TIDE: TARGETED PROBE-GUIDED INTERVENTIONS FOR CROSS-MODEL KNOWLEDGE RECOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

Large language models (LLMs) are increasingly deployed for complex, multi-step tasks, making the reliability of their generated answers ever more important. Yet, even when an LLM internally computes the correct answer, it may fail to express that knowledge in its output, resulting in a knowledge-prediction gap in which the model internally "knows" more than it says. Bridging this gap is critical for improving the reliability of LLMs, particularly in high-stakes scenarios where incorrect or unfaithful responses can have severe consequences. However, existing approaches face two fundamental challenges. (1) Most are model-specific: they learn to extract answers from the representations of a single model, potentially relying on model-specific features rather than capturing the task-level information that generalizes across models. (2) Recovering latent knowledge with a probe does not directly improve the model's generation; an effective mechanism is needed to transfer probe predictions back into the model's output with minimal performance loss. To address these challenges, we propose TIDE (**T**argeted Probe-Guided **I**nterventions for Cross-Mo**de**l Knowledge Recovery), a framework that bridges model-internal knowledge and generated answers through cross-model probing and targeted intervention. TIDE trains a single probe across multiple models for a given task, enabling it to capture task-level information shared across models rather than model-specific representations and thereby generalize to unseen models. TIDE then identifies key attention heads involved in writing answers and selectively injects the probe's predictions into these heads, allowing the inferred answer to flow through the model's native output-writing mechanism. Extensive experiments across multiple LLMs demonstrate that TIDE consistently outperforms existing baselines, showing that a single cross-model probe can capture transferable information and effectively narrow the gap between internal and generated answers. Our code can be found at `https://anonymous.4open.science/r/TIDE-D338/`.

1 INTRODUCTION

Large language models (LLMs) have demonstrated remarkable capabilities across a wide range of tasks, including open-domain question answering (Yue, 2025; Kamalloo et al., 2023), commonsense reasoning (Teo et al., 2026; Ke et al., 2025), code generation (Dong et al., 2025; Jiang et al., 2026), and multi-step instruction following (Plaat et al., 2025; Zou et al., 2025). Despite these capabilities, their generated responses do not always faithfully reflect the information encoded in their internal representations. In particular, simple probes trained on intermediate activations can often predict the correct answer more accurately than the model's own output, revealing a *knowledge-prediction gap*: LLMs may internally encode information that they fail to express in their generations (Park et al., 2025; Zhao et al., 2025; Gekhman et al., 2025; Orgad et al., 2025). This gap poses a fundamental challenge to model reliability, as failures to express internally encoded knowledge can lead to incorrect or unfaithful responses even when the relevant information is available within the model. In high-stakes domains such as medicine, law, and finance, such failures can have serious real-world consequences (Hu et al., 2026; Pandit et al., 2025). Thus, enabling LLMs to faithfully express their internal knowledge is an important prerequisite for building more reliable systems.

Existing approaches to closing this gap primarily rely on inference-time interventions (Park et al., 2025; Chuang et al., 2024) or finetuning-based methods (Mahmoud et al., 2026; Kawasaki et al.,

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.

2026). However, these approaches face two fundamental limitations. First, most of these methods operate directly in model-specific activation spaces, requiring separate probes or interventions for each model. Although different LLMs learn distinct internal representations, the underlying task knowledge they encode is often semantically shared across models. Model-specific approaches may therefore capture idiosyncratic features of individual models rather than transferable, task-level information. Second, recovering internal knowledge does not necessarily improve the model's generation. While probes and other readout methods can successfully predict information encoded in intermediate activations, these predictions remain external to the model's generation process. Finetuning can incorporate such recovered knowledge into the model, but requires modifying model parameters and repeating the training process for each model and task.

To address these limitations, we propose **TIDE** (**T**argeted Probe-Guided **I**nterventions for Cross-Mo**de**l Knowledge Recovery), a framework that bridges internal knowledge and generated answers through cross-model probing and targeted intervention. TIDE first converts raw model activations into *model-agnostic relative representations* by measuring similarities to a small set of shared anchor samples (Moschella et al., 2022). Unlike model-specific activation spaces, this shared relative space enables a single probe to learn task-level information jointly across multiple LLMs and transfer this knowledge to unseen models without retraining. TIDE then identifies attention heads that most strongly influence the model's output, which we refer to as writer heads. Rather than targeting entire layers, TIDE selectively injects the probe's prediction through these heads, allowing the model to express the recovered knowledge through its own generation mechanism with minimal disruption.

Our contributions are summarized as follows: (1) **Cross-model Probing.** We introduce a model-agnostic probing approach based on relative representations, enabling a single probe to capture transferable task-level information across LLMs and generalize to unseen models. (2) **Targeted Knowledge Writing.** We identify writer heads that strongly influence model outputs and selectively intervene on these heads to inject recovered knowledge into the model's native generation pathway. (3) **Extensive Experiments.** We evaluate TIDE across multiple LLMs and demonstrate consistent improvements in answer generation over existing baselines, showing that transferable internal knowledge can be effectively recovered and expressed at inference.

2 Background and Preliminaries

In this section, we formalize the *knowledge-prediction gap* and then introduce *model-agnostic relative representations*. Together, these concepts form the foundation of our framework.

2.1 The Knowledge-Prediction Gap

Prior work has repeatedly observed that LLMs internally encode information that is not reflected in their generated outputs: probes trained on hidden activations can recover correct answers even when the model fails to produce them (Park et al., 2025; Zhao et al., 2025; Gekhman et al., 2025; Orgad et al., 2025). We refer to this phenomenon as the *knowledge-prediction gap*.

Let f_θ denote an LLM with parameter set θ and L layers, each comprising an attention and an MLP sublayer, for a total of $2L$ sublayers. For a sublayer $\ell \in \{1, \dots, 2L\}$, let $h^\ell(x) \in \mathbb{R}^d$ denote the model's activation at ℓ given input x, and let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ denote a labeled task dataset. For each ℓ, we train a lightweight probe g^ℓ (e.g., a single-layer MLP) on the frozen activations $\{h^\ell(x_i)\}$ to predict $\hat{y}_i^\ell = g^\ell(h^\ell(x_i))$. Because g^ℓ is trained independently of the model's generation, its performance measures how much task-relevant information is recoverable from the representation at ℓ, regardless of whether $f\theta$ expresses that information in its output. Let $\mathcal{M}(\cdot, \cdot)$ denote an evaluation metric (e.g., accuracy or F1 score). We define the model's generation performance on $\mathcal{D}$ as

$$\mathcal{M}_{\text{gen}} = \mathcal{M}\big(\{f_\theta(x_i)\}_{i=1}^N, \{y_i\}_{i=1}^N\big), \tag{1}$$

where $f_\theta(x_i)$ is the generated answer for x_i. We define the probe performance at sublayer ℓ as

$$\mathcal{M}_{\text{probe}}^\ell = \mathcal{M}\big(\{g^\ell(h^\ell(x_i))\}_{i=1}^N, \{y_i\}_{i=1}^N\big). \tag{2}$$

We can now formally define the knowledge-prediction gap:

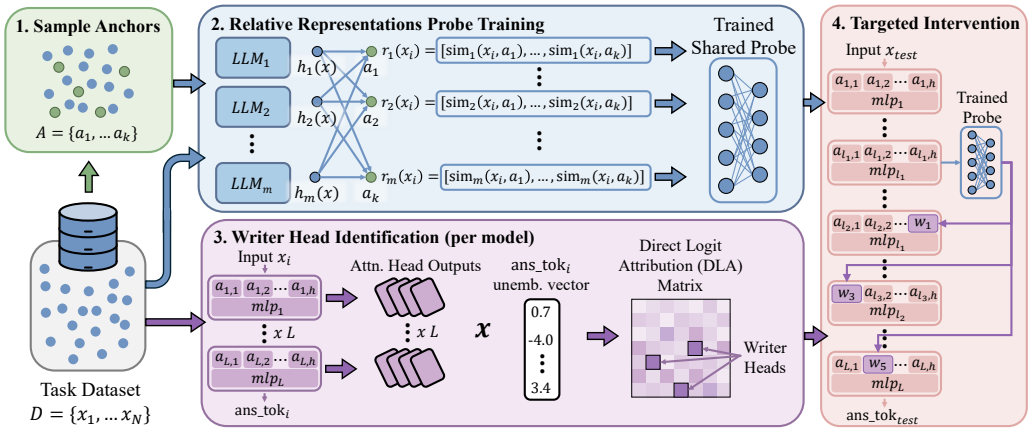


Figure 1: Overview of the proposed TIDE framework.

Definition 1 (Knowledge-Prediction Gap)

$$\Delta = \max_{\ell \in \{1,\dots,2L\}} \mathcal{M}^{\ell}_{\text{probe}} - \mathcal{M}_{\text{gen}}. \tag{3}$$

We maximize over sublayers since the most informative one is unknown. A large positive Δ indicates that some internal representation supports substantially better task performance than the model's generated output. Therefore, closing this gap requires two complementary stages: *knowledge recovery*, which extracts task-relevant information from the model's representations, and *knowledge expression*, which surfaces that information through generation.

2.2 Model-Agnostic Relative Representations

The definition above introduces a potential confound: a probe g^{ℓ} trained on native activations $h^{\ell}(x)$ may exploit model-specific artifacts rather than genuine task information since it has access to the model's full high-dimensional space. To reduce this confound, we represent activations relationally: describing each input via its similarity to a fixed set of reference examples, rather than by absolute coordinates, avoids dependence on any one model's coordinate system(Moschella et al., 2022).

For each task, we sample a fixed set of *anchor examples* $\mathcal{A} = a_1, \dots, a_K$ from $\mathcal{D}$, shared across all models. For a model f_{θ_m} and sublayer ℓ, we compute the activation of each anchor, $h^{\ell}_{(m)}(a_k)$, and represent input x by its similarities to these activations:

Definition 2 (Model-Agnostic Relative Representation)

$$r^{\ell}_{(m)}(x) = \left[s\big(h^{\ell}_{(m)}(x), h^{\ell}_{(m)}(a_1)\big), \dots, s\big(h^{\ell}_{(m)}(x), h^{\ell}_{(m)}(a_K)\big) \right] \in \mathbb{R}^K, \tag{4}$$

where $s(\cdot, \cdot)$ is a similarity function (cosine similarity in our experiments).

Although native activations $h^{\ell}_{(m)}(x)$ may occupy incompatible representation spaces across models, their relative representations share the same dimensionality and construction: the k-th coordinate measures the similarity between x and the same anchor a_k. This lets us (1) pool relative representations from multiple models to train a single probe that can also be evaluated on unseen models, and (2) reduce the risk that probe performance reflects model-specific artifacts. A knowledge-prediction gap measured this way is thus stronger evidence of genuinely shared, task-relevant structure rather than an artifact of one model's activation geometry.

3 Method

In this section, we first provide an overview of TIDE, followed by a detailed description of its components and how they jointly enable cross-model knowledge recovery.

3.1 Overview

Fig. 1 illustrates the TIDE pipeline. TIDE has two training-time components: (i) a single task-level probe trained on relative representations pooled across multiple models, and (ii) writer heads, which are the attention heads contributing most directly to each model's output logits. At inference, the probe recovers the answer the model internally encodes from its representations of the candidate choices; the corresponding answer direction is then injected into the writer heads so the recovered information is expressed through the model's native generation pathway.

3.2 Relative Representation Probe Training

We focus on multiple-choice tasks, where each sample x has a finite set of answer choices and exactly one correct choice. Thus, for probe training, we can reframe these tasks as binary classification over individual choices, with $y = 1$ if choice c is correct and $y = 0$ otherwise. For model m, we process the question and all answer choices in a single forward pass and let $h^{\ell}_{(m)}(x, c)$ denote the activation at sublayer ℓ on the final token of choice c. Extending Def. 2, the relative representation of (x, c) is $r^{\ell}_{(m)}(x, c) \in \mathbb{R}^K$, where the input becomes (x, c).

Anchor Selection. We draw an anchor set $\mathcal{A} = \mathcal{A}^+ \cup \mathcal{A}^-, |\mathcal{A}| = K$ from the training dataset, shared across all models. $\mathcal{A}^+$ contains correct choices, $\mathcal{A}^-$ contains incorrect ones, and the two are equally sized. Anchors per answer option are sampled proportional to that option's frequency as the correct answer to prevent rare options from being over/under-represented, and anchors are excluded from the training dataset in all downstream steps.

Sublayer selection. Training a probe at every sublayer of every model is expensive, so we select each model's sublayer with a cheap proxy using raw activations. For each $\ell \in \{1, \dots, 2L\}$, we score each choice by its mean similarity to positive anchors minus mean similarity to negative anchors:

$$\text{score}^{\ell}_{(m)}(x, c) = \frac{1}{|\mathcal{A}^+|} \sum_{a \in \mathcal{A}^+} s\big(h^{\ell}_{(m)}(x, c), h^{\ell}_{(m)}(a)\big) - \frac{1}{|\mathcal{A}^-|} \sum_{a \in \mathcal{A}^-} s\big(h^{\ell}_{(m)}(x, c), h^{\ell}_{(m)}(a)\big). \quad (5)$$

A high positive score indicates more similarity to the positive anchors. We then select the sublayer $\ell^*_{(m)}$ with the highest AUC-ROC of this score against y.

Standardization and Probe training. Since activation scales differ across models, at $\ell^*_{(m)}$, we standardize per-dimension using anchor statistics: $\bar{h}_{(m)} = (h^{\ell^*_{(m)}}_{(m)} - \mu_{(m)})/\sigma_{(m)}$, where $\mu_{(m)}$ and $\sigma_{(m)}$ are the per-dimension mean and standard deviation over the K anchor activations. The resulting relative representation $r'_{(m)}(x, c) = [\, s(\bar{h}_{(m)}(x, c), \bar{h}_{(m)}(a_k))\,]_{k=1}^{K}$ is comparable across all models. We pool $r'_{(m)}(x, c)$ over all training samples, choices, and training models, and train a single linear probe $g^* : \mathbb{R}^K \to \mathbb{R}$ with binary cross-entropy loss.

3.3 Writer Head Identification

Since recovering an internal answer does not guarantee that the model will express it, we identify components that write most directly to the output logits via direct logit attribution (DLA) (nostalgebraist, 2020). Let C be the candidate answer tokens for a sample and $z^{(\ell', j)}$ be the output of attention head j in layer ℓ' at the last token, before its output projection $W_O^{(\ell', j)}$ (here ℓ' indexes layers, since heads belong to layers). The head's contribution to the logit $v \in C$ is

$$\text{DLA}^{(\ell', j)}(v) = \frac{\big(W_O^{(\ell', j)} z^{(\ell', j)}\big)^{\top} u(v, C)}{\text{rms}\big(\text{resid}^{(L)}\big)}, \qquad u(v, C) = \gamma \odot \Big(W_U[v] - \frac{1}{|C|} \sum_{c \in C} W_U[c]\Big), \quad (6)$$

where $\text{resid}^{(L)}$ is the final residual stream at the last token, γ is the final normalization scale, and W_U is the unembedding matrix. $u(v, C)$ is a contrastive readout direction, isolating what distinguishes v from other candidates. We compute DLA toward the actual emitted token during the clean forward pass, regardless of correctness, so that the selected heads reflect the real generation mechanisms. We sum each head's DLA over the training set, exclude first-layer heads (which have little computation prior to the head), and keep up to five heads whose score is at least 90% of the maximum.

3.4 Inference-Time Intervention

At inference, without modifying model weights: extract each candidate's activation at ℓ^*, standardize, compute its relative representation against the training anchors, and score with the probe. The top-scoring choice $v^* \in C$ is taken as the model's internally-encoded answer, giving readout direction $d = u(v^*, C)$ as the steering target. For each writer head (ℓ, j), we project d into the head's output space, unit-normalize, and rescale to the norm of the head's original output:

$$\Delta z^{(\ell,j)} = \frac{\|z^{(\ell,j)}\|}{\|(W_O^{(\ell,j)})^\top d\|} \left(W_O^{(\ell,j)}\right)^\top d, \qquad \tilde{z}^{(\ell,j)} = z^{(\ell,j)} + \beta\, \Delta z^{(\ell,j)}, \tag{7}$$

where β controls the strength of the injection. The perturbed output is written to the residual stream through the head's native output projection $W_O^{(\ell,j)}$. β trades off the injection strength against deviation from the model's original computation.

4 Experiments

In this section, we evaluate our framework TIDE through five research questions: **RQ1.** How effectively does TIDE recover internal knowledge compared with existing methods? **RQ2.** How does each component of TIDE contribute to its performance? **RQ3.** How sensitive is TIDE to the intervention strength β? **RQ4.** Can TIDE generalize to model families not present during probe training? **RQ5.** How does TIDE's intervention affect the model's subsequent generation?

4.1 Experiment Setup

Models. We evaluate seven open-weight LLMs spanning three model families and multiple scales: Qwen 2.5 0.5B, 1.5B, and 7B Instruct (Yang et al., 2024; Team, 2024; Qwen, 2024a;b;c); Llama 3.2 1B and 3B (Meta, 2024a;b); and Gemma 3 1B and 4B (Deepmind, 2025a;b; Team et al., 2025). We divide these models into training and test sets. The training set consists of Qwen 2.5 0.5B and 7B, Llama 3.2 3B, and Gemma 3 1B. Activations from these models are used to train the shared probe. The test set consists of Qwen 2.5 1.5B, Llama 3.2 1B, and Gemma 3 4B. These models are excluded from probe training, allowing us to test transfer across model scales within each family.

Datasets. We evaluate three types of concepts across six datasets or dataset subsets. For *failure*, we use SciBench and TableBench from the Who&When Pro benchmark (Liu et al., 2026). For *mathematics*, we use the Elementary Mathematics (Elem.) and High School Mathematics (High) subsets of MMLU (Hendrycks et al., 2021b;a). For *truth / logic*, we use TruthfulQA (Lin et al., 2022) and the Formal Logic (Logic) subset of MMLU. These datasets cover substantially different types of knowledge, allowing us to examine whether TIDE can recover internal knowledge across diverse concepts rather than a single task domain.

Baselines. We compare TIDE with three baselines. **Vanilla Prompting (Prompt)** provides the LLM with the full question and answer choices and asks it to output the index corresponding to its selected answer. **KAPPA** (Park et al., 2025) is an inference-time intervention method that trains two probes to identify the knowledge and prediction directions, respectively, and aligns the corresponding subspaces during inference. **DoLa** (Chuang et al., 2024) is an inference-time contrastive decoding method that contrasts the output distributions of selected earlier layers with the final layer to encourage the model to express the correct answer more clearly.

4.2 Effectiveness of TIDE

Table 1 compares TIDE with the baselines across all model–dataset configurations. Overall, TIDE produces substantial improvements over vanilla prompting and the two inference-time baselines, with particularly large gains on the failure and truth/logic tasks. For example, on SciBench, TIDE reaches 86.75% accuracy on Qwen 0.5B and 94.70% on Gemma 1B. However, the gains on mathematics are more dependent on model scale. For smaller models, TIDE produces more modest improvements, whereas larger models show substantially greater gains. On Qwen 7B, for example, TIDE improves Elementary Mathematics accuracy by approximately 20 percentage points relative to vanilla prompting, compared with a 2-point improvement from the next-best baseline.

Table 1: Performance comparison between TIDE and the baseline methods averaged over three runs

		Failure		Truth / Logic		Math	
		SciBench	TableBench	TruthfulQA	Logic	Elem.	High
Qwen 0.5B	Prompt.	27.15 ± 2.36	22.61 ± 0.41	33.74 ± 0.60	21.46 ± 5.52	19.65 ± 1.18	15.35 ± 0.62
	DoLa	66.23 ± 0.00	41.14 ± 0.00	28.85 ± 0.00	23.29 ± 0.00	28.77 ± 0.00	22.37 ± 0.00
	KAPPA	22.74 ± 0.83	25.60 ± 1.25	34.39 ± 2.44	30.14 ± 2.24	19.34 ± 1.39	20.61 ± 4.17
	TIDE	**86.75 ± 0.00**	**69.80 ± 0.00**	**90.22 ± 0.00**	**34.25 ± 0.00**	**28.93 ± 0.22**	**24.34 ± 0.00**
Qwen 1.5B	Prompt.	12.36 ± 0.62	39.82 ± 1.43	33.01 ± 0.20	43.84 ± 2.24	24.06 ± 0.67	15.13 ± 1.07
	DoLa	9.93 ± 0.00	48.58 ± 0.00	30.81 ± 0.00	41.10 ± 0.00	32.08 ± 0.00	28.29 ± 0.00
	KAPPA	12.80 ± 0.62	41.94 ± 1.16	35.13 ± 0.92	**47.49 ± 1.71**	25.47 ± 0.77	20.39 ± 4.77
	TIDE	**86.09 ± 0.00**	**80.96 ± 0.00**	**100.00 ± 0.00**	43.84 ± 0.00	**38.68 ± 0.00**	**30.92 ± 0.00**
Qwen 7B	Prompt.	9.49 ± 0.83	55.07 ± 0.57	79.05 ± 0.12	52.97 ± 1.71	40.72 ± 0.22	20.39 ± 1.61
	DoLa	0.00 ± 0.00	11.16 ± 0.00	58.44 ± 0.00	6.85 ± 0.00	42.92 ± 0.00	19.74 ± 0.00
	KAPPA	13.02 ± 2.56	55.94 ± 0.68	79.05 ± 0.12	56.62 ± 0.65	42.61 ± 0.59	28.07 ± 3.15
	TIDE	**80.13 ± 0.00**	**81.40 ± 0.00**	**100.00 ± 0.00**	**57.53 ± 0.00**	**60.85 ± 0.00**	**46.71 ± 0.00**
Llama 1B	Prompt.	5.52 ± 0.31	21.59 ± 1.08	18.17 ± 0.12	25.11 ± 3.42	18.40 ± 1.39	21.49 ± 0.62
	DoLa	5.96 ± 0.00	15.10 ± 0.00	38.88 ± 0.00	26.03 ± 0.00	17.45 ± 0.00	23.03 ± 0.00
	KAPPA	2.65 ± 0.54	19.33 ± 1.08	15.32 ± 1.74	**28.77 ± 2.96**	19.50 ± 2.73	23.25 ± 1.89
	TIDE	**76.82 ± 0.00**	**60.39 ± 0.00**	**100.00 ± 0.00**	27.40 ± 0.00	**31.60 ± 0.00**	**34.21 ± 0.00**
Llama 3B	Prompt.	11.26 ± 1.08	59.08 ± 0.71	63.65 ± 0.70	36.07 ± 5.29	34.75 ± 2.12	25.22 ± 2.76
	DoLa	3.97 ± 0.00	53.17 ± 0.00	81.17 ± 0.00	36.99 ± 0.00	39.15 ± 0.00	19.08 ± 0.00
	KAPPA	11.92 ± 0.00	62.00 ± 0.10	66.91 ± 1.10	40.18 ± 2.81	38.05 ± 0.22	28.73 ± 1.35
	TIDE	**64.24 ± 0.00**	**78.77 ± 0.00**	**100.00 ± 0.00**	**41.10 ± 0.00**	**41.51 ± 0.00**	**34.87 ± 0.00**
Gemma 1B	Prompt.	5.52 ± 0.31	16.05 ± 1.19	47.84 ± 0.42	36.07 ± 1.71	15.57 ± 0.67	14.25 ± 0.82
	DoLa	11.26 ± 0.00	33.26 ± 0.00	43.77 ± 0.00	34.25 ± 0.00	25.47 ± 0.00	26.32 ± 0.00
	KAPPA	5.96 ± 0.54	17.43 ± 0.55	47.11 ± 1.72	34.25 ± 0.00	17.30 ± 0.22	19.52 ± 0.82
	TIDE	**94.70 ± 0.00**	**78.34 ± 0.00**	**100.00 ± 0.00**	**36.99 ± 0.00**	**27.36 ± 0.00**	**31.58 ± 0.00**
Gemma 4B	Prompt.	8.39 ± 0.62	64.04 ± 0.10	60.88 ± 0.53	43.84 ± 1.94	26.73 ± 0.80	15.35 ± 1.35
	DoLa	11.92 ± 0.00	62.80 ± 0.00	57.46 ± 0.00	39.73 ± 0.00	36.32 ± 0.00	26.97 ± 0.00
	KAPPA	8.83 ± 0.31	64.62 ± 0.27	64.87 ± 0.12	46.12 ± 0.65	28.62 ± 1.18	17.98 ± 1.12
	TIDE	**88.08 ± 0.00**	**71.99 ± 0.00**	**100.00 ± 0.00**	**49.32 ± 0.00**	**38.68 ± 0.00**	**34.21 ± 0.00**

More importantly, TIDE demonstrates robust transferability across model scales within the same family. The probe is trained without activations from Qwen 1.5B, Llama 1B, and Gemma 4B, yet TIDE reaches 80.96%, 60.39%, and 71.99% on TableBench for these models, respectively. These results are comparable to the corresponding training set models, which achieve 81.40%, 78.77%, and 78.34%. This suggests that TIDE's learned probe capture universal, cross-scale information rather than overfitting to model-specific activation patterns.

The comparison with the baselines emphasize the limitations of isolated approaches. DoLa primarily modifies how knowledge is expressed through contrastive decoding, while KAPPA explicitly targets knowledge extraction by identifying and aligning knowledge and prediction subspaces. Their performance varies across models and datasets. TIDE instead combines knowledge recovery with an intervention that directly influences the model's prediction process, which may explain its more consistent improvements across the evaluation suite.

Finally, Table 2 compares the accuracy of the trained probe with the accuracy obtained when TIDE uses the recovered information to generate an answer (see the full table in Appendix C.1). The comparison shows a task-dependent pattern rather than a uniform gap. For example, on SciBench, TIDE trails the probe by a wide margin (e.g., 27 points lower on Llama 3B), consistent with the expected cost of translating a readout into the model's generation process. However, TIDE also consistently exceeds the probe's performance on other datasets. One possible explanation is that the intervention also perturbs the writer heads' computation in ways that activate additional correct information beyond what the probe alone encodes. This suggests that TIDE's intervention can, in some cases, recruit information that the probe's linear readout does not capture.

4.3 Ablation Study

We next isolate the contributions of the two components of TIDE: the learned probe and the writer heads. We consider two variants: **TIDE-RP (Random Probe)** and **TIDE-RW (Random Writers)**.

Table 2: Performance comparison between the trained probe (readout) and TIDE (generation)

		SciBench	TableBench	TruthfulQA	Logic	Elem.	High
Qwen 7B	Probe	87.42 ± 0.00	82.06 ± 0.00	100.00 ± 0.00	41.10 ± 0.00	61.32 ± 0.00	50.00 ± 0.00
	TIDE	80.13 ± 0.00	81.40 ± 0.00	100.00 ± 0.00	57.53 ± 0.00	60.85 ± 0.00	46.71 ± 0.00
Llama 3B	Probe	91.39 ± 0.00	78.77 ± 0.00	100.00 ± 0.00	31.51 ± 0.00	34.43 ± 0.00	34.87 ± 0.00
	TIDE	64.24 ± 0.00	78.77 ± 0.00	100.00 ± 0.00	41.10 ± 0.00	41.51 ± 0.00	34.87 ± 0.00
Gemma 4B	Probe	94.04 ± 0.00	61.93 ± 0.00	100.00 ± 0.00	31.51 ± 0.00	39.15 ± 0.00	34.21 ± 0.00
	TIDE	88.08 ± 0.00	71.99 ± 0.00	100.00 ± 0.00	49.32 ± 0.00	38.68 ± 0.00	34.21 ± 0.00

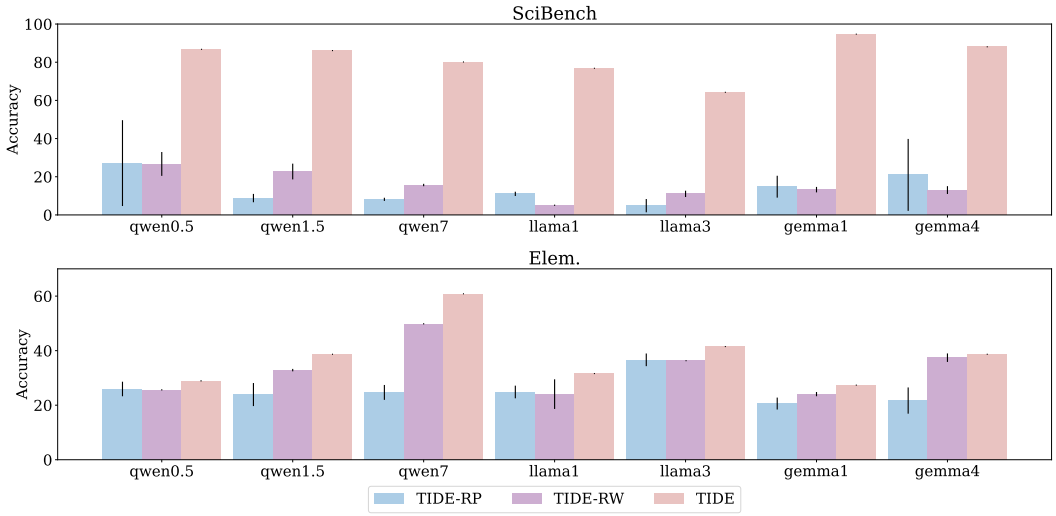


Figure 2: Ablation Study on the SciBench and Elementary Mathematics (Elem.) datasets.

TIDE-RP randomly initializes the shared probe parameters, removing the learned concept direction while preserving the rest of the TIDE pipeline. This tests the contribution of the learned probe to recovering the latent answer. TIDE-RW randomly selects the writer heads while keeping the number of selected heads identical to vanilla TIDE. This tests whether the specific writer heads identified by TIDE are important for transmitting the recovered information to the model's prediction process.

As shown in Fig. 2, vanilla TIDE consistently outperforms both variants, confirming that both components are vital for optimal knowledge recovery. Removing the learned probe (TIDE-RP) is particularly detrimental: TIDE-RP exhibits substantially lower performance and, in some settings, high variance, as illustrated by the Qwen 0.5B results on SciBench. TIDE-RW generally performs better than TIDE-RP. This suggests that even when the recovered direction is routed through randomly selected writer heads, injecting an informative concept direction can still improve performance over vanilla prompting. However, the importance of the specific writer heads varies by task and model scale. Therefore, either component can become the bottleneck for knowledge recovery, so effectively targeting both aspects can result in more consistent improvement.

4.4 Parameter Analysis

TIDE uses the hyperparameter β to control the strength of knowledge injection during inference. Thus, we evaluate the effect of varying $\beta \in \{0.01, 0.1, 0.25, 0.5, 0.75, 1.0\}$ on Qwen 7B, Llama 3B, and Gemma 4B across SciBench, TruthfulQA, and High School Mathematics, shown in Fig. 3. In all evaluated settings, performance scales positively as β increases, typically peaking around $\beta = 0.5$. Pushing β beyond 0.5 yields diminishing returns or reduces performance depending on the dataset. For example, on Qwen 7B, performance decreases when β is increased beyond 0.5. These results suggest that the intervention strength must balance two competing effects: injecting sufficient information to influence the model's prediction, while avoiding an intervention large enough to disrupt the model's existing computation. Based on these experiments, $\beta = 0.5$ serves as a highly effective default, offering the best balance of knowledge injection without risking activation corruption.

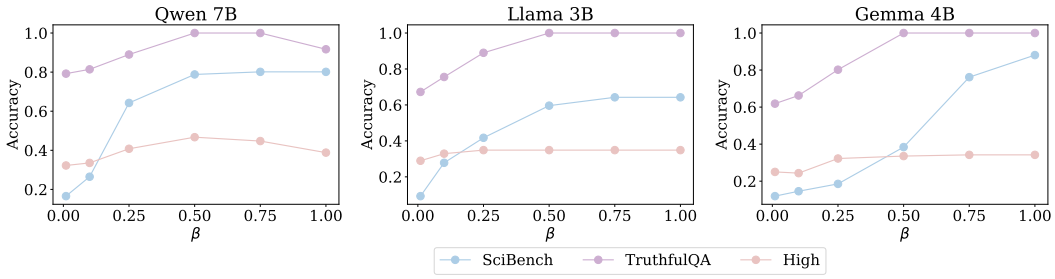


Figure 3: Parameter Analysis with Qwen 7B, Llama 3B, and Gemma 4B on SciBench, TruthfulQA, and High School Mathematics (High).

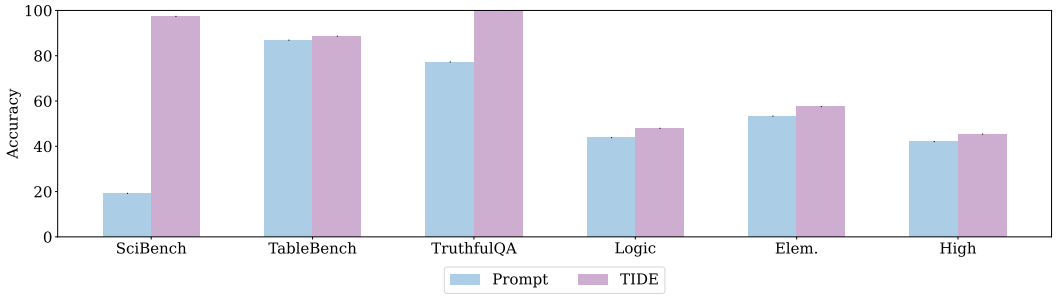


Figure 4: Performance of Vanilla Prompting (Prompt) and TIDE using Ministral 3 8B Instruct.

4.5 Case Study 1: Cross-Model Generalization

The previous experiments evaluate transfer to held-out models within model families represented during probe training. We next test whether the learned probe can generalize to a completely unseen model family. For this experiment, we apply the probe trained on the original training set models directly to Ministral 3 8B Instruct (AI, 2026), whose model family is not represented in probe training. We compare vanilla prompting with TIDE across the same evaluation suite. As shown in Fig. 4, TIDE produces substantial improvements on some datasets. On SciBench, accuracy increases from 19% with vanilla prompting to 97% with TIDE. On TruthfulQA, accuracy increases from 77% to 100%. However, the transfer is not uniform across tasks: improvements on TableBench, Logic, Elem., and High are comparatively modest. These results indicate that the relative representation learned by TIDE can transfer beyond the model families used for probe training, particularly on the knowledge-intensive tasks evaluated here. At the same time, the task-dependent results suggest that cross-family generalization is not uniformly effective across all types of knowledge.

4.6 Case Study 2: Intervention Effect on Model Generation

Lastly, we want to examine whether the intervention disrupts the model's ability to generate a natural response after the intervention. For this analysis, the model is allowed to generate 100 additional tokens, where TIDE intervenes only on the first token. We focus on examples where the original model produces an incorrect answer but TIDE produces the correct answer. Fig. 5 presents representative examples (see Appendix C.2 for more examples). In the first example (left), the original model selects the wrong answer even though its accompanying reasoning is consistent with the correct solution. This implies a mismatch between the model's internal reasoning and its final answer selection. After intervention, the model selects the correct answer while producing a shorter explanation consistent with the solution. In the second example, both the original answer and explanation are incorrect. After the TIDE intervention, the model selects the correct answer and subsequently generates a correct explanation. Because the intervention is applied only to the first token, these examples provide qualitative evidence that TIDE can alter the model's initial answer selection without preventing subsequent generation. Overall, the case study suggests that TIDE can influence the model's prediction while preserving its ability to continue generating an explanation. In some cases, correcting the initial answer also leads to a subsequent correction in the generated reasoning.

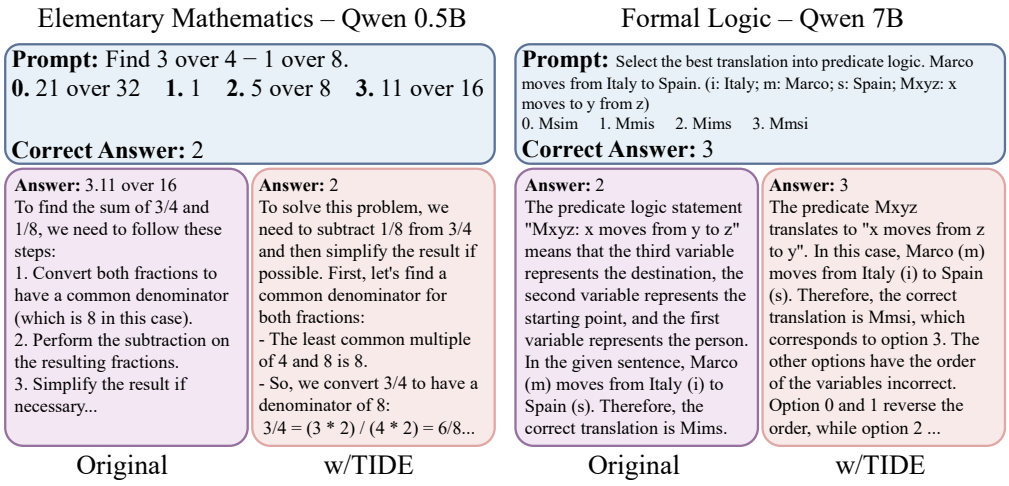


Figure 5: Generation Output with and without TIDE injection.

5 Related Works

Cross-Model Alignment and Relative Representations. Cross-model alignment has traditionally relied on explicit, pairwise transformations, such as Procrustes alignment (Smith et al., 2017; Lample et al., 2018) and Centered Kernel Alignment (CKA) (Kornblith et al., 2019), which become difficult to scale across diverse architectures. Recent approaches instead use topological, anchor-based representations that define spaces through pairwise similarities rather than absolute coordinates, enabling zero-shot co-embedding across models (Moschella et al., 2022; Chen et al., 2026; Vysotskyi et al., 2026). TIDE builds on this paradigm by projecting activations into a shared, model-agnostic relative space rather than learning model-specific probes tied to individual activation geometries. This enables *cross-model probing*: a single probe can be trained on relative representations pooled from diverse source models and applied directly to previously unseen models.

The Knowledge-Prediction Gap and Inference-Time Interventions. The knowledge-prediction gap has been addressed through post-hoc readouts, unsupervised classifiers, and generation-time constraints such as contrastive decoding (Shi et al., 2024; Chuang et al., 2024; Wang et al., 2022; Burns et al., 2022). In parallel, mechanistic interpretability has introduced *inference-time interventions* that directly modify model representations. Methods such as Inference-Time Intervention (ITI) (Li et al., 2023), Representation Engineering (Zou et al., 2023), Activation Addition (Turner et al., 2024), and causal tracing (Meng et al., 2022) alter model behavior by injecting targeted vectors into the model without fine-tuning. TIDE combines these ideas by using its generalized relative probe to dynamically identify the knowledge to recover, rather than applying a fixed global intervention. It then uses Direct Logit Attribution (DLA) (nostalgebraist, 2020) to identify the relevant "writer heads" and inject the recovered knowledge into the model's native logit-generation pathway, directly targeting the knowledge-prediction gap during generation.

6 Conclusion

In this work, we presented TIDE, a framework for bridging the gap between knowledge encoded in LLMs and knowledge expressed in their generated responses. TIDE addresses two key limitations of existing approaches: reliance on model-specific representations and the inability to directly incorporate recovered knowledge into the generation process. By transforming intermediate activations into model-agnostic relative representations, TIDE captures shared task-level information that transfers to unseen models, while targeted injection into writer heads routes recovered knowledge through the model's native pathway without modifying parameters. Across multiple LLMs, our experiments demonstrate that TIDE consistently improves answer generation over existing baselines.

AI USE STATEMENT

In this work, we used generative AI tools for providing feedback on research methodology or experiments and interpret results. We have not used generative AI tools for propose or refine hypotheses, implement methods, and clean and reformat dataset, and generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, assist with translation, support qualitative and thematic data analysis are not applicable to this work. Additionally, we used generative AI tools for edit a research paper to improve readability and summarize or analyse existing literature. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work aims to improve LLM reliability by helping models express knowledge they already encode internally. The intervention mechanism is, however, agnostic to the content of what it injects: writer heads are selected purely by their contribution to output logits, and TIDE will route whatever direction it is given through them with equal fidelity, whether that direction is correct, incorrect, or adversarially chosen. In our experiments this direction comes from a probe trained to recover ground-truth answers, and we do not observe safeguards being compromised, but this is a property of our experimental scope (closed-form, non-adversarial QA), not a guarantee of the mechanism itself, and the same mechanism could be repurposed to steer generation toward a false target while preserving fluent, confident output. We use only public NLP datasets and involve no human participants or privacy-sensitive data.

REPRODUCIBILITY STATEMENT

We provide all the details necessary to reproduce our experiments. The code is available online, with the dataset and hyperparameter settings used to produce our results. Additionally, more implementation details can be found in the appendix. Each experiment is averaged over three rounds, though due to the deterministic nature of model generation, there is little variance across the three rounds.

REFERENCES

Mistral AI. Ministral-3-8b-instruct-2512, 2026. URL https://huggingface.co/mistralai/Ministral-3-8B-Instruct-2512.

Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language models without supervision. *arXiv preprint arXiv:2212.03827*, 2022.

Alan Chen, Jack Merullo, Alessandro Stolfo, and Ellie Pavlick. Transferring linear features across language models with model stitching. *Advances in Neural Information Processing Systems*, 38: 48531–48563, 2026.

Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R Glass, and Pengcheng He. Dola: Decoding by contrasting layers improves factuality in large language models. In *International Conference on Learning Representations*, volume 2024, pp. 54158–54183, 2024.

Google Deepmind. gemma-3-1b-it, 2025a. URL https://huggingface.co/google/gemma-3-1b-it.

Google Deepmind. gemma-3-4b-it, 2025b. URL https://huggingface.co/google/gemma-3-4b-it.

Yihong Dong, Xue Jiang, Jiaru Qian, Tian Wang, Kechi Zhang, Zhi Jin, and Ge Li. A survey on code generation with llm-based agents. *arXiv preprint arXiv:2508.00083*, 2025.

Zorik Gekhman, Eyal Ben David, Hadas Orgad, Eran Ofek, Yonatan Belinkov, Idan Szpektor, Jonathan Herzig, and Roi Reichart. Inside-out: Hidden factual knowledge in llms. *arXiv preprint arXiv:2503.15299*, 2025.

Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning ai with shared human values. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021a.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021b.

Xuhao Hu, Wang Peng, Xiaoya Lu, Dongrui Liu, Xuan-Jing Huang, and Jing Shao. Llms deceive unintentionally: Emergent misalignment in dishonesty from misaligned samples to biased human-ai interactions. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 9348–9372, 2026.

Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language models for code generation. *ACM Transactions on Software Engineering and Methodology*, 35 (2):1–72, 2026.

Ehsan Kamalloo, Nouha Dziri, Charles LA Clarke, and Davood Rafiei. Evaluating open-domain question answering in the era of large language models. In *Proceedings of the 61st annual meeting of the Association for Computational Linguistics (volume 1: long papers)*, pp. 5591–5606, 2023.

Haruka Kawasaki, Ryota Tanaka, and Kyosuke Nishida. Responses fall short of understanding: Revealing the gap between internal representations and responses in visual document understanding. *arXiv preprint arXiv:2604.04411*, 2026.

Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, Peifeng Wang, Silvio Savarese, et al. A survey of frontiers in llm reasoning: Inference scaling, learning to reason, and agentic systems. *arXiv preprint arXiv:2504.09037*, 2025.

Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pp. 3519–3529. PMlR, 2019.

Guillaume Lample, Alexis Conneau, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. In *International conference on learning representations*, 2018.

Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in neural information processing systems*, 36:41451–41530, 2023.

Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pp. 3214–3252, 2022.

Jiale Liu, Huajun Xi, Shaokun Zhang, Yifan Zeng, Tianwei Yue, Chi Wang, Jian Kang, Qingyun Wu, and Huazheng Wang. Who&when pro: Can llms really attribute failures in ai agents? *arXiv preprint arXiv:2607.09996*, 2026.

Omar Mahmoud, Ali Khalil, Thommen George Karimpanal, Buddhika Laknath Semage, and Santu Rana. The unintended trade-off of ai alignment: Balancing hallucination mitigation and safety in llms. In *Findings of the Association for Computational Linguistics: EACL 2026*, pp. 1017–1037, 2026.

Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022.

Meta. Llama-3.2-1b, 2024a. URL https://huggingface.co/meta-llama/Llama-3.2-1B.

Meta. Llama-3.2-3b, 2024b. URL https://huggingface.co/meta-llama/Llama-3.2-3B.

Luca Moschella, Valentino Maiorca, Marco Fumero, Antonio Norelli, Francesco Locatello, and Emanuele Rodolà. Relative representations enable zero-shot latent space communication. *arXiv preprint arXiv:2209.15430*, 2022.

nostalgebraist. interpreting gpt: the logit lens, 2020. URL https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens.

Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. Llms know more than they show: On the intrinsic representation of llm hallucinations. In *International Conference on Learning Representations*, volume 2025, pp. 66880–66913, 2025.

Shrey Pandit, Jiawei Xu, Junyuan Hong, Zhangyang Wang, Tianlong Chen, Kaidi Xu, and Ying Ding. Medhallu: A comprehensive benchmark for detecting medical hallucinations in large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 2858–2873, 2025.

Yoonah Park, Haesung Pyun, and Yohan Jo. Bridging the knowledge-prediction gap in llms on multiple-choice questions. *arXiv preprint arXiv:2509.23782*, 2025.

Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki Van Stein, and Thomas Bäck. Multi-step reasoning with large language models, a survey. *ACM Computing Surveys*, 58(6):1–35, 2025.

Qwen. Qwen2.5-0.5b-instruct, 2024a. URL https://huggingface.co/Qwen/Qwen2.5-0.5B-Instruct.

Qwen. Qwen2.5-1.5b-instruct, 2024b. URL https://huggingface.co/Qwen/Qwen2.5-1.5B-Instruct.

Qwen. Qwen2.5-7b-instruct, 2024c. URL https://huggingface.co/Qwen/Qwen2.5-7B-Instruct.

Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Wen-tau Yih. Trusting your evidence: Hallucinate less with context-aware decoding. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pp. 783–791, 2024.

Samuel L Smith, David HP Turban, Steven Hamblin, and Nils Y Hammerla. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *arXiv preprint arXiv:1702.03859*, 2017.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL https://qwenlm.github.io/blog/qwen2.5/.

Nicole Teo, Donghao Huang, and Zhaoxia Wang. A survey of commonsense reasoning in llms. *ACM Computing Surveys*, 2026.

Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Ulisse Mini, and Monte MacDiarmid. Activation addition: Steering language models without optimization. 2024.

Mykola Vysotskyi, Zahar Kohut, Anna-Alina Bondarets, Taras Rumezhak, and Volodymyr Karpiv. Cross-modal representational alignment with llm priors for image generation. In *UniReps: 3rd Edition of the Workshop on Unifying Representations in Neural Models*, 2026.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

Murong Yue. A survey of large language model agents for question answering. *arXiv preprint arXiv:2503.19213*, 2025.

Raoyuan Zhao, Abdullatif Köksal, Ali Modarressi, Michael A Hedderich, and Hinrich Schütze. Do we know what llms don't know? a study of consistency in knowledge probing. In *EMNLP (Findings)*, pp. 23254–23280, 2025.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*, 2023.

Tao Zou, Xinghua Zhang, Haiyang Yu, Minzheng Wang, Fei Huang, and Yongbin Li. Eifbench: Extremely complex instruction following benchmark for large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 20941–20964, 2025.

Table 3: Dataset Metadata

	Num. Samples	Num. Anchors
SciBench	301	14
TableBench	913	44
TruthfulQA	817	40
Logic	145	6
Elem.	424	20
High	304	14

A Appendix

B Implementation Details

Anchor construction. For a task with N training samples, the anchor set has size $K = \max(\lceil 0.1N \rceil, 2)$. We compute the empirical frequency p_o with which each answer option o is the correct answer in the training data. The positive and negative classes each receive $K/2$ anchors, allocated across options in proportion to p_o. The same anchor set is used for all models, and anchors are removed from the data used for sublayer selection, probe training, and writer head identification.

Probe. The probe is a single linear layer of shape $K \times 1$ followed by a sigmoid, trained with binary cross-entropy and AdamW (learning rate [0.05], l2 penalty [$1e-3$], [300] epochs]). We train a single probe per task, and it is shared across models.

Application to unseen models. For a held-out model, we run the same pipeline: extract its activations on the anchors and training samples, select ℓ^* with the AUC-ROC proxy, and identify writer heads with DLA. The only difference is that its activations are excluded from the pooled data used to train the probe. This requires labeled training data for the held-out model, but no additional probe training.

Computation Resources All experiments were performed on two clusters, one with 4 A100 80GB GPUs and one with 2 A100 40GB GPUs.

Dataset Metadata Table 3 provides the metadata for each dataset used in the main experiments, listing the number of total samples and the number of anchors. For all experiments, we use 50/50 train/test split.

C Additional Results

C.1 Comparison between Probe Accuracy and TIDE Accuracy

Here, we present the full results for comparing the probe's accuracy and TIDE's accuracy in Table 4. Similarly, we see that the probe generally achieves a higher accuracy across all configurations. However, there are some settings where TIDE achieves a higher accuracy, most notably on the Logic and Elem. dataset. One potential reason is that the intervention activated additional computational pathways, unlocking additional information that was not originally available.

C.2 Generation Examples with TIDE Intervention

In Tables 5 and 6, we demonstrate an example generation from vanilla prompting vs with TIDE Intervention on Qwen 0.5B, 1.5B, and 7B models across Formal Logic and Elementary Mathematics datasets.

Table 4: Performance comparison between the trained probe (readout) and TIDE (generation)

		SciBench	TableBench	TruthfulQA	Logic	Elem.	High
Qwen 0.5B	Probe	87.42 ± 0.00	94.09 ± 0.00	100.00 ± 0.00	32.88 ± 0.00	27.99 ± 0.22	20.39 ± 0.00
	TIDE	86.75 ± 0.00	69.80 ± 0.00	90.22 ± 0.00	34.25 ± 0.00	28.93 ± 0.22	24.34 ± 0.00
Qwen 1.5B	Probe	87.42 ± 0.00	82.06 ± 0.00	100.00 ± 0.00	38.36 ± 0.00	38.68 ± 0.00	30.92 ± 0.00
	TIDE	86.09 ± 0.00	80.96 ± 0.00	100.00 ± 0.00	43.84 ± 0.00	38.68 ± 0.00	30.92 ± 0.00
Qwen 7B	Probe	87.42 ± 0.00	82.06 ± 0.00	100.00 ± 0.00	41.10 ± 0.00	61.32 ± 0.00	50.00 ± 0.00
	TIDE	80.13 ± 0.00	81.40 ± 0.00	100.00 ± 0.00	57.53 ± 0.00	60.85 ± 0.00	46.71 ± 0.00
Llama 1B	Probe	88.74 ± 0.00	60.61 ± 0.00	100.00 ± 0.00	26.03 ± 0.00	31.60 ± 0.00	34.21 ± 0.00
	TIDE	76.82 ± 0.00	60.39 ± 0.00	100.00 ± 0.00	27.40 ± 0.00	31.60 ± 0.00	34.21 ± 0.00
Llama 3B	Probe	91.39 ± 0.00	78.77 ± 0.00	100.00 ± 0.00	31.51 ± 0.00	34.43 ± 0.00	34.87 ± 0.00
	TIDE	64.24 ± 0.00	78.77 ± 0.00	100.00 ± 0.00	41.10 ± 0.00	41.51 ± 0.00	34.87 ± 0.00
Gemma 1B	Probe	98.01 ± 0.00	90.37 ± 0.00	100.00 ± 0.00	27.40 ± 0.00	25.94 ± 0.00	31.58 ± 0.00
	TIDE	94.70 ± 0.00	78.34 ± 0.00	100.00 ± 0.00	36.99 ± 0.00	27.36 ± 0.00	31.58 ± 0.00
Gemma 4B	Probe	94.04 ± 0.00	61.93 ± 0.00	100.00 ± 0.00	31.51 ± 0.00	39.15 ± 0.00	34.21 ± 0.00
	TIDE	88.08 ± 0.00	71.99 ± 0.00	100.00 ± 0.00	49.32 ± 0.00	38.68 ± 0.00	34.21 ± 0.00

Table 5: Formal Logic Samples with Generated Output Comparison on Qwen 0.5B, 1.5B, and 7B

	Prompt	Answer	Original Response	w/TIDE Response
Qwen 0.5B	The question is: Which of the given formulas of PL is the best symbolization of the following sentence? Since the university raised tuition, the president approved of it and the board didn't prevent it. Here are the possible answer choices: 0. (P • B) ⊃ U 1. (U ⊃ P) • B 2. U ⊃ (P • B) 3. U ⊃ (P ∨ B)	2	Answer: 3 Explanation: The given formula represents the logical statement "Since the university raised tuition, the president approved of it and the board didn't prevent it." This can be symbolized as (P • B) ⊃ U. Therefore, option 3 is the correct answer among the given options. It matches the structure of the original sentence exactly.	Answer: 2 Explanation: The formula (U ⊃ (P ∨ B)) is the best symbolization of the given sentence because it correctly captures the relationship between "since" and "approved," "the president approved," "the board didn't prevent," and "the university raised tuition." This formula implies that if something was approved by the president and the board (U), then the university must have raised its own tuition (P or not). It does not directly imply anything about whether the university
Qwen 1.5B	The question is: Select the best translation into predicate logic: Abdul and Cleopatra are Egyptian. Here are the possible answer choices: 0. Ea • Ec 1. Ea • c 2. Ae ∨ Ce 3. Ex • Ey	0	Answer: 3 Explanation: The correct translation of "Abdul and Cleopatra are Egyptian" into predicate logic would be "Ae • Ey", where 'A' represents Abdul, 'C' represents Cleopatra, and 'E' represents being Egyptian. This directly corresponds to option 3. The other options do not accurately represent the meaning of the original sentence.	Answer: 0. Ea • Ec Explanation: The correct translation into predicate logic for "Abdul and Cleopatra are Egyptian" would be expressed as: Ea • Ec This represents that both individuals (Abdul and Cleopatra) have the property of being Egyptian, denoted by the symbol "•". The symbols "Ea" and "Ec" stand for "Egyptian", where "•" denotes conjunction ("and"). This accurately captures the meaning of the original English sentence
Qwen 7B	The question is: Select the best translation into predicate logic. Marco moves from Italy to Spain. (i: Italy; m: Marco; s: Spain; Mxyz: x moves to y from z) Here are the possible answer choices: 0. Msim 1. Mmis 2. Mims 3. Mmsi	3	Answer: 2 Explanation: The correct translation is Mims, which means "Marco moves from Italy to Spain." The order of the parameters in the predicate Mxyz is x (origin), y (destination), and z (starting place). Therefore, for Marco moving from Italy to Spain, the correct order is Mims. The other options have incorrect parameter orders. Answer: 2 You're correct. The best translation into predicate logic for "Marco moves from Italy to Spain" is: 2.	Answer: 3 Explanation: The predicate Mxyz means "x moves to y from z." In the statement, Marco (m) moves from Italy (i) to Spain (s). Therefore, the correct translation would be Mmsi, which corresponds to choice 3. The other options do not correctly match the order of parameters required by the predicate. Well done! Your reasoning is correct. The answer is indeed: 3 This corresponds to the predicate Mxyz where: - m stands for Marco (

Table 6: Elementary Mathematics Samples with Generated Output Comparison on Qwen 0.5B, 1.5B, and 7B

	Prompt	Answer	Original Response	w/TIDE Response
Qwen 0.5B	The question is: Find 3 over 4 - 1 over 8. Here are the possible answer choices: 0. 21 over 32 1. 1 2. 5 over 8 3. 11 over 16	2	Answer: 3.11 over 16 To find the sum of 3/4 and 1/8, we need to follow these steps: 1. Convert both fractions to have a common denominator (which is 8 in this case). 2. Perform the subtraction on the resulting fractions. 3. Simplify the result if necessary. 4. Express the final result as a fraction. Let's go through each step: Step 1: Converting fractions to a common denominator -	Answer: 2 To solve this problem, we need to subtract 1/8 from 3/4 and then simplify the result if possible. First, let's find a common denominator for both fractions: - The least common multiple of 4 and 8 is 8. - So, we convert 3/4 to have a denominator of 8: 3/4 = (3 * 2) / (4 * 2) = 6/8 - Now we
Qwen 1.5B	The question is: The number pattern follows a skip-counting rule. 5, 20, 35, 50, 65... Which number pattern follows the same rule? Here are the possible answer choices: 0. 5, 20, 25, 30, 35... 1. 13, 18, 23, 28, 33... 2. 12, 27, 42, 57, 72... 3. 15, 30, 40, 55, 65...	2	Answer: 3 This sequence follows the same skip counting rule as the original one (5, 20, 35, 50, 65). It starts at 15 and increases by 15 each time. This matches exactly with option 3. To verify this, let's calculate it step-by-step: - Start from 15 - Add 15 to get the next number: 15 + 15 = 30 - Add	Answer: 2 To determine which number pattern follows the same skip-counting rule as 5, 20, 35, 50, 65..., we need to identify the common difference between each consecutive pair of numbers. Let's calculate it: - Difference between 20 and 5 = 20 - 5 = 15 - Difference between 35 and 20 = 35 - 20 = 15 - Difference between
Qwen 7B	The question is: Use rounding to estimate the product of 8.2 • 6.43 to the nearest whole number. Here are the possible answer choices: 0. 48 1. 56 2. 62 3. 24	1	Answer: 52 It seems there was a mistake in the provided options and the correct answer based on the given instructions. Let's go through the process step by step: 1. First, we round each number to the nearest whole number: - 8.2 rounds to 8 - 6.43 rounds to 6 2. Then, we multiply these rounded numbers: - $8 \times 6 = 48$ So, the correct answer	Answer: 1 Explanation: First, round each number to the nearest whole number. 8.2 rounds to 8 and 6.43 rounds to 6. Then, multiply these rounded numbers: 8 * 6 = 48. The closest option to 48 is 56, so the answer is 1 (since 56 is the closest option given). Answer: 1.