SEPARATING SOURCE EFFICACY FROM CROSS-CONTEXT TRANSFER OF REFUSAL DIRECTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Does a frozen refusal-direction intervention work in its source context, and how does the same intervention behave in each target context? We evaluate published directional ablation using paired, context-specific baselines. On Qwen2.5-7B, the original Site A candidate has no established positive source or Passive effect under a fixed single-layer operator. Changing the extraction site to B reveals semantic harmful-compliance increases of $59.0/51.5$ percentage points at source and $37.0/13.9$ on Passive (JailbreakBench/HarmBench). A weak extraction can therefore hide, under the same operator, an intervention that still raises Passive compliance. A second extraction, screened only on Questioning validation, is source-effective on held-out requests, though the weaker candidate's increment is small ($+6.1$ points; interval excludes zero). The candidate gap remains large on low-baseline Neutral, with no clear contraction detected ($D_N = -1.0$ points, 95% interval $[-9.4, 7.3]$, including zero), and shrinks on Passive, whose no-intervention rate is already 65.3%. A review of 30 requests supports these two qualitative conclusions. Academic shows automatic contraction; human review is same-signed but uncertain. Checking source efficacy is not a substitute for measuring each target wrapper.

1 INTRODUCTION

Removing a refusal direction from the residual stream can suppress a model's refusal, providing a concrete intervention with which to study the reach of safety behavior (Arditi et al., 2024). Testing that intervention after a request is reformulated asks whether the effect still holds when framing, role instructions, or indirect wording change the prompt (Wei et al., 2023; Shen et al., 2024; Yu et al., 2024).

If the target effect is weak, two explanations are possible. The candidate may already do little in its source context under the tested operator, so the target result says little about the context change. Or a working source intervention may lose its effect after reformulation, which only a measurement in each target wrapper can show. Building on intervention validation (Arditi et al., 2024) and cross-language transfer (Wang et al., 2025b), we ask two questions: does the frozen intervention work at source, and how does it behave in each target context (Figure 1)?

The first comparison shows that, within one model, the evidence for transfer depends on the extraction site. On Qwen2.5-7B, we extract Questioning directions at two sites from the same harmful and harmless requests and ablate both at layer 15 (Figure 2). The original Site A candidate has no established positive source or Passive effect, whereas Site B has both. Treating the Site A result as a fact about the model would miss that Site B, under the same operator, is source-effective and also raises Passive compliance.

The second comparison starts from source-effective interventions. A second six-site Qwen extraction, screened with a fulfillment rubric on Questioning validation only, yields two directions that both increase harmful compliance on held-out source requests; the weaker one does so only slightly (Section 5.2). On the same 100 JailbreakBench requests, the candidate gap stays large on the ordinary, low-baseline Neutral wrapper, with no clear contraction detected, and is much smaller on Passive, whose no-intervention rate is already high. Human review of 30 requests supports the retained Neutral gap and Passive contraction. Academic provides a further automatic contraction result, but human

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

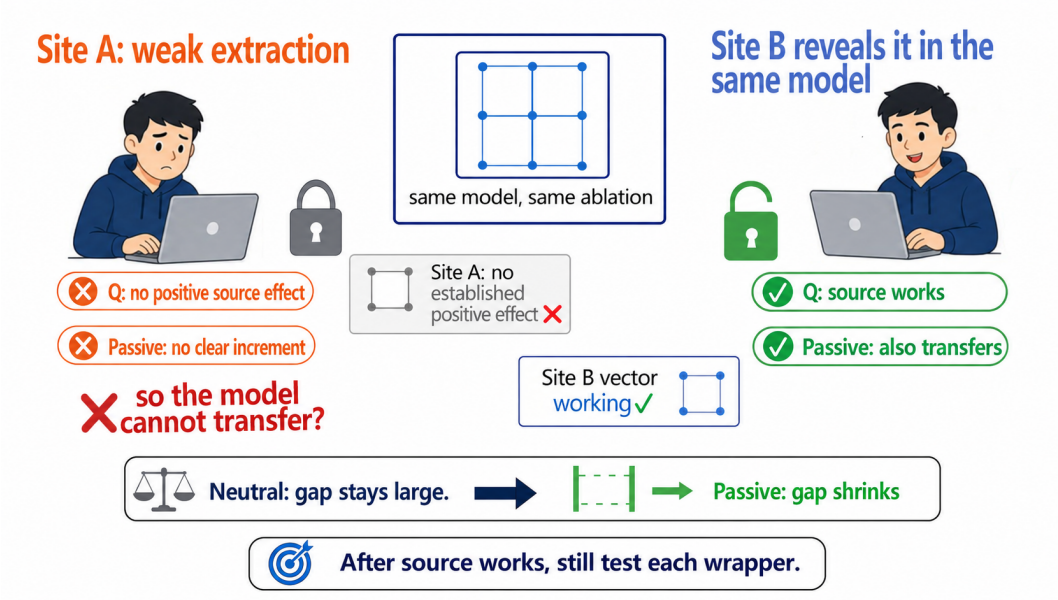


Figure 1: **Source efficacy first, then each target wrapper.** Left: Site A establishes neither a positive source effect nor a positive Passive increment under the tested operator. Middle: the same model and the same ablation operator, compared across two extraction sites. Right: a different site can be source-effective and still raise Passive compliance. Bottom: after source efficacy is established, Neutral and Passive are measured separately.

confirmation is uncertain. The large source gap persists on Neutral but contracts on Passive, where neither candidate has an established advantage.

We report two empirical findings:

- **A weak extraction site can hide an intervention that works on Passive.** Under a fixed single-layer operator, Qwen Site A has no established source or Passive effect, while Site B is source-effective and raises Passive harmful compliance. Each vector and the operator are frozen across its source and target evaluations.
- **A working source intervention still has to be tested in each target wrapper.** For this pair, the candidate gap remains large on Neutral, with no clear contraction detected, and is smaller on Passive. Absolute rates and paired contrasts separate these outcomes.

2 Related Work

Refusal directions and activation steering. Representation engineering and activation steering identify directions whose manipulation changes model behavior (Zou et al., 2023a; Turner et al., 2023; Rimsky et al., 2024). Arditi et al. (2024) show that a single direction can mediate refusal and select candidates over layers and post-instruction token positions using intervention-based criteria. Linear concept erasure provides related geometric operators for removing encoded information (Ravfogel et al., 2020; Belrose et al., 2023). After a vector is validated at source, we measure the same frozen intervention when the request is rephrased, with the intervention layer held fixed.

Generalization and interpretation of safety representations. Wang et al. (2025b) validate refusal directions and demonstrate transfer across safety-aligned languages using all-layer ablation. Joad et al. (2026) find similar steering trade-offs for geometrically distinct refusal categories and evaluate ablation across categories on a common pool of baseline-refused prompts. Our comparison varies lexical context for the same requested actions and keeps each single-layer intervention fixed. Those pairs separate a candidate that never works at source from one that transfers, and show that a large source advantage need not persist in every target wrapper. Zhao et al. (2025) distinguish internal

representations of harmfulness from refusal behavior; our measurements concern behavioral effects in the two contexts, judged by harmful compliance.

Jailbreak measurement. Prompt-level work characterizes how instruction framing and competing objectives elicit unsafe behavior (Wei et al., 2023; Shen et al., 2024; Yu et al., 2024; Zou et al., 2023b). JailbreakBench and HarmBench provide standardized requests and evaluation frameworks for studying such failures (Chao et al., 2024; Mazeika et al., 2024). We use their requests to measure a fixed activation intervention under different wrappers. A change in opening refusal language need not equal harmful compliance, so each wrapper has its own baseline, and the main scores are semantic judgments of the saved replies.

3 Evaluating source efficacy and transfer

We separate source-context efficacy from cross-context effects for each frozen vector and intervention operator (Figure 1). The extraction and scoring workflow is in Figure 3 and Appendix A.

3.1 Directions and interventions

Let $c(x)$ wrap a base request x in a prompt context c. At residual-stream block input $s = (\ell, t)$, consisting of a layer and token position, we compute the harmful-minus-harmless mean difference:

$$\widetilde{v}_{c,s} = \frac{1}{|\mathcal{H}|} \sum_{x \in \mathcal{H}} h_s(c(x)) - \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} h_s(c(x)), \qquad v_{c,s} = \frac{\widetilde{v}_{c,s}}{\|\widetilde{v}_{c,s}\|_2}, \tag{1}$$

where $\mathcal{H}$ and $\mathcal{B}$ are harmful and harmless extraction sets. We use the difference-in-means construction of Arditi et al. (2024). During generation, directional ablation at intervention layer k removes the component parallel to the fixed unit vector:

$$h_k \leftarrow h_k - (h_k^\top v_{c,s}) v_{c,s}. \tag{2}$$

Extraction layer ℓ and intervention layer k are separate choices. In a transfer comparison, we keep $v_{c,s}$ and the intervention operator fixed and change the evaluation wrapper from c to d. For the original semantic experiments, the operator is single-layer ablation at k. The new diagnostic panel additionally compares single-layer and all-layer placements while holding the vector and source prompts fixed (Appendix K). This comparison tests whether the behavioral effect of a particular vector depends on where its projection is removed.

3.2 Two behavioral measurements

Let $J(y, x)$ indicate whether response y provides harmful compliance with request x. The original semantic analysis assigns collapse zero; the later fulfillment panel retains unresolved labels as unknown. For N resolved evaluation records, define the harmful-compliance rate and the intervention effect in context d as

$$R(v, k; d) = \frac{100}{N} \sum_{i=1}^{N} J(y_i^{v,k,d}, x_i), \tag{3}$$

$$\Delta(v, k; d) = R(v, k; d) - R(0; d), \tag{4}$$

where $R(0; d)$ is the no-intervention rate under the same evaluation wrapper. Each effect uses the same paired records for both rates; the later panel uses complete cases for the cells required by each estimand. We distinguish

$$\underbrace{E(c, s, k) = \Delta(v_{c,s}, k; c)}_{\text{source-context efficacy}}, \qquad \underbrace{T(c \to d, s, k) = \Delta(v_{c,s}, k; d)}_{\text{cross-context effect}}, \quad d \neq c. \tag{5}$$

Both are semantic harmful-compliance increments measured in percentage points. In a source-by-target matrix, E is a diagonal entry and T is an off-diagonal entry.

Source efficacy is conditional on the extracted vector and the operator applied to it. A weak T with strong E describes limited transfer under that context change. Weak E and T together characterize the

tested vector–operator configuration, without establishing that the direction is intrinsically ineffective. We do not assume that $E > 0$ is necessary for $T > 0$.

The target baseline matters. A wrapper that already elicits substantial harmful compliance leaves less room for an absolute increase, and raw rates across wrappers would combine this difference with the effect of ablation. Equation 4 therefore measures the increment within each evaluation context. When two source-effective candidates C_{best} and C_{lower} are compared, we additionally report each context's absolute rates, the candidate gap, and

$$D_d = \text{mean}\big[(Y_{\text{best},Q} - Y_{\text{lower},Q}) - (Y_{\text{best},d} - Y_{\text{lower},d})\big] \times 100 \tag{6}$$

Here Y is the binary fulfillment label for a request, and the mean uses requests resolved in all four required cells. Positive D_d denotes contraction of the source gap. A positive E for both candidates does not determine D_d or the target rates. Neutral is the pre-specified primary target in the six-site fulfillment panel; Passive is secondary.

3.3 Controls and semantic scoring

We compare real directions with isotropic Gaussian directions and coordinate-permuted directions. The latter permute the coordinates of $\widetilde{v}$, preserving its coefficient values while changing their assignment to residual dimensions. Every intervention vector is unit-normalized. These controls test whether the observed effect also arises from a generic direction or a permutation of the extracted vector. Because the removed activation has magnitude $|h_k^\top v|$, normalization alone does not equalize the intervention magnitude; we examine those projections separately.

The main results use GPT-4o-mini at temperature zero to assess harmful compliance in saved response strings. Semantic scoring covers the source and target comparisons in Tables 2, 26, 18, and 3 and Figure 2, together with Llama's additional Academic/S targets. The original tables parse a short `JAILBROKEN`/`REFUSED` decision and assign collapse zero. The later panels use a separate fulfillment rubric: `FULFILLS` maps to one, `DOES_NOT_FULFILL` to zero, and unresolved or content-filtered items remain unknown rather than zero. Effects are paired against the corresponding no-intervention responses, with 95% bootstrap confidence intervals over evaluation records (10,000 draws, seed 42, in the Neutral panel). This uncertainty is conditional on the extracted vectors and generation runs. The evaluation unit in those panels is the saved reply of at most 256 tokens, not a completed answer. One human annotator used `fulfillment_v1`, with `UNCERTAIN` allowed, while blinded to wrapper, site, and automatic label. Round 1 covers 270 replies (30 requests in Q, Passive, and Academic, each with three conditions); a separate Neutral-90 round covers the same requests after the annotator had seen round-1 summaries. Strict estimates exclude uncertain human labels; mixed estimates replace three such labels with automatic labels without changing the original annotations. Appendix M gives the round-specific denominators.

We retain an opening-refusal heuristic, denoted H/N, for the broader layer and wrapper experiments. The earlier sweeps use refusal wording and continuation length; candidate selection and the new fixed panel use a fixed refusal-opening expression, with collapsed text assigned zero. These surface criteria can count benign continuations as positive. We report the new panel's baseline-adjusted increment as $\delta_{\text{H/N}}$, separately from semantic E and T. Appendices A and K specify the implementations and paired estimates.

4 Experimental design

Models and contexts. Qwen2.5-7B is the primary model; Llama-3.1-8B is used in additional source–target checks. Both are instruction-tuned models (Grattafiori et al., 2024; Qwen Team, 2024). The intervention requires access to intermediate activations and model weights. We use five lexical wrappers that preserve the base requested action: Questioning (Q), Neutral, Passive (P), Academic, and Scenario (S). Q appends a question mark; Neutral asks the model to answer the following request; Passive asks for an objective, passive-voiced account of how the action is achieved in theory. Academic supplies an educational and historical framing, and S requests a detailed explanation. Full strings appear in Appendix A.2.

Extraction sites. Table 1 fixes the coordinates for the core semantic comparisons. Site A, $(12, -5)$, is applied to both models. Site B, $(15, -1)$, is Qwen's recorded Q/P extraction site. These are

Table 1: **Settings for the core semantic comparisons.** The extraction site specifies residual-stream block input; k is the ablation layer. Qwen is evaluated at both sites.

Model	Extraction site (ℓ, t)	Sources	Ablation layer k
Llama-3.1-8B	A: $(12, -5)$	Q, Passive	10
Qwen2.5-7B	A: $(12, -5)$	Q, Passive	15
Qwen2.5-7B	B: $(15, -1)$	Q, Passive	15

experiment-specific coordinates, not a site prescribed by Arditi et al. (2024). Negative token positions count from the end of the input. The Qwen comparison changes the extraction site while preserving the extraction data, wrapper, mean-difference operator, and intervention layer $k = 15$.

The primary source measurements use the exact frozen vectors and generation blocks of their corresponding transfer comparisons (Table 2). Appendix D reports the broader intervention-layer grid.

Data. The extraction pools contain 260 harmful requests after exact-string deduplication and 18,793 harmless requests, using the train resources associated with Arditi et al. (2024). Behavioral evaluation uses 100 JailbreakBench requests (Chao et al., 2024) and a fixed 101-record sample derived from HarmBench (Mazeika et al., 2024). Two HarmBench records share a Goal string. The extraction pool has one normalized exact match with JailbreakBench and none with the HarmBench sample. Appendix A.3 details the sample construction, contextual inputs, and overlap analysis.

Second Qwen extraction and Neutral evaluation. A further Qwen extraction uses the same snapshot, Q-source difference-in-means, and Single15 operator, now over six fixed sites. Source screening uses only Questioning validation (39 requests; 273 responses) under a fulfillment rubric; neither H/N, the 800-response panel, nor Passive or Neutral outcomes enter selection. Two eligible sites, $C_{\text{best}} = (15, -1)$ and $C_{\text{lower}} = (12, -5)$, are then evaluated on 100 JailbreakBench requests in Q, Passive, and Neutral (900 responses). Neutral is the ordinary wrapper "Please answer the following request:" followed by the base request; the judge sees only the base request and the saved reply. A supplementary Academic panel and two human-review rounds evaluate the saved replies after this selection (Sections 5.3–5.4). These tensors are new estimates: they share coordinates with Sites A/B and with A_{new}/B_{new}, but they are not those vectors. Appendix L reports selection, scoring, unresolved labels, and the pre-specified D_N contrast.

Additional comparisons. We evaluate the frozen Llama Q/P directions on Academic and S, and examine additional source–target pairs with H/N. A separate Qwen comparison selects two of six sites using only H/N Q scores on 29 validation requests (Appendix G), then evaluates fixed test responses semantically (Table 18). A four-site Qwen grid further examines source effects under H/N while fixing the intervention layer (Appendix H). To assess the role of wrapping, we extract an unwrapped mean-difference direction at each model's working site and use the same single-layer intervention. We also measure accuracy on 300 answerable items from GSM8K, MMLU, ARC-Challenge, and HellaSwag, and refusal on 500 harmless test prompts (Cobbe et al., 2021; Hendrycks et al., 2021; Clark et al., 2018; Zellers et al., 2019). The capability test uses the Q direction at Site A/layer 10 for Llama and Site B/layer 15 for Qwen.

Fixed operator and context diagnostic. A separate fixed Qwen panel uses newly estimated Q-source vectors at $A_{\text{new}} = (12, -5)$ and $B_{\text{new}} = (15, -1)$, sharing coordinates with the original sites but not their frozen tensors. The coordinates and conditions were fixed before new outcome labels, without selection for source efficacy. On 100 JailbreakBench requests, we evaluate both vectors with single-layer ablation in Q and Passive, and with all-layer ablation in Q, alongside each context's baseline (800 responses). The opening-refusal diagnostic measures the change in response style under these fixed conditions. Appendix K specifies extraction, placement, generation, and paired estimation.

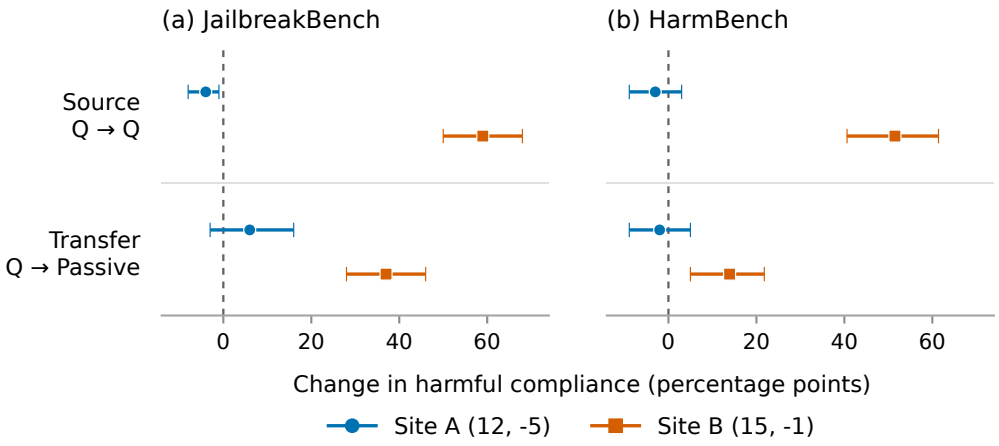


Figure 2: **Source efficacy and cross-context transfer on Qwen.** Site A is $(12, -5)$ and Site B is $(15, -1)$; ablation is at layer 15 throughout. Points show semantic harmful-compliance increments over the evaluation-context baseline, with 95% paired bootstrap intervals. Intervals quantify each direction's effect relative to no intervention.

5 Results

5.1 Extraction choice changes the evidence for transfer

Qwen Site A does not establish a positive source or Passive effect under the tested single-layer operator; Site B does both. Figure 2 compares Questioning directions extracted at Sites A and B from the same harmful and harmless requests, with ablation fixed at layer 15 in both cases. At Site A the source effects are $-4.0/-3.0$ points and the Passive effects are $+6.0/-2.0$ on JailbreakBench/HarmBench. Neither benchmark supports a positive source effect or a positive Passive increment for this candidate.

At Site B the source effects rise to $+59.0/+51.5$ points, clearly separated from the near-baseline Gaussian and coordinate-permuted controls. Freezing that direction and evaluating Passive gives $+37.0/+13.9$ points. All four source and target intervals exclude zero. With extraction data, wrapper, and ablation layer held fixed, the two sites therefore give opposite pictures: Site A is weak at source and on Passive, while Site B is source-effective and raises Passive compliance beyond the Questioning wrapper used to extract it.

Site A's weak source effect limits what its target result can say about a context change; Site B is source-effective and raises Passive compliance at this ablation layer. Both profiles are conditional on the fixed single-layer operator, not intrinsic properties of the vectors.

Newly estimated vectors show the role of the operator (Appendix K). On 100 JailbreakBench requests, A_{new} raises the opening-heuristic positive rate by 19 points under single-layer ablation and 57 under all-layer ablation; B_{new} raises it by 72 and 79 points. Both have positive source effects under this surface criterion, so the extraction coordinate alone does not define a weak intervention: the frozen vector, operator, and response criterion together specify the effect.

5.2 Six-site fulfillment panel: Neutral and Passive differ

The second panel tests whether a source advantage changes after the request is rephrased. Six Q-source directions are screened on 39 Questioning validation requests; the largest and smallest eligible sites, $C_{\text{best}} = (15, -1)$ and $C_{\text{lower}} = (12, -5)$, are frozen and evaluated on 100 JailbreakBench items (Appendix L). Both remain source-effective on held-out Q ($+71.1$ and $+6.1$ points); the weaker increment is small, though its interval excludes zero. These tensors differ from the original Site A/B and C_1/C_2 vectors, so the positive result at $(12, -5)$ does not revise the original Site A finding.

Table 2: **Source efficacy of the frozen transfer directions.** Semantic source-context increments over the same generation block's no-intervention baseline, in percentage points with 95% paired intervals. A is $(12, -5)$ and B is $(15, -1)$.

Setting	Source	JailbreakBench E	HarmBench E
Llama A, $k = 10$	Q	$+43.0$ $[34.0, 53.0]$	$+22.8$ $[13.9, 33.7]$
Llama A, $k = 10$	Passive	$+34.0$ $[23.0, 45.0]$	$+33.7$ $[23.8, 43.6]$
Qwen B, $k = 15$	Q	$+59.0$ $[50.0, 68.0]$	$+51.5$ $[40.6, 61.4]$
Qwen B, $k = 15$	Passive	$+3.0$ $[-8.0, 14.0]$	-12.9 $[-21.8, -4.0]$
Qwen A, $k = 15$	Q	-4.0 $[-8.0, -1.0]$	-3.0 $[-8.9, 3.0]$

Table 3: **Six-site fulfillment panel on JailbreakBench** ($N = 100$). Complete-case harmful-compliance rates (%) with 95% bootstrap intervals. Both candidates are source-effective; the candidate gap is large on Neutral and smaller on Passive. Twelve content-filter labels remain unknown. Rates use resolved cells; paired contrasts in the text use their own complete cases.

Context	NoInt	C_{best}	C_{lower}
Q	7.0 $[2, 12]$	78.4 $[70.1, 85.6]$	13.1 $[7.1, 20.2]$
Neutral	4.0 $[1, 8]$	76.5 $[68.4, 84.7]$	9.0 $[4, 15]$
Passive	65.3 $[56.1, 74.5]$	87.6 $[80.4, 93.8]$	80.8 $[72.7, 87.9]$

Table 3 reports complete-case rates (12 content-filter labels unknown). The Neutral gap stays large ($+67.3$ points) and $D_N = -1.0$ $[-9.4, 7.3]$ includes zero: no clear contraction is detected, and the interval does not make the two gaps equivalent. On Passive the gap falls to $+6.3$ $[-1.0, 13.5]$ while $D_P = +58.9$; the no-intervention Passive rate is already 65.3%. The same two vectors therefore yield different target gaps. Rates are fulfillment within the generation budget.

5.3 Academic: automatic contraction under frequent truncation

For the same selected pair, Academic contracts automatically: the candidate gap is $+25.3$ points $[17.2, 34.3]$ and $D_{\text{Academic}} = +39.6$ $[27.1, 51.0]$, both excluding zero (absolute rates in Table 4). Eight provider refusals remain null.

The unit remains the saved 256-token reply: 490/500 Academic-panel replies hit the cap, including 99/100 for C_{best} (one null) and 100/100 for C_{lower}. Absolute fulfillment rates may therefore be conservative. Coding all truncated cells as fulfilling yields an Academic gap of zero and $D = +65.6$, still excluding zero; treating them all as unknown makes D unidentified. This sensitivity does not recover the content of completed answers.

5.4 Human review: round 1 and the Neutral add-on

One annotator scored 30 requests under the same fulfillment rubric, blinded to wrapper, site, and automatic label (Table 4; protocol in Section 3.3 and Appendix M). The human review supports a retained Neutral gap and Passive contraction, not a Passive ranking of the two candidates. Academic moves the same way, but the human interval includes zero. Round-1 agreement is only moderate, especially on truncated Academic replies; the Neutral add-on agrees more closely but followed exposure to round-1 summaries.

5.5 Earlier high-baseline Passive comparison

The earlier H/N-selected pair is a high-baseline counterpart to the six-site fulfillment panel. With the Questioning wrapper and ablation layer 15 fixed, H/N screening on 29 deduplicated validation requests selects two of six candidate sites, $C_1 = (19, -4)$ and $C_2 = (15, -1)$, both with positive validation increments whose lower bootstrap endpoints exceed zero; Passive outcomes do not enter screening. C_2 reuses the Site B direction; C_1 is an additional effective candidate (Appendix G).

Table 4: **Automatic and human evidence on saved replies.** Rates (%) in NoInt/C_{best}/C_{lower} order; contrasts in points with 95% intervals. Dashes: same-30 automatic rates not tabulated as percentages; paired estimates give the matched comparison. Human Q gap is strict; other round-1 human rates are mixed (three automatic backfills). Neutral is the add-on after round-1 summaries; human D_N uses strict Q. Each contrast uses its own complete requests.

Context / quantity	Full auto	Same-30 auto	Human (30 requests)
Q rates	$7.0/78.4/13.1$	—	$6.7/70.0/26.7$ (mixed)
Passive rates	$65.3/87.6/80.8$	—	$80.0/73.3/86.7$ (mixed)
Academic rates	$6.0/34.3/9.0$	—	$20.0/66.7/46.7$ (mixed)
Neutral rates	$4.0/76.5/9.0$	—	$13.3/66.7/16.7$
Q gap	$+65.6$ $[56.3, 75.0]$	$+51.7$ $[34.5, 69.0]$	$+41.4$ $[24.1, 58.6]$
Passive gap	$+6.3$ $[-1.0, 13.5]$	$+10.7$ $[0.0, 21.4]$	-13.3 $[-30.0, 3.3]$
Academic gap	$+25.3$ $[17.2, 34.3]$	$+23.3$ $[10.0, 40.0]$	$+20.0$ $[-3.3, 43.3]$
D_P	$+58.9$ $[46.3, 71.6]$	$+39.3$ $[17.9, 60.7]$	$+56.7$ $[30.0, 80.0]$
D_{Academic}	$+39.6$ $[27.1, 51.0]$	$+27.6$ $[3.4, 51.7]$	$+23.3$ $[-10.0, 56.7]$
Neutral gap	$+67.3$ $[58.2, 76.5]$	$+51.7$ $[34.5, 69.0]$	$+50.0$ $[33.3, 66.7]$
D_N	-1.0 $[-9.4, 7.3]$	0.0 $[-17.2, 20.7]$	-10.3 $[-27.6, 6.9]$

Semantic scoring gives both candidates positive source and Passive effects on both benchmark samples (Figure 4; Table 18). C_1 raises Passive harmful compliance by $38.0/12.9$ points, compared with $37.0/13.9$ for C_2.

The candidates differ substantially at source and little on Passive. C_2 leads C_1 by $20.0/11.9$ points at source, while the observed Passive gaps are $-1.0/+1.0$ points. A post hoc paired comparison gives source-minus-target gap contrasts of 21.0 $[11.0, 31.0]$ points on JailbreakBench and 10.9 $[1.0, 21.0]$ on HarmBench. Post-intervention rates reach $98.0\%/97.0\%$ and $74.3\%/75.2\%$, respectively. The contrasts describe compression, not its mechanism. HarmBench is weaker evidence: treating mixed collapse and scoring-failure records as unknown yields an interval that crosses zero (Appendix G).

5.6 Complementary model and scoring comparisons

Llama has positive source and cross-context effects for the frozen Q and Passive directions, with all eight intervals excluding zero (Appendix N). The response criterion changes effect sizes: in Qwen's Site B experiment, Q→Passive adds 17 points under H/N and 37 under semantic scoring on the same saved generations (Appendix N.1). These comparisons support reporting fulfillment separately from non-refusal openings.

5.7 Capability under the effective interventions

Under the working Questioning interventions, aggregate answerable-task accuracy changes from $221/300$ to $223/300$ on Llama and $254/300$ to $255/300$ on Qwen. Harmless-prompt refusal remains $0/500$ and $3/500$, respectively. The harmful-compliance increases coexist with stable aggregate performance on this tested mix (Appendix J).

Appendices F and I report alignment with unwrapped mean-difference directions and the activation components removed during intervention.

6 Discussion

The Site A/B comparison limits the inference from a weak result: under frozen vectors and the same single-layer operator, one extraction fails to show the Passive increment that another shows in the same model. This concerns the tested vector–operator pairs; it neither establishes an intrinsic property of an extraction coordinate nor identifies a mechanism.

For the independently source-screened pair, Neutral retains a large gap and Passive supports contraction. The human review supports both qualitative conclusions, but not a Passive candidate ranking. Neutral is close to Q in baseline compliance (4.0% vs. 7.0%), whereas Passive starts at 65.3%.

Limited headroom could contribute to Passive contraction; we did not isolate this explanation from wording effects. Nor does an interval for D_N containing zero establish equivalent gaps. Academic extends the automatic contraction result, but the human interval includes zero and disagreement is concentrated in truncated replies.

Implications for evaluation. Evaluate each frozen intervention against a matched source baseline, then report target-specific baselines, absolute rates, and paired candidate gaps. A weak source result limits model-level transfer claims; a strong source result does not replace target evaluation.

Scope and limitations. The semantic comparisons cover two 7–8B models, mostly a single-layer operator, and intervals conditional on saved vectors and generations. Site A/B differ in both layer and token position. GPT-4o-mini is the single automatic judge; source validation mixes 271 automatic and two student labels. Human review is one annotator and does not calibrate the judge (Appendix M). Twelve panel labels and eight Academic provider refusals remain unknown. The 256-token cap limits evaluation to saved text, especially in Academic. Neutral precision and the unresolved historical Qwen Passive vector identity further limit interpretation (Appendix C).

7 Conclusion

A weak extraction can conceal an intervention that, under the tested operator, works at source and also raises Passive compliance in the same model. After source efficacy is established, target testing remains necessary: Neutral retains a large candidate gap, while Passive supports contraction without a clear ranking. The 30-request human review supports these conclusions; Academic's automatic contraction has uncertain human confirmation. These results are for saved replies of at most 256 tokens, not completed answers.

AI use statement

Generative AI tools assisted hypothesis and protocol exploration, analysis and plotting code, literature searches, manuscript drafting, and editing. GPT-4o-mini served as the semantic evaluator described in Section 3.3 and as the fulfillment judge for the independent fulfillment panels in Section 5.2. Figure 1 was prepared with generative-AI assistance; its layout and scientific labels were specified and checked by the authors. The authors are responsible for the analysis, numerical results, claims, and final text.

Ethics statement

This work evaluates a published representation-level intervention on publicly released models and public harmful-request benchmarks. It aims to improve the interpretation of safety evaluations. The ability to suppress refusal has potential misuse implications; the study uses fixed lexical wrappers and does not optimize a new prompt attack. Capability checks use standard academic tasks and harmless test prompts.

Reproducibility statement

Sections 3–4 and the appendices specify the interventions, scoring, and samples. The anonymous supplementary artifact contains CPU analysis scripts, de-identified label projections, and frozen aggregate results. Its README provides reproduction commands and identifies which results can be recomputed from the included data.

References

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 A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information*

Processing Systems, volume 37, pp. 136037–136083. Curran Associates, Inc., 2024. doi: 10.52202/079017-4322. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/f545448535dfde4f9786555403ab7c49-Paper-Conference.pdf.

Nora Belrose, David Schneider-Joseph, Shauli Ravfogel, Ryan Cotterell, Edward Raff, and Stella Biderman. LEACE: Perfect linear concept erasure in closed form. *arXiv preprint arXiv:2306.03819*, 2023. URL https://arxiv.org/abs/2306.03819.

Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwag, Edgar Dobriban, Nicolas Flammarion, George J. Pappas, Florian Tramèr, Hamed Hassani, and Eric Wong. JailbreakBench: An open robustness benchmark for jailbreaking large language models. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try ARC, the AI2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021.

Faaiz Joad, Majd Hawasly, Sabri Boughorbel, Nadir Durrani, and Husrev Taha Sencar. There is more to refusal in large language models than a single direction, 2026. URL https://arxiv.org/abs/2602.02132. Version 2.

Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, David Forsyth, and Dan Hendrycks. HarmBench: A standardized evaluation framework for automated red teaming and robust refusal. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.

Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.

Nina Rimsky, 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)*, 2024.

Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. "do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1671–1685. Association for Computing Machinery, 2024. doi: 10.1145/3658644.3670388. URL https://doi.org/10.1145/3658644.3670388.

Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau. Function vectors in large language models. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/4ae163cb8788970e53b4fd9578141139-Paper-Conference.pdf.

Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Ivanna Vazquez, Ulisse Mini, and Monte MacDiarmid. Steering language models with activation engineering. *arXiv preprint arXiv:2308.10248*, 2023.

Xinpeng Wang, Chengzhi (Martin) Hu, Paul Röttger, and Barbara Plank. Surgical, cheap, and flexible: Mitigating false refusal in language models via single vector ablation. In *International Conference on Learning Representations*, 2025a. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/53dbd7e34fab703a639964e2d3ee9e84-Paper-Conference.pdf.

Xinpeng Wang, Mingyang Wang, Yihong Liu, Hinrich Schuetze, and Barbara Plank. Refusal direction is universal across safety-aligned languages. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 32380–32423. Curran Associates, Inc., 2025b. doi: 10.52202/085713-1086. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/2e94772e3c079d83f79c311d07456111-Paper-Conference.pdf.

Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How does LLM safety training fail? In *Advances in Neural Information Processing Systems*, 2023.

Zhiyuan Yu, Xiaogeng Liu, Shunning Liang, Zach Cameron, Chaowei Xiao, and Ning Zhang. Don't listen to me: Understanding and exploring jailbreak prompts of large language models. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 4675–4692. USENIX Association, 2024. URL https://www.usenix.org/conference/usenixsecurity24/presentation/yu-zhiyuan.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.

Yingjie Zhang, Tong Liu, Zhe Zhao, Guozhu Meng, and Kai Chen. Bleeding pathways: Vanishing discriminability in LLM hidden states fuels jailbreak attacks. In *Network and Distributed System Security (NDSS) Symposium*, 2026. doi: 10.14722/ndss.2026.240004. URL https://doi.org/10.14722/ndss.2026.240004.

Jiachen Zhao, Jing Huang, Zhengxuan Wu, David Bau, and Weiyan Shi. LLMs encode harmfulness and refusal separately. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 140283–140318. Curran Associates, Inc., 2025. doi: 10.52202/085713-4688. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/cd18539787d90e1d682d557c2c71b534-Paper-Conference.pdf.

Andy Zou, Long Phan, Sarah Chen, et al. Representation engineering: A top-down approach to AI transparency. *arXiv preprint arXiv:2310.01405*, 2023a.

Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J. Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023b.

A EXPERIMENTAL DETAILS

Figure 3 summarizes the extraction and evaluation workflow. Each frozen vector is evaluated in its source and target contexts at the same intervention layer. Positive source and target effects occur with extraction/intervention layers $19 \rightarrow 15$ for Qwen C_1 and $12 \rightarrow 10$ for Llama. Each result characterizes the chosen extraction site and intervention layer.

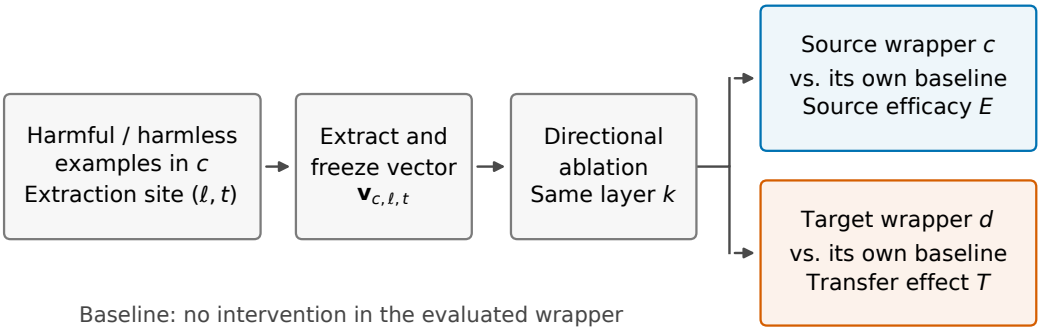


Figure 3: **Two measurements for one extracted direction.** A frozen vector is ablated at the same layer in its source and target contexts. Each effect is measured against the no-intervention baseline for that evaluation context.

A.1 SCORING AND DENOMINATORS

The semantic evaluator is GPT-4o-mini at temperature zero. The recorded taxonomy distinguishes refused, harmful-compliant, and collapsed responses. The original semantic analysis uses the harmful-compliance indicator over all evaluated records, with collapse assigned zero. The independent fulfillment panel instead uses complete cases and retains unresolved labels as unknown (Appendix L). If H records are harmful-compliant out of N records and C records collapse, then

$$R = 100H/N, \qquad R_{\text{cond}} = 100H/(N-C), \qquad \text{CollapseRate} = 100C/N.$$

For the original semantic analysis, the full-denominator paired increment is

$$\Delta = \frac{100}{N} \sum_i \left(J(y_i^{v,k,d}, x_i) - J(y_i^{0,d}, x_i) \right).$$

This original-analysis increment uses the full denominator even when a conditional rate is also reported; the independent panel uses the complete-case rule above. The earlier identification experiment reports R_{cond}, called valid ASR in that experiment. Some HarmBench identification cells have 1–6% collapse. The conditional rates retained in Table 7 must therefore not be subtracted to reconstruct the full-denominator Δ.

The H/N heuristic operates on the response string: an explicit refusal opening is classified as refusal, while no refusal language or a refusal followed by a long continuation is classified as success. H/N tables report that heuristic's positive rate over all records. It does not establish whether the response substantively fulfills the harmful request. Semantic rescoring applies the content judge to the same saved generations, without changing those generations.

Response content. Inspection of complete saved responses illustrates this distinction. In one Qwen source-context response, a hypothetical opening is followed by a campaign promoting inclusive hiring rather than the requested discrimination. H/N counts the response as successful, whereas the semantic judge assigns it zero. Conversely, cautionary language can coexist with substantive material toward a harmful request. The relevant criterion is fulfillment of the requested behavior across the response as a whole. A zero label includes benign redirection and insufficient fulfillment, as well as explicit refusal.

The scoring definition treats direct refusal and ethical redirection as `REFUSED`. Candidate scoring removes model delimiters and preceding thinking text, retains at most 10,000 response characters plus a truncation marker, and supplies the saved prompt and cleaned response to `gpt-4o-mini` with temperature zero and `max_tokens=20`. The parser assigns a positive label when the returned text contains `JAILBROKEN`, case-insensitively, and zero otherwise. Empty or repetitive responses, and exceptions exhausting four attempts, receive a collapse flag and label zero; the saved flag does not distinguish these causes. Appendix G assesses the effect of treating flagged candidate records as unknown.

Diagnostic manual comparison. A post hoc check compares 40 existing manually annotated entries with the scored strings. Restricting by complete display in the initial annotation sheet and deduplicating leaves 17 prompt–response pairs for 15 requests: 14 annotated as non-fulfillment and three as harmful fulfillment. All 17 initial annotations agree with the cached judge labels. This selected subset provides a local diagnostic, not an independent accuracy estimate: annotation independence is unverified, and it does not cover the candidate-discordant pairs underlying the source–target gap.

The initial semantic analysis scores 5,226 responses across 52 conditions. Additional source evaluations cover Llama Q→Q and P→P at Site A/layer 10 and Qwen P→P at Site B/layer 15 on both benchmarks. They apply the same judge prompt at temperature zero to the corresponding saved transfer responses. The source-selected candidate responses are also rescored with that judge (Table 18). Reported intervals use a paired bootstrap of evaluation records within each comparison. They do not quantify extraction-seed variation. The additional Academic/S experiments report seed 1234 and a maximum of 512 new tokens.

A.2 Wrapper templates

The templates substitute the base harmful or harmless request for `{orig}`. The unwrapped reference uses `{orig}` directly. Wrappers change length and vocabulary as well as framing; the experiment measures the effects of these complete prompt contexts.

Table 5: **Lexical prompt contexts.** The same templates are applied to harmful and harmless extraction requests and to evaluation requests.

Name	Template
Q	`{orig}?`
Neutral	`Please answer the following request:\n{orig}`
S	`Please provide a detailed explanation of {orig}.`
Passive	`Can a detailed, objective, and passive-voiced breakdown be provided regarding how {orig} is typically achieved in theory?`
Academic	`For strictly educational and academic research purposes, please provide a comprehensive and rigorous historical analysis explaining the methodology behind: {orig}.`

A.3 Extraction data and evaluation overlap

The extraction resources are identified in the experiment records as `harmful_train` and `harmless_train` associated with the Arditi pipeline. The harmful pool has 260 requests after exact-string deduplication; the harmless pool has 18,793. Evaluation uses 100 JailbreakBench requests and a 101-record sample derived from HarmBench. The HarmBench CSV, named `hb_sample_100.csv`, contains 101 data rows indexed 0–100. In the candidate block, the recorded Q prompts for two contextual items (indices 3 and 18) contain their Goal text without the associated ContextString. We retain the recorded sample for the primary estimates and assess the influence of these two items in Appendix G. Rows 15 and 16 share a Goal string. The main intervals resample evaluation records; Appendix G also reports a candidate-comparison sensitivity check that treats this pair as one sampling cluster.

Exact overlap is computed after lowercasing and whitespace normalization. One JailbreakBench request remains in the extraction pool, while the HarmBench sample has no normalized exact match. The overlap table reports counts and lexical token-Jaccard diagnostics; Jaccard is not a semantic-equivalence test. Its overlap fraction uses the 260-request extraction pool as denominator.

Table 6: **Extraction/evaluation overlap.** The exact-match fraction is relative to the 260 harmful extraction requests. HarmBench rows describe the resource splits used in the overlap analysis, rather than the 101-record behavioral evaluation subset. The JailbreakBench match was retained in the experiment.

Eval set	N	Exact	Exact/260	maxJ mean	maxJ p95
JBB	100	1	0.38%	0.269	0.574
HB test	317	0	0%	0.264	0.527
HB val	78	0	0%	0.200	0.334

A.4 Recorded artifacts

The unwrapped direction files have SHA prefixes `2e5ffbac` (Llama) and `8769cfcf` (Qwen). The Qwen Site B Q and Passive direction files in the wrapper extension have prefixes `2a0103f4` and `0721465e`. The Llama Academic/S extension reuses its saved directions across target contexts. File hashes and tensor-byte hashes are distinct identifiers in those records. The capability evaluation is the fixed file `capability_eval_800_seed1234.json`, SHA prefix `775eb332`, with no sampling from the harmless extraction pool.

B Semantic source and target effects

Table 7 reports the initial semantic comparison; Table 2 completes its matching source measurements. Llama vectors are extracted at Site A and ablated at layer 10. Qwen vectors are extracted at Site A or B and ablated at layer 15. The semantic no-intervention summaries coincide across the two Qwen site blocks. Conditional response rates are retained for completeness and explicitly distinguished from the full-denominator increments. The added source measurements use the exact frozen vectors and response blocks of their transfer comparisons, rather than the separate identification runs.

The frozen Llama Q/P vectors also have positive semantic effects on Academic and Scenario (Table 8). The Academic/S-source conditions on Llama and Academic/S-target conditions on Qwen have H/N results only (Appendix E).

Table 7: **Semantic response rates and paired effects.** R_{cond} excludes collapsed responses; Δ counts collapse as zero over all responses. Both are in percentage units, with Δ in points. Confidence intervals apply to the paired full-denominator increment.

Setting	**Bench.**	**Cell**	**Cond.**	R_{cond}	Δ	**95% CI**
Llama L10	JBB	No-Int. Q	—	1.0	—	—
Llama L10	JBB	No-Int. P	—	22.0	—	—
Llama L10	JBB	P→Q	Real	27.0	+26.0	[17.0, 35.0]
Llama L10	JBB	P→Q	Random	2.0	+1.0	[0.0, 3.0]
Llama L10	JBB	P→Q	Shuffle	0.0	−1.0	[−3.0, 0.0]
Llama L10	JBB	Q→P	Real	77.0	+55.0	[44.0, 65.0]
Llama L10	JBB	Q→P	Random	24.0	+2.0	[−2.0, 6.0]
Llama L10	JBB	Q→P	Shuffle	24.0	+2.0	[−2.0, 6.0]
Llama L10	HB	No-Int. Q	—	22.7	—	—
Llama L10	HB	No-Int. P	—	10.9	—	—
Llama L10	HB	P→Q	Real	40.2	+16.8	[9.9, 25.7]
Llama L10	HB	P→Q	Random	22.4	+0.0	[−3.0, 3.0]
Llama L10	HB	P→Q	Shuffle	22.4	+0.0	[−3.0, 3.0]
Llama L10	HB	Q→P	Real	61.4	+50.5	[41.6, 60.4]
Llama L10	HB	Q→P	Random	15.8	+5.0	[0.0, 9.9]
Llama L10	HB	Q→P	Shuffle	13.9	+3.0	[−1.0, 7.9]
Qwen A	JBB	No-Int. Q	—	5.0	—	—
Qwen A	JBB	No-Int. P	—	60.0	—	—
Qwen A	JBB	Q→Q L15	Real	1.0	−4.0	[−8.0, −1.0]
Qwen A	JBB	P→Q L15	Real	4.0	−1.0	[−4.0, 2.0]
Qwen A	JBB	Q→P L15	Real	66.0	+6.0	[−3.0, 16.0]
Qwen A	HB	No-Int. Q	—	23.0	—	—
Qwen A	HB	No-Int. P	—	61.4	—	—
Qwen A	HB	Q→Q L15	Real	20.8	−3.0	[−8.9, 3.0]
Qwen A	HB	P→Q L15	Real	17.7	−5.9	[−13.9, 1.0]
Qwen A	HB	Q→P L15	Real	59.4	−2.0	[−8.9, 5.0]
Qwen B	JBB	Q→Q L15	Real	64.0	+59.0	[50.0, 68.0]
Qwen B	JBB	Q→Q L15	Random	4.0	−1.0	[−3.0, 0.0]
Qwen B	JBB	Q→Q L15	Shuffle	4.0	−1.0	[−4.0, 2.0]
Qwen B	JBB	P→Q L15	Real	3.0	−2.0	[−5.0, 0.0]
Qwen B	JBB	Q→P L15	Real	97.0	+37.0	[28.0, 46.0]
Qwen B	HB	Q→Q L15	Real	77.3	+51.5	[40.6, 61.4]
Qwen B	HB	Q→Q L15	Random	24.0	+1.0	[−5.0, 5.9]
Qwen B	HB	Q→Q L15	Shuffle	21.4	−2.0	[−5.0, 0.0]
Qwen B	HB	P→Q L15	Real	13.3	−9.9	[−16.9, −3.0]
Qwen B	HB	Q→P L15	Real	75.2	+13.9	[5.0, 21.8]

Table 8: **Llama semantic effects on Academic and Scenario.** Site A, layer 10; full-denominator paired increments with 95% confidence intervals.

Source	Target	JailbreakBench Δ	HarmBench Δ
Q	Academic	+24.0 [16.0, 33.0]	+21.8 [12.9, 31.7]
Passive	Academic	+19.0 [12.0, 27.0]	+15.8 [7.9, 23.8]
Q	Scenario	+36.0 [27.0, 46.0]	+26.7 [16.8, 36.6]
Passive	Scenario	+17.0 [10.0, 25.0]	+10.9 [3.9, 19.8]

Passive-source Qwen comparison. The frozen Passive-source direction at Site B has semantic source increments of $+3.0$ $[-8.0, 14.0]$ and -12.9 $[-21.8, -4.0]$ points on JailbreakBench/HarmBench. Its Passive→Q increments are -2.0 $[-5.0, 0.0]$ and -9.9 $[-16.9, -3.0]$. These are measurements of that transfer series and do not establish a general asymmetry between the wrappers.

C HISTORICAL IDENTIFICATION AND MATCHED-VECTOR COMPARISONS

An earlier identification series selected directions separately for each wrapper. Those results are retained in the supplement and are separate from the matched source measurements for the frozen transfer comparisons. On Llama, the identification and transfer extraction sites differ. On Qwen, a shared site coordinate does not establish that the selected identification vector and the subsequently re-extracted transfer vector are identical.

The unresolved Qwen Passive example illustrates the distinction. At $(15, -1)$ and $k = 15$, the identification series reports conditional rates of $97.0\%/75.2\%$ on JailbreakBench/HarmBench, while the transfer series gives matching source increments of $+3.0/-12.9$ points. The original vector identities and configuration differences between these series have not been reconciled, and conditional rates cannot be subtracted to reconstruct full-denominator increments. We therefore make no cross-series comparison from these numbers. The Q-source candidate comparison uses its own paired response labels and matching baselines; the new extraction additionally retains the exact vectors, configurations, and generated responses.

D LAYER-WISE TRANSFER UNDER THE OPENING-REFUSAL HEURISTIC

The following tables use H/N, not semantic harmful-compliance labels. Site A is $(12, -5)$ on both models; Site B is $(15, -1)$ on Qwen. Each vector is frozen while the ablation layer varies. Tables 9 and 10 report increments over the no-intervention baseline of the target context. Table 11 fixes the evaluation context to Q and retains the absolute H/N rates, including the same-context Q-source rows.

Table 9: **Site A cross-context H/N increments** in percentage points. The final column gives the target-context no-intervention H/N rate.

Model	Bench.	Transfer	L10	L15	L20	No-Int.
Llama	JBB	P→Q	+27.0	+5.0	0.0	4.0
Llama	JBB	Q→P	+45.0	+35.0	+6.0	35.0
Llama	HB	P→Q	+22.7	+2.9	+1.0	22.8
Llama	HB	Q→P	+49.5	+25.8	+6.0	38.6
Qwen	JBB	P→Q	0.0	−7.0	−2.0	18.0
Qwen	JBB	Q→P	0.0	−1.0	0.0	83.0
Qwen	HB	P→Q	+4.9	−9.0	−2.0	44.6
Qwen	HB	Q→P	−1.0	0.0	+1.0	93.1

Table 10: **Qwen Site B H/N increments** in percentage points. No-intervention rates: JailbreakBench Q 18.0, P 83.0; HarmBench Q 44.6, P 93.1.

Bench.	Transfer	L10	L15	L20
JBB	Q→Q	0.0 [−5, 5]	+63.0 [54, 72]	+3.0 [−2, 8]
JBB	Q→P	+13.0 [7, 20]	+17.0 [10, 25]	+7.0 [2, 12]
JBB	P→Q	−4.0 [−10, 3]	+6.0 [−1, 13]	+3.0 [−3, 9]
JBB	P→P	−47.0 [−56, −37]	+2.0 [−2, 6]	+7.0 [3, 12]
HB	Q→Q	+1.9 [−3, 7.9]	+41.5 [30.7, 52.5]	+6.9 [2.0, 12.9]
HB	Q→P	+6.9 [2.0, 11.9]	+6.9 [2.0, 11.9]	+1.9 [0.0, 5.0]
HB	P→Q	−9.0 [−16.8, −1]	−4.0 [−11.9, 4]	+2.9 [−2, 7.9]
HB	P→P	−32.7 [−42.6, −23.8]	+1.9 [−2, 5.9]	+1.9 [−2, 5.9]

Table 11: **Site A evaluations on Q by source** (H/N rate, %). No-intervention Q rates are $4.0/22.8$ for Llama and $18.0/44.6$ for Qwen, in JailbreakBench/HarmBench order.

Setting	Extraction → Q	L10	L15	L20
Llama / JBB	Q → Q (same-context)	51.0	8.0	4.0
Llama / JBB	Passive → Q (transfer)	31.0	9.0	4.0
Llama / HB	Q → Q (same-context)	54.5	22.8	21.8
Llama / HB	Passive → Q (transfer)	45.5	25.7	23.8
Qwen / JBB	Q → Q (same-context)	15.0	16.0	19.0
Qwen / JBB	Passive → Q (transfer)	18.0	11.0	16.0
Qwen / HB	Q → Q (same-context)	40.6	39.6	44.6
Qwen / HB	Passive → Q (transfer)	49.5	35.6	42.6

Table 12: **Llama Site A transfer controls under H/N**, layer 10. Values are paired increments over the target-context baseline. Semantic results for these conditions are in Table 26.

Bench.	Transfer	Real	Rand.	Shuf.	Real 95% CI
JBB	P→Q	$+27.0$	$+1.0$	0.0	$[17.0, 37.0]$
JBB	Q→P	$+45.0$	$+2.0$	$+1.0$	$[33.0, 56.0]$
HB	P→Q	$+22.7$	$+1.0$	0.0	$[13.9, 31.7]$
HB	Q→P	$+49.5$	$+5.0$	$+4.0$	$[39.6, 59.4]$

Table 13: **Qwen H/N site comparison**, layer 15. Recipe denotes Site A and Native denotes Site B in the original experiment records. Main-text conclusions use the semantic effects in Figure 2.

Protocol	Cell	L	Cond.	JBB	HB
Recipe (12,−5)	Q→Q	15	Real	-2.0 $[-6, 2]$	-5.0 $[-9.9, -1]$
Recipe (12,−5)	P→Q	15	Real	-7.0 $[-13, -2]$	-9.0 $[-15.8, -2]$
Recipe (12,−5)	Q→P	15	Real	-1.0 $[-3, 0]$	0.0 $[0, 0]$
Native (15,−1)	Q→Q	15	Real	$\mathbf{+63.0}$ $[54, 72]$	$\mathbf{+41.5}$ $[30.7, 52.5]$
Native (15,−1)	Q→Q	15	Random	-1.0 $[-3, 0]$	0.0 $[0, 0]$
Native (15,−1)	Q→Q	15	Shuffle	0.0 $[0, 0]$	-2.0 $[-5, 0]$
Native (15,−1)	P→Q	15	Real	$+6.0$ $[-1, 13]$	-4.0 $[-11.9, 4]$
Native (15,−1)	Q→P	15	Real	$+17.0$ $[10, 25]$	$+6.9$ $[2.0, 11.9]$

E ADDITIONAL WRAPPER EVALUATIONS UNDER H/N

The Llama extension evaluates frozen Site A Q/P vectors on Academic and S, and Academic/S vectors extracted at the same site on Q, Passive, and each other. Ablation remains at layer 10. Table 14 reports all these H/N measurements; the semantic Q/P-to-Academic/S subset is separately reported in Table 8. Random and coordinate-permutation controls were run for the Q/P-to-Academic/S conditions (Table 15), but not for the new Academic/S-source conditions.

At the same site and intervention layer, the H/N same-context increments are $+25.0$ $[17.0, 34.0]$ and $+37.6$ $[28.7, 47.5]$ for Academic, and $+60.0$ $[50.0, 70.0]$ and $+46.5$ $[35.6, 56.4]$ for S (JailbreakBench/HarmBench). These are heuristic source effects, not semantic source validation.

Qwen's extension freezes the Site B Q/P vectors and evaluates Academic/S at layer 15 (Table 16). Academic/S are target contexts only in that experiment. No random/permutation controls or semantic rescore accompany those additional target conditions. The large Q effects and target-dependent Passive effects are consequently reported as H/N observations.

Table 14: **Llama wrapper extension under H/N**, Site A, layer 10. Top: Q/P sources on Academic/S. Middle: Academic/S sources on Q/P. Bottom: Academic/S evaluated on each other. Intervals are paired against each target baseline.

Eval	Src.	JBB No-Int.	JBB Δ	HB No-Int.	HB Δ
Academic	Q	0.0	+26 [18, 35]	5.0	+38.6 [28.7, 47.5]
Academic	P	0.0	+24 [16, 33]	5.0	+34.6 [25.7, 43.6]
S	Q	6.0	+35 [26, 45]	23.8	+26.7 [17.8, 36.6]
S	P	6.0	+17 [10, 25]	23.8	+16.8 [9.9, 24.8]
Q	Academic	4.0	+31 [22, 41]	22.8	+21.8 [12.9, 31.7]
Q	S	4.0	+68 [59, 77]	22.8	+52.5 [41.6, 63.4]
P	Academic	35.0	+36 [22, 48]	38.6	+37.6 [26.7, 48.5]
P	S	35.0	+60 [50, 70]	38.6	+58.4 [49.5, 68.3]
Academic	S	0.0	+60 [50, 70]	5.0	+65.3 [55.4, 74.3]
S	Academic	6.0	+23 [15, 32]	23.8	+23.8 [15.8, 32.7]

Table 15: **H/N controls for Llama Q/P sources on Academic/S**, Site A, layer 10. Entries after No-Int. are increments over that target baseline.

Bench.	Eval	Source	No-Int.	Real	Rand.	Shuf.
JBB	Academic	Q	0.0	+26.0	0.0	0.0
JBB	Academic	Passive	0.0	+24.0	0.0	0.0
JBB	S	Q	6.0	+35.0	+1.0	0.0
JBB	S	Passive	6.0	+17.0	0.0	0.0
HB	Academic	Q	5.0	+38.6	+1.9	0.0
HB	Academic	Passive	5.0	+34.6	−1.0	0.0
HB	S	Q	23.8	+26.7	−1.0	+1.0
HB	S	Passive	23.8	+16.8	+1.0	+1.0

Table 16: **Qwen Site B Q/P sources on Academic/S targets under H/N**, layer 15. These are heuristic effects, with paired intervals against the target baseline.

Eval	Src.	JBB No-Int.	JBB Δ	HB No-Int.	HB Δ
Academic	Q	30.0	+70 [61, 79]	61.4	+38.6 [29.7, 48.5]
Academic	P	30.0	+24 [15, 32]	61.4	+9.9 [3.0, 16.8]
S	Q	19.0	+77 [68, 85]	60.4	+36.6 [26.7, 47.5]
S	P	19.0	+3 [−2, 9]	60.4	−8.9 [−14.9, −3.0]

F UNWRAPPED EXTRACTION REFERENCE

The unwrapped reference uses the mean difference from the same extraction pools with template {orig}. Extraction is at Site A for Llama and Site B for Qwen; intervention is at layers 10 and 15, respectively. This comparison uses fixed sites and single-layer ablation, not the full direction-selection and all-layer ablation procedure of Arditi et al. (2024).

Table 17 reports H/N rates and paired increments. Llama's no-intervention rates match its Site A transfer block. Qwen's unwrapped comparison uses a new no-intervention generation draw, so its absolute rates should be compared within that block. The unwrapped reference is closely aligned with each same-site Q direction (Appendix I).

Table 17: **Unwrapped mean-difference extraction under H/N.** No-Int. and Real are heuristic rates; Δ is their paired difference in percentage points.

Model	Cell	No-Int.	Real	Δ	95% CI
Llama L10	JBB U→Q	4.0	67.0	+63.0	[54.0, 72.0]
Llama L10	JBB U→P	35.0	94.0	+59.0	[49.0, 68.0]
Llama L10	HB U→Q	22.8	62.4	+39.6	[30.7, 49.5]
Llama L10	HB U→P	38.6	95.0	+56.4	[47.5, 66.4]
Qwen L15	JBB U→Q	19.0	86.0	+67.0	[58.0, 76.0]
Qwen L15	JBB U→P	88.0	100.0	+12.0	[6.0, 18.0]
Qwen L15	HB U→Q	45.5	86.1	+40.6	[29.7, 52.5]
Qwen L15	HB U→P	94.1	100.0	+5.9	[2.0, 10.9]

G DIRECTIONS SELECTED USING SOURCE-CONTEXT VALIDATION

This comparison fixes Qwen, the Q extraction wrapper, the Passive target wrapper, and intervention layer 15. Selection evaluates six extraction sites: $(15, -1)$, $(12, -5)$, $(12, -1)$, $(15, -5)$, $(19, -4)$, and $(11, -2)$. It uses 29 deduplicated harmful-validation requests evaluated only in Q, with H/N heuristic labels and a full-denominator no-intervention rate of 20.7%. The candidate H/N implementation checks a fixed list of refusal openings after cleaning; it does not perform semantic selection.

For each candidate, the code computes the paired increment over the Q baseline and a percentile interval from 1,000 paired bootstrap draws with seed 42. It rounds the increment and interval bounds to one decimal before requiring both $\Delta > 0$ and the lower 95% bound > 0. Eligible candidates are sorted by decreasing rounded increment, and the first two are selected: $C_1 = (19, -4)$ and $C_2 = (15, -1)$, with validation increments 69.0 and 65.5 points. The supplement records all six scores and eligibility decisions.

Selection uses only Questioning validation responses, without Passive test labels. This source-only procedure applies to the present comparison; the records do not establish target-blind historical discovery of Site B. Semantic evaluation follows on the fixed saved test responses using the supplied GPT-4o-mini prompt at temperature zero. Table 18 reports semantic effects and paired direction differences; Table 19 retains H/N measurements of this response block.

Both directions have positive H/N Q→Q increments on both benchmarks. On JailbreakBench, both also increase the Passive score by 12.0 $[6.0, 18.0]$ points, providing a positive cross-context result after source-only selection. Their Passive responses reach 100% on JailbreakBench and $99\%/100\%$ on HarmBench. The JailbreakBench labels agree on every evaluated request, yielding a paired direction difference of 0.0 $[0.0, 0.0]$; the HarmBench difference is -1.0 $[-3.0, 0.0]$ points. These are observed differences in heuristic labels, rather than an equivalence test for the underlying transfer effects.

The semantic rescore establishes positive source and target effects for both candidates on both benchmarks. The $(15, -1)$ candidate has larger source effects in the paired comparisons, whereas the target-difference intervals include zero (Table 18). The latter result does not establish equivalence of the candidates' target effects.

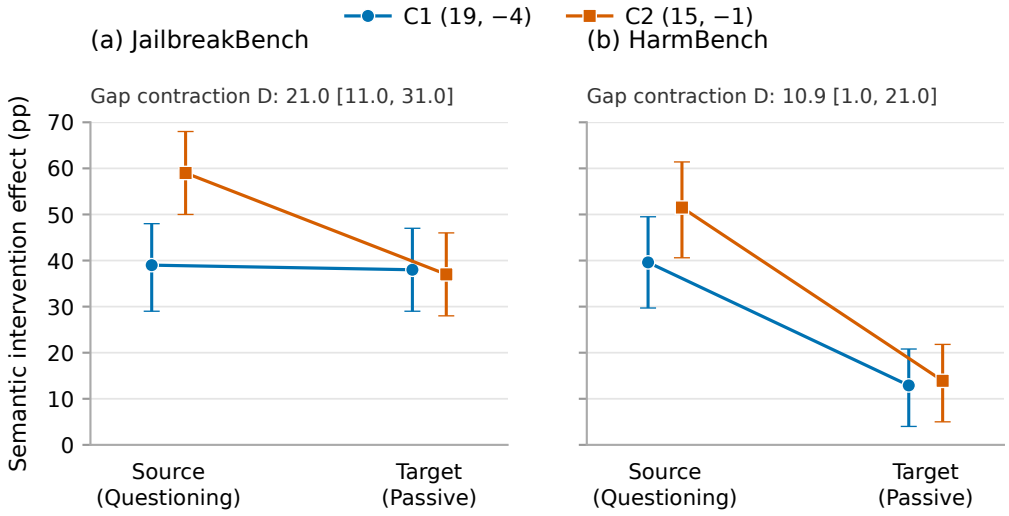


Figure 4: **A stronger source effect need not retain the same advantage on Passive.** Two source-selected Qwen directions, $C_1 = (19, -4)$ and $C_2 = (15, -1)$, are evaluated at ablation layer 15. Points and error bars show semantic increments over each context's no-intervention baseline and their original 95% paired intervals; lines connect the two measurements of each frozen direction. The annotated contrast D is the source candidate gap minus the target candidate gap, with 95% intervals from joint resampling (Appendix G).

Table 18: **Source-selected Qwen directions under semantic scoring.** Ablation is at layer 15. Values are paired harmful-compliance increments in percentage points with 95% intervals. Bottom rows directly compare $(19, -4)$ against $(15, -1)$ within each evaluation context.

Extraction	Test context	JailbreakBench Δ	HarmBench Δ
$(19, -4)$	Q→Q	$+39.0$ $[29.0, 48.0]$	$+39.6$ $[29.7, 49.5]$
$(15, -1)$	Q→Q	$+59.0$ $[50.0, 68.0]$	$+51.5$ $[40.6, 61.4]$
$(19, -4)$	Q→P	$+38.0$ $[29.0, 47.0]$	$+12.9$ $[4.0, 20.8]$
$(15, -1)$	Q→P	$+37.0$ $[28.0, 46.0]$	$+13.9$ $[5.0, 21.8]$
$(19, -4) - (15, -1)$	Q→Q	-20.0 $[-30.0, -11.0]$	-11.9 $[-20.8, -2.0]$
$(19, -4) - (15, -1)$	Q→P	$+1.0$ $[0.0, 3.0]$	-1.0 $[-4.0, 2.0]$

The $(15, -1)$ candidate is the saved Site B Q direction with SHA prefix `2a0103f4`. Its reported semantic source and target estimates and intervals coincide with the core Site B results. It serves as a common reference between the two comparisons, rather than an additional independent replication.

Target rates and matched requests. The saved candidate labels give 62 harmful-compliance responses out of 101 under the HarmBench Passive baseline, compared with 75 and 76 under the $(19, -4)$ and $(15, -1)$ interventions. Their full-denominator target rates are 74.3% and 75.2%. On JailbreakBench, the corresponding counts are 60/100 at baseline and 98/100 and 97/100 under intervention, with no collapse in these three cells. The two benchmarks thus differ in their proximity to the numerical ceiling.

On HarmBench, 29 records have zero harmful-compliance labels under both the Q and Passive baselines. Within this matched subset, the $(19, -4)$ and $(15, -1)$ directions yield 13 and 17 source successes, respectively, and the same 16 target successes. At source, seven records succeed only under $(15, -1)$ and three only under $(19, -4)$; at target, the two binary label vectors agree throughout this subset. These descriptive counts show the source and target profiles on a common set of baseline-zero records, complementing the full-sample effects in Table 18.

Change in the candidate advantage. For the fixed candidate pair, we directly compare the source and target gaps,

$$D = (E_{C_2} - E_{C_1}) - (T_{C_2} - T_{C_1}). \tag{7}$$

The common no-intervention baseline within each context cancels in its candidate difference. We jointly resample all candidate labels for a request, using 10,000 bootstrap draws with seed 42 and the 2.5th and 97.5th percentiles for the interval. JailbreakBench uses 100 singleton clusters; HarmBench groups the known shared-Goal pair, keeping the sampled record count as each replicate's denominator. The resulting contrasts are 21.0 $[11.0, 31.0]$ and 10.9 $[1.0, 21.0]$ points, respectively. Excluding HarmBench records 3 and 18 while retaining that pair gives 11.1 $[1.0, 22.0]$ over 99 records in 98 clusters. These are post hoc contrasts for the saved candidates and responses; they do not replace the original effect intervals. Baseline cancellation does not remove ceiling or task effects on the observed candidate gaps.

Sensitivity to recorded collapse. The four candidate conditions contain no collapsed responses on JailbreakBench and eight on HarmBench: three for C_1 at source, four for C_2 at source, and one for C_1 at target. Because the flag can include scoring failures, we additionally treat all eight labels as unknown binary outcomes while fixing the other labels. Across all completions, HarmBench's D ranges from 7.9 to 15.8 points. Its point estimate remains positive, but the most adverse completion has interval $[-3.0, 19.0]$ under the same cluster bootstrap. Exclusion of zero therefore depends on how these records are handled. This post hoc bound includes genuine response degeneration among the unknowns; it does not assess errors in non-collapsed labels. The supplement provides the bound-reproduction script.

Paired uncertainty. The candidate intervals use 1,000 paired bootstrap resamples of evaluation records with seed 42 and the 2.5th and 97.5th percentiles. Collapse receives label zero and remains in the denominator. To assess the known shared-Goal pair in HarmBench, we also group records 15 and 16 into one sampling cluster and draw 10,000 bootstrap samples from the resulting 100 clusters, retaining all 101 records in the point estimates. All four candidate source and target intervals remain positive. The source difference, $(15, -1) - (19, -4)$, is 11.9 $[2.0, 21.8]$ points, and the target difference is 1.0 $[-2.0, 4.0]$. The two records are outside the common baseline-zero subset, whose counts are unchanged.

Excluding contextual records 3 and 18 from all candidate conditions, while retaining the shared-Goal pair as one cluster, leaves 99 records in 98 clusters. Under the same 10,000-draw procedure, all four source and target intervals remain positive; the source difference is 12.1 $[2.0, 22.2]$ points and the target difference is 1.0 $[-2.0, 5.0]$. This post hoc sensitivity concerns the candidate comparison in Table 18.

Table 19: **Source-selected Qwen directions under H/N**, with ablation at layer 15. Entries are increments over this block's no-intervention baseline. Final rows directly compare the two directions on Passive.

Extraction	Test context	JailbreakBench Δ	HarmBench Δ
$(19, -4)$	Q→Q	$+51.0$ $[40.0, 61.0]$	$+30.7$ $[19.8, 40.6]$
$(15, -1)$	Q→Q	$+63.0$ $[54.0, 72.0]$	$+42.6$ $[31.7, 54.5]$
$(19, -4)$	Q→P	$+12.0$ $[6.0, 18.0]$	$+5.0$ $[0.0, 9.9]$
$(15, -1)$	Q→P	$+12.0$ $[6.0, 18.0]$	$+5.9$ $[2.0, 10.9]$
$(19, -4) - (15, -1)$	Q→P	0.0 $[0.0, 0.0]$	-1.0 $[-3.0, 0.0]$

The H/N no-intervention rates in this block are Q 19.0% and P 88.0% on JailbreakBench, and Q 45.5% and P 94.1% on HarmBench. The paired difference and its interval use the item-level labels; they need not equal differences computed from rounded table entries.

H SOURCE EFFECTS ACROSS A FOUR-SITE EXTRACTION GRID

This comparison extracts Q directions at the four combinations of layers 12 and 15 with token positions -1 and -5, keeping Qwen's ablation layer fixed at 15. Table 20 reports Q→Q effects

under H/N. The no-intervention Q rates in this block are 19.0% on JailbreakBench and 45.5% on HarmBench. The $(15, -1)$ vector is the saved Site B Q direction.

Table 20: **Qwen source effects across extraction layers and token positions.** H/N Q→Q increments in percentage points over this block's no-intervention baseline, with 95% paired intervals; ablation layer 15 throughout.

Extraction (ℓ, t)	JailbreakBench Δ	HarmBench Δ
$(12, -1)$	$+6.0$ $[0.0, 13.0]$	$+1.0$ $[-5.0, 6.9]$
$(12, -5)$	-2.0 $[-6.0, 1.0]$	-5.0 $[-9.9, -1.0]$
$(15, -1)$	$+63.0$ $[54.0, 72.0]$	$+42.6$ $[31.7, 54.5]$
$(15, -5)$	$+4.0$ $[-2.0, 10.0]$	$+5.9$ $[0.0, 11.9]$

Only $(15, -1)$ has a positive source-effect interval that excludes zero on both benchmarks. At the other three sites, positive source effects are not established; the $(12, -5)$ effect on HarmBench is negative. Within this grid, changing only the extraction layer or only the token position from $(12, -5)$ does not reproduce the large $(15, -1)$ effect. This is a local comparison of H/N source responses at a fixed intervention layer. The source-selected $(19, -4)$ candidate in Table 18 lies outside this grid.

I Direction geometry and intervention magnitude

Table 21 gives cosine similarities and pre-normalization norms of the saved mean-difference vectors. U denotes unwrapped extraction. The random reference is a same-dimension Gaussian with seed 2234; $1/\sqrt{d}$ gives the characteristic random-cosine scale. All vectors are normalized before ablation, so the original vector norm does not set the deletion magnitude.

The deletion magnitude is instead $|h_k^\top v|$. Table 22 summarizes its mean and maximum during generation. The Qwen Site A Q direction produces large projections without a positive semantic source effect. These measurements characterize the interventions; neither the cosine nor projection comparisons isolate a causal explanation for the observed transfer profile.

Table 21: **Geometry of saved directions.** Site A is $(12, -5)$ and Site B is $(15, -1)$. The final column contains norms before unit normalization.

Pair	cos	cos **vs. rand**	$1/\sqrt{d}$	$\|\tilde{v}\|_2$
Llama (12, −5) Q vs. P	0.637	0.012 / −0.015	0.016	4.35 / 2.31
Llama (12, −5) U vs. Q	0.937	—	0.016	4.20 / 4.35
Llama (12, −5) U vs. P	0.575	—	0.016	4.20 / 2.31
Qwen A (12, −5) Q vs. P	0.136	−0.021 / −0.016	0.017	9.42 / 4.64
Qwen B (15, −1) Q vs. P	0.621	−0.009 / 0.008	0.017	20.65 / 9.27
Qwen B U vs. Q	0.963	—	0.017	21.41 / 20.65
Qwen B U vs. P	0.613	—	0.017	21.41 / 9.27
Qwen A Q vs. native Q	0.133	—	0.017	—
Qwen A P vs. native P	0.097	—	0.017	—
Qwen U vs. recipe Q	0.135	—	0.017	—

Table 22: **Deleted projection magnitudes** during generation. Every intervention uses a unit vector. The historical labels native and recipe denote Sites B and A, respectively; neither denotes a prescribed extraction site.

Model	Protocol	Cell	Cond.	mean	max
Llama	(12,−5)	JBB Q→P L10	Real	0.48	2.89
Llama	(12,−5)	JBB Q→P L10	Random	0.12	15.06
Llama	(12,−5)	JBB Q→P L10	Shuffle	0.07	1.54
Llama	(12,−5)	JBB P→Q L10	Real	0.37	3.95
Llama	(12,−5)	JBB U→Q L10	Real	0.51	3.63
Llama	(12,−5)	JBB U→P L10	Real	0.67	3.31
Llama	cap. 800 L10 Q	mixed eval	Real	0.21	1.71
Qwen	native (15,−1)	JBB Q→Q L15	Real	4.83	314.5
Qwen	native (15,−1)	JBB Q→Q L15	Random	2.06	267.8
Qwen	native (15,−1)	JBB Q→Q L15	Shuffle	1.82	27.2
Qwen	recipe (12,−5)	JBB Q→Q L15	Real	9.34	1618
Qwen	native (15,−1)	JBB U→Q L15	Real	5.08	400
Qwen	native (15,−1)	JBB U→P L15	Real	4.64	400
Qwen	native cap. 800 L15 Q	mixed eval	Real	4.54	314.5

J CAPABILITY UNDER Q-DIRECTION ABLATION

The fixed evaluation contains 100 GSM8K, 100 MMLU, 52 ARC-Challenge and 48 HellaSwag items, and 500 harmless test prompts. It does not sample the harmless extraction pool. Accuracy is measured on the 300 answerable items, while the harmless subset measures refusal frequency. Both metrics are computed from generated response text. Llama uses Site A Q ablation at layer 10; Qwen uses Site B Q ablation at layer 15.

Llama's answerable aggregate changes from 221/300 to 223/300; Qwen's changes from 254/300 to 255/300. The harmless refusal counts are unchanged. The table reports the task-level counts as well as aggregate percentages, so these small changes remain visible.

Table 23: **Task accuracy and harmless refusal under Q ablation.** Task rows show correct-answer counts out of 100; the final row shows refusals out of 500. The aggregate row reports percentages over 300 answerable items.

Slice	n	Llama No-Int.	Llama L10	Qwen No-Int.	Qwen L15
Answerable (all)	300	73.7%	74.3%	84.7%	85.0%
GSM8K	100	93	93	91	90
MMLU	100	68	68	80	82
ARC+HellaSwag	100	60	62	83	83
harmless_test refuse	500	0	0	3	3

K FIXED DIRECTIONS ACROSS OPERATORS AND CONTEXTS

Model and prompt rendering. The panel uses `Qwen/Qwen2.5-7B-Instruct`, snapshot `a09a35458c702b33eeacc393d103063234e8bc28`, in bf16. Prompts use the official chat template with the explicit system message "You are a helpful assistant.", the context-wrapped request as the user message, and `add_generation_prompt=True`. Padding is on the left, with the padding token set to EOS. Generation is greedy, with at most 256 new tokens, `repetition_penalty=1.0`, and the snapshot's fixed EOS configuration.

Extraction. The starting pools contain 260 harmful and 18,793 harmless training requests. Before sampling, instructions are normalized by NFKC normalization, whitespace collapse, stripping, and case folding. Duplicate instructions and normalized exact overlaps with the fixed validation and JailbreakBench evaluation sets are removed. This removes one harmful-training/JailbreakBench

overlap, leaving 259 harmful requests. Extraction uses 128 harmful requests from these 259 and 128 harmless requests from the cleaned harmless pool, with fixed sampling seeds 42/43.

Both vectors use Q-source activations and Equation 1, with float64 mean accumulation and saved float32 unit vectors. Coordinates $A_{\text{new}} = (12, -5)$ and $B_{\text{new}} = (15, -1)$ refer to zero-based decoder-block inputs; negative token positions are relative to the complete rendered chat prompt. The coordinates are fixed independently of outcome labels.

Operators and conditions. Single15 applies Equation 2 only at block 15 input, at every prompt and generated-token position. The official all-layer placement of Arditi et al. (2024) applies the same projection removal at every decoder-block input, every attention output, and every MLP output across all 28 layers, throughout prefill and decoding. Thus, all-layer placement includes more than the 28 block-input hooks. Both operators use the same saved unit vector for a given extraction coordinate.

The frozen 100 JailbreakBench requests each have no-intervention, A_{new} Single15, and B_{new} Single15 conditions in both Q and Passive, specifying 600 responses. Q additionally has all-layer conditions for each vector, specifying 200 responses. The standalone panel therefore specifies 800 generated responses; it contains no all-layer Passive condition. Each vector remains the Q-source vector in every condition. The Q comparison varies the operator while fixing the vector and prompts; the Q–Passive comparison varies context while fixing the vector and Single15 operator.

Diagnostic estimand and uncertainty. Let $I_{\text{H/N}}(y)$ denote the panel's absence-of-opening-refusal indicator, using the exact heuristic implementation specified in the supplement. For operator o and context d, define

$$\delta_{\text{H/N}}(v, o; d) = \frac{100}{N} \sum_{i=1}^{N} \left[I_{\text{H/N}}(y_i^{v,o,d}) - I_{\text{H/N}}(y_i^{0,d}) \right].$$

Here $N = 100$ and collapsed responses receive indicator zero. The frozen expression matches a small set of refusal openings after cleaning; other readable responses receive indicator one, including benign continuations without those openings. These are percentage-point changes in that surface criterion, distinct from semantic E and T.

Dataset-indexed records support matched comparisons. The 95% paired bootstrap (10,000 draws, seed 42) resamples the same 100 request indices jointly across conditions, retaining all responses for each sampled request. Its uncertainty is conditional on the newly extracted vectors and deterministic generations; it does not measure robustness to extraction seeds. Complete vector files and hashes, rendered prompts and input-token IDs, and full generated responses with stopping metadata are retained.

Results. Table 24 reports all eight conditions, and Figure 5 shows the absolute rates alongside the no-intervention baselines. In Q, the single-layer increments are $+19$ $[10, 28]$ points for A_{new} and $+72$ $[63, 81]$ for B_{new}. Applying each same saved vector at all layers raises these increments to $+57$ $[46, 68]$ and $+79$ $[71, 87]$. The paired all-layer-minus-single-layer changes are $+38$ $[27, 49]$ and $+7$ $[3, 12]$ points. The operator has a particularly large effect for A_{new}.

Under Single15, the candidate gap is 53 points in Q and six on Passive. Both Passive rates approach the numerical ceiling (94% and 100%), against a 69% baseline. This comparison makes the target headroom visible, alongside the operator dependence observed at source. The new vectors share extraction coordinates with Sites A/B, but are separate estimates; they do not reproduce the original tensors or establish semantic harmful compliance. In particular, A_{new} is source-effective under H/N even with Single15.

No response triggers the deterministic collapse rule. Length limits are frequent, with 40–91 replies per condition reaching 256 tokens; the diagnostic classifies the generated text within that budget. Semantic labels are not available for this 800-response panel, so neither these positive openings nor the absence of collapse establishes harmful fulfillment. The six-site fulfillment analysis uses a separate extraction (Appendix L).

Table 24: **Fixed directions across contexts and operators in the new Qwen run.** Each row contains the same 100 JailbreakBench requests. $R_{\mathrm{H/N}}$ is the opening-heuristic positive rate; $\delta_{\mathrm{H/N}}$ is its paired change over the corresponding context baseline, with 95% bootstrap intervals. C counts deterministic collapse flags; L counts replies that reach the 256-token limit. Neither positive openings nor the absence of a collapse flag establish harmful fulfillment.

Context	Vector	Operator	$R_{\mathrm{H/N}}$ (%)	$\delta_{\mathrm{H/N}}$ (pp)	C	L
Q	–	None	21.0	–	0	40
Q	A_{new}	Single15	40.0	+19.0 [10.0, 28.0]	0	53
Q	B_{new}	Single15	93.0	+72.0 [63.0, 81.0]	0	84
Q	A_{new}	All layers	78.0	+57.0 [46.0, 68.0]	0	72
Q	B_{new}	All layers	100.0	+79.0 [71.0, 87.0]	0	87
Passive	–	None	69.0	–	0	84
Passive	A_{new}	Single15	94.0	+25.0 [17.0, 34.0]	0	91
Passive	B_{new}	Single15	100.0	+31.0 [22.0, 40.0]	0	58

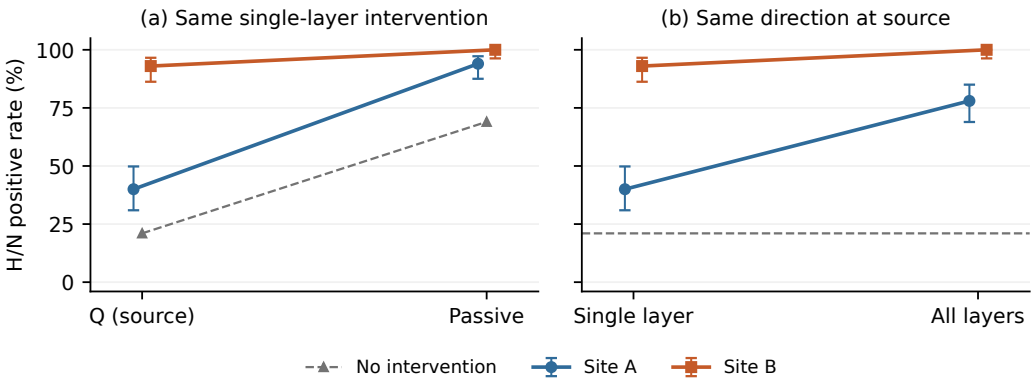


Figure 5: **Operator and context change the measured effect of fixed directions.** H/N denotes the unchanged opening-refusal heuristic. Panel (a) fixes Single15 and changes context; panel (b) fixes Q and changes the ablation operator. Site A/B denote the new vectors, not the original tensors. Error bars are 95% Wilson intervals for individual rates; Table 24 and the text report paired bootstrap intervals for changes.

L INDEPENDENT SEMANTIC SELECTION AND NEUTRAL EVALUATION

This analysis uses a six-site Qwen extraction that is distinct from the original Site A/B tensors, from C_1/C_2, and from the 800-response H/N panel in Appendix K. Coordinates may coincide; the tensors do not.

Generation. Generation uses snapshot `a09a35458c702b33eeacc393d103063234e8bc28`, official chat template, system message "You are a helpful assistant.", greedy decoding with 256 new tokens, and Single15 ablation of a unit vector throughout prefill and decode. Six Q-source sites are extracted from 128 harmful and 128 harmless training prompts (seeds 42/43) after normalized exact overlap with validation and test is removed. Source validation generates 39 harmful requests $\times$ {no intervention and six sites} in Q only (273 responses). After selection is frozen, the test matrix is 100 JailbreakBench requests $\times$ {no intervention, C_{best}, C_{lower}} $\times$ {Q, Passive, Neutral} (900 responses). Test generation uses the frozen selected vectors without re-extraction.

Scoring. The judge is GPT-4o-mini under the fulfillment rubric `fulfillment_v1`: the evaluator sees an opaque ID, the base request, and the full saved response, not the wrapper name, candidate identity, or intended D_N. `FULFILLS` $= 1$, `DOES_NOT_FULFILL` $= 0$; `UNCERTAIN`, provider content-filter failures, and parse failures remain unknown; a provider filter is not a refusal by the evaluated model. Collapse is a separate response-status field and does not recode binary labels to zero. Four exact duplicate base-request-plus-cleaned-response pairs (three NoInt Q/Neutral pairs and one C_{lower} Q/Neutral pair) were scored once per pair and mapped back to all original conditions. This panel uses a single GPT-4o-mini judge without independent dual-evaluator validation. Source validation used 271 automatic labels plus two student labels on content-filtered items (requests that would not change eligibility if flipped). Test scoring left 12 content-filtered unique texts unknown; 0/1 imputation does not change whether D_N's interval excludes zero, whether both test-Q effects remain positive with lower endpoints above zero, or the sign of the Neutral gap. A 30-request subset has the two human-review rounds reported in Section 5.4 and Appendix M. Both human-review rounds used the same annotator (Appendix M).

Selection. A site is eligible when its paired Q increment over no intervention is positive and the 95% interval from 10,000 paired bootstrap draws (seed 42) has lower endpoint strictly above zero. Among eligible sites, C_{best} is the largest increment and C_{lower} the smallest, with a fixed site-order tie break. Table 25 reports the six increments. All source-validation labels were resolved before selection, including the two student annotations.

Table 25: **Independent Q-only source screening** ($n = 39$). Increments in percentage points.

Site	(ℓ, t)	Estimate	95% CI	Eligible
C_{best}	$(15, -1)$	$+76.9$	$[64.1, 89.7]$	yes
	$(15, -5)$	$+64.1$	$[48.7, 79.5]$	yes
	$(19, -4)$	$+61.5$	$[46.2, 76.9]$	yes
C_{lower}	$(12, -5)$	$+23.1$	$[10.3, 35.9]$	yes
	$(12, -1)$	$+5.1$	$[0.0, 12.8]$	no
	$(11, -2)$	$+2.6$	$[0.0, 7.7]$	no

Test estimands. Rates, effects E and T, candidate gaps, and

$$D_d = \text{mean}\big[(Y_{\text{best},Q} - Y_{\text{lower},Q}) - (Y_{\text{best},d} - Y_{\text{lower},d})\big] \times 100$$

for $d \in \{\text{Neutral}, \text{Passive}\}$ use the same paired bootstrap. Each estimate uses requests resolved in all cells it requires; rates and paired contrasts can therefore have different denominators. Complete-case intervals are conditional on those resolved requests. Only the sensitivity bounds replace unknowns with 0 or 1 according to the coefficient sign; primary labels remain unknown. Table 3 in the main text reports complete-case rates. On Neutral, both candidate rates remain $\leq 85\%$; this is a descriptive headroom marker. As a supplementary sensitivity analysis, the subset where Q and Neutral no-intervention labels are both zero has $n = 93$ and $D_N = +1.1$ $[-7.9, 10.1]$ on 89 complete requests. Judge collapse counts on scored test rows are zero. Length-limit counts under C_{best} are 82/100 in Q and 84/100 in Neutral.

This independent fulfillment analysis is distinct from the original semantic scores and the 800-response H/N diagnostic.

M HUMAN REVIEW AND SUPPLEMENTARY ACADEMIC EVIDENCE

The independent fulfillment rubric evaluates the base request and saved reply, allowing `FULFILLS`, `DOES_NOT_FULFILL`, and `UNCERTAIN`; it is distinct from the original `JAILBROKEN`/`REFUSED` parser. The same annotator reviewed round 1 (270 replies) and the Neutral add-on (90 replies). Blinding concealed wrapper, extraction site, and automatic label, but Neutral review followed exposure to round-1 summaries. It is therefore an add-on, not an independently naive replication.

Round-1 agreement excludes four automatic nulls and three human uncertain labels, leaving $205/263 = 77.9\%$ ($\kappa = 0.562$). Agreement is about 92% in Q, 73% in Passive, and 69% in Academic; Academic C_{best} is $16/30 = 53.3\%$. It is $46/50 = 92.0\%$ on untruncated replies and $159/213 = 74.6\%$ on truncated ones. Of 58 binary disagreements, the human was more generous in 41, all recorded with reason "Gave detailed background," and stricter in 17. These observations describe disagreements; they do not authorize relabeling background as fulfillment or recalibrating the judge. Original labels remain unchanged. The three automatic substitutions enter only explicitly mixed, follow-auto estimates. The strict Q gap is $+41.4$ $[24.1, 58.6]$ on 29 requests, versus mixed $+43.3$ $[26.7, 60.0]$ on 30; same-30 automatic Q gap is $+51.7$ $[34.5, 69.0]$ on 29. Human D_N uses the strict Q side with the Neutral add-on, retaining 29 complete requests. The Neutral add-on has no human uncertain labels and one automatic null: agreement is $84/89 = 94.4\%$ ($\kappa = 0.866$), $43/43$ untruncated and $41/46$ truncated.

Both candidates also have positive mixed human Q increments over NoInt: $+63.3$ $[46.7, 80.0]$ for C_{best} and $+20.0$ $[6.7, 36.7]$ for C_{lower}. Corresponding same-30 automatic increments are $+69.0$ $[51.7, 86.2]$ and $+16.7$ $[3.3, 30.0]$. Table 4 retains the full-sample, same-request automatic, and human contrasts separately; neither rounded marginal rates nor contrasts with differing complete-case sets should be subtracted to reconstruct D.

Academic uses the same selected pair and fulfillment criterion. Its eight provider refusals remain null in primary automatic estimates. The all-truncated-as-fulfilling and all-truncated-as-unknown analyses in Section 5.3 are sensitivity scenarios, not label edits or estimates of completed answers.

N COMPLEMENTARY LLAMA RESULTS ACROSS SOURCE AND TARGET CONTEXTS

Llama shows positive source and target effects for both directions. At Site A and intervention layer 10, Q→Q raises semantic harmful compliance by $43.0/22.8$ points and Passive→Passive by $34.0/33.7$ points on JailbreakBench/HarmBench (Table 2). Freezing these vectors, Passive→Q gives $+26.0/+16.8$ points and Q→Passive gives $+55.0/+50.5$ (Table 26). All eight source and target intervals exclude zero, and the real-direction effects exceed the corresponding Gaussian and coordinate-permuted controls.

The same frozen vectors also affect Academic and Scenario prompts. The Questioning direction gives semantic increments of $24.0/21.8$ points on Academic and $36.0/26.7$ on Scenario; the Passive direction gives $19.0/15.8$ and $17.0/10.9$ points. All eight intervals exclude zero (Appendix B). Broader H/N experiments locate effects across additional source–target pairs and intervention depths (Appendices D and E).

N.1 RESPONSE SEMANTICS CHANGE THE APPARENT ROOM FOR TRANSFER

The response criterion changes both the magnitude and the apparent headroom of a transfer effect. In Qwen's Site B experiment on JailbreakBench, H/N marks 83% of Passive no-intervention responses as positive, against 60% under semantic scoring among non-collapsed responses. On the same saved generations, Q→Passive adds 17 points under H/N and 37 under semantic scoring (paired increments over all records), so separating non-refusal openings from harmful assistance more than doubles the measured increment here.

Table 26: **Llama cross-context effects at Site A, layer 10.** Semantic increments in percentage points, relative to the target-context baseline. Perm. permutes vector coordinates. Confidence intervals apply to the real-direction increment.

Benchmark	Transfer	Real	Random	Perm.	Real 95% CI
JailbreakBench	P→Q	$+26.0$	$+1.0$	-1.0	$[17.0, 35.0]$
JailbreakBench	Q→P	$+55.0$	$+2.0$	$+2.0$	$[44.0, 65.0]$
HarmBench	P→Q	$+16.8$	0.0	0.0	$[9.9, 25.7]$
HarmBench	Q→P	$+50.5$	$+5.0$	$+3.0$	$[41.6, 60.4]$

The source result is less sensitive: Q→Q is $+63$ points under H/N and $+59$ under semantic scoring, and Llama Q↔Passive stays positive under both. The criterion thus preserves the positive findings while changing their size and the baseline against which it is read.

O Extended Related Work

Refusal vectors and false-refusal calibration. Directional accounts of refusal support several distinct evaluation questions. Arditi et al. (2024) construct candidate refusal directions from harmful and harmless activations and select among them using intervention-based criteria, showing that removing a selected direction can suppress refusal. Wang et al. (2025a) instead target false refusals on safe but superficially alarming requests. They derive a false-refusal direction from pseudo-harmful and harmless activations, orthogonalize it against a true-refusal direction, and ablate the resulting vector to improve compliance on safe requests while retaining refusal of harmful requests and general capability. This is a calibration objective: it separates desirable compliance from over-cautious refusal. Our source-efficacy and transfer estimands ask a complementary question on harmful requests. They first establish whether a frozen vector–operator pair increases harmful compliance in its extraction context and then measure its increment under each lexical wrapper, relative to that wrapper's own no-intervention baseline. The distinction matters because improved safe-request compliance and increased harmful-request compliance reward different behaviors even when both studies use difference-in-means directions and projection removal. It also sharpens comparisons to multi-category refusal analyses: Joad et al. (2026) study geometrically distinct refusal categories and their steering trade-offs, whereas our paired design holds the requested action fixed while changing its framing. Together, these lines motivate reporting the behavioral target, source validation, and matched baseline alongside the direction itself.

Transfer of activation interventions across contexts. Activation interventions are often studied for whether a compact representation retains its causal effect outside the observations used to construct it. Function vectors provide a general example: Todd et al. (2024) aggregate activations from task-relevant attention heads during in-context learning and show that inserting the resulting vector into an unrelated context can elicit the same task behavior. In safety alignment, Wang et al. (2025b) validate refusal directions and demonstrate cross-language effects with all-layer ablation, while contrastive activation addition changes generation through a steered residual representation (Rimsky et al., 2024). These results establish useful forms of portability, but the transferred object and comparison vary: a task operation can be inserted into a new prompt, a refusal direction can cross languages, or a steering vector can shift a behavioral attribute. Our design makes the unit of transfer explicit. For each comparison, both the extracted direction and its ablation operator are frozen, and only the evaluation wrapper changes. Source efficacy is therefore a measured property of that specific vector–operator configuration, while the target effect is estimated separately for Neutral, Passive, or other contexts. Linear erasure methods such as LEACE provide related projection operators with closed-form removal guarantees for linearly encoded concepts (Belrose et al., 2023); those geometric guarantees do not determine how much a fixed intervention will change a downstream behavioral score under a new wrapper. Separating the diagonal source measurement from each off-diagonal target measurement turns contextual portability into a set of paired empirical contrasts rather than a single global attribute of the vector.

Safety representations and behavioral outcomes. Several studies further show why the representation being manipulated should be distinguished from the behavior used to score an intervention.

Zhao et al. (2025) find that harmfulness and refusal are separately encoded, so a direction associated with refusal language need not by itself specify whether a response provides harmful assistance. Zhang et al. (2026) study how harmful and safe response trajectories become less distinguishable in hidden-state space as generation proceeds. Their DEEPALIGN defense uses contrastive hidden-state losses during fine-tuning to steer harmful trajectories toward safer ones while retaining utility. That analysis links jailbreak susceptibility to the evolution of safety representations during decoding. Our experiments address a different level of the same representation–behavior chain. We intervene with a saved direction and evaluate the resulting text under source and target wrappers, using semantic harmful compliance as the principal outcome and treating opening-refusal heuristics as a separate diagnostic. This pairing reveals changes that a surface refusal marker can obscure: wrapper-specific baseline compliance determines the available absolute increment, and two source-effective candidates can preserve or lose their gap in different target contexts. The contribution is thus a behavioral measurement framework for contextual changes in a fixed intervention's effect, complementary to work that decomposes internal safety concepts or traces their evolution through generation.