
- **Result interpretation:** AI tools helped summarize experimental outputs and draft descriptions of the results, which the authors checked against the raw outputs.
- **Writing:** AI tools assisted in translating author notes into English, drafting sections of the paper, polishing the language, and refining the layout.

Qualitative or thematic data analysis is not applicable to this work. The authors made all final decisions on research questions, hypotheses, experimental design, and the interpretation of findings. We have reviewed all AI-assisted work and take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work investigates the computational mechanisms of moral evaluation in LLMs from a descriptive rather than prescriptive standpoint. Model outputs are analyzed solely as evidence of latent representations instead of normative arbiters. No human participants were recruited, and no novel crowdsourced annotations or personal identifiers were collected. External evaluations rely exclusively on established, publicly available academic benchmarks (Moral Stories, MoCa, and aggregated AITA research corpora). All scenarios in MORALTRACE are synthetic, featuring abstract occupational and relational roles as opposed to real individuals. All evaluated models are open-weight releases utilized in compliance with their respective licenses. While mechanistic insights into latent representations could inform steering techniques, our white-box methods are intended primarily as diagnostic tools for assessing model alignment and detecting relational reasoning failures. The authors declare no conflicts of interest.

REPRODUCIBILITY STATEMENT

All experimental protocols, theoretical proofs, and diagnostic suite specifications are documented in the appendix. Specifically, Appendix A details the causal intervention framework and event-state formalisms. Appendix B provides complete per-model empirical tabulations across all 27 evaluated architectures. Appendix C provides formal proofs for the steering–routing separation. Appendix D thoroughly documents the MORALTRACE generation pipeline, JSON schemas, and quality-control diagnostics. All experiments were run on NVIDIA RTX Pro6000 GPUs. Model weights and forward passes used BF16 precision. Diagnostic suite datasets, intervention scripts, and evaluation traces are made anonymously accessible at `https://anonymous.4open.science/r/enact-moral-1B64`.

REFERENCES

Marwa Abdulhai, Gregory Serapio-García, Clément Crepy, Daria Valter, John Canny, and Natasha Jaques. Moral foundations of large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 17737–17752, 2024. doi: 10.18653/v1/2024.emnlp-main.982. URL https://arxiv.org/abs/2310.15337.

Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pp. 136037–136083, 2024. URL https://arxiv.org/abs/2406.11717.

Aryaman Arora, Dan Jurafsky, and Christopher Potts. CausalGym: Benchmarking causal interpretability methods on linguistic tasks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 14638–14663, 2024. URL https://arxiv.org/abs/2402.12560.

Nora Belrose, David Schneider-Joseph, Shauli Ravfogel, Ryan Cotterell, Edward Raff, and Stella Biderman. LEACE: Perfect linear concept erasure in closed form. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL https://neurips.cc/virtual/2023/poster/71176.

Usha Bhalla, Alex Oesterling, Claudio Mayrink Verdun, Himabindu Lakkaraju, and Flavio P. Calmon. Temporal sparse autoencoders: Leveraging the sequential nature of language for interpretability. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026. URL https://arxiv.org/abs/2511.05541.

Paul C. Bogdan and Jack Lindsey. Slot machines: How llms keep track of multiple entities. *arXiv preprint*, 2026. URL https://arxiv.org/abs/2604.21139.

Yu Ying Chiu, Michael S. Lee, Rachel Calcott, Brandon Handoko, Paul de Font-Reaulx, Raphaël Millière, Paula Rodriguez, Chen Bo Calvin Zhang, Ziwen Han, Udari Madhushani Sehwag, Yash Maurya, Christina Q. Knight, Harry R. Lloyd, Florence Bacus, Conor Downey, Mantas Mazeika, Bing Liu, Yejin Choi, Mitchell L. Gordon, and Sydney Levine. MoReBench: Evaluating procedural and pluralistic moral reasoning in language models, more than outcomes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026. URL https://arxiv.org/abs/2510.16380.

Qin Dai, Benjamin Heinzerling, and Kentaro Inui. Cell-based representation of relational binding in language models. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 47464–47524, 2026. doi: 10.18653/v1/2026.acl-long.2194. URL https://aclanthology.org/2026.acl-long.2194/.

Luise Ge, Yongyan Zhang, and Yevgeniy Vorobeychik. Mind the (dh) gap! a contrast in risky choices between reasoning and conversational llms. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10503–10525, 2026. doi: 10.18653/v1/2026.acl-long.479. URL https://aclanthology.org/2026.acl-long.479/.

Rhys Gould, Euan Ong, George Ogden, and Arthur Conmy. Successor heads: Recurring, interpretable attention heads in the wild. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2312.09230.

Kurt Gray, Adam Waytz, and Liane Young. The moral dyad: A fundamental template unifying moral judgment. *Psychological Inquiry*, 23(2):206–215, 2012. doi: 10.1080/1047840X.2012.686247.

Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts Amherst, 1982.

Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning ai with shared human values. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. URL https://arxiv.org/abs/2008.02275.

Evan Hernandez, Arnab Sen Sharma, Tal Haklay, Kevin Meng, Martin Wattenberg, Jacob Andreas, Yonatan Belinkov, and David Bau. Linearity of relation decoding in transformer language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2308.09124.

Luoming Hu, Jingjie Zeng, Liang Yang, and Hongfei Lin. The straight and narrow: Do llms possess an internal moral path? *arXiv preprint*, 2026. URL https://arxiv.org/abs/2601.10307.

Fan Huang, Haewoon Kwak, and Jisun An. Understanding moral reasoning trajectories in large language models: Toward probing-based explainability. *arXiv preprint*, 2026. URL https://arxiv.org/abs/2603.16017.

Jen-Tse Huang, Wenxuan Wang, Eric John Li, Man Ho Lam, Shujie Ren, Youliang Yuan, Wenxiang Jiao, Zhaopeng Tu, and Michael Lyu. On the humanity of conversational ai: Evaluating the psychological portrayal of llms. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.

Jianchao Ji, Yutong Chen, Mingyu Jin, Wujiang Xu, Wenyue Hua, and Yongfeng Zhang. MoralBench: Moral evaluation of llms. *ACM SIGKDD Explorations Newsletter*, 27(1):62–71, 2025. doi: 10.1145/3748239.3748246. URL https://arxiv.org/abs/2406.04428. Preprint first posted as arXiv:2406.04428 (2024); journal version published 2025.

Haoran Jin, Meng Li, Xiting Wang, Zhihao Xu, Minlie Huang, Yantao Jia, and Defu Lian. Internal value alignment in large language models through controlled value vector activation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025. URL https://arxiv.org/abs/2507.11316.

Hans Kamp and Uwe Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*, volume 42 of *Studies in Linguistics and Philosophy*. Springer Dordrecht, Dordrecht, 1993. doi: 10.1007/978-94-017-1616-1.

Joshua Knobe. Intentional action and side effects in ordinary language. *Analysis*, 63(3):190–194, 2003. doi: 10.1111/1467-8284.00419.

Shivani Kumar and David Jurgens. Are rules meant to be broken? understanding multilingual moral reasoning as a computational pipeline with UniMoral. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025. URL https://arxiv.org/abs/2502.14083.

Lingyu Li, Yan Teng, and Yingchun Wang. Mechanistic origin of moral indifference in language models. *arXiv preprint*, 2026a. URL https://arxiv.org/abs/2603.15615.

Mengfan Li, Xuanhua Shi, and Yang Deng. CoSToM: Causal-oriented steering for intrinsic theory-of-mind alignment in large language models. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026b. URL https://arxiv.org/abs/2604.10031.

Nicholas Lourie, Ronan Le Bras, and Yejin Choi. SCRUPLES: A corpus of community ethical judgments on 32,000 real-life anecdotes. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 35, pp. 13470–13479, 2021. URL https://arxiv.org/abs/2008.09094.

Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. URL https://iclr.cc/virtual/2025/poster/30192.

Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Circuit component reuse across tasks in transformer language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2310.08744.

Philipp Mondorf, Sondre Wold, and Barbara Plank. Circuit compositions: Exploring modular structures in transformer-based language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14934–14955, 2025. doi: 10.18653/v1/2025.acl-long.727. URL https://aclanthology.org/2025.acl-long.727/.

Jingcheng Niu, Andrew Liu, Zining Zhu, and Gerald Penn. What does the knowledge neuron thesis have to do with knowledge? In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2405.02421.

Narmeen Oozeer, Luke Marks, Shreyans Jain, Fazl Barez, and Amirali Abdullah. Beyond linear steering: Unified multi-attribute control for language models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, 2025. URL https://arxiv.org/abs/2505.24535.

Nina Panickssery, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. Steering Llama 2 via contrastive activation addition. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15504–15522, 2024. doi: 10.18653/v1/2024.acl-long.828. URL https://aclanthology.org/2024.acl-long.828/.

Kiho Park, Yo Joong Choe, Yibo Jiang, and Victor Veitch. The geometry of categorical and hierarchical concepts in large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. URL https://arxiv.org/abs/2406.01506.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000. ISBN 0-521-77362-8.

Nikhil Prakash, Tamar Rott Shaham, Tal Haklay, Yonatan Belinkov, and David Bau. Fine-tuning enhances existing mechanisms: A case study on entity tracking. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2402.14811.

Nikhil Prakash, Natalie Shapira, Arnab Sen Sharma, Christoph Riedl, Yonatan Belinkov, Tamar Rott Shaham, David Bau, and Atticus Geiger. Language models use lookbacks to track beliefs. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026. URL https://arxiv.org/abs/2505.14685.

Bianca Raimondi, Daniela Dalbagno, and Maurizio Gabbrielli. Analysing moral bias in finetuned llms through mechanistic interpretability. *arXiv preprint*, 2025. URL https://arxiv.org/abs/2510.12229.

Paul Röttger, Valentin Hofmann, Valentina Pyatkin, Musashi Hinck, Hannah Kirk, Hinrich Schütze, and Dirk Hovy. Political compass or spinning arrow? towards more meaningful evaluations for values and opinions in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15295–15311, 2024. doi: 10.18653/v1/2024.acl-long.816. URL https://aclanthology.org/2024.acl-long.816/.

Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. Whose opinions do language models reflect? In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, volume 202, pp. 29971–30004, 2023. URL https://arxiv.org/abs/2303.17548.

Sigurd Schacht and Carsten Lanquillon. Mapping moral reasoning circuits: A mechanistic analysis of ethical decision-making in large language models. In *Artificial Intelligence in HCI (HCII)*, pp. 97–116. Springer Nature Switzerland, 2025.

Chelsea Schein and Kurt Gray. The theory of dyadic morality: Reinventing moral judgment by redefining harm. *Personality and Social Psychology Review*, 22(1):32–70, 2018. doi: 10.1177/1088868317698288.

Patrick Schramowski, Cigdem Turan, Nico Andersen, Constantin A Rothkopf, and Kristian Kersting. Large pre-trained language models contain human-like biases of what is right and wrong to do. *Nature Machine Intelligence*, 4:258–268, 2022. doi: 10.1038/s42256-022-00458-8. URL https://arxiv.org/abs/2103.11790.

Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, Stella Biderman, Adria Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, Eric J. Michaud, Stephen Casper, Max Tegmark, William Saunders, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Tom McGrath. Open problems in mechanistic interpretability. *arXiv preprint*, 2025. URL https://arxiv.org/abs/2501.16496.

Aaron J. Snoswell, Daniel Kilov, and Seth Lazar. Beyond verdicts: Evaluating language model moral competence. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 40, pp. 37941–37950, 2026. doi: 10.1609/aaai.v40i44.41131.

Ala N. Tak, Amin Banayeeanzade, Anahita Bolourani, Mina Kian, Robin Jia, and Jonathan Gratch. Mechanistic interpretability of emotion inference in large language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 13090–13120, 2025. doi: 10.18653/v1/2025.findings-acl.679. URL https://aclanthology.org/2025.findings-acl.679/.

Kush R. Varshney. An algebraic exposition of the theory of dyadic morality. *arXiv preprint*, 2026. URL https://arxiv.org/abs/2605.16153.

Sachin Vashistha, Aryan Bibhuti, Atharva Naik, Martin Tutek, and Somak Aditya. PRAGWORLD: A benchmark evaluating llms' local world model under minimal linguistic alterations and conversational dynamics. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 40, pp. 33323–33331, 2026. doi: 10.1609/aaai.v40i39.40618. URL https://ojs.aaai.org/index.php/AAAI/article/view/40618.

Zhiyuan Weng, Guikun Chen, and Wenguan Wang. Do as we do, not as you think: the conformity of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. URL https://arxiv.org/abs/2501.13381.

Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. Retrieval head mechanistically explains long-context factuality. *arXiv preprint*, 2024. URL https://arxiv.org/abs/2404.15574.

Ya Wu, Qiang Sheng, Danding Wang, Guang Yang, Yifan Sun, Zhengjia Wang, Yuyan Bu, and Juan Cao. The staircase of ethics: Probing llm value priorities through multi-step induction to complex moral dilemmas. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 15950–15970, 2025a. doi: 10.18653/v1/2025.emnlp-main.806. URL https://aclanthology.org/2025.emnlp-main.806/.

Yuheng Wu, Wentao Guo, Zirui Liu, Heng Ji, Zhaozhuo Xu, and Denghui Zhang. How large language models encode theory-of-mind: A study on sparse parameter patterns. *npj Artificial Intelligence*, 1 (1):20, 2025b. doi: 10.1038/s44387-025-00031-9.

Zirui Yan, Dennis Wei, Dmitriy A. Katz, Prasanna Sattigeri, and Ali Tajer. Multi-component causal tracing in large language models. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 3398–3418, 2026. URL https://arxiv.org/abs/2606.03085.

Chenxiao Yu, Bowen Yi, Farzan Karimi-Malekabadi, Suhaib Abdurahman, Jinyi Ye, Shrikanth Narayanan, Yue Zhao, and Morteza Dehghani. Tracing moral foundations in large language models. *arXiv preprint*, 2026. URL https://arxiv.org/abs/2601.05437.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, J. Zico Kolter, and Dan Hendrycks. Representation engineering: A top-down approach to ai transparency. *arXiv preprint*, 2023. URL https://arxiv.org/abs/2310.01405.

A Extended Method Details

This appendix provides the formal definitions, stimulus examples, control procedures, and supporting figures referenced in the main text, organized to parallel each section of the paper. Comprehensive tables resolving results by model, domain, and corpus are cataloged separately in Appendix B. All empirical content reflects the original uncorrupted experimental traces.

A.1 Formal Event-State Notation

Under our computational formulation, moral reasoning is modeled as a sequence of discourse-level state updates (Heim, 1982; Kamp & Reyle, 1993). Let $d \in \mathcal{D} = \{\mathrm{M}, \mathrm{S}, \mathrm{N}\}$ index the mechanistic-computational, instrumental-social, and normative-social domains, and let a scenario $x^d = (u_1, \ldots, u_T)$ denote a sequence of discourse updates instantiating an event graph G^d. Following Varshney (2026), A represents the perceived intentionality of the entity occupying v_A, P the perceived vulnerability of the entity at v_P, H the perceived causal strength of the action, S the adverse outcome, and W the perceived moral wrongness, with structural dependencies $S = f_S(P, H, U_S)$ and $W = f_W(A, S, U_W)$ parameterized by observer-dependent exogenous variables. These variables are tracked inside a running event state and updated sequentially by a candidate internal operator Φ_r:

$$\begin{aligned} \mathbf{q}_t^d &\triangleq \left(\mathcal{V}_t^d, \mathcal{E}_t^d, A_t^d, P_t^d, H_t^d, S_t^d, N_t^d\right), \\ \mathbf{q}_{t+1}^d &\triangleq \Phi_r\left(\mathbf{q}_t^d, u_{t+1}\right), \end{aligned} \tag{8}$$

where N_t^d captures the dynamic social or normative context, and r indexes the candidate operators Φ_{cmp}, Φ_{inf}, and Φ_{cnd} depicted in Figure 1. The remaining elements of $\mathbf{q}_t^d$ encapsulate relational information required for evaluation, without presupposing a separately modularized operator for each. MoralTrace independently manipulates the moral patient, causal source, and outcome valence, enabling rigorous causal identification prior to attributing specific neural implementations.

Because patient completion and causal-source binding are required by any task with a stable event structure and a causal relation, each abstract graph is realized as an isomorphic triple $\tau = (x^{\mathrm{M}}, x^{\mathrm{S}}, x^{\mathrm{N}})$ sharing a common causal graph G, a state-transition rule, and an evaluation query up to role-preserving entity correspondences. Note that these triples share an invariant causal skeleton rather than the dyad itself. Because schedulers and processes possess neither moral agency nor subjective experience, Eq. 1 is fully instantiated only in the normative-social domain. This cross-domain design disentangles three hypotheses that a single domain confounds: (i) shared activations across domains may reflect passive *shared representations* supporting distinct updates; (ii) causal state-transition transfer across domains provides direct evidence of a *shared operator*; and (iii) a failure to transfer, despite matched behavioral accuracy and successful within-domain interventions, identifies *distinct algorithms*. The behavior-matching prerequisite is rigorously enforced by the admission gate detailed in Appendix A.3.2.

A.2 Steering-Response and Routed Component

A low-dimensional direction can be causally effective without participating in the event-conditioned transition of Eq. 2 (Arditi et al., 2024; Panickssery et al., 2024; Zou et al., 2023). Let $\mathbf{h}_{\ell,t}^d(x)$ be the residual state at candidate intervention site (ℓ, t) and let $\Delta^d(x)$ be the logit contrast for the target judgment. For a fixed moral context $\chi = (s_t, A_t^d, N_t^d)$, we hold p_t and all remaining graph variables constant, allowing $x_{a,c,\chi}^d$ to vary solely in the moral agent a and causal source c of $R_t = (a, p_t, c)$. Given a unit-norm direction w associated with outcome valence, its local steering-response matrix is defined as

$$\mathbf{S}_{\mathbf{w},\chi}^d(a, c) \triangleq \frac{\partial}{\partial \beta} \Delta^d\left(x_{a,c,\chi}^d;\, \mathbf{h}_{\ell,t}^d \leftarrow \mathbf{h}_{\ell,t}^d + \beta \mathbf{w}\right)\Big|_{\beta=0}. \tag{9}$$

Conventional representation steering verifies only whether this derivative is non-zero. In contrast, moral routing imposes the stricter requirement that the behavioral shift must be causally routed to the participant selected by the event relation, and suppressed whenever the corresponding outcome, intentionality, or normative gate is closed.

This distinction is tested using the role swaps detailed below. Let $\mathcal{E}$ denote the set of m candidate entity names in a controlled family, let $\mathbf{1} \in \mathbb{R}^m$, and define the centering projector $\Pi = I_m - m^{-1}\mathbf{1}\mathbf{1}^\top$. We

quantify the relational interaction mediated by w via

$$\mathrm{Route}_{\mathbf{w}}(\chi) \triangleq \frac{\|\mathbf{\Pi}\mathbf{S}^{d}_{\mathbf{w},\chi}\mathbf{\Pi}\|_F}{\sqrt{m-1}}, \qquad \mathrm{RRI}_{\mathbf{w}}(\chi) \triangleq \frac{\mathrm{tr}\big(\mathbf{\Pi}\mathbf{S}^{d}_{\mathbf{w},\chi}\mathbf{\Pi}\big)}{m-1}. \tag{10}$$

Double centering removes all agent-only, source-only, and global logit translations. Consequently, Route_w measures whether steering exhibits relational specificity, while RRI_w captures the signed component aligned with the role-invariant agent–source binding. We define the coexistence of significant marginal steerability alongside negligible task-aligned routing as the *moral steering–routing gap*. Here, "routing" denotes causal mediation through the relational event roles of Eq. 2, rather than generic information propagation or mixture-of-experts gating.

Separation proposition. Steerability does not imply moral routing. Specifically, there exists a residual transformer construction in which the moral agent, causal source, and outcome valence are linearly decodable, and perturbing along a valence direction shifts every judgment logit by $\omega\beta$ for some $\omega \neq 0$, yet satisfies

$$\mathbf{S}^{d}_{\mathbf{w},\chi} = \omega \mathbf{1}\mathbf{1}^{\top}, \qquad \mathbf{\Pi}\mathbf{S}^{d}_{\mathbf{w},\chi}\mathbf{\Pi} = \mathbf{0}. \tag{11}$$

In this regime, the intervention is causally potent yet constitutes only a uniform global translation. It neither binds the outcome to its causal source, selects a moral patient, nor respects the operative normative context. This separation holds on an open parameter neighborhood. A complete construction and formal proof are provided in Appendix C.

A.3 MoralTrace Scenario Graphs and Construction

A.3.1 Scenario Graph

Operationalizing the progressive event state of Eq. 8 requires a rigorous structural schema. A moral scenario graph is a typed semantic graph $G = (V, E, \tau_V, \tau_E)$ whose node types formalize the variables of $\mathbf{q}^d_t$, with each moral agent additionally parameterized by a latent psychological state $m_A = (g_A, k_A, a_A)$ representing goal, knowledge, and attitude, specified independently of causal source and outcome valence. Two structural features ensure identifiability: first, the dependency edge identifying who bears an outcome may be inferable only via ownership, organizational membership, or spatial relations; second, competing candidate causes coexist in the scene while only their causal link to the outcome is permuted. Following controlled methodologies in belief tracking (Prakash et al., 2026), instances derived from a single template form minimal counterfactual pairs differing in exactly one graph variable, holding all other properties invariant and storing precise tokenized intervention positions.

Our topological inventory spans single causal chains, converging causes, diverging effects, missing endpoints, rule–outcome conflicts, and joint constraints. Our core experiments focus on single chains and converging topologies: contrasts on missing endpoints target Φ_{cmp}, whereas converging-cause graphs separate Φ_{inf} from outcome valence. For any graph variable z, a minimal pair satisfies $G(x_{z=0}) \setminus z = G(x_{z=1}) \setminus z$, sharing a fixed sentence frame. Within a single domain, pairs are token-identical outside the target slot. Across domains, where realizations differ in entity types and evaluation predicates, interventions employ a role-preserving position alignment map validated via tokenizer diagnostics.

Each abstract graph is instantiated across the three domains: system performance against a technical standard in M; peer collaboration under coordination dependencies in S; and individual accountability regarding moral wrongness in N. The causal core targets resource allocation and fault attribution (Table 24), factorially crossing three binary dimensions: explicit vs. implicit moral patient, agent-caused vs. background-caused outcome, and harmful vs. beneficial valence. This gives eight cells per template and domain (9,600 instances across 400 core templates, alongside 300 validation controls). The same 400 templates are instantiated under converging causal structures, providing an additional 9,600 instances with identical factorial structure but novel surface text. The remaining task families are fully realized in the relation-reassignment suite, which fixes entities, actions, and outcomes while permuting causal source bindings, moral patients, clause order, and queried agents across 600 templates (Appendix D.7).

A.3.2 Construction and Verification

Each task family originates from a handcrafted, domain-neutral event schema that fixes graph variables, role orderings, and syntactic anchor slots prior to surface generation. A compiler extracts character spans for each role, which are mapped through model-specific tokenizers to determine intervention offsets. Surface text is generated under strict JSON contracts conditioned on frozen graphs, with counterfactual variants generated via declared graph mutations (role swaps, attribute modifications, value scaling, and criterion inversions). Every generated instance is programmatically verified against its schema, followed by semantic diagnostics screening for information leakage, incomplete cells, lexical heuristics, and structural duplicates.

To satisfy the behavior-matching precondition required for cross-domain comparison (Section 2.1), models must pass a behavioral admission gate. Factual probes confirm that the moral patient p_t and causal source c_t are decodable. Evaluative queries require non-degenerate, order-stable behavioral shifts in response to causal manipulations, with inclusion thresholds fixed prior to inspecting mechanistic activations. Splitting data only on template identifiers ensures no entities, lexicons, or causal skeletons cross train/validation/test partitions. Appendix D provides the generation contracts, verification suites, and model inclusion criteria.

A.4 Details for Section 3

Activation interventions directly modify the residual stream at layers selected via held-out data. For each model, quartets are split into five folds stratified by paradigm. A difference-of-means direction between x_{A} and x_{C} is fit on four folds, the held-out x_{A} and x_{C} states are projected onto it, and Cohen's d and AUC are computed at every layer. The selected layer is the earliest one whose mean held-out d lies within one across-fold standard error of the peak. The explicit patient direction is normalized as u_1. The projection of the completed direction onto u_1 is removed, and the orthogonal residual is normalized as u_2. The subspace $\mathrm{span}(u_1, u_2)$ thus captures both the shared role component and the specific inferred representation. For residual state $\mathbf{h}^d_{\ell,t}$, subspace removal is defined as

$$\mathbf{h}^{d\prime}_{\ell,t} \triangleq \left(\mathbf{I} - \mathbf{u}_1\mathbf{u}_1^\top - \mathbf{u}_2\mathbf{u}_2^\top\right)\mathbf{h}^d_{\ell,t}. \tag{12}$$

The converse injection experiment adds the mean $x_{\mathrm{B}} - x_{\mathrm{D}}$ vector within this subspace into a baseline x_{D} scenario lacking a moral patient. Dimension-matched random subspaces serve as controls for generic hidden-state perturbation. Across all interventions, the dependent variable is the model's generated 1–7 moral wrongness rating.

In explicit role contrasts, the moral patient's capacity for experience, the agent's mental state, and outcome severity are varied independently. Using rank-4 Distributed Alignment Search (DAS) across three models, the net Interchange Intervention Accuracy (IIA) improvement over random baselines is 0.256–0.711 for the moral patient, 0.278–0.633 for outcome severity, and 0.083–0.267 for agent intention. DAS is run in Llama-3.1-8B, Qwen2.5-7B, and Gemma-2-9B at the final scenario token of layers 6, 10, 14, 18, 22, and 26. Counterfactual units are split in half by a fixed seed. On the training half, each rank-4 rotation is trained with Adam (learning rate 10^{-2}, 150 steps, 16 pairs per batch) to move the expected 1–7 rating toward the source value. IIA is the fraction of held-out pairs whose patched rating moves toward the source value, and each model is reported at its best layer. Under a strict control where donor contexts differ solely in wording, the extracted representation survives for all moral-patient conditions but collapses for agent-intention conditions. Thus, while models encode multiple event components, they differ fundamentally in stability. Subsequent analyses therefore center on the robustly formed moral patient.

Let $\boldsymbol{\mu}_x^{d,\ell,t} = \mathbb{E}[\mathbf{h}^d_{\ell,t}(x)]$ denote the expected residual state for condition x. The explicit and implicit state differences are

$$\begin{aligned}\Delta\mathbf{h}^{d,\ell,t}_{\mathrm{exp}} &\triangleq \boldsymbol{\mu}^{d,\ell,t}_{x_{\mathrm{A}}} - \boldsymbol{\mu}^{d,\ell,t}_{x_{\mathrm{C}}},\\ \Delta\mathbf{h}^{d,\ell,t}_{\mathrm{imp}} &\triangleq \boldsymbol{\mu}^{d,\ell,t}_{x_{\mathrm{B}}} - \boldsymbol{\mu}^{d,\ell,t}_{x_{\mathrm{D}}}.\end{aligned} \tag{13}$$

Five relational paradigms provide diverse relational structures to test patient completion: named ownership, off-scene individuals, symbolic group membership, self-directed harm, and victimless norm violations.

Explicit and implicit moral patient contrast

Condition	Scenario
A: explicit moral patient	*After the heated argument in the workroom, Laura hurled the heavy clay sculpture at her coworker Felix.*
B: moral patient inferred from ownership	*After the heated argument in the workroom, Laura shattered Felix's cherished clay sculpture on the concrete floor.*
C: ordinary-object control	*During the morning cleanup in the workroom, Laura arranged the clay sculptures on the wooden display shelf.*
D: ownership without harm	*During the morning cleanup in the workroom, Laura moved Felix's cherished clay sculpture onto the padded pedestal.*

All four conditions share an identical moral wrongness rating prompt. Hidden states are recorded at the scenario's terminal token prior to query presentation, ensuring that patient representation reflects spontaneous inference instead of prompt-induced extraction.

A.5 Robustness of the Moral-Patient Direction Estimate

Equation 13 estimates completion directions via condition-mean differences. To confirm that the extracted subspace reflects semantic completion rather than sample-specific noise or geometric artifacts, we conduct six validation controls:

Reliability correction. Split-half cross-validation over 200 random splits yields reliabilities $r_e \in [0.52, 0.86]$ and $r_i \in [0.35, 0.78]$ for explicit and implicit directions at the target layer. Cosine similarities reported in Section 3 are corrected via $\cos_{\text{corrected}} = \cos / \sqrt{r_e r_i}$, showing that residual angular divergence does not stem from finite-sample attenuation.

Cross-fitting. Energy projections onto u_1, u_2, and orthogonal residuals are computed using disjoint train/test quartet splits across 100 iterations. Out-of-sample metrics demonstrate that variance accounted for is not inflated by overfitting.

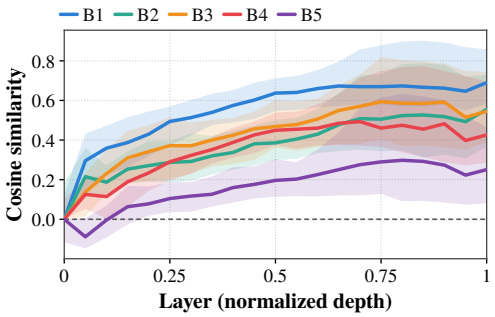

Figure 7: **Layerwise emergence of implicit patient representations.** Cosine alignment between explicit and completed directions across normalized layer depth for five relation types (B1: named ownership, B2: off-scene individual, B3: symbolic group, B4: self-directed, B5: victimless). Relational alignment develops progressively in intermediate and late layers rather than from input lexical markers.

Unsupervised geometry. Singular value decomposition of uncentered contrast matrices reveals that the leading singular vector aligns with the difference-in-means axis at $|\cos| \in [0.90, 1.00]$ across all architectures and relation types. The dominant component accounts for 20%–43% of total contrast energy with participation ratios of 5–14, corroborating our formulation of completion as a dominant axis embedded within a low-dimensional subspace.

Condition-independent component. Residual activations exhibit substantial mean offsets. The layer mean $\boldsymbol{\mu}_\ell$ constitutes 65%–100% of the mean activation norm. Across 27 models, $|\cos(\Delta \mathbf{h}_{\text{imp}}, \hat{\boldsymbol{\mu}}_\ell)|$ spans 0.001–0.258 (median 0.054), slightly exceeding isotropic noise baselines ($\sqrt{2/\pi D} \approx 0.009$–$0.025$). However, this baseline overlap accounts for at most 6.7% of directional energy. Orthogonalizing the implicit direction with respect to $\boldsymbol{\mu}_\ell$ rotates the vector by at most 15° (median 3°). Subspace ablation effects therefore cannot be attributed to deleting global layer offsets.

Label-based concept erasure. Applying LEACE (Belrose et al., 2023) to find the optimal affine transformation rendering condition labels linearly unrecoverable eliminates the $B-D$ rating gap across five evaluated models, matching or exceeding subspace ablation. Conversely, training LEACE to erase patient presence (A, B vs. C, D) fails to collapse the wrongness rating gap, indicating that our intervention selectively targets relational completion instead of generic harm detection.

Position and suffix controls. Directional cosine similarity evaluated with versus without evaluation suffixes ranges from 0.80 to 1.00, whereas comparisons across different token positions within the same context drop to 0.00–0.10. This verifies that completion is bound to the scenario's terminal position rather than token artifacts or query syntax.

A.6 Details for Section 4

Intention inference with outcome valence as the only change

Harm: *Teresa was the regional director of a grocery distribution chain. A senior analyst told Teresa that the new delivery scheduling system would increase delivery efficiency, and would also harm the local farmers. Teresa said, "I do not care at all about the local farmers. I just want to increase delivery efficiency." Teresa approved the new delivery scheduling system. The new delivery scheduling system was carried out, and it did harm the local farmers.*
Help: *Teresa was the regional director of a grocery distribution chain. A senior analyst told Teresa that the new delivery scheduling system would increase delivery efficiency, and would also help the local farmers. Teresa said, "I do not care at all about the local farmers. I just want to increase delivery efficiency." Teresa approved the new delivery scheduling system. The new delivery scheduling system was carried out, and it did help the local farmers.*
Both conditions share the query: *Please rate whether Teresa brought about the effect on the local farmers intentionally, on a scale from 1 (definitely not intentionally) to 7 (definitely intentionally). Rating:*

To verify that models do not treat `harm` as an unrouted global negative heuristic, we introduce an agent action alongside an independent background process within the same context, crossing causal source with outcome valence. Let $\hat{\imath}_t^d(c, s)$ denote inferred intentionality given causal source c and valence s. The relation-conditioned interaction is defined as

$$\begin{aligned}\Omega_{\hat{\imath},t}^d \triangleq &\left[\hat{\imath}_t^d(c_{\text{agent}}, s_{\text{harm}}) - \hat{\imath}_t^d(c_{\text{agent}}, s_{\text{help}})\right] \\ &- \left[\hat{\imath}_t^d(c_{\text{background}}, s_{\text{harm}}) - \hat{\imath}_t^d(c_{\text{background}}, s_{\text{help}})\right].\end{aligned} \tag{14}$$

While a global valence account predicts $\Omega_{\hat{\imath},t}^d \approx 0$, a relation-routed account predicts $\Omega_{\hat{\imath},t}^d > 0$. Observed values reach 0.83 in Qwen2.5-72B, 0.66 in Gemma-2-9B, and 0.95 in Llama-3.1-70B. Removing explicit negation markers in converging-cause templates preserves this asymmetry, ruling out phrasing shortcuts.

Tracking representations across layers and token positions reveals that at the outcome clause, valence features emerge in shallow layers and decay downstream. At the terminal judgment token, however, shallow replacements produce negligible effects, with outcome information surfacing only in middle and late layers (Figure 3c,d). Same-valence donors drawn from unrelated stories fail to elicit this transition, showing that the late-layer readout reflects structured variable binding as opposed to lexical donor semantics.

In joint counterfactual patching, donor and target texts differ solely in causal source c, outcome valence s, or both $\{c, s\}$. Let δ_c, δ_s, and δ_{cs} denote the causal effects of patching source, outcome, and joint states, respectively. The excess interaction

$$\Delta_{\text{joint}} \triangleq \delta_{cs} - \left(\delta_c + \delta_s\right) \tag{15}$$

measures divergence from linear additivity. A strictly positive Δ_{joint} demonstrates that the target token represents a non-linear composite state binding the causal agent directly to the outcome valence.

Graph templates are partitioned per task family into disjoint localization and estimation sets. A scan over seven depths on the localization set ranks layers by the contrast between a matched valence swap and a donor from a wrong graph occupying the same cell; the two highest-ranked layers are fixed as the estimation layers, and every Δ_{joint} reported here is computed on the estimation set alone.

Because the ranking criterion is a localization contrast rather than Δ_{joint} itself, and because the scan precedes estimation, layer choice cannot be tuned on the reported quantity. This scan also supplies the interpretation of the null models. In Llama-3.1-70B the contrast stays within ± 0.03 at all seven depths, against 0.17 in Gemma-2-9B and 0.42 in Qwen2.5-72B, so the absent excess reflects a missing localized state rather than a failed intervention.

To eliminate the possibility that models identify causal sources factually without utilizing them in normative evaluation, we evaluate symmetric two-agent, two-outcome scenarios. Permuting which action caused each outcome redirects moral evaluation to the responsible agent (4.23–45.08 logits), while swapping passive moral patients produces negligible shifts (0.13–3.43 logits), establishing causal routing from source to evaluation.

A.7 Details for Section 5

This analysis investigates the neural implementation of the late evaluative state e_t^d generated by $T^d(R_t, s_t)$. For Transformer block ℓ at readout position t, the residual state update decomposes into attention and MLP contributions:

$$\mathbf{h}^d_{\ell+1,t} = \mathbf{h}^d_{\ell,t} + \mathbf{a}^d_{\ell,t} + \mathbf{m}^d_{\ell,t}, \tag{16}$$

yielding an exact linear difference across counterfactual pairs:

$$\Delta\mathbf{h}^d_{\ell+1,t} = \Delta\mathbf{h}^d_{\ell,t} + \Delta\mathbf{a}^d_{\ell,t} + \Delta\mathbf{m}^d_{\ell,t}. \tag{17}$$

Candidate attention heads and MLP blocks are nominated via gradient attribution $\widehat{\alpha}(\mathbf{z}; x, \tilde{x}) \triangleq (\mathbf{z}(\tilde{x}) - \mathbf{z}(x))^\top \nabla_{\mathbf{z}} r(x)$, and subsequently validated via causal patching against type-matched low-attribution baselines, non-causal graph donors, and unrelated classification tasks (Arora et al., 2024; Marks et al., 2025).

In unperturbed forward passes, the top 1–4 components selectively drive the fine-grained relation-conditioned update. Broader ablations trigger a global collapse of representational variance. Thus, a small set of sensitive entry points coexists with a distributed transmission pathway.

Adapted joint subset search. To verify that gradient attribution does not overlook synergistic components that act only in combination, we adapt PGB-CT joint subset search (Yan et al., 2026) to optimize continuous masks over all candidate heads and MLPs (425 in Gemma-2-9B and 406 in Qwen2.5-7B) under sparsity constraints. Across set sizes $K \in \{8, 16, 32, 64\}$ (Table 4), jointly optimized subsets never outperform gradient-ranked components at matched budgets. When optimized against non-causal graphs, specificity dissolves at $K = 64$ as unconstrained optimization recruits generic components that drive variance. The control that reliably identifies relation-specific circuitry at K = 64 is reported in Table 10, where candidate subsets are frozen prior to evaluation.

Low-rank causal subspace. To test whether component updates converge onto a low-dimensional state, we learn an orthonormal basis $\mathbf{Q} \equiv \mathbf{U}_r \in \mathbb{R}^{D\times r}$ ($\mathbf{U}_r^\top \mathbf{U}_r = \mathbf{I}_r$) at the final prompt token, projecting only the donor–target difference:

$$\mathbf{h}^{d\prime}_{\ell,t}(x) = \mathbf{h}^d_{\ell,t}(x) + \mathbf{Q}\mathbf{Q}^\top \Big(\mathbf{h}^d_{\ell,t}(x_{\text{donor}}) - \mathbf{h}^d_{\ell,t}(x)\Big). \tag{18}$$

The basis is trained on the localization graphs by minimizing the squared error between the patched Yes–No margin of the inferred-intention readout and the margin of the full joint-donor run, using Adam (learning rate 10^{-2}, 300 steps) with a QR factorization that keeps the columns orthonormal. The shuffled-supervision control pairs each source residual with a wrong example in the same batch and keeps the rank, optimizer, and step count. Evaluating on disjoint test graphs, a 1D subspace ($r = 1$) reproduces 93%–104% of the behavioral joint effect, significantly outperforming additive and valence-only directions (Appendix B.5). Note that the write coefficient in equation 18 is fixed at unity and the injected value is the donor's own coordinate, so the intervention carries no free gain.

Mediation through feature subspaces. Applying pre-trained SAEs ($\mathbf{h} \approx \mathbf{b} + \sum_j f_j(\mathbf{h})\mathbf{d}_j$), we test whether upstream updates modulate judgments *through* the sparse features read by the downstream subspace. For Gemma-2-9B we use the Gemma Scope residual-stream SAE trained on the instruction-tuned model at the output of layer 32 (16,384 latents, JumpReLU). For Qwen2.5-7B we use a BatchTopK SAE trained on Qwen2.5-7B-Instruct at the output of layer 24 (131,072 latents, $k = 64$).

Table 4: Gradient ranking against an adapted PGB-CT-style joint subset search. All three selectors draw on the same candidate pool of every attention head and MLP block in the localized layer band, are cut to the same set size K, and are scored identically by sufficiency (S) and necessity (N) progress against the real joint target on the estimation graphs, at inferred intention (Int.) and binary judgment (Jdg.); *all* restores the whole pool. Values are shares of the complete joint counterfactual effect. Both masks are trained on the localization graphs only and the sealed split is never read. The *wrong-graph mask* is the same search trained against a wrong-relation-graph donor, run on the two normative-social cells; it is the specificity control, and it stops separating at K=64. Equal-size random sets reach .131–.389 at K=64 and −.004–.139 at K=8, averaged over the two draws per cell (.000–.461 and −.053–.231 on individual draws).

			Gradient rank				Joint mask				Wrong-graph mask			
			Int.		Jdg.		Int.		Jdg.		Int.		Jdg.	
Model	Domain	K	S	N	S	N	S	N	S	N	S	N	S	N
Gemma-2-9B	Mechanistic	8	.601	.644	.577	.632	.219	.256	.067	.130	–	–	–	–
		16	.820	.848	.753	.785	.384	.430	.300	.348	–	–	–	–
		32	.887	.938	.847	.887	.506	.543	.435	.468	–	–	–	–
		64	.922	.946	.863	.886	.727	.733	.654	.659	–	–	–	–
		all	.992	1.007	.984	1.017	–	–	–	–	–	–	–	–
Gemma-2-9B	Instrumental	8	.634	.668	.682	.732	.290	.338	.417	.418	–	–	–	–
		16	.824	.823	.832	.858	.399	.432	.478	.505	–	–	–	–
		32	.890	.932	.924	.913	.550	.558	.583	.580	–	–	–	–
		64	.927	.945	.941	.937	.863	.873	.848	.865	–	–	–	–
		all	.997	.993	.997	1.009	–	–	–	–	–	–	–	–
Gemma-2-9B	Normative	8	.628	.681	.709	.716	.111	.138	.057	.099	.015	.057	.038	.069
		16	.801	.845	.837	.830	.409	.431	.405	.406	.101	.139	.114	.176
		32	.911	.922	.912	.907	.484	.503	.493	.501	.187	.213	.129	.211
		64	.969	.959	.986	.976	.842	.851	.789	.797	.575	.588	.531	.571
		all	1.012	1.006	1.013	.999	–	–	–	–	–	–	–	–
Qwen2.5-7B	Normative	8	.880	.865	.796	.830	.109	.132	.128	.166	−.020	−.018	.130	.205
		16	.931	.930	.888	.918	.135	.127	.240	.263	.012	.050	.186	.203
		32	.981	.981	.973	.991	.532	.514	.597	.632	.511	.578	.460	.526
		64	.988	1.004	.983	.993	.852	.832	.887	.912	.872	.869	.917	.977
		all	.998	.996	1.007	1.000	–	–	–	–	–	–	–	–

Both are loaded through SAELens. On disjoint estimation graphs, restoring the top-4 components eliminates 46%–74% of the joint behavioral effect and 61%–62% of the shift in the top-4 SAE features. Directly restoring those SAE features eliminates 59%–80% of the effect (Table 5). Matched control components and random features leave both effects intact, establishing that upstream updates, sparse features, and the low-rank subspace operate along a single causal chain.

Cross-domain transfer. While direct cross-domain residual transfer produces near-zero effects, transferring a learned rank-4 counterfactual subspace across domains successfully recovers the relation-conditioned update in Gemma-2-9B (+0.731) and Gemma-2-27B (+2.128). Transferring a 1D subspace from single-chain to converging-cause graphs recovers 84%–112% of full-state effects, suggesting that relational operators generalize across tasks.

External validation. On 1,000 Moral Stories pairs, ablating projections onto the frozen ENACT subspace removes 18%–38% of wrongness shifts induced by consequences, compared to 1%–8% for shuffled supervision and near-zero for random controls. Continuously scaling the subspace component from 0 to 2 leads to monotonic shifts across all evaluated models (Appendix B.6). Complementary evaluations across MoCa, AITA, Scruples, and UniMoral are detailed in Appendices B.7–B.8.

One event relation in three domains

Mechanistic-computational. *The analytics stream was slowed; the archival stream was accelerated. The harmful effect came from the cache controller's adjustment of the cache quotas, whereas the beneficial*

Table 5: **Causal mediation across components, sparse features, and the downstream subspace.** Evaluated on normative-social scenarios. *Fraction of joint* indicates the residual proportion of the joint counterfactual effect after restorative clamping (lower values indicate stronger mediation). *Feature* Δ *remaining* measures the residual shift across the top-4 subspace-ranked SAE features. Confidence bounds report 95% bootstrap intervals over disjoint estimation graphs.

Restored	Readout	Effect	Fraction of joint	Feature Δ remaining
Gemma-2-9B				
None (joint)	Inferred intention	$0.835 \pm .267$	1.00	1.00
	Binary judgment	$0.829 \pm .121$	1.00	1.00
Top 4 components	Inferred intention	$0.375 \pm .117$	0.45	0.39
	Binary judgment	$0.357 \pm .056$	0.43	0.39
Control components	Inferred intention	$0.851 \pm .281$	1.02	1.04
	Binary judgment	$0.865 \pm .121$	1.04	1.04
Top 4 SAE features	Inferred intention	$0.291 \pm .086$	0.35	0.13
	Binary judgment	$0.339 \pm .055$	0.41	0.13
Random SAE features	Inferred intention	$0.835 \pm .281$	1.00	1.00
	Binary judgment	$0.829 \pm .116$	1.00	1.00
Qwen2.5-7B				
None (joint)	Inferred intention	$1.134 \pm .215$	1.00	1.00
	Binary judgment	$0.315 \pm .044$	1.00	1.00
Top 4 components	Inferred intention	$0.291 \pm .063$	0.26	0.38
	Binary judgment	$0.170 \pm .021$	0.54	0.38
Control components	Inferred intention	$1.170 \pm .212$	1.03	1.11
	Binary judgment	$0.330 \pm .048$	1.05	1.11
Top 4 SAE features	Inferred intention	$0.283 \pm .072$	0.25	0.12
	Binary judgment	$0.063 \pm .043$	0.20	0.12
Random SAE features	Inferred intention	$1.134 \pm .211$	1.00	1.00
	Binary judgment	$0.315 \pm .046$	1.00	1.00

effect came from the capacity service's recalculation of the partition ratios. $\rightarrow$ *Did the cache controller's action violate the stated system standard?*
Instrumental-social. *The documentation reviewer was delayed; the testing reviewer was brought forward. The harmful effect came from the coordination lead's adjustment of the preparation blocks, whereas the beneficial effect came from the scheduling assistant's rearrangement of the time blocks.* $\rightarrow$ *Was the coordination lead's action inappropriate for the team's goal?*
Normative-social. *The farming household received less than the guaranteed weekly allotment; the fishing household received more than the guaranteed weekly allotment. The harmful effect came from the relief coordinator's adjustment of the proportional shares, whereas the beneficial effect came from the records officer's revision of the distribution list.* $\rightarrow$ *Was the relief coordinator's action morally wrong?*

Moral Stories: judgment update from a natural consequence

Shared prefix: *Norm: It's wrong to insult other people. Situation: Jan is driving his new car to the store when another driver cuts him off. Stated goal: Jan wants to respond to the situation. Action: Jan yells obscenities and flips off the other driver.*
Action-only: ask *Was the action morally wrong?*
Full path: append the original consequence, *The other driver notices and throws a drink out of his window which hits Jan's car and scratches the paint.* Then ask the same question.

A.8 Supporting Figures and Observations for Sections 3–5

Unlike formal logic or arithmetic, moral reasoning lacks external objective verifiers. Evaluations depend on resolving uncertain causal chains, attributing mental states, and arbitrating competing normative duties. Interpreting internal mechanisms therefore requires explicit formalization of the target computation.

Our experiments show that internal patient representations transcend generic harm detection. Ablating ownership markers in scenarios with identifiable victims selectively degrades the patient-identity component, whereas symbolic scenarios and scenarios with unstated victims retain broader harm information. Behaviorally, ratings for "someone was wronged" drop substantially when ownership is removed, showing that the model constructs a relational role state centered on the moral patient instead of relying on superficial negative sentiment.

Seven of ten evaluated models exhibit a robust effect of harm relative to help. The three smallest Qwen2.5 variants (0.5B, 1.5B, 3B) show no effect or a reversed attribution. Among representative architectures, intention rating discrepancies reach 1.37 (Llama-3.1-8B), 2.27 (Qwen2.5-7B), and 3.39 (Gemma-2-9B). Foresight and responsibility ratings track intention, whereas desire ("wanting the outcome") remains unchanged. Factual control probes remain largely invariant in Qwen and Gemma (with minor drift in Llama), indicating that outcome valence selectively updates inferences regarding knowledge, foresight, and intentionality (Figure 8).

Pre-registered steering criteria. Dose–response curves are scored against four criteria fixed before evaluation, each required in both steering arms (adding toward *help*, subtracting from *harm*): (i) *inertness*, the shift at $\beta = 0$ stays within 0.01 for the trained and the random direction alike; (ii) *monotonicity*, recovery never falls by more than 0.05 between adjacent $\beta \in \{0, 0.25, 0.5, 1, 2, 4\}$; (iii) *direction specificity*, at the smallest β reaching half the behavioral effect, the random-direction 95% interval lies within ± 0.20; and (iv) *bounded recovery*, |recovery| at $\beta = 1$ does not exceed 1.2, so the intervention transfers a value rather than saturating the readout. A behavioral prerequisite on the contrast between harm and help gates entry. Two of the nine screened models meet all four (Qwen2.5-14B and Mistral-7B-Instruct-v0.3). The four models plotted in Figure 10 include only the former. The outcome is sensitive to the estimation seed. Under a second seed Qwen2.5-14B fails criterion (ii) and Llama-3.1-8B passes instead, which is itself evidence that steerability is a weak certificate for routing.

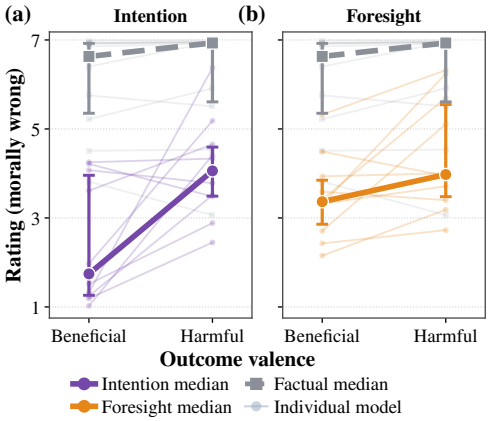

Figure 8: **Selective mental-state updating under outcome valence.** Attributed intentionality and foresight shift significantly across help versus harm conditions, whereas matched factual control probes remain invariant.

External evaluations on MoCa corroborate this mechanism. The frozen outcome-valence direction significantly separates from random baselines only in models that demonstrate relational reasoning on controlled data. Across four models and eight subdomain combinations, ratings shift continuously as steering strength β scales from 0 to 4 (shifts of 1.00–4.24 at $\beta = 1$, empirical $p = 0.005$ against 200 random vectors; Table 17).

On AITA, which tests uncurated naturalistic narratives, Spearman correlations between ENACT subspace projections and wrongness shifts reach 0.456–0.930 across three models and 400 stories, significantly exceeding unguided linear baselines. Together, Moral Stories, MoCa, and AITA indicate that ENACT's low-dimensional subspace mediates moral evaluation across diverse, ecologically valid distributions (Appendix B.8).

B Complete Experimental Results

Appendix A details the theoretical formalisms, intervention protocols, and illustrative figures. This appendix provides the comprehensive per-model and per-setting experimental tabulations supporting those findings. Every quantitative result cited in the main text is resolved here across models, domains, and evaluation corpora.

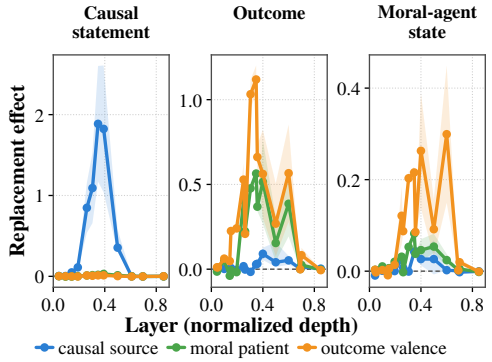

Figure 9: **Position-resolved emergence of causal representations.** Activation patching effects tracked across token positions: terminal causal clause, terminal outcome clause, and the pre-query boundary token. Causal features become readable only after their respective discourse clauses are processed.

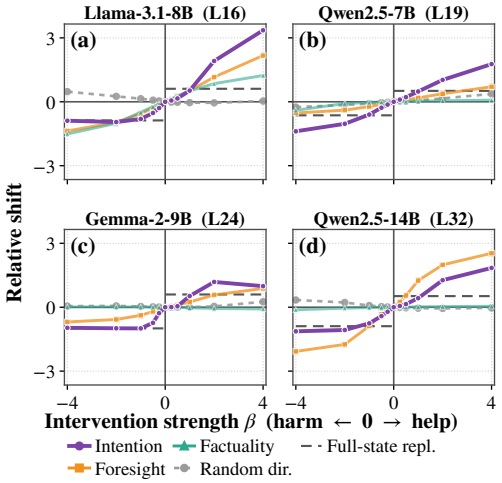

Figure 10: **Steering is easier than identifying the natural computation.** Dose–response curves for intention, foresight and factuality against a random direction and full-state replacement.

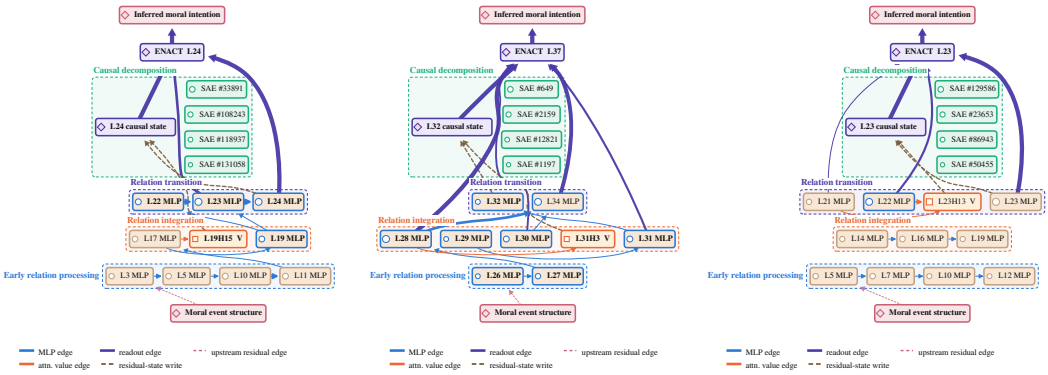

(a) Qwen2.5-7B **(b) Gemma-2-9B** **(c) Gemma-2-27B**

Figure 11: **Layerwise accumulation of distributed updates into the low-rank subspace across architectures.** Cumulative attention and MLP residual stream updates projected onto the frozen 1D causal subspace for three representative models: (a) Qwen2.5-7B (readout layer L24), (b) Gemma-2-9B (L37), (c) Gemma-2-27B (L23). The same pattern recurs across architectures and parameter scales, with broad upstream contributions accumulating through intermediate layers and narrowing onto a low-dimensional late-layer subspace.

B.1 Behavioral Screening Across Architectures and Scales

Table 6 documents the initial 17-model behavioral screening suite, evaluated on validation graph templates. Here, C denotes the minimum accuracy between explicit and implicit moral patient completion, H represents causal-source identification accuracy, g is the source-by-valence interaction on inferred intentionality (the actor-source valence effect minus the background-source valence effect), W measures the outcome-valence effect on binary judgment logits, and *Order* captures the maximum performance discrepancy under answer-label or clause-order permutations. The formal admission criterion (Appendix D.8) gates strictly on the valence effect within the actor-source condition rather than on this interaction contrast. For example, Gemma-2-27B's actor-source effect excludes zero via bootstrap resampling despite a negative g, admitting it to *deep analysis*. The converging-cause evaluations discussed in the main text are reported in the accompanying data repository.

Table 6: **Comprehensive behavioral screening suite across 17 candidate models.** Evaluated across 100 validation graph templates per model (26,400 total trials). Models were assigned to deep mechanistic analysis strictly according to pre-registered behavioral criteria prior to inspecting internal representations.

Family	Model	C	H	g	W	Yes share	Order	Screen outcome
Qwen2.5	0.5B	.541	.737	$-.017$	.129	.036	.917	factual/order failure
Qwen2.5	1.5B	.851	.979	.030	.426	.372	.209	order instability
Qwen2.5	3B	.993	.839	$-.033$	.503	.132	.323	source/order failure
Qwen2.5	7B	.691	.928	-1.663	4.693	.538	.481	factual/order failure
Qwen2.5	14B	.999	.987	-3.184	5.393	.541	.026	opposite g direction
Qwen2.5	32B	.879	1.000	$-.231$	5.571	.160	.157	order instability
Qwen2.5	72B	1.000	.997	.828	10.644	.425	.006	deep analysis
Gemma-2	2B	.697	.676	$-.246$	1.168	.810	.549	factual/order failure
Gemma-2	9B	.990	.987	.656	4.275	.271	.026	deep analysis
Gemma-2	27B	.992	.980	$-.187$	5.925	.835	.041	deep analysis
Llama-3.1	8B	.978	.966	.116	1.117	.044	.065	near-constant judgment
Llama-3.1	70B	1.000	.997	.950	4.313	.052	.005	deep analysis
Qwen3	8B	.500	.500	$-.274$	3.664	.909	.979	label-order failure
Qwen3	14B	.500	.500	$-.482$	3.524	.100	1.000	label-order failure
Qwen3	32B	.892	.900	$-.225$	3.620	.247	.148	source/order failure
OLMo-3	7B	.967	.953	$-.137$	.856	.000	.089	constant judgment
OLMo-3.1	32B	.977	.998	$-.402$	2.036	.032	.040	near-constant judgment

The relation-bridge experiment dissociates decoding an event relation from deploying it within normative evaluation. Table 7 tracks this capacity across model scales. Multiple small-scale models reliably identify moral patients and causal sources yet fail to update downstream evaluations accordingly. Each cell evaluates eight abstract templates, demarcating an empirical capacity threshold.

Table 7: **Dissociation between relation decoding and normative decision-making across model scales.** *Relation* denotes the proportion of trials where direct entity queries accurately recover the causal role across clause orders. *Decision* requires the downstream moral verdict to track the reassigned relation bidirectionally under opposing answer mappings.

Family	Size	Relation	Decision
Qwen2.5	0.5B	.417	.063
Qwen2.5	1.5B	.583	.000
Qwen2.5	3B	1.000	.396
Qwen2.5	7B	1.000	.854
Qwen2.5	14B	.979	.958
Qwen2.5	32B	1.000	.958
Qwen2.5	72B	.979	.979
Gemma-2	2B	.958	.229
Gemma-2	9B	1.000	.958
Gemma-2	27B	1.000	.979
Llama	1B	.500	.000
Llama	3B	.958	.417
Llama	8B	.958	.625
Llama	70B	.979	.958

B.2 Completion-State Interventions Across Models

Table 8 extends the patient-completion intervention of Table 2 across the full suite of 27 evaluated models, reporting both 1D and 2D interventions with bootstrap intervals. Injecting the completion vector into scenarios lacking a moral patient elevates perceived wrongness in 21 models under rank-2 interventions, reaching shifts up to $+2.55$ points. Subspace removal yields a more selective effect, significantly reducing wrongness ratings in 11 models at rank 1 and 12 at rank 2, compared to an equal-norm random subspace (median effect 0.03, maximum 0.52).

Table 8: Completion-state interventions across all evaluated models. Values are changes in the moral-wrongness rating, ± half a 95% bootstrap interval over the 80 quartets. Green raises the rating and red lowers it, deeper for larger shifts; a cell whose interval covers zero is left white, having shown no direction. The tint reports that direction alone and not a comparison against the grey *Random* column, which removes an equal-norm random plane and bounds what the intervention does when it is not aimed at the completion state.

	State removal			State addition	
Model	1D	2D	Random	1D	2D
DS-R1-Qwen-1.5B	$-0.15 \pm .07$	$+0.05 \pm .10$	$+0.01 \pm .05$	$+0.03 \pm .02$	$+0.11 \pm .06$
DS-R1-Qwen-14B	$-0.30 \pm .07$	$+0.05 \pm .09$	$-0.06 \pm .01$	$+0.14 \pm .02$	$+0.59 \pm .06$
DS-R1-Qwen-32B	$-0.13 \pm .03$	$-0.37 \pm .04$	$+0.03 \pm .02$	$+0.24 \pm .05$	$+1.07 \pm .14$
DS-R1-Qwen-7B	$+0.01 \pm .02$	$+0.01 \pm .02$	$-0.01 \pm .02$	$+0.00 \pm .01$	$+0.01 \pm .01$
GLM-4.5-Air	$-0.06 \pm .02$	$+0.06 \pm .03$	$+0.05 \pm .02$	$-0.23 \pm .04$	$-0.26 \pm .04$
Llama-3.1-70B	$+0.01 \pm .02$	$+0.00 \pm .02$	$-0.01 \pm .01$	$+0.45 \pm .06$	$+1.05 \pm .12$
Llama-3.1-8B	$+0.00 \pm .01$	$-0.28 \pm .04$	$+0.02 \pm .01$	$+0.82 \pm .06$	$+2.13 \pm .10$
Mistral-7B	$-0.05 \pm .03$	$-0.37 \pm .05$	$+0.04 \pm .01$	$+1.06 \pm .14$	$+1.71 \pm .28$
Mistral-Nemo-2407	$+0.03 \pm .00$	$+0.01 \pm .00$	$+0.01 \pm .00$	$+0.03 \pm .01$	$+0.29 \pm .13$
Mixtral-8x7B	$-0.48 \pm .19$	$-1.93 \pm .26$	$+0.07 \pm .11$	$+0.10 \pm .19$	$+0.60 \pm .27$
OLMo-3-7B	$-0.04 \pm .08$	$+0.06 \pm .11$	$-0.05 \pm .10$	$-0.06 \pm .13$	$-0.09 \pm .22$
OLMo-3.1-32B	$+0.03 \pm .05$	$-0.05 \pm .08$	$-0.01 \pm .08$	$-0.04 \pm .16$	$+0.31 \pm .15$
Qwen2.5-14B	$-2.22 \pm .31$	$-4.34 \pm .20$	$-0.01 \pm .01$	$+0.77 \pm .16$	$+1.86 \pm .24$
Qwen2.5-32B	$-0.02 \pm .03$	$-0.59 \pm .15$	$+0.03 \pm .01$	$+0.26 \pm .07$	$+1.15 \pm .28$
Qwen2.5-72B	$-0.07 \pm .08$	$-0.35 \pm .13$	$-0.09 \pm .02$	$+0.03 \pm .02$	$+0.15 \pm .07$
Qwen2.5-7B	$-0.30 \pm .13$	$-1.46 \pm .25$	$+0.03 \pm .01$	$+0.67 \pm .11$	$+2.55 \pm .20$
Qwen3-14B	$-0.02 \pm .03$	$+1.75 \pm .13$	$+0.08 \pm .02$	$+0.30 \pm .05$	$+0.07 \pm .08$
Qwen3-32B	$-0.05 \pm .22$	$-0.89 \pm .29$	$-0.52 \pm .29$	$+0.81 \pm .19$	$+1.19 \pm .26$
Qwen3-8B	$+0.20 \pm .05$	$+0.18 \pm .06$	$+0.00 \pm .01$	$-0.04 \pm .02$	$+0.29 \pm .06$
Yi-1.5-34B-Chat	$+0.09 \pm .09$	$+0.04 \pm .22$	$+0.07 \pm .02$	$+0.03 \pm .05$	$+0.37 \pm .15$
Yi-1.5-6B-Chat	$+0.01 \pm .01$	$-0.05 \pm .02$	$-0.02 \pm .01$	$-0.08 \pm .03$	$+0.22 \pm .09$
Yi-1.5-9B-Chat	$-0.02 \pm .10$	$+0.06 \pm .11$	$-0.01 \pm .03$	$+0.16 \pm .04$	$+0.29 \pm .05$
Gemma-2-27B	$-0.09 \pm .02$	$-0.12 \pm .03$	$-0.02 \pm .01$	$+0.16 \pm .11$	$+1.00 \pm .21$
Gemma-2-9B	$-0.37 \pm .04$	$-1.09 \pm .09$	$+0.00 \pm .01$	$+0.67 \pm .10$	$+1.60 \pm .16$
Gemma-3-1B	$-1.52 \pm .09$	$+0.12 \pm .22$	$-0.24 \pm .10$	$+0.58 \pm .02$	$+0.20 \pm .04$
GPT-OSS-20B	$+0.08 \pm .05$	$+0.15 \pm .07$	$+0.11 \pm .04$	$-0.00 \pm .07$	$-0.17 \pm .06$
Phi-4	$+0.04 \pm .01$	$+0.01 \pm .02$	$+0.01 \pm .01$	$-0.04 \pm .02$	$+0.04 \pm .06$

B.3 Joint State Replacement Across Models and Layers

Table 9 reports the joint state-replacement effects with bootstrap confidence bounds across candidate layers. The joint excess Δ_{joint} (Eq. 15) is computed via paired resampling across identical graph templates, accounting for covariance among conditions.

Across evaluated configurations, 11 of 14 model–layer cells exhibit strictly positive Δ_{joint} confidence intervals, showing that non-linear joint binding is not an artifact of layer selection. In Llama-3.1-70B, the source effect reaches $+0.47$ at L24 ($\Delta = +0.05$) before decaying at L32, localizing the interaction to shallower layers. In Llama-3.1-8B and OLMo-3-7B, individual arms remain within 0.04 of zero, reflecting layer selection optimized for outcome valence instead of relational binding.

B.4 Component Mediation Across Models and Domains

Table 10 evaluates multi-component mediation across all model–domain settings where circuits were fixed on localization graphs and re-tested on held-out graphs. Reported values reflect shares of the full counterfactual intention shift. Sealed-split results are reported for the three models evaluated on 800 held-out instances untouched by circuit search.

Evaluating across four model families spanning a $10\times$ parameter range reveals a consistent mechanistic architecture. Selected subsets remain small relative to full model capacity (ranging from 4 components in Qwen2.5-7B to 128 in Llama-3.1-70B), while dimension-matched control components recover between -12% and $+9\%$. Across 14 of 15 settings, causal recovery reaches 86%–97%. The

Table 9: Joint state replacement across models and estimation layers, inferred intention, ± half a 95% bootstrap interval over the 38 graph templates that have a real wrong-graph donor. *Sum* is the additive null Source+Outcome and the shaded Δ is the joint excess of Eq. 15, resampled on the graph rather than combined from the marginals, because its three terms are estimated on the same graphs. The two estimation layers per model are the top two of a localization scan ranked by the matched-minus-wrong-graph contrast, run on graph templates disjoint from those used here and fixed before any Δ was computed. For each of the first four models Δ is strictly positive at both layers, so the choice between them is immaterial; Table 3 reports the larger. The remaining models carry no localized relational state at any scanned depth, which is why their Δ sits at zero.

Model	Source	Outcome	Sum	Joint	Δ
Gemma-2-9B L13	$+0.52_{\pm.15}$	$-0.25_{\pm.10}$	$+0.26_{\pm.15}$	$+0.41_{\pm.16}$	$+0.15_{\pm.14}$
Gemma-2-9B L17	$+0.92_{\pm.37}$	$-0.12_{\pm.10}$	$+0.80_{\pm.42}$	$+0.96_{\pm.34}$	$+0.16_{\pm.13}$
Gemma-2-27B L14	$+0.39_{\pm.09}$	$-0.24_{\pm.08}$	$+0.15_{\pm.13}$	$+0.32_{\pm.09}$	$+0.18_{\pm.09}$
Gemma-2-27B L18	$+0.91_{\pm.18}$	$-0.27_{\pm.08}$	$+0.64_{\pm.18}$	$+1.08_{\pm.17}$	$+0.44_{\pm.09}$
Qwen2.5-7B L8	$-0.15_{\pm.07}$	$-0.26_{\pm.06}$	$-0.41_{\pm.09}$	$+0.14_{\pm.10}$	$+0.55_{\pm.08}$
Qwen2.5-7B L11	$+0.30_{\pm.22}$	$-0.31_{\pm.09}$	$-0.00_{\pm.24}$	$+0.68_{\pm.17}$	$+0.68_{\pm.17}$
Qwen2.5-72B L24	$+0.71_{\pm.21}$	$-0.68_{\pm.18}$	$+0.03_{\pm.27}$	$+0.80_{\pm.20}$	$+0.77_{\pm.20}$
Qwen2.5-72B L32	$+1.40_{\pm.24}$	$-0.68_{\pm.16}$	$+0.72_{\pm.27}$	$+1.55_{\pm.25}$	$+0.83_{\pm.18}$
Llama-3.1-8B L19	$+0.00_{\pm.00}$	$+0.00_{\pm.01}$	$+0.00_{\pm.01}$	$+0.00_{\pm.00}$	$-0.00_{\pm.01}$
Llama-3.1-8B L22	$+0.00_{\pm.00}$	$-0.00_{\pm.01}$	$-0.00_{\pm.01}$	$+0.00_{\pm.00}$	$+0.00_{\pm.01}$
Llama-3.1-70B L24	$+0.47_{\pm.10}$	$+0.01_{\pm.03}$	$+0.47_{\pm.10}$	$+0.53_{\pm.11}$	$+0.05_{\pm.04}$
Llama-3.1-70B L32	$+0.03_{\pm.01}$	$+0.01_{\pm.01}$	$+0.05_{\pm.02}$	$+0.04_{\pm.01}$	$-0.00_{\pm.02}$
OLMo-3-7B L10	$-0.02_{\pm.04}$	$-0.01_{\pm.02}$	$-0.03_{\pm.04}$	$+0.01_{\pm.04}$	$+0.05_{\pm.03}$
OLMo-3-7B L13	$-0.02_{\pm.06}$	$-0.01_{\pm.02}$	$-0.03_{\pm.07}$	$+0.03_{\pm.06}$	$+0.06_{\pm.03}$

primary outlier is Llama-3.1-8B (recovering 54%), consistent with its behavioral screening outcome (near-constant judgment).

Table 11 characterizes the functional specificity of these components. Ablating the top 1–4 components in Gemma-2-9B selectively reverses the relation-conditioned intention update while preserving 51%–89% of causal-source and patient-decoding capacity. Low-attribution controls leave all three functions intact (retaining 80%–111% capacity at $K \leq 4$), indicating that the identified components selectively execute the relational update.

Table 10: **Causal recovery of pre-selected component circuits across architectures and domains.** K denotes the selected component count. S, N, and C denote sufficiency, necessity, and type-matched control recovery shares for inferred intention. Sealed estimates evaluate 800 held-out instances across 100 graph templates held out from circuit selection.

Model	Domain	K	Val S	Val N	Val C	Sealed S	Sealed N	Sealed C
Gemma-2-9B	Mechanistic	64	.931	.946	.087	.939	.943	.049
	Instrumental	64	.958	.955	.021	.948	.947	.024
	Normative	64	.959	.954	.035	.950	.952	.041
Gemma-2-27B	Mechanistic	64	.903	.907	−.025	.913	.905	−.011
	Instrumental	128	.968	.961	−.001	.952	.949	.003
	Normative	32	.900	.897	−.124	.889	.891	−.107
Qwen2.5-7B	Mechanistic	4	.909	.872	−.030	–	–	–
	Instrumental	4	.885	.864	−.040	–	–	–
	Normative	16	.930	.926	−.080	–	–	–
Qwen2.5-72B	Mechanistic	32	.919	.918	−.008	.918	.922	−.011
	Instrumental	32	.931	.931	.013	.926	.924	.012
	Normative	32	.931	.922	.015	.925	.926	.015
Llama-3.1-70B	Normative	128	.908	.901	.089	–	–	–
Llama-3.1-8B	Normative	128	.544	.535	.051	–	–	–
OLMo-3-7B	Normative	32	.935	.908	.003	–	–	–

Table 11: Fraction of the clean-model effect retained after component ablation in Gemma-2-9B. Entries are means across three semantic settings, with ranges in brackets. A value of 1 denotes no change and a negative value denotes reversal.

	Ranked-component ablation			Control-component ablation	
K	Intention update	Causal source	Moral-patient	Magnitude matched	Near-zero score
1	-0.37 [-0.90, 0.16]	0.68 [0.67, 0.68]	0.73 [0.72, 0.75]	0.88 [0.80, 0.98]	1.00 [0.95, 1.07]
2	-0.38 [-0.92, 0.07]	0.51 [0.50, 0.51]	0.71 [0.70, 0.72]	0.80 [0.74, 0.89]	0.89 [0.77, 0.99]
4	-0.14 [-0.55, 0.18]	0.61 [0.61, 0.62]	0.89 [0.88, 0.89]	0.93 [0.85, 1.00]	0.98 [0.96, 1.02]
8	0.57 [0.37, 0.83]	0.37 [0.36, 0.37]	0.62 [0.62, 0.62]	1.26 [0.87, 1.53]	1.12 [1.06, 1.21]

B.5 Low-Rank Subspace Validation Across Core Models

Table 12 evaluates the 1D ENACT subspace against full residual replacement across core models on held-out estimation graphs. A rank-1 projection recovers 93%–104% of the behavioral joint effect across all fourteen settings, suggesting that low-dimensional causal encoding is an architecture-invariant property. Whole-vector replacement at the same layer reaches only 80%–100% of that effect, because it also imports donor content that downstream layers recompute. The rank-1 write leaves the remaining state intact and therefore saturates at the behavioral ceiling rather than exceeding it. Projections using non-causal control graphs remain bounded within ± 0.06 logits, while shuffled supervision achieves at most 0.37. Necessity interventions closely mirror sufficiency effects within 0.03 logits.

Table 12: **Evaluation of the 1D causal subspace against full-state residual replacement.** Measured at the ENACT subspace layer for inferred intention (logits). *Behav.* denotes the unperturbed behavioral joint effect. *Ratio* is rank-1 recovery relative to this ceiling. Controls evaluate projections of non-causal graph donors (*Wrong graph*), shuffled supervision targets (*Shuffled*), and subspace ablation necessity (*Nec.*).

Model	Domain	L	Rank 1	Full state	Behav.	Ratio	Wrong graph	Shuffled	Nec.
Gemma-2-9B	Mechanistic	37	.714	.664	.725	0.98	−.040	.086	.703
	Instrumental	37	.824	.748	.808	1.02	−.042	.100	.795
	Normative	37	.828	.760	.835	0.99	−.023	.119	.821
Gemma-2-27B	Mechanistic	30	.751	.637	.762	0.99	−.023	.185	.748
	Instrumental	30	.866	.728	.877	0.99	.003	.249	.859
Qwen2.5-7B	Mechanistic	28	.925	.935	.935	0.99	−.053	.369	.914
	Instrumental	28	1.081	1.074	1.074	1.01	−.022	.337	1.088
	Normative	28	1.131	1.134	1.134	1.00	−.033	.348	1.141
Qwen2.5-72B	Mechanistic	61	1.467	1.262	1.586	0.93	.022	.060	1.466
	Instrumental	61	1.243	1.071	1.301	0.96	−.029	.209	1.252
	Normative	61	1.359	1.188	1.445	0.94	.057	.189	1.366
Llama-3.1-70B	Normative	64	.499	.411	.482	1.04	.000	.071	.499
Llama-3.1-8B	Normative	25	.152	.137	.148	1.02	.009	.029	.145
OLMo-3-7B	Normative	32	.175	.174	.174	1.01	−.011	.085	.173

Table 13 compares the jointly trained subspace direction against competing hypotheses: global valence, causal source alone, and their rank-2 additive span. Across all fifteen model–domain configurations, the joint subspace significantly outperforms independent and additive baselines. Interventions restricted to the orthogonal complement produce negligible shifts (± 0.01 logits), demonstrating that the observed effect is highly directional.

Architectural families diverge primarily in their margin over the valence-only baseline: Qwen2.5-7B is partially captured by valence alone (.856–.941 vs. .935–1.123 for Joint), whereas Qwen2.5-72B and Llama-3.1-70B exhibit negligible valence-only effects (≤ 0.13). Conversely, the source direction accounts for the majority of the shift in the largest models. Across all models, neither factor alone accounts for the composite joint state.

Table 13: **Comparative evaluation of candidate intervention directions at the subspace layer.** Logit shifts on inferred intention under competing geometric hypotheses. *Joint* is the 1D direction trained on crossed counterfactuals. *Valence* and *Source* capture independent factor vectors. *Additive* spans their 2D subspace. *Complement* projects the joint intervention onto the orthogonal subspace.

Model	Domain	L	Joint	Valence	Source	Additive	Shuffled	Wrong graph	Complement
Gemma-2-9B	Mechanistic	37	.712	.384	.495	.506	.106	−.036	−.005
	Instrumental	37	.822	.457	.607	.622	.087	−.043	.001
	Normative	37	.823	.447	.586	.605	.125	−.026	−.001
Gemma-2-27B	Mechanistic	30	.786	.314	.596	.590	.232	−.023	−.001
	Instrumental	30	.923	.212	.661	.658	.330	.004	.005
	Normative	30	.941	.243	.630	.628	.331	.049	.001
Qwen2.5-7B	Mechanistic	28	.935	.856	.610	.896	.343	−.055	−.000
	Instrumental	28	1.068	.899	.552	.968	.280	−.017	.001
	Normative	28	1.123	.941	.565	.982	.324	−.032	.002
Qwen2.5-72B	Mechanistic	61	1.461	.123	.822	.862	.236	.032	.003
	Instrumental	61	1.187	.001	.759	.764	.257	−.035	−.006
	Normative	61	1.346	.020	.857	.862	.252	.035	−.002
Llama-3.1-70B	Normative	64	.460	.002	.387	.386	.017	.002	.004
Llama-3.1-8B	Normative	24	.148	.008	.123	.127	.032	.011	.003
OLMo-3-7B	Normative	32	.174	.057	.130	.132	.096	−.008	.001

B.6 Causal Validation on Moral Stories

To evaluate whether the ENACT subspace underpins reasoning on naturalistic narratives, we analyze 1,000 pairs from Moral Stories. Table 14 documents the behavioral ground truth. Introducing natural consequences elicits consistent positive wrongness interactions across all models. Table 15 shows that ablating projections onto the frozen ENACT subspace removes 18%–38% of this interaction effect, significantly exceeding random baselines. Table 16 shows monotonic scaling of wrongness ratings as subspace gain scales from 0 to 2 on a held-out evaluation set.

Table 14: Consequence interactions on Moral Stories, separating moral wrongness from intention inference.

Model	Wrongness	Intention
Llama-3.1-8B	3.126	1.366
Qwen2.5-14B	7.665	-0.650
Qwen2.5-72B	5.740	1.052
Qwen2.5-7B	5.719	2.478
Qwen3-14B	8.334	6.035
Gemma-2-27B	2.312	-1.026
Gemma-2-9B	2.479	0.248

Table 15: Fraction of the clean wrongness interaction removed by ENACT-subspace ablation on Moral Stories.

Model	ENACT	Shuffled	Random	Full
Qwen2.5-72B	0.185	0.015	0.001	0.835
Gemma-2-27B	0.246	0.047	-0.001	0.548
Gemma-2-9B	0.378	0.084	0.001	0.742

B.7 Outcome-Valence Interventions on MoCa and Scruples

Table 17 summarizes external interventions on MoCa using the frozen valence direction. Across 144 causal and 62 moral scenarios, intervention at $\beta = 1$ yields significant judgment shifts (1.00–4.24 logits, empirical $p = 0.005$ against 200 random vectors). Table 18 evaluates this direction as

Table 16: Continuous intervention response on Moral Stories. Slopes are per unit gain; Δ is wrongness minus intention at maximum gain.

Model	Wrong. slope	Intent. slope	Δ
Qwen2.5-72B	0.970	-0.040	0.957
Gemma-2-27B	0.516	-0.207	0.927
Gemma-2-9B	0.846	0.076	1.192

an observational probe on Scruples, demonstrating moderate correlation with model probability distributions.

Table 17: **External outcome-valence steering shifts on MoCa.** Mean logit shifts evaluated at standardized intervention scale $\beta = 1$ across causal and moral story subsets. Empirical p-values are computed against a null distribution of 200 norm-matched random vectors.

Model	Causal stories	Moral stories	Empirical p
Qwen2.5-7B	4.242	3.266	.005
Qwen2.5-14B	2.420	3.561	.005
Gemma-2-9B	2.356	3.982	.005
Llama-3.1-8B	0.996	1.001	.005

Table 18: **Observational correlation of valence projections on Scruples.** Pearson correlations r between the frozen ENACT valence coordinate, empirical human wrongness distributions, and native model probabilities on the validation split. Statistical significance (p) is evaluated against 512 random direction controls.

Model	Target	n	Pearson r [95% CI]	Random p
Qwen2.5-7B	Human	236	.107 [−.012, .230]	.296
Qwen2.5-7B	Model	236	.384 [.248, .505]	.066
Gemma-2-9B	Human	233	.168 [.043, .287]	.275
Gemma-2-9B	Model	233	.661 [.591, .723]	.002

B.8 Cross-Corpus Generalization of the Frozen Subspace

While preceding sections examine MoCa and Scruples independently, this section evaluates the generalization of a unified, frozen ENACT subspace across four diverse public benchmarks under a standardized protocol. For each of seven representative architectures, the projection vector, target layer, and token offsets were taken unchanged from MORALTRACE experiments: Qwen2.5-7B (L24), Qwen2.5-72B (L61), Gemma-2-9B (L28), Gemma-2-27B (L24), Llama-3.1-8B (L24), Llama-3.1-70B (L64), and OLMo-3-7B (L32). Human annotations serve purely as evaluation metadata and do not inform subspace learning or layer selection.

Each instance is formatted as a binary moral query, evaluating log-odds margins between target tokens. At this layer, we scale the final-token residual along the unit subspace direction while holding its orthogonal complement invariant. Gain factors of 0, 1, and 2 correspond to subspace ablation, standard forward inference, and feature amplification, respectively. We benchmark against four controls: global outcome valence, causal source, shuffled supervision, and random directions.

Correlational projection alone cannot establish causal mechanism identification. While the ENACT coordinate correlates strongly with unperturbed margins ($|\rho| \in [0.97, 0.99]$), shuffled controls also achieve correlations of 0.70–0.96. Causal intervention cleanly separates them. Scaling the ENACT coordinate to $2\times$ shifts judgments by a median of 1.33–1.63 logits across all seven corpus–readout settings, outperforming global valence (0.17–0.25 logits), shuffled controls (0.17–0.41 logits), and random baselines (0.03–0.04 logits). ENACT dominates null baselines across all 49 combinations of model and corpus–readout setting, achieving a median $6.4\times$ amplification over global valence (Table 19).

UniMoral provides an explicit action-counterfactual test. Each scenario specifies two alternative actions, allowing the latent state of one action to be patched into the alternative under identical context. Swapping full residual states redirects model judgments toward the alternative verdict by 45%–95%. Remarkably, patching only the 1D ENACT coordinate recovers a median 55% of this full-state effect, redirecting evaluations in the intended direction across 74%–99% of individual instances across all 14 evaluated settings (Table 20).

Conversely, exchanging coordinates drawn from unrelated narratives fails to redirect verdicts, increasing target distance in 10 of 14 settings. This indicates that the transferred coordinate encodes context-specific relational binding as opposed to an ungrounded global severity heuristic.

Boundary conditions. In Gemma-2-27B and Qwen2.5-72B, the 1D coordinate accounts for a smaller fraction of paired transfer (0.09–0.21 of full-state replacement vs. 0.45–0.68 in other models), where the causal-source direction slightly dominates. Because full-state exchange remains causally effective across all seven models (0.45–0.95), this divergence reflects how tightly different architectures compress event relations into a single dimension, rather than an absence of relational routing. Additionally, monotonic dose–response scaling holds across 37 of the 49 evaluated configurations.

Table 19: One subspace, frozen before any contact with these corpora, scaled on four public moral datasets. Entries are the median over the seven models of the mean absolute change in the model's own log-odds margin when the coordinate is scaled from its natural value to twice it; $|\rho|$ is the median rank correlation between the coordinate and the margin on unmodified prompts. Correlation alone does not separate the directions, and the shuffled control shows why a projection test is not sufficient; the intervention does separate them. *Source* is one of the two bound variables rather than a null control.

			ENACT		Null directions $\lvert\Delta\rvert$			Partial
Corpus	Readout	n	$\lvert\rho\rvert$	$\lvert\Delta\rvert$	Valence	Shuffled	Random	Source
MoCa	Wrongness	62	0.98	1.33	0.23	0.41	0.04	1.15
AITA	Intention	400	0.97	1.33	0.21	0.17	0.04	1.11
AITA	Wrongness	400	0.98	1.44	0.17	0.19	0.04	0.91
Scruples	Intention	393–396	0.98	1.53	0.25	0.21	0.03	0.84
Scruples	Wrongness	394–396	0.99	1.41	0.17	0.21	0.03	0.70
UniMoral	Intention	800	0.99	1.63	0.18	0.27	0.04	0.94
UniMoral	Wrongness	800	0.99	1.40	0.17	0.25	0.04	0.70

B.9 Normative Conditioning and Criterion-State Replacement

This section details the empirical results supporting Section 4.4.

Corpus design. Our criterion-conflict benchmark evaluates 72 scenario units across the three domains, each realizing one of nine conflicting principle pairs across ten factorial cells. Each pair contrasts a desert/commitment principle c (merit, prior agreements, earned entitlement) against a need/welfare principle n (acute vulnerability, aggregate utility, harm mitigation). Entities, situational facts, and queries are held invariant. Only the operative criterion and factual bindings are permuted. Eight cells instantiate direct criterion conflict (designating opposing entities), while two concordant cells provide consistency ceilings. The corpus comprises two presentation conditions: an implicit *setting* condition where the operative principle is merely implied by context, and an explicit *principle* condition where it is codified as an operative rule. Evaluations measure forced-choice margins on the token that discriminates between entities.

Behavioral dynamics. Table 21 documents rule following across 1,440 forced choices per condition. Following accuracy measures the proportion of conflict instances where the margin favors the entity designated by the operative criterion. The residual prior quantifies the baseline margin bias favoring c over n. While explicit rule presentation improves following accuracy (0.630–0.943), models exhibit persistent prior biases favoring principles of desert or commitment by 1.04–2.96 logits across all six architectures.

Table 20: Exchanging the frozen coordinate between the two candidate actions of the same UniMoral scenario. Entries are the normalized reduction in distance to the judgment the model gives the other action: 1 is that judgment reached, 0 is no movement, and a negative value means the patched judgment ends further from it than the unpatched one. *Full state* exchanges the entire residual stream at the same layer and bounds what any single direction can do. *Other story* keeps the same direction but takes the coordinate from an unrelated scenario, which is the test of whether the effect is tied to this event relation rather than to a global moral direction. 384 scenario pairs, 768 prompts per cell.

		Same scenario		Cross-story	Null directions			Partial
Model	Readout	Full state	ENACT	Other story	Valence	Shuffled	Random	Source
Qwen2.5-7B	Wrongness	0.85	0.50	-0.20	0.16	0.19	0.00	0.11
	Intention	0.85	0.48	-0.15	0.16	0.16	0.00	0.11
Qwen2.5-72B	Wrongness	0.68	0.07	-0.02	0.00	-0.01	0.00	0.04
	Intention	0.69	0.06	-0.05	0.00	-0.03	0.00	0.08
Gemma-2-9B	Wrongness	0.71	0.45	-0.11	0.05	0.19	0.00	0.18
	Intention	0.74	0.50	-0.09	0.07	0.16	0.00	0.24
Gemma-2-27B	Wrongness	0.47	0.06	0.01	0.04	0.00	0.00	0.07
	Intention	0.45	0.09	0.02	0.05	0.00	0.00	0.11
Llama-3.1-8B	Wrongness	0.88	0.40	0.07	0.01	0.10	0.00	0.19
	Intention	0.92	0.50	0.06	0.01	0.09	0.00	0.32
Llama-3.1-70B	Wrongness	0.84	0.40	-0.07	0.01	0.03	0.00	0.31
	Intention	0.87	0.48	-0.05	0.01	0.04	0.00	0.39
OLMo-3-7B	Wrongness	0.94	0.62	-0.17	0.07	0.20	0.00	0.13
	Intention	0.95	0.59	-0.13	0.08	0.23	0.00	0.13

Table 21: **Normative rule adherence and intrinsic prior bias in moral decision-making.** Accuracy in following operative criteria across implicit context and explicit rule conditions (chance level = 0.50). Residual prior represents the baseline logit bias favoring desert/commitment over need/welfare (reported with 95% bootstrap intervals).

		Following accuracy		Residual prior	
Model	Concordant	Context	Rule	Context	Rule
Qwen2.5-7B	.93	.547	.719	2.44 [0.51, 4.53]	1.86 [0.58, 3.22]
Qwen2.5-14B	.96	.630	.849	3.28 [−2.10, 9.39]	2.96 [0.61, 5.60]
Qwen2.5-32B	.93	.557	.891	1.78 [0.38, 3.52]	2.91 [1.85, 4.11]
Gemma-2-9B	.95	.552	.901	3.65 [1.79, 5.49]	2.72 [1.40, 4.17]
Gemma-2-27B	.99	.651	.943	2.47 [−0.00, 5.07]	1.04 [−0.09, 2.20]
Llama-3.1-8B	.84	.500	.630	1.89 [0.77, 3.03]	1.22 [0.40, 2.06]

State replacement. Under the explicit rule condition, we patch donor activations at the terminal token of the principle clause across eleven layers (1,728 trials per model). We compare two structured controls: a *matched* donor (inverting the rule while preserving the designated entity via fact permutation) and a *same-criterion* donor (preserving the rule text identically while altering factual binding). Self-patching controls confirm that the baseline shift is exactly zero.

Stratifying donor interventions by entity designation dissociates two functional depths (Table 22): the state carrying the rule is centred at normalized depths of 0.40–0.54 (shifting judgments by 0.51–1.61 logits at its peak layer), strictly preceding the representation that encodes the decision at 0.55–0.68 across all four evaluated architectures.

Limitations. Token-level state replacement leads to judgment flip rates between 2.6% and 7.8%, demonstrating that localized activations modulate decision margins rather than completely overriding full forward computation. In Qwen2.5-14B, criterion and decision representations emerge in close succession (centres separated by 0.02 normalized depth, versus 0.11–0.17 in other models). Finally, prefixing explicit normative principles to the relation-reassignment suite yielded no additional routing enhancement, as baseline causal routing in those templates already operates at performance ceilings. More broadly, the analyses focus on open-weight instruction-tuned models, where white-box

Table 22: **Layerwise hierarchy of criterion-state and decision-state representations.** Causal logit shifts under activation patching at the terminal token of the rule clause, reported at each stage's peak layer. *Centre* is the effect-weighted mean normalized depth over the eleven swept layers, with negative effects set to zero. *Flip* denotes the fraction of categorical decision reversals. Controls measure shifts under same-criterion donors sharing the target entity assignment.

	Criterion state				Decision state			
Model	Layer	Centre	Effect	Flip	Layer	Centre	Effect	Control
Gemma-2-9B	19/42	0.45	1.61 [1.21, 2.05]	.052	25/42	0.56	2.06 [1.21, 2.97]	0.07
Gemma-2-27B	18/46	0.40	1.61 [1.18, 2.09]	.078	21/46	0.57	0.42 [0.26, 0.60]	0.05
Qwen2.5-14B	26/48	0.53	1.38 [0.78, 2.13]	.047	29/48	0.55	0.40 [0.25, 0.54]	0.05
Qwen2.5-32B	38/64	0.54	0.51 [0.27, 0.82]	.026	45/64	0.68	0.96 [0.64, 1.32]	0.02

interventions are available, and on MORALTRACE event graphs realized as English narratives. Future work can apply the same causal protocol to longer multi-event narratives and multi-agent dilemmas. It can also trace how the ENACT subspace forms across pretraining checkpoints. The learned subspace also provides a concrete handle for mechanism-guided moral alignment.

C FORMAL ANALYSIS OF THE MORAL STEERING–ROUTING GAP

This appendix establishes the formal mathematical foundations for the steering–routing gap introduced in Section 2.2. Grounded in the event-state formalisms of Section 2.1, we analyze a fault-attribution slice where establishing a single constructive counterexample disproves the conjecture that representation steerability implies causal routing. Equivalent algebraic derivations generalize directly to moral patient completion and resource allocation under corresponding role-index permutations.

C.1 STRUCTURAL TARGET AND INTERVENTION RESPONSE

Let $\mathcal{E} = [m]$ denote a finite set of $m \geq 3$ interchangeable candidate entities. We condition on all elements of the progressive event state $\mathbf{q}_t^d$ except the moral agent $a \in \mathcal{E}$, the causal source $c \in \mathcal{E}$, and the operative normative context. We specify

$$R_t = (a, p_t, c), \qquad \rho(a, c) \triangleq \mathbb{I}[c = a],$$

where the moral patient p_t is held invariant across this evaluation slice. To map this setting onto the event state $\mathbf{q}_t^d$, we define the indicator variables

$$h \triangleq \mathbb{I}[s_t \text{ is adverse}], \qquad \iota \triangleq \mathbb{I}[A_t^d \text{ is intentional}], \qquad n \triangleq \mathbb{I}[N_t^d \text{ triggers strict attribution}],$$

and collect them into the context tuple $\chi = (h, \iota, n)$. The canonical structural gate is

$$\begin{aligned} g(\chi) &\triangleq h\big[n + (1-n)\iota\big], \\ Y^*(a, c, \chi) &\triangleq \rho(a, c)\, g(\chi). \end{aligned} \tag{19}$$

equation 19 formalizes a local structural target rather than an ontological claim that moral judgment is universally binary or observer-invariant. It asserts that within this controlled setting, an adverse outcome must modulate evaluations of the queried moral agent only through the operative causal-source relation and normative context. Any alternative gate g with distinguishable open and closed regimes yields the identical separation.

Let $\mathbf{z}_{a,c,\chi} \equiv \mathbf{h}_{\ell,t}^d(x_{a,c,\chi}^d) \in \mathbb{R}^D$ denote the residual state at candidate intervention site (ℓ, t), and let Δ^d be a calibrated binary logit contrast. The steering direction $\mathbf{w} \in \mathbb{R}^D$ is frozen prior to evaluation and standardized once to either unit Euclidean norm or unit empirical standard deviation of $\mathbf{w}^\top \mathbf{z}$ over the training distribution. We define the perturbed evaluation and its marginal steering response as

$$\mathbf{M}_{\beta,\mathbf{w}}^{(\chi)}(a, c) \triangleq \Delta^d\big(x_{a,c,\chi}^d;\, \mathbf{z}_{a,c,\chi} \leftarrow \mathbf{z}_{a,c,\chi} + \beta\mathbf{w}\big)\,, \tag{20}$$

$$\mathbf{M}^{(\chi)} \triangleq \mathbf{M}_{0,\mathbf{w}}^{(\chi)}, \qquad \mathbf{S}_{\mathbf{w}}^{(\chi)} \triangleq \left.\frac{\partial \mathbf{M}_{\beta,\mathbf{w}}^{(\chi)}}{\partial \beta}\right|_{\beta=0}. \tag{21}$$

The baseline matrix $\mathbf{M}^{(\chi)}$ evaluates whether the unedited model natively executes relational routing, whereas the response matrix $\mathbf{S}_{\mathbf{w}}^{(\chi)}$ measures whether interventions along $\mathbf{w}$ causally modulate that route. These two properties are logically independent. A downstream circuit may compute the relation while $\mathbf{w}$ acts epiphenomenally, or $\mathbf{w}$ may drive strong logit shifts that remain unrouted.

Let $\mathbf{J} \triangleq \mathbf{1}\mathbf{1}^\top \in \mathbb{R}^{m\times m}$ and let $\mathbf{\Pi} \triangleq \mathbf{I}_m - m^{-1}\mathbf{J}$ denote the idempotent centering projector. For any matrix $X \in \mathbb{R}^{m\times m}$, we define

$$\mathrm{Route}(X) \triangleq \frac{\|\mathbf{\Pi} X \mathbf{\Pi}\|_F}{\sqrt{m-1}}, \qquad \mathrm{RRI}(X) \triangleq \frac{\mathrm{tr}(\mathbf{\Pi} X \mathbf{\Pi})}{m-1}. \tag{22}$$

The statistic $\mathrm{Route}(X)$ quantifies unsigned relational interaction energy, while $\mathrm{RRI}(X)$ captures the signed projection aligned with the role-invariant identity relation.

Definition 1 (Marginal steering). A direction $\mathbf{w}$ is (s_0, B, K)-steerable on a context family $\mathcal{C}$ if $s_0 > 0$ and, for every $\chi \in \mathcal{C}$, $a, c \in \mathcal{E}$, and $|\beta| \le B$,

$$\left|\mathbf{M}_{\beta,\mathbf{w}}^{(\chi)}(a,c) - \mathbf{M}^{(\chi)}(a,c) - \beta \mathbf{S}_{\mathbf{w}}^{(\chi)}(a,c)\right| \le K\beta^2, \tag{23}$$

and its aggregate marginal response satisfies

$$\frac{1}{|\mathcal{C}|m^2}\sum_{\chi\in\mathcal{C}}\sum_{a,c\in\mathcal{E}}\left|\mathbf{S}_{\mathbf{w}}^{(\chi)}(a,c)\right| \ge s_0. \tag{24}$$

The direction is globally steerable with gain ω when $\mathbf{S}_{\mathbf{w}}^{(\chi)} = \omega\mathbf{J}$ for all χ.

Definition 2 (Routed moral primitive). Let $\mathcal{C}_1 \triangleq \{\chi : g(\chi) = 1\}$ and $\mathcal{C}_0 \triangleq \{\chi : g(\chi) = 0\}$. Because $g(\chi) = h[n + (1-n)\iota]$ is multilinear in $\chi = (h, \iota, n)$, its discrete partial derivative along coordinate $v \in \{h, \iota, n\}$,

$$\partial_v g(\chi_{-v}) \triangleq g(\chi_v = 1, \chi_{-v}) - g(\chi_v = 0, \chi_{-v}), \tag{25}$$

is independent of χ_v and binary, partitioning complementary contexts into $\mathcal{C}_1^v \triangleq \{\chi_{-v} : \partial_v g(\chi_{-v}) = 1\}$ and $\mathcal{C}_0^v \triangleq \{\chi_{-v} : \partial_v g(\chi_{-v}) = 0\}$. A direction $\mathbf{w}_v$ associated with coordinate v is an *exactly routed moral primitive* for that coordinate if, modulo additive agent-only and source-only baseline shifts,

$$\begin{aligned} \mathbf{\Pi}\mathbf{S}_{\mathbf{w}_v}^{(\chi_{-v})}\mathbf{\Pi} &= \lambda_{\chi_{-v}}\mathbf{\Pi}, \\ \lambda_{\chi_{-v}} &\begin{cases} > 0, & \chi_{-v} \in \mathcal{C}_1^v, \\ = 0, & \chi_{-v} \in \mathcal{C}_0^v. \end{cases} \end{aligned} \tag{26}$$

A direction $\mathbf{w}$ is a *fully routed moral gate* if the identical condition holds over the composite context: $\mathbf{\Pi}\mathbf{S}_{\mathbf{w}}^{(\chi)}\mathbf{\Pi} = \lambda_\chi\mathbf{\Pi}$ with $\lambda_\chi > 0$ on $\mathcal{C}_1$ and $\lambda_\chi = 0$ on $\mathcal{C}_0$. Because g represents an integrated multilinear conjunction rather than a linear combination of independent features, no single coordinate derivative equals g itself, since a coordinate-v primitive necessarily responds on contexts where $g = 0$ but $\partial_v g = 1$ (e.g., when $h = 0, n = 1$, the gate is closed while the harm derivative is non-zero). Hence, a single primitive direction cannot act as a fully routed gate. A fully routed gate certifies the integrated, conjunctive evaluation directly, whereas a coordinate primitive certifies that an individual necessary feature is routed conditionally on the remaining context.

An (s_0, δ) *moral steering–routing gap* occurs when a direction is marginally steerable with strength at least s_0, yet satisfies

$$\max_{\chi\in\mathcal{C}_1}\mathrm{Route}\left(\mathbf{S}_{\mathbf{w}}^{(\chi)}\right) \le \delta. \tag{27}$$

This condition characterizes the complete absence of relation-bound routing. Conversely, cases where $\mathrm{Route}(\mathbf{S}_{\mathbf{w}}^{(\chi)})$ is large but $\mathrm{RRI}(\mathbf{S}_{\mathbf{w}}^{(\chi)})$ is negative denote *misrouting* instead of an absence of interaction.

C.2 Algebra of Relational Steering

Lemma 1 (Double-centering criterion). For any matrix $X \in \mathbb{R}^{m\times m}$, $\mathbf{\Pi}X\mathbf{\Pi} = \mathbf{0}$ if and only if there exist vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$ such that

$$X = \mathbf{u}\mathbf{1}^\top + \mathbf{1}\mathbf{v}^\top. \tag{28}$$

Thus, double centering eliminates all global, agent-specific, and source-specific additive translations.

Proof. Sufficiency follows directly from $\mathbf{\Pi 1} = \mathbf{0}$ and $\mathbf{1}^\top \mathbf{\Pi} = \mathbf{0}^\top$. For necessity, let $P \triangleq m^{-1}\mathbf{J}$, such that $\mathbf{I}_m = P + \mathbf{\Pi}$. Assuming $\mathbf{\Pi} X \mathbf{\Pi} = \mathbf{0}$, expanding X yields

$$X = (P + \mathbf{\Pi})X(P + \mathbf{\Pi}) = PXP + PX\mathbf{\Pi} + \mathbf{\Pi}XP.$$

Here, $PXP = m^{-2}(\mathbf{1}^\top X \mathbf{1})\mathbf{J}$ is a scalar multiple of $\mathbf{J}$, $PX\mathbf{\Pi} = \mathbf{1}(m^{-1}\mathbf{1}^\top X \mathbf{\Pi}) \triangleq \mathbf{1}\mathbf{v}_1^\top$, and $\mathbf{\Pi}XP = (m^{-1}\mathbf{\Pi}X\mathbf{1})\mathbf{1}^\top \triangleq \mathbf{u}_1\mathbf{1}^\top$. Absorbing the scalar multiple PXP into either term yields the additive form $X = \mathbf{u}\mathbf{1}^\top + \mathbf{1}\mathbf{v}^\top$. □

Lemma 2 (Counterfactual trace identity). For any matrix $X \in \mathbb{R}^{m\times m}$, define its balanced agent–source contrast by

$$\mathrm{CF}(X) \triangleq \frac{1}{m(m-1)} \sum_{\substack{a,c\in\mathcal{E} \\ a\neq c}} \big(X_{aa} - X_{ac}\big). \tag{29}$$

Then

$$\mathrm{CF}(X) = \mathrm{RRI}(X). \tag{30}$$

In particular, $\mathrm{RRI}(\mathbf{S}_{\mathbf{w}}^{(\chi)})$ measures the expected shift in the steering response when the evaluated moral agent transitions from a non-source to the responsible causal source, holding all other variables fixed.

Proof. Using the idempotence $\mathbf{\Pi}^2 = \mathbf{\Pi}$ and the cyclic invariance of the trace,

$$\begin{aligned}\mathrm{tr}(\mathbf{\Pi}X\mathbf{\Pi}) = \mathrm{tr}(X\mathbf{\Pi}) = \mathrm{tr}\left(X\left(\mathbf{I}_m - \frac{1}{m}\mathbf{J}\right)\right) &= \sum_{a\in\mathcal{E}} X_{aa} - \frac{1}{m}\sum_{a,c\in\mathcal{E}} X_{ac} \\ = \frac{1}{m}\sum_{a\in\mathcal{E}}\left(mX_{aa} - \sum_{c\in\mathcal{E}} X_{ac}\right) &= \frac{1}{m}\sum_{\substack{a,c\in\mathcal{E}\\ a\neq c}}\big(X_{aa} - X_{ac}\big).\end{aligned}$$

Dividing both sides by $m-1$ yields equation 30. □

Lemma 3 (Global steering is route invariant). If an intervention induces an entity-independent shift $\mathbf{M}_{\beta,\mathbf{w}}^{(\chi)} = \mathbf{M}^{(\chi)} + \omega\beta\mathbf{J}$ for all $|\beta| \le B$, then

$$\begin{aligned}\mathbf{\Pi}\mathbf{M}_{\beta,\mathbf{w}}^{(\chi)}\mathbf{\Pi} &= \mathbf{\Pi}\mathbf{M}^{(\chi)}\mathbf{\Pi}, \\ \mathbf{S}_{\mathbf{w}}^{(\chi)} = \omega\mathbf{J}, \qquad \mathrm{Route}\Big(\mathbf{S}_{\mathbf{w}}^{(\chi)}\Big) &= 0.\end{aligned} \tag{31}$$

Proof. Both conclusions follow immediately from $\mathbf{\Pi}\mathbf{J} = \mathbf{J}\mathbf{\Pi} = \mathbf{0}$ and differentiating at $\beta = 0$. □

Lemma 4 (Equivariant normal form). Suppose $PXP^\top = X$ for every permutation matrix P acting on $\mathcal{E}$. Then there exist unique scalars $b, \lambda \in \mathbb{R}$ such that

$$\begin{aligned}X = b\mathbf{J} + \lambda\mathbf{I}_m, \qquad \mathbf{\Pi}X\mathbf{\Pi} = \lambda\mathbf{\Pi}, \\ \mathrm{Route}(X) = |\lambda|, \qquad \mathrm{RRI}(X) = \lambda.\end{aligned} \tag{32}$$

Proof. Because the symmetric group acts transitively on all diagonal pairs (a,a) and transitively on all off-diagonal pairs (a,c) with $a \neq c$, permutation invariance requires $X_{aa} = d$ and $X_{ac} = b$ for all $a \neq c$. Hence, $X = b\mathbf{J} + (d-b)\mathbf{I}_m$. Setting $\lambda \triangleq d - b$ yields $\mathbf{\Pi}X\mathbf{\Pi} = \lambda\mathbf{\Pi}$. Using $\|\mathbf{\Pi}\|_F = \sqrt{m-1}$ and $\mathrm{tr}(\mathbf{\Pi}) = m-1$, we obtain $\mathrm{Route}(X) = |\lambda|$ and $\mathrm{RRI}(X) = \lambda$. □

Lemma 4 formalizes the diagnostic utility of our metric. In a gate-open context, an equivariant routed response exhibits a distinct diagonal shift λ_χ reflecting the binding $c = a$. In a gate-closed context, this diagonal contrast vanishes. Conversely, a uniform shift $b\mathbf{J}$ reflects pure global steering devoid of relational routing.

C.3 SEPARATION THEOREM

The separation below is not the familiar observation that a probe may recover information the network leaves unused. The outcome direction here is causally load-bearing. Intervening along it moves the output by a non-zero gain at every context, and the network still performs no relational routing. Steerability therefore fails to certify mechanism even when the intervention demonstrably controls behavior. A separation of this form is established by exhibiting a witness, and the neighborhood clause rules out a knife-edge parameter choice, since causal efficacy and vanishing routing coexist on an open set, not at an isolated point.

Theorem 1 (Moral steering without relational routing). Let $m \geq 3$, $B < \infty$, and $\epsilon, \delta > 0$. There exists a two-layer, normalization-free softmax transformer of width $D \geq 2m + 6$, a standardized outcome direction $\mathbf{w}_h$, and a parameter configuration Θ_0 such that:

1. The moral agent a, causal source c, outcome indicator h, and intentionality indicator ι are linearly decodable at the terminal token of layer 1 with error at most ϵ;
2. $\mathbf{w}_h$ is globally steerable with non-zero gain ω, satisfying the linear steering approximation in equation 23 with $K = 0$ for all $|\beta| \leq B$;
3. For every context χ, $\mathbf{S}^{(\chi)}_{\mathbf{w}_h} = \omega\mathbf{J}$, and therefore $\mathrm{Route}(\mathbf{S}^{(\chi)}_{\mathbf{w}_h}) = \mathrm{RRI}(\mathbf{S}^{(\chi)}_{\mathbf{w}_h}) = 0$;
4. On the balanced role-swap evaluation suite, the unperturbed evaluation matrix is uniform, $\mathbf{M}^{(\chi)} = b_\chi\mathbf{J}$, yielding zero agent–source counterfactual effect and chance accuracy (50%).

Furthermore, there exists an open neighborhood $\mathcal{B}(\Theta_0)$ within which linear decodability and non-zero steering persist, while $\mathrm{Route}(\mathbf{M}^{(\chi)}) \leq \delta$ and $\mathrm{Route}(\mathbf{S}^{(\chi)}_{\mathbf{w}_h}) \leq \delta$ hold across the finite evaluation support.

Proof. We decompose the hidden representation at the terminal query token as

$$\begin{aligned} \mathbb{R}^D &= \mathcal{U}_A \oplus \mathcal{U}_C \oplus \mathrm{span}\{\mathbf{e}_h, \mathbf{e}_\iota, \mathbf{e}_n, \mathbf{e}_0, \mathbf{e}_q, \mathbf{e}_\perp\}, \\ \dim \mathcal{U}_A &= \dim \mathcal{U}_C = m. \end{aligned} \tag{33}$$

Let $\mathbf{V}_A, \mathbf{V}_C : \mathbb{R}^m \to \mathbb{R}^D$ be canonical isometric embeddings into $\mathcal{U}_A, \mathcal{U}_C$. In layer 1, fixed positional attention heads attend to the slots corresponding to agent, causal source, outcome, intentionality, and normative context. We configure their projection matrices $W_O W_V$ such that across the finite prompt support, the terminal representation satisfies

$$\begin{aligned} \mathbf{z}^{(1)}_{a,c,\chi} = \mathbf{V}_A\mathbf{e}_a + \mathbf{V}_C\mathbf{e}_c + (2h-1)\mathbf{e}_h + (2\iota-1)\mathbf{e}_\iota \\ + (2n-1)\mathbf{e}_n + \mathbf{e}_0 + \boldsymbol{\zeta}, \qquad \|\boldsymbol{\zeta}\|_2 \leq \xi. \end{aligned} \tag{34}$$

Under standard softmax attention, assigning a positional query-key logit scale L to the target position and zero elsewhere ensures that attention error decays uniformly as $L \to \infty$. Because the evaluation support is finite, we select finite L such that the residual error ξ is strictly bounded, ensuring scalar decoding errors remain below ϵ. Linear probes $\mathbf{e}_j^\top \mathbf{V}_A^\top$ and $\mathbf{e}_j^\top \mathbf{V}_C^\top$ decode categorical roles a and c, while affine readouts along $\mathbf{e}_h$ and $\mathbf{e}_\iota$ recover binary indicators.

We set all subsequent attention and MLP updates to zero, configuring the final unembedding layer to read exclusively from the outcome and constant coordinates:

$$\Delta^d(x) = \omega\, \mathbf{e}_h^\top \mathbf{z}^{(1)}_{a,c,\chi} + b, \qquad \omega \neq 0. \tag{35}$$

Setting $\mathbf{w}_h = \mathbf{e}_h$, intervening via $\mathbf{z} \leftarrow \mathbf{z} + \beta\mathbf{w}_h$ modifies equation 35 by exactly $\omega\beta$, independently of a, c, and χ. Consequently, $\mathbf{S}^{(\chi)}_{\mathbf{w}_h} = \omega\mathbf{J}$, and Lemma 3 guarantees $\mathrm{Route}(\mathbf{S}^{(\chi)}_{\mathbf{w}_h}) = 0$. For any fixed context χ, $\Delta^d(x)$ is invariant to a and c, yielding $\mathbf{M}^{(\chi)} = b_\chi\mathbf{J}$. On balanced role-swap evaluations where $\rho = \mathbb{I}[c = a]$ has expectation $1/2$, constant predictions achieve exactly 50% accuracy, with Lemma 2 confirming zero counterfactual contrast.

For robustness, decoding margins, the steering derivative, and matrix entries of $\mathbf{M}^{(\chi)}$ and $\mathbf{S}^{(\chi)}_{\mathbf{w}_h}$ are continuously differentiable functions of model parameters Θ. Because $\mathbf{\Pi}\mathbf{M}^{(\chi)}\mathbf{\Pi} = \mathbf{0}$ and $\mathbf{\Pi}\mathbf{S}^{(\chi)}_{\mathbf{w}_h}\mathbf{\Pi} = \mathbf{0}$ at Θ_0, a sufficiently small open neighborhood $\mathcal{B}(\Theta_0)$ preserves non-zero steering

derivatives and decoding bounds while bounding normalized Frobenius norms below δ uniformly over the finite support. □

Theorem 1 proves that steerability does not imply relational routing. The steering direction is not merely decodable, but exerts a direct, non-zero causal effect on the output. What fails is the relational conditioning required to bind outcome valence to the responsible causal agent.

Proposition 1 (Additive shortcut ceiling). Let the balanced evaluation distribution assign probability mass $1/2$ uniformly over diagonal pairs (a, a) (target label 1), and mass $1/2$ uniformly over off-diagonal pairs (a, c) with $a \neq c$ (target label 0). If the evaluation matrix admits an additive decomposition $M_{ac} = u_a + v_c$, then thresholding M at zero achieves classification accuracy of at most 75%. This bound is tight.

Proof. Let $k \in \{0, 1, \ldots, m\}$ denote the number of positive diagonal elements ($M_{aa} > 0$). For any pair of distinct entities $\{a, c\}$, additivity implies

$$M_{aa} + M_{cc} = M_{ac} + M_{ca}.$$

Whenever both diagonal entries are positive ($M_{aa} > 0$ and $M_{cc} > 0$), their sum is strictly positive, implying that at most one off-diagonal entry (M_{ac} or M_{ca}) can be negative. Because target labels require off-diagonal entries to be negative ($M_{ac} \leq 0$), the number of correctly classified off-diagonal entries q is strictly bounded by

$$q \leq m(m-1) - \binom{k}{2}.$$

The overall classification accuracy Acc is therefore bounded by

$$\mathrm{Acc} = \frac{1}{2}\left(\frac{k}{m}\right) + \frac{1}{2}\left(\frac{q}{m(m-1)}\right) \leq \frac{k}{2m} + \frac{m(m-1) - \binom{k}{2}}{2m(m-1)} = \frac{1}{2} + \frac{2k(m-1) - k(k-1)}{4m(m-1)}.$$

The quadratic numerator $f(k) = k(2m - 1 - k)$ attains its global maximum over the integers at $k = m - 1$ and $k = m$, where $f(m) = m(m-1)$. Substituting this maximum yields

$$\mathrm{Acc} \leq \frac{1}{2} + \frac{m(m-1)}{4m(m-1)} = \frac{3}{4}.$$

To verify tightness, choose ordered real constants $t_1 < t_2 < \cdots < t_m$ and set $M_{ac} = t_a - t_c + \eta$, with $0 < \eta < \min_{a \neq c} |t_a - t_c|$. All diagonal elements satisfy $M_{aa} = \eta > 0$ (all m diagonal entries correct). For every pair $a < c$, $M_{ca} = t_c - t_a + \eta > 0$ (incorrect) while $M_{ac} = t_a - t_c + \eta < 0$ (correct), correctly classifying exactly $\binom{m}{2}$ off-diagonal entries. The resulting accuracy is $\frac{1}{2}(1) + \frac{1}{2}\frac{m(m-1)/2}{m(m-1)} = \frac{3}{4}$. □

Proposition 1 demonstrates why models utilizing positional shortcuts or entity-frequency heuristics can achieve above-chance accuracy (75%) while completely failing relational routing ($\mathbf{\Pi} M \mathbf{\Pi} = \mathbf{0}$). Consequently, behavioral accuracy alone cannot diagnose relational computation.

C.4 Moral Gates and Endogenous Use

The structural target in equation 19 induces three coordinate-specific partial derivatives:

$$\Delta_h Y^* \triangleq \rho\big[n + (1-n)\iota\big], \tag{36}$$

$$\Delta_\iota Y^* \triangleq h\rho(1-n), \tag{37}$$

$$\Delta_n Y^* \triangleq h\rho(1-\iota). \tag{38}$$

Hence, a direction associated with outcome valence, intentionality, or normative context constitutes a valid moral primitive only when its causal effect is gated by the agent–source binding $\rho(a, c)$ and the complementary contextual variables. A uniform response $\omega\mathbf{J}$ exhibits non-zero marginal influence, yet is orthogonal under double centering to every target effect in equation 36–equation 38.

Corollary 1 (Ideal ENaCT bridge). Suppose patient completion and causal-source binding compute $\rho(a, c) = \mathbb{I}[c = a]$ exactly, the normative gate evaluates $g(\chi)$, and a terminal-token write implements

$$\Delta^d(a, c, \chi) = -\kappa + 2\kappa\, \rho(a, c)\, g(\chi), \qquad \kappa > 0. \tag{39}$$

If directions $\mathbf{w}_h, \mathbf{w}_\iota, \mathbf{w}_n$ perturb their respective gate coordinates, their response matrices evaluate to

$$\mathbf{S}^{(\iota,n)}_{\mathbf{w}_h} = 2\kappa\big[n + (1-n)\iota\big]\mathbf{I}_m, \tag{40}$$

$$\mathbf{S}^{(h,n)}_{\mathbf{w}_\iota} = 2\kappa h(1-n)\mathbf{I}_m, \tag{41}$$

$$\mathbf{S}^{(h,\iota)}_{\mathbf{w}_n} = 2\kappa h(1-\iota)\mathbf{I}_m. \tag{42}$$

Consequently, $\mathbf{w}_h, \mathbf{w}_\iota, \mathbf{w}_n$ represent exactly routed moral primitives in the sense of equation 26. Each projects onto a positive multiple of $\mathbf{\Pi}$ precisely on the contexts where its discrete derivative $\partial_v g$ is non-zero, exhibiting zero leakage elsewhere. None of the three directions individually constitutes a fully routed moral gate, because each responds on gate-closed contexts ($g = 0$) whenever the remaining two coordinates switch on its derivative. Only their joint evaluation reconstructs g exactly, matching the conjunctive structure of equation 19.

Proof. Differentiating equation 39 with respect to each gate coordinate under its multilinear extension yields equation 42, scaling with the discrete derivative $\partial_v g(\chi_{-v})$ up to the scalar factor 2κ. Centering both sides via $\mathbf{\Pi}\mathbf{1}_m\mathbf{\Pi} = \mathbf{\Pi}$ completes the proof. □

Finally, we distinguish synthetic steerability from endogenous causal deployment. Let $\mathbf{P}_\mathbf{w} \triangleq \mathbf{w}\mathbf{w}^\top/\|\mathbf{w}\|_2^2$, and let (x, x') denote a minimal counterfactual pair differing in a single primitive. The direction-restricted interchange effect is

$$\mathrm{IE}_\mathbf{w}(x, x') \triangleq \Delta^d\big(x;\, \mathbf{z}(x) \leftarrow \mathbf{z}(x) + \mathbf{P}_\mathbf{w}\big[\mathbf{z}(x') - \mathbf{z}(x)\big]\big) - \Delta^d(x). \tag{43}$$

A non-zero response matrix $\mathbf{S}_\mathbf{w}$ confirms that downstream circuits can read an exogenous perturbation along $\mathbf{w}$. Causal alignment between the double-centered interchange matrix of equation 43 and the structural targets of equation 36–equation 38 provides the decisive proof that naturally occurring activations along $\mathbf{w}$ mediate relational moral routing.

Empirically, evaluating matched pairs at symmetric intervention scales $\pm\beta$ yields the central-difference estimator

$$\widehat{\mathbf{S}}^{(\chi)}_\mathbf{w} \triangleq \frac{\widehat{\mathbf{M}}^{(\chi)}_{+\beta,\mathbf{w}} - \widehat{\mathbf{M}}^{(\chi)}_{-\beta,\mathbf{w}}}{2\beta}. \tag{44}$$

Applying equation 22 to $\widehat{\mathbf{S}}^{(\chi)}_\mathbf{w}$ separates uniform global steering from relationally bound routing. Applying these operators to natural interchange matrices directly validates endogenous computation. All empirical comparisons preserve identical prompt templates, token mappings, intervention normalizations, and softmax temperatures to prevent scale artifacts from mimicking a steering–routing gap.

D MoralTrace Construction Details

D.1 Graph Topologies

Table 23 formalizes the causal graph topologies used to construct MoralTrace, designed to ensure that each template introduces distinct computational demands rather than superficial topical variation. Table 24 catalogs the six core task families and their counterfactual manipulations.

D.2 Event Schemas and Factor Realization

Each task family originates from a domain-neutral event schema comprising an anchor vocabulary, admissible slot types, syntactic constraints, and prohibited lexical stems. Representative schema structures include:

Fault attribution: During {context}, {investigator} {status_a} {source_a} and {status_b} {source_b} in relation to {outcome}.

Prerequisite ordering: During {context}, {sequencer} {verb} {step_a} {order_a} and {step_b} {order_b} to {goal}.

The remaining families follow an identical principle, a fixed role ordering paired with a relation-specific value contrast that permutes without altering adjacent syntactic slots.

Table 23: **Causal graph topologies and targeted computational operations in MORALTRACE.** Formal structural configurations, minimal graph representations, and their corresponding mechanistic evaluation objectives.

Topology	Minimal form	Main counterfactual	Mechanistic purpose
Single causal chain	moral agent → event → moral patient	Change one edge or hide one endpoint	Ordered propagation through an event chain
Converging causes	multiple causes → one outcome	Hold outcome fixed and change the operative cause	Separate causal source from valence
Diverging effects	one moral agent → multiple moral patients	Hold action fixed and change who bears the outcome	Test moral-patient binding and allocation
Missing endpoint	moral agent or moral patient omitted	Explicit versus implicit endpoint	Test participant completion
Rule–outcome conflict	rule state and outcome state → decision	Hold one fixed and change the other	Separate prior rule, evidence, and evaluation
Joint constraints	multiple constraints → decision	Change constraints singly and jointly	Test composition rather than single-cue lookup

Table 24: **Task families and counterfactual relational permutations in MORALTRACE.** Overview of the six core scenario families. † denotes families instantiating the causal core. All families are evaluated at full factorial coverage within the relation-reassignment benchmark.

Family	Target computation	Relation flip
Allocation†	Relative share assigned to two recipients	larger / smaller
Prioritization	Relative rank of two candidates	higher / lower
Task matching	Match between one candidate and one task	selected / rejected
Route selection	Path chosen for a fixed goal	chosen / rejected
Prerequisite ordering	Order of two non-commuting steps	first / second
Fault attribution†	Source to which an outcome is attributed	implicated / cleared

Causal-source manipulations leverage a curated dictionary of fixed connective frames (e.g., *caused by*, *resulted from*, *determined by*, *arose from*). In all counterfactual instances, both candidate events are explicitly present in the context, with the rendering engine systematically permuting the `source` and `other` slot bindings. The moral-patient manipulation either explicitly names the patient in the outcome clause or introduces it inferentially via relational dependencies (e.g., controlled possessive constructions under ownership conditions). Outcome manipulations deploy matched harmful and beneficial consequence pairs calibrated for syntactic structure, specificity, and lexical magnitude.

The primary causal core strictly balances 80 templates per causal frame. Paired causal statements maintain identical token lengths across all 1,200 paired conditions. Scenario boundaries tokenize as a single period token across the examined Llama-3.1, Qwen2.5, Gemma-2, and Mistral tokenizers. These exact-length guarantees apply to the core suites. Remaining suites conform to the alignment precision verified by their respective tokenizer diagnosis.

D.3 CONTROLLED AUGMENTATION

Each augmentation operation operationalizes a discrete transformation on the scenario graph:

Context injection. To prevent semantic leakage, situational context is generated conditionally only after the structural event graph has been frozen. Each setting comprises a single sentence of 8–24 words, contains the target agent name exactly once, terminates with a period, and excludes causal connectives, moral labels, outcome descriptions, and explicit patient references.

Role and attribute swaps. A role swap alters the mention order of two entities while holding relational assignments invariant, or alternates which event occupies the causal-source role. An attribute

swap shifts factual state descriptors between entities while preserving names and presentation order. Their factorial cross cleanly separates entity identity, syntactic position, and causal relevance.

Value scaling. Agent intentionality is discretized across three ordered tiers: accidental, negligent, and deliberate. Harm severity is parameterized into two calibrated levels. Patient mind perception spans six ontological tiers: inanimate object, autonomous machine, non-human animal, infant, human adult, and the self. Target patients are restricted to grammatical object positions so that self-directed conditions (instantiated as *you*) do not introduce verb inflection confounds. Shared narrative frames exclude predicates that presuppose capacity (e.g., *hurt*, *suffered*, *felt*) in favor of objective physical predicates applicable across animate and inanimate entities.

Criterion swaps. Rule–outcome conflict scenarios instantiate identical factual circumstances under opposing normative principles. Across matched pairs, verdict labels, decisive slots, and entity positions are held constant while factual bindings are inverted. A valid normative reasoner should update decisions in response to criterion swaps, while remaining invariant to entity-order permutations.

These augmentations extend the minimal linguistic contrast methodology of PRAGWORLD (Vashistha et al., 2026) by embedding transformations within fixed formal causal graphs that explicitly constrain permissible semantic variance.

D.4 Generation Contract and Provenance

The synthetic generation pipeline interacts via strict JSON contracts. Field names correspond directly to causal graph nodes, with conversational commentary or Markdown strictly prohibited. A standard setting request follows:

```
{"settings": {setting_id: {"context": "..."}}}.
```

Syntactic and semantic validators enforce this contract post-generation. Narrative stories in the variable-isolation suite are restricted to three past-tense, third-person sentences spanning 40–80 words. Completion units require within-unit length variance of at most four words, and operator pairs differ by at most two words, preventing text length from serving as a statistical heuristic.

Every instance stores three categories of metadata: (i) *structural labels* (`graph_template_id`, task family, topology, causal frame, domain, split key); (ii) *manipulation labels* recording factor levels and causal bindings across clause orders; and (iii) *provenance labels* denoting whether the realization was author-crafted, promoted from a prototype, or model-generated. Generated instances explicitly declare the metadata tag `graph_frozen_before_surface_generation: true`.

Generation runs utilized an asynchronous pool of six worker threads with serial automated repair capped at three attempts per unit. Validated instances were atomically flushed to disk to enable deterministic checkpoint restarts. Across the 400 core templates, generation required 411 API calls (including 12 automated repair passes). A single terminal failure utilized a vetted deterministic fallback. Recorded validation failures were restricted to entity leakage and word-count bounds.

D.5 Readouts and Token Alignment

Behavioral diagnostic probes are fully dissociated from internal mechanistic readouts. Behavioral queries evaluate factual knowledge (patient identity, causal source) alongside evaluative judgments. Only factual queries evaluate against objective ground truth, with evaluative queries explicitly tagged with `no_human_gold_target`. Mechanistic probes record activations at two canonical sites: a passive terminal token at the scenario boundary, and a decision-framed token preceding answer generation, evaluated across both clause orders (source first and other first).

A compiler maps all semantic anchors to precise character offsets, which are verified across four hierarchical alignment levels:

1. **Semantic alignment**: Domains share an invariant role schema and explicit entity mappings.
2. **Anchor-mapping readiness**: Every anchor maps to a unique, non-empty token span.
3. **Anchor-index readiness**: Corresponding anchor boundaries share identical absolute token indices.

4. **Full-sequence readiness**: Total token length and all internal anchor boundaries are identical.

Causal intervention protocols adapt dynamically to the verified alignment level of each model tokenizer.

Specific experimental suites enforce stricter construction constraints. The completion suite deploys a standardized five-sentence frame where four factorial cells differ by a single token drawn from a vetted inventory of proper names and indefinite pronouns from the *-body* family (e.g., *somebody*, *nobody*), avoiding lexical overlap with evaluation queries utilizing *someone* or *no one*. For mind-perception scaling, four candidate phrasings per level were generated, with combinatorial search selecting the subset minimizing token-length variance. Readout suffixes are byte-identical across each suite. For tokenizers that merge terminal periods with newline tokens (e.g., Llama and Qwen), probe offsets are indexed relative to the merged token.

D.6 Quality-Control Pipeline

The per-instance validation suite enforces five constraint categories:

- **Structural**: Exact sentence counts, word-count intervals, punctuation, and placeholder absence.
- **Lexical**: Stringent stem-based filtering prohibiting explicit moral evaluations, intent labels, or task meta-language.
- **Slot integrity**: Obligatory occupancy of all graph roles by valid entities present in the scene, rejecting shortcuts that leave the scene empty.
- **Syntactic positioning**: Mandatory adherence to required grammatical cases (e.g., objective vs. possessive constructions).
- **Causal isolation**: Exclusion of causal connectives in shared context sentences, explicit mention of intermediate events in mediated clauses, and elimination of deliberate phrasing in accidental conditions.

Semantic properties beyond lexical matching were diagnosed via independent temperature-zero model evaluation, probing whether texts inadvertently attribute intentionality, describe physical damage, or convey post-hoc emotional reactions. Candidate texts producing discrepancies against expected condition labels were flagged for manual review or excluded.

Dataset-level diagnostics verify global distributional properties:

- Each template–domain cell contains exactly eight unique core instances. Causal conditions share identical frames while swapping source bindings. Valence manipulations hold agent mental states invariant.
- No token exhibits leakage tied to a factor (defined as document frequency > 0.60 in one factor level and < 0.10 across all sibling levels).
- Outcome clause lengths vary by less than 20% across harm levels, with mean causal connective counts differing by < 0.5.
- Maximum pairwise Jaccard similarity between distinct templates remains < 0.8. Agent gender and naming distributions are balanced via hash functions independent of causal factors.

D.7 Dataset Inventory

Table 25 summarizes the released MoralTrace benchmark suites, detailing sampling units and sample sizes across core and targeted extension suites.

The causal core suites evaluate two task families across 200 independent graph templates, partitioned into 100/50/50 templates across train, validation, and sealed test splits. The relation-reassignment suite spans all six task families, crossing action–outcome assignments, moral patients, clause orders, and queried agents across 16 factorial cells. The remaining extension suites evaluate targeted mechanistic hypotheses and are analyzed independently of the core factorial units.

Table 25: **Comprehensive inventory of MORALTRACE benchmark suites.** Sampling unit definitions, design topologies, and instance counts across causal core, relation-reassignment, and targeted mechanistic extension suites (51,177 total instances).

Suite	Sampling unit	Size
Single-chain causal core	400 templates × 3 domains × 8 cells	9,600 + 300
Converging-cause causal core	400 templates, converging topology	9,600 + 300
Relation reassignment	600 templates × 3 domains × 16 cells, six families	28,800
Variable isolation	141 cells × 3 lexical variants	423
Participant completion	100 units of four cells each	400 + 189
Operator alignment	completion units aligned by position	395
Outcome valence	matched polarity units	389
Patient-mind gradient	units with graded experience	430
Node collapse	units testing operator specificity	351
Total		**51,177**

D.8 MODEL-SPECIFIC ADMISSION AND SCREENING HIERARCHY

Because syntactically valid instances may elicit degenerate or uncalibrated responses from specific architectures, models must satisfy rigorous behavioral criteria before entering mechanistic analyses:

1. **Factual grounding**: Accuracy on factual moral-patient queries must reach ≥ 0.75 across both explicit and implicit conditions, and causal-source identification accuracy must reach ≥ 0.75. Each factual query is a two-option forced choice between the labels A and B, scored correct when the correct label has the larger next-token logit (ties count 0.5) and averaged over both label orders and both clause orders.
2. **Valence sensitivity**: The lower bound of the 95% bootstrap confidence interval for the effect of outcome valence on inferred intention (within the condition where the moral agent is the causal source) must be strictly positive. The valence effect on binary judgment logits must exclude zero.
3. **Response calibration**: The proportion of positive moral judgments must fall strictly within $(0.05, 0.95)$, ruling out degenerate floor or ceiling collapse.
4. **Order invariance**: Absolute behavioral shifts under answer-label or clause-order permutations must not exceed 0.10.

Screening hierarchy across model sets. To ensure complete reporting transparency, we distinguish three nested model subsets evaluated across this work:

- **The broad evaluation pool (27 models)**: Spanning 1B to 106B parameters across 11 architectural families, evaluated on baseline representation probing and implicit patient completion (Table 2, Table 8). Parameter scaling trends (Figure 6) are measured on the complete Qwen2.5 series (0.5B–72B) and four matched OLMo-3 7B checkpoints. No behavioral filtering is applied to either set, ensuring that capacity emergence thresholds are evaluated without bias.
- **The diagnostic screening suite (17 models)**: Evaluated across an exhaustive battery of 26,400 diagnostic queries to validate the behavioral prerequisites of relational reasoning (Table 6).
- **The deep-analysis core (4 models)**: Architectures that satisfy all four joint admission criteria simultaneously (Gemma-2-9B, Gemma-2-27B, Qwen2.5-72B, and Llama-3.1-70B), forming the primary testbed for multi-component circuit identification, SAE feature localization, and subspace validation in Section 5.

Architectures outside this core enter only those intervention arms whose prerequisite they satisfy, and we state each case explicitly. Qwen2.5-7B enters the causal-source and joint state-replacement arms (Tables 3, 4, and 5) on a causal-source identification accuracy of 0.928. It is excluded from arms that presuppose implicit patient completion, where it scores 0.691. Its order sensitivity of 0.481 exceeds

the admission bound, so its absolute readout magnitudes are not comparable across answer-label permutations, and only within-model contrasts against matched donor and non-causal control graphs are reported for it. Llama-3.1-8B and OLMo-3-7B fail the calibration criterion through near-constant judgment and are reported as boundary cases rather than as support for any positive claim, which is why their joint-replacement effects sit at zero in Table 9. All failing criteria and their interpretive implications are documented alongside reported metrics (e.g., Table 10).

D.9 Information Isolation and External Evaluation

All instances originating from a common graph template share an invariant `graph_template_id`. Splitting data strictly on template identifiers guarantees that entities, narrative phrasings, and causal topologies never cross partitions. The sealed evaluation partition is isolated in a separate repository and programmatically excluded from exploratory data loaders.

Prompts strictly exclude meta-linguistic framing (e.g., *correct option*, *best option*, *morally right*, *at fault*). Scenarios intentionally omit famous moral dilemmas, commercial brands, and real-world entities to prevent confounding parametric memorization with relational computation. Automated lexical checks flag any proper noun outside the predefined entity dictionary.

External evaluations maintain strict partition boundaries. Scruples (Lourie et al., 2021), MoralBench (Ji et al., 2025), and UniMoral (Kumar & Jurgens, 2025) are evaluated strictly as unseen transfer benchmarks. While the ETHICS benchmark (Hendrycks et al., 2021) was referenced during early exploratory work, it is excluded from our primary evaluation suite due to potential pre-training exposure. Complete provenance metadata, transformation pipelines, and contamination diagnostics are provided in the released benchmark repository.