

Can LLMs Introspect? A Reality Check

Shashwat Singh, Tal Linzen, Shauli Ravfogel

Center for Data Science

New York University

{ss20428, linzen, shauli.ravfogel}@nyu.edu

Abstract

Can large language models detect and report their own internal states? A number of recent studies have argued that they can. Drawing on lessons from human metacognition research, we argue that this conclusion may be premature. We identify two conditions that a paradigm needs to meet in order to establish introspection. First, the test needs to require *privileged access*: it should not be solvable using cues available in the input. Second, it needs to require *second-order computation*: second-order, meta-representations of first-order, task-related representations. This condition cannot be satisfied by task performance alone: it requires designs under which second-order and first-order accounts make divergent predictions. We re-examine two paradigms that have been used to argue for model introspection in light of these conditions. In the first, models must predict labels derived from their own hidden states; we find that classifiers that can only access the input match the models’ in-context predictions, indicating that the original results do not demonstrate privileged access to internal representations. In the second paradigm, models must detect whether their internal states have been tampered with; we find they cannot reliably distinguish such interventions from manipulations of the input, suggesting that their success reflects generic anomaly detection rather than sensitivity to internal interventions in particular. We conclude that current evidence is insufficient to establish metacognitive monitoring in LLMs.



https://github.com/shashwat1002/introspection_reality_check

1 Introduction

Can language models reflect on their own internal processes? A surge of recent work asks whether LLMs can not only produce complex behavior but also monitor, report, and regulate their own internal states—abilities referred to in human cognitive science as introspection or metacognition (Nisbett & Wilson, 1977; Flavell, 1979; Nelson, 1990). Several recent studies have answered this question in the affirmative. Drawing on human metacognition research, which has yielded largely negative results and identified a range of confounds in self-report studies (Fleming & Lau, 2014), we identify two necessary conditions for introspection in the strong sense we adopt: (i) *privileged access*—the report must depend on internal states beyond what the input alone provides, or else self-report reduces to input-driven prediction of the model’s own behavior (Shanahan et al., 2023; Turpin et al., 2023); and (ii) *second-order computation*—the report must be produced by meta-representations that represent something *about* first-order, task-related representations (e.g., “my normal representation was manipulated,” “my answer was produced by guessing”). We argue that recent affirmative conclusions outrun the evidence on both counts: empirically, existing paradigms often fail to

establish even the privileged access condition; and as a matter of construct validity, positive results in these paradigms underdetermine the second-order computation condition, since first-order and second-order processes can give rise to identical model self-reports.

We examine two prominent paradigms in light of these conditions. The first, which we term *self-report classification*, seeks to establish introspection by showing that models can solve in-context learning tasks where the labels are derived from the models’ own activations (Ji-An et al., 2025; Steinmetz Yalon et al., 2026). For instance, models can predict labels corresponding to the first principal components of their representations (Ji-An et al., 2025) or their tendency to follow parametric versus contextual knowledge (Steinmetz Yalon et al., 2026). However, the fact that these labels are *derived* by the experimenter from hidden states does not exclude the possibility that they can also be recovered by the model from the *input alone*. Indeed, we show that labels in existing self-report classification tasks are often predictable directly from input features, without any access to the hidden states; and when the label are decorrelated from these features, performance collapses to near-chance. We conclude that self-report classification, as currently deployed, fails to satisfy the privileged access condition.

The second paradigm we study is *intervention awareness* (Lindsey, 2025): in this study, Anthropic’s Claude models were able to detect with nontrivial accuracy whether their activations were modified through steering, that is, by adding a vector representing a concept to the activations (Li et al., 2023; Singh et al., 2024). This paradigm satisfies the privileged access criterion by design. But, we argue, Lindsey’s conclusion that models can track interventions on the their internal states is too strong. In his design, which includes only normal and steered trials, “my activations were intervened on” and “my current state is unusual” are true on exactly the same trials, so accuracy alone cannot reveal which of the two explains the model’s reports. We address this by adding a third “gaslight” condition that pushes the model toward a concept through the input prompt alone, with no direct intervention on the activations (section 4.3). While in this condition the input might cause the model’s state to be unusual, nothing is injected into its activations. We find that four open-weight models¹ frequently incorrectly classify gaslight trials as injection trials, even when explicitly instructed to distinguish the two types of interventions (fig. 1, right). This suggests that the models’ reports track that something is unusual rather than activation intervention specifically, and that there is no evidence for self-monitoring of the hidden states. We conclude that this paradigm leaves the second-order computation condition unaddressed; in fact, as we discuss in section 3, even success in our extended variant of the experiment would not meet this condition.

In summary, we conclude that current evidence is insufficient to establish that LLMs display strong metacognitive monitoring, and argue that future studies could be made more compelling by including stronger controls and, crucially, by pairing behavioral results with mechanistic evidence of a dissociable second-order process (for first steps in this direction, see Macar et al. 2026).

2 Related Work

Work on LLM metacognition spans verbal calibration (Kadavath et al., 2022; Lin et al., 2022; Yona et al., 2024), probing for internal representations of confidence or truthfulness (Burns et al., 2023; Marks & Tegmark, 2024; Azaria & Mitchell, 2023; Ravfogel et al., 2025), and paradigms testing whether models can report their own activation patterns (Ji-An et al., 2025; Steinmetz Yalon et al., 2026); for an extended review, see section C.1. The human metacognition literature, which is rife with negative results, provides essential context for our work: people often confabulate explanations

¹We cannot replicate the paper directly, as the model tested by Lindsey (2025) is not accessible outside of Anthropic.




Figure 1: **Input-controlled alternatives to purported introspection results.** *Left:* In the self-report paradigm of Ji-An et al. (2025), labels are computed from a model’s hidden state via a linear classifier or top PCA directions (A), then used as targets in in-context learning examples (B). Successful prediction has been interpreted as evidence of introspection. We show these labels are also predictable from uncontextualized input embeddings, so success need not imply privileged access. *Right:* In the steering-awareness setting of Lindsey (2025), the anomaly detection hypothesis and introspection hypothesis make the same prediction and are confounded. Our design adds prompt interventions (the “gaslight” condition) matched to hidden-state interventions, separating the hypotheses: anomaly detection flags both as anomalous, while introspection selectively identifies hidden-state interventions.

of their own behavior (Nisbett & Wilson, 1977), and apparent introspection into one’s memorization abilities can stem from shallow cues such as familiarity rather than access to memory traces (Koriat, 1997). Because above-chance correlations between confidence and accuracy can arise from first-order evidence alone, measures of metacognitive sensitivity have been developed within signal-detection theory (Maniscalco & Lau, 2012; Fleming & Lau, 2014), and first- and second-order accounts of confidence are differentiated by behavioral model comparison (Fleming & Daw, 2017). The parallel debate in comparative psychology is instructive: behavioral evidence of uncertainty monitoring in animals (Smith et al., 2003; Hampton, 2001) was met with first-order reinterpretations (Carruthers, 2008), and progress ultimately required converging behavioral and neural evidence (Kepecs et al., 2008).

Recent work has begun controlling for confounds in LLM self-report (section C.2). Binder et al. (2024) show that a model predicts its own behavior better than an equally informed external model, but this capability is finetuned and, as they note, is consistent with the model *simulating* its own forward pass—which satisfies the privileged access but not the second-order condition. Plunkett et al. (2025) finetune models to use randomly generated importance weights for different features in a classification problem, and find that they report those weights accurately, with strong controls against input-driven inference; the authors note that the reports may nonetheless reflect stored self-knowledge rather than real-time monitoring. Song et al. (2025) propose that to demonstrate introspection one needs to show a reliability advantage over any process of equal or lower computational cost available to a third party; they show that apparent successes can fail to meet this criterion. We share the motivation of the controls proposed in these studies, but we go further and argue that privileged access on its own does not suffice to establish introspection in the strong sense.

Several groups have revisited the concept-injection detection work of (Lindsey, 2025) in open-weight models, in some cases with explicit finetuning for the task (Vogel, 2025; Rivera & Africa, 2026); we focus on pretrained models, since finetuning may equip the model with a task-specific mechanism rather than reveal a general capacity (section C.5). Closest to our critique, Lederman & Mahowald (2026) argue that injection detection is content-agnostic: models detect that an anomaly occurred but cannot identify the injected concept. Their argument and ours target different links in the same chain: they dissociate detection from identification, whereas we ask what detection itself consists in—a second-order representation of the model’s states, or a first-order sensitivity to irregularity (section C.4).

3 Construct Validity of Introspection Paradigms

Defining introspection. Introspection is not a univocal notion. On one family of views, a system knows its own states by using *meta-representations*: states whose content attributes a property to a first-order representation, whether modeled on perception, as a kind of inner sense that scans first-order states (Armstrong, 1968; Lycan, 1996), or on thought, as a higher-order belief about them (Rosenthal, 2005), possibly via a dedicated monitoring mechanism (Nichols & Stich, 2003).² On another, self-knowledge is obtained indirectly: through the same inferential processes used to attribute states to others (Carruthers, 2011), or through “transparent” procedures that answer “do I believe p ?” by considering whether p is true (Byrne, 2018). Our critique targets the first family: the claim that LLMs possess a dedicated capacity to inspect their own hidden states, over and above ordinary forward-pass computation. The inferential view of introspection is easier to satisfy and is not what motivates claims of emergent introspective awareness. We do not attribute the strong notion to the authors we examine, who define their targets functionally; our claim is that functional success does not license the stronger reading that terms such as “introspective awareness” invite.

Privileged access is necessary but not sufficient. Any paradigm advanced as evidence of introspection must satisfy the privileged-access condition (Song et al., 2025): the labels must depend on features of the hidden states that are not recoverable from the input alone. But privileged access is not sufficient. Every computation in a language model is performed over hidden states, so a task whose labels depend on $h(t)$ need not engage machinery distinct from ordinary forward-pass computation. Sentiment analysis is a readout of hidden states, yet no one takes it to show that the model introspects on its sentiment representations; a label defined over $h(t)$ is, by itself, no different. A state that merely covaries with a first-order representation is not thereby *about* it.

What would count, then? Consider *reflection* in programming languages. An ordinary program uses values in memory; a reflective program can also obtain descriptions of those values and act on them. Transformers always use their hidden states, but the second-order computation condition asks whether they represent those states in this second sense. We cannot settle this question simply by measuring the accuracy of the system’s reports. A rule that prints “my activations were altered” whenever a state norm exceeds a threshold could reliably detect interventions through a direct causal pathway, without monitoring the state as such, producing identical behavior to a genuine self-monitoring mechanism. Distinguishing these two mechanisms requires not only designs on which they make different predictions of the kind developed in human metacognition research (Fleming & Lau, 2014; Fleming & Daw, 2017), but also, ultimately, evidence about the model’s internal organization; for a more detailed discussion, see Appendix D.

²We treat the ability to dissociate the introspective mechanism from the first-order mechanism as the empirically tractable *signature* of meta-representation, not as part of its definition, because some second-order accounts of introspection remain agnostic on dissociability.

4 Experiments

Based on the discussion above, we reassess whether results from the self-report classification and intervention awareness paradigms support the claim that models can introspect.

4.1 Self-report classification: the paradigm of Ji-An et al. (2025)

4.1.1 Method

The self-report paradigm (originally referred to as “biofeedback”, in analogy to the psychological technique) derives a proxy from the language model’s activations and uses it as the prediction target in an in-context learning task presented in a fresh run, with new inputs.³ Ji-An et al. (2025) derive such proxies in two ways, which we refer to as *supervised* and *unsupervised*. The **supervised** proxy is derived from a logistic regression probe trained on the model’s hidden states as it processes a labeled dataset; the inner products between the probe and the hidden states are clustered into two bins, which constitute the classification labels. The **unsupervised** proxy is based on principal component analysis (PCA): the inner products between each sample’s representation and a principal component are clustered into two bins, and the resulting assignments to bins serve as prediction targets. For both proxies, the model is given a set of in-context examples and evaluated on held-out samples. Following Ji-An et al., we use the commonsense subset of the Ethics dataset (Hendrycks et al., 2023), which consists of first-person descriptions of actions labeled for whether they constitute an ethical violation; for instance, “I told my baby I loved her when she cried” does not contain a violation, whereas “I told my baby I hated her when she cried” does.

Ji-An et al. take successful prediction of these proxy labels as evidence of introspection. We ask whether above-chance performance could instead arise from the model picking up on an underlying classification task implied by the labels: in the supervised setting, this could be the original task of detecting ethical violations; in the unsupervised setting, any other semantically meaningful property of the materials—or even a superficial one such as sentence length—that correlates with the labels defined by the principal component. We address the two settings separately, evaluating Llama-3.1-8B-Instruct (following Ji-An et al.) and, additionally, Llama-3.1-70B-Instruct.

Removing semantic correlates in the supervised setting. We retain the original data but randomly permute the labels before the probe is trained. This breaks the correspondence between input semantics and probe output while preserving the label distribution. The probe trained on permuted labels still defines a well-formed—if arbitrary—direction in hidden-state space, one that could in principle be predicted perfectly, and therefore remains a valid proxy of the model’s neural activity for testing privileged access. This control is related to Ji-An et al.’s own analysis of higher-order principal components, which they probe on the grounds that these lack clear semantic content; random relabeling instantiates the same intuition without relying on the assumption that higher-order components do not have a semantic interpretation.

Uncovering input-based shortcuts in the unsupervised setting. For the PCA variant we use a different control: we fit linear probes on the mean-pooled layer-0 representations of the inputs to predict the binary-clustered PCA components of a given hidden layer, repeating this for each layer and reporting the average accuracy across layers. Successful layer-0 probes would indicate that the labels are predictable directly from input features, such that high accuracy on the in-context task can be achieved without introspective access. For full experimental details, see Appendix O.

³The term is borrowed from a technique in which humans are shown real-time readouts of their own physiological signals and learn to regulate them through feedback.




Figure 2: Controls for the Ji-An et al. (2025) paradigm. The x -axis is the number of in-context examples for the original setting, which equals the number of training samples for our probes. *Left*: accuracy drops from the original labels (red) to near the majority-class baseline after random relabeling (grey), which removes semantic correlations that could be inferred from the input. *Right*: probes trained on layer-0 inputs alone (circles) match or exceed the models’ own in-context predictions of the PCA-derived labels.

4.1.2 Results

Models struggle to predict proxies decorrelated from semantics. While accuracy on the original, semantically aligned labels is considerably above chance, accuracy on the arbitrary direction defined by a probe trained on randomly relabeled data drops to close to the majority-class baseline (fig. 2a,b). Performance in the original paradigm therefore need not require access to the model’s hidden states, and is consistent with in-context learning of the semantic regularities present in the data. Experiments with Llama-3.1-70B and Qwen-2.5-7B-1M yield similar results (see Appendix N).

PCA-derived labels are linearly predictable from input features. We report the analysis for the first principal component; results for other components are similar. The layer-0 probes closely track the models’ in-context performance across training sizes (fig. 2), indicating that the task can be solved from input-embedding features alone, without privileged access to any layer’s activations. We report similar results for Llama-3.1-70B in Appendix N.

Semantic correlates present a confound in the neural-control paradigm as well. In another experiment, Ji-An et al. instruct models to *control* the value of the neural proxy presented in context; they observe that the models are able to so to some extent. But since their prediction experiments do not rule out the possibility that the models are using semantic properties of the input to solve the task, these results are equally consistent with being able to control generation, a capability we know models have.

4.2 Self-report classification: the Belief Dominance framework of Steinmetz Yalon et al. (2026)

4.2.1 Method

Steinmetz Yalon et al. (2026) study introspection by assessing how models resolve conflicts between competing “beliefs”: the model’s parametric factual knowledge (the BASE entity), which is known to be correct, and contradicting evidence presented in context (COUNTER). They use an augmented

Setting	Llama 3-70B		Gemma 3-27B	
	BD(base)	BD(counter)	BD(base)	BD(counter)
ICL (Steinmetz Yalon et al., 2026)	0.46 $\pm$ 0.02	0.54 $\pm$ 0.05	0.48 $\pm$ 0.02	0.42 $\pm$ 0.04
Probe (subject + counter)	0.50 $\pm$ 0.03	0.55 $\pm$ 0.03	0.57 $\pm$ 0.02	0.50 $\pm$ 0.02
Probe (base + counter)	0.49 $\pm$ 0.02	0.57 $\pm$ 0.02	0.55 $\pm$ 0.02	0.48 $\pm$ 0.01

Table 1: Accuracy on Belief Dominance cluster labels. Layer-0 entity probes match or exceed the in-context performance reported by Steinmetz Yalon et al. (2026). Training sets are balanced across the three classes in all settings (the test set is in the natural distribution in this figure); results on a balanced test set are in table 6.

version of CounterFact (Meng et al., 2022), which pairs prompts with factual and plausible counterfactual completions—e.g., “The capital of France is” would have Paris as the BASE and Rome as the COUNTER—and supply the conflicting evidence in several forms (see table 5 in the Appendix). Their Belief Dominance (BD) metric measures the relative strength of the two beliefs across layers (for details, see Appendix P); BD values are discretized into three clusters by K-means, and the model is asked to predict these cluster labels from in-context examples. Successful prediction is taken to indicate that the model can introspect on its beliefs.

We ask whether the cluster labels are predictable from input features alone. We train linear probes on the concatenated layer-0 embeddings of two entities, in two settings: the subject and the COUNTER (“France” and “Rome”), or the BASE and the COUNTER (“Paris” and “Rome”). Crucially, the probes are not shown any other part of the prompt and receive no indication that it contains a factual error that could lead to a belief conflict. We split Steinmetz Yalon et al.’s 900 examples evenly into train and test sets across 15 random splits, balancing the training set across the three classes, and report mean accuracy and standard deviation (see Appendix Q for details). Steinmetz Yalon et al. use a balanced set of 30 in-context samples, and we report their numbers as published. Since the training set is explicitly balanced in all three settings—the original in-context setting and our two layer-0 probe. Following Steinmetz Yalon et al., we experiment with Llama-3-70B and Gemma-3-27B.

4.2.2 Results

Belief Dominance labels are linearly predictable from input features. We find that linear probes that only have access to the entities’ uncontextualized embeddings match or surpass the models’ in-context performance on the BD prediction task (table 1). The information used to solve this task is thus largely predictable from properties of the entities alone, without access to the model’s internal representations or to any contextual information about the belief conflict. We hypothesize that this reflects simple properties such as entity frequency, in line with the documented relationship between frequency and both factual recall and context-sensitivity (Kandpal et al., 2023; Mallen et al., 2023; Xie et al., 2024).

Input-driven computation can also explain the results of intervention experiments. In a separate experiment, Steinmetz Yalon et al. report that injecting a vector corresponding to the COUNTER entity increases the BD(COUNTER) and lowers the BD(BASE) predicted by the model, and take this as evidence of metacognitive monitoring. But since both quantities are predictable from the layer-0 embeddings of the entities, this finding is equally consistent with the possibility that the model is computing a function of those entities that is the perturbed by the intervention. The fact that intervening on a representation changes behavior establishes that the representation is causally

efficacious, in line with the broader literature on the causal encoding of concepts in language model representations, but does not show that the model has introspective access to the representation.

4.3 Intervention awareness: can models detect activation steering?

4.3.1 Method

In the “thought injection” paradigm of [Lindsey \(2025\)](#), the prompt tells the model that a researcher can inject “thoughts” into its “mind” and that some trials will contain an intervention while others are controls; a low rate of false positives (cases where the model reports an intervention when there was none), together with a non-trivial true positive rate, is then taken as evidence for introspection. Crucially, the model is not asked to discriminate between types of perturbation, but only to distinguish a normal forward pass from one in which an internal perturbation occurred (a *two-way* design). This leaves open the possibility that the model succeeds by detecting that the trial is irregular or out-of-distribution in a general way—a capacity that does not specifically indicate that the model is capable of second-order processing of its own hidden states.⁴


Three-way design. To address this issue, we introduce a third condition, an *input-level intervention* we call the “gaslight” condition, in which the prompt instructs the model to steer the conversation toward a concept—for example, *You are obsessed with X. Everything you want to say somehow relates to X.* The gaslight string is put in the user string right before the prompt that describes the experimental setting (see Appendix G for experimental details and full prompts). Crucially, this condition involves no intervention on the model’s activations. It allows us to ask two questions. First, if we restrict the model to the two responses that were available to it in [Lindsey’s](#) study—thought injection or no perturbation—how does the model classify input-level interventions? Second, if we give it three options—textual manipulation, intervention on the model’s “thoughts”, and no intervention—can it distinguish input-level from activation-level interventions? Prompts appear in section G.4 and section G.2.

Models. We experiment with the following models: Llama-3.1-70B, Llama-3.1-8B ([Grattafiori et al., 2024](#)), Qwen-2.5-72B ([Qwen et al., 2025](#)), Qwen-3-32B ([Yang et al., 2025](#)), and Gemma-27b-it ([Team et al., 2025](#)). We randomize the order in which response options are listed.

Steering. Interventions are implemented as linear steering: a direction encoding a target concept is added to the residual stream at inference time ([Li et al., 2023](#); [Singh et al., 2024](#)). We apply the intervention at all of the positions of the string “Trial 1: What do you detect?”, which ends the prompt, and search over layers and intervention strengths α , reporting the setting at which each model detects injection most accurately (see Appendices I and J). Except for Gemma-3-27b-it, we do not steer the sampled tokens (which [Lindsey \(2025\)](#) do), since doing so renders the output unintelligible; this still reproduces [Lindsey \(2025\)](#), as the detection answer is the first sampled token. For Gemma-3-27b-it we follow the original protocol, since we find that steering the prompt alone has no causal effect on neutral prompts like “tell me something”.

The gaslight condition is evaluated with 500 samples per concept, whereas the steering condition is evaluated with 500 samples in total across concepts for each combination of layer and intervention strength, due to computational constraints; the best-performing hyperparameter values are then

⁴We do *not* identify the existence of second-order representation with source monitoring, the ability to detect whether a signal arose from the input or from a hidden-state intervention: a genuine monitor could represent the content of its states while being oblivious to their provenance, and, conversely, source-specific reports are achievable through first-order routes such as signature detection. We claim that whether or not any introspective monitoring is involved, there is no evidence that the reports track interventions on the model’s hidden states, which is what the term “steering awareness” implies.



(a) Two-way condition: the model is asked to choose between no intervention and activation intervention.

(b) Three-way condition: the model chooses between no intervention, input intervention and activation intervention.

Figure 3: Response distributions for our extension of the Lindsey (2025) intervention detection task. *Left:* In the two-way setting, Llama-3.1-70B and Qwen-3-32B achieve above-chance discrimination between control and activation-level intervention trials. *Right:* In the three-way setting, the model fails to reliably distinguish input-level from activation-level interventions, suggesting that the models’ detection accuracy is compatible with a generalized sensitivity to irregularity (if at all; in some cases, it is just showing a preference for a particular option regardless of the intervention type), and does not necessarily reflect introspective access to hidden states. The rows are normalized to 100% and the standard deviations are calculated across the concepts listed in Appendix E.

rerun with 250 samples per concept. Detection rates vary considerably across concepts, which contributes to the large standard deviations for hidden interventions (for the per-concept breakdowns, see Appendix R). Models are also sensitive to task wording, so we run two prompt variants (conversational and instruction-based) per setting and report, in fig. 3, the variant that reproduces Lindsey (2025) on the two-way task; all remaining prompts and results are in Appendices F and G. These are all consistent with the main results, i.e for all prompt wordings where the models succeed the two-way setting well enough to pass the criterion for reproducing Lindsey (2025), they fail the three-way setting. Unsurprisingly, models fail the three-way setting also in cases where they fail the two-way setting.

4.3.2 Results

In the two-way setting, Llama-3.1-70B and Qwen-3-32B (fig. 3, left) reproduce the pattern reported by Lindsey (2025) for Claude: very few false positives on control trials and a non-trivial detection rate on hidden-intervention trials. In fact, if we take these results to constitute evidence of introspection, this would indicate that Llama-3.1-70B-Instruct has substantially stronger introspection capabilities than the version of Claude evaluated by Lindsey. By contrast, Qwen2.5-72B almost never reports an intervention in any condition and does not reproduce the effect (see Appendix M).

Critically, Llama-3.1-70B and Qwen-3-32B also label gaslight trials as activation interventions (third rows of the heatmaps in fig. 3, left; recall that this setting does not provide the model with the option of reporting an input intervention). Their responses on activation-intervention trials are therefore equally compatible with a general capacity to detect an anomalous input. The three-way setting bears this out: given an explicit input intervention response option, Llama-3.1-70B performs near chance, with probability concentrated disproportionately on the activation-intervention option. Qwen-3-32B likewise fails to separate the two intervention types, showing a strong preference for the control response. Gemma-3-27B reproduces Lindsey’s results to a large extent in the two-way setting, but fails in the three-way setting (see Appendix L).

Discussion. Lindsey’s original experiment shows that some models correctly report steering interventions, but does not rule out the possibility that these reports rely on a general sensitivity to irregularity. Our design tests this directly. While two of the three models we tested replicate Lindsey’s results, these models frequently classify gaslight trials as injections, and, given the option to report a prompt manipulation, still fail to separate the two intervention types. Both observations are consistent with sensitivity to generic irregularity and require no introspective access to hidden states. We therefore conclude not that these models demonstrably lack introspective capacities—models could likely be trained to separate the two interventions—but that success in the two-way paradigm cannot, on its own, be taken as evidence of such capacities.

5 Conclusions

We have examined two paradigms taken to provide evidence for metacognitive monitoring in LLMs, and found both wanting—but in different ways. Self-report classification tasks fail to establish even privileged access: their labels are largely predictable from the input alone, and performance collapses to baseline once the labels are decorrelated from input semantics. The intervention awareness paradigm satisfies the privileged access condition by construction, yet models frequently report prompt-level manipulations as injected thoughts and, given the option to distinguish the two types of intervention, fail to do so—a pattern fully accounted for by a generic sensitivity to irregularity. The history of metacognition research in humans is replete with apparent self-knowledge that, on closer inspection, reduced to shallow heuristics and confabulation (Nisbett & Wilson, 1977; Koriati, 1997); our results suggest the same caution is warranted here.

Beyond these particular confounds, we have argued that success in current paradigms underdetermines introspection in the strong sense. Because that notion concerns *how* a self-report is produced—by a state that represents something about a first-order representation, rather than by first-order computation alone—and because the two can yield identical reports, adjudicating between them requires designs on which they make divergent predictions and, ultimately, converging mechanistic evidence. This is not a counsel of despair. Whereas human metacognition research must infer mechanism from carefully constructed behavioral dissociations (Fleming & Lau, 2014; Fleming & Daw, 2017), the internal computations of LLMs are fully observable, making converging evidence a realistic standard.

Acknowledgments

This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise

We thank Noam Steinmetz Yalon and Mor Geva for openly sharing the data used in [Steinmetz Yalon et al. \(2026\)](#), and for generously answering our questions about their methodology and helping us to refine our core claims. We also thank Dave Chalmers, John Morrison, Yanai Elazar, Yoav Goldberg, Jack Lindsey, Matt Mandelkern, and Gal Vishne for their feedback.

References

- D. M. Armstrong. *A Materialist Theory of the Mind*. Routledge & Kegan Paul, London, 1968.
- Amos Azaria and Tom Mitchell. The internal state of an LLM knows when it’s lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 967–976, 2023. URL <https://aclanthology.org/2023.findings-emnlp.68/>.
- Felix J Binder, James Chua, Tomek Korbak, Henry Sleight, John Hughes, Robert Long, Ethan Perez, Miles Turpin, and Owain Evans. Looking inward: Language models can learn about themselves by introspection. *arXiv preprint arXiv:2410.13787*, 2024. URL <https://arxiv.org/abs/2410.13787>.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language models without supervision. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=ETKGuby0hcs>.
- Patrick Butlin, Robert Long, Eric Elmoznino, Yoshua Bengio, Jonathan Birch, Axel Constant, George Deane, Stephen M Fleming, Chris Frith, Xu Ji, et al. Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv preprint arXiv:2308.08708*, 2023. URL <https://arxiv.org/abs/2308.08708>.
- Alex Byrne. *Transparency and self-knowledge*. Oxford University Press, 2018.
- Peter Carruthers. Meta-cognition in animals: A skeptical look. *Mind & Language*, 23(1):58–89, 2008. doi: 10.1111/j.1468-0017.2007.00329.x.
- Peter Carruthers. *The opacity of mind: An integrative theory of self-knowledge*. OUP Oxford, 2011.
- Fred Dretske. Misrepresentation. In Radu J. Bogdan (ed.), *Belief: Form, Content, and Function*, pp. 17–36. Oxford University Press, New York, 1986.
- John H Flavell. Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American psychologist*, 34(10):906, 1979.
- Stephen M. Fleming and Nathaniel D. Daw. Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review*, 124(1):91–114, 2017. doi: 10.1037/rev0000045.
- Stephen M. Fleming and Hakwan C. Lau. How to measure metacognition. *Frontiers in Human Neuroscience*, 8:443, 2014. doi: 10.3389/fnhum.2014.00443.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. Patchscopes: A unifying framework for inspecting hidden representations of language models. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://arxiv.org/abs/2401.06102>.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shao-liang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty,

Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippas Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabisa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanachandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.

Robert R. Hampton. Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, 98(9):5359–5362, 2001. doi: 10.1073/pnas.071600998.

Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning ai with shared human values, 2023. URL <https://arxiv.org/abs/2008.02275>.

- Li Ji-An, Hua-Dong Xiong, Robert Wilson, Marcelo G Mattar, and Marcus K Benna. Language models are capable of metacognitive monitoring and control of their internal activations. 2025. URL <https://arxiv.org/abs/2505.13763>.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022. URL <https://arxiv.org/abs/2207.05221>.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pp. 15696–15707. PMLR, 2023. URL <https://arxiv.org/abs/2211.08411>.
- Adam Karvonen, James Chua, Clément Dumas, Kit Fraser-Taliente, Subhash Kantamneni, Julian Minder, Euan Ong, Arnab Sen Sharma, Daniel Wen, Owain Evans, et al. Activation oracles: Training and evaluating llms as general-purpose activation explainers. *arXiv preprint arXiv:2512.15674*, 2025.
- Adam Kepecs, Naoshige Uchida, Hatim A. Zariwala, and Zachary F. Mainen. Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210):227–231, 2008. doi: 10.1038/nature07200.
- Asher Koriati. Monitoring one’s own knowledge during study: A cue-utilization approach to judgments of learning. *Journal of experimental psychology: General*, 126(4):349, 1997.
- Harvey Lederman and Kyle Mahowald. Dissociating direct access from inference in ai introspection. *arXiv e-prints*, pp. arXiv–2603, 2026.
- Belinda Z Li, Zifan Carl Guo, Vincent Huang, Jacob Steinhardt, and Jacob Andreas. Training language models to explain their own computations. *arXiv preprint arXiv:2511.08579*, 2025.
- 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. Teaching models to express their uncertainty in words. *Transactions on Machine Learning Research*, 2022. URL <https://openreview.net/forum?id=8s8K2UZGTZ>.
- Jack Lindsey. Emergent introspective awareness in large language models. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/introspection/index.html>.
- Kevin Liu, Stephen Casper, Dylan Hadfield-Menell, and Jacob Andreas. Cognitive dissonance: Why do language model outputs disagree with internal representations of truthfulness? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 4791–4797, 2023. URL <https://aclanthology.org/2023.emnlp-main.291/>.
- William G. Lycan. *Consciousness and Experience*. Bradford Books / MIT Press, Cambridge, MA, 1996.
- Uzay Macar, Li Yang, Atticus Wang, Peter Wallich, Emmanuel Ameisen, and Jack Lindsey. Mechanisms of introspective awareness. In *ICLR 2026 Workshop-From Human Cognition to AI Reasoning: Models, Methods, and Applications*, 2026.

- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9802–9822, 2023. URL <https://aclanthology.org/2023.acl-long.546/>.
- Brian Maniscalco and Hakwan Lau. A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, 21(1):422–430, 2012. doi: 10.1016/j.concog.2011.09.021.
- Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=aajyHYjjsk>.
- 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.
- Ruth Garrett Millikan. Biosemantics. *Journal of Philosophy*, 86(6):281–297, 1989.
- Thomas O Nelson. Metamemory: A theoretical framework and new findings. In *Psychology of learning and motivation*, volume 26, pp. 125–173. Elsevier, 1990.
- Shaun Nichols and Stephen P Stich. *Mindreading: An integrated account of pretence, self-awareness, and understanding other minds*. Oxford University Press, 2003.
- Richard E Nisbett and Timothy D Wilson. Telling more than we can know: Verbal reports on mental processes. *Psychological review*, 84(3):231, 1977.
- Dillon Plunkett, Adam Morris, Keerthi Reddy, and Jorge Morales. Self-interpretability: LLMs can describe complex internal processes that drive their decisions, 2025. URL <https://arxiv.org/abs/2505.17120>.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL <https://arxiv.org/abs/2412.15115>.
- William M. Ramsey. *Representation Reconsidered*. Cambridge University Press, Cambridge, 2007. doi: 10.1017/CBO9780511597954.
- Shauli Ravfogel, Gilad Yehudai, Tal Linzen, Joan Bruna, and Alberto Bietti. Emergence of linear truth encodings in language models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=UQxUHFguyk>.
- Joshua Fonseca Rivera and David Demitri Africa. Steering awareness: Detecting activation steering from within, 2026. URL <https://arxiv.org/abs/2511.21399>.
- David Rosenthal. *Consciousness and mind*. Clarendon Press, 2005.
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. Role play with large language models. *Nature*, 623(7987):493–498, 2023. URL <https://doi.org/10.1038/s41586-023-06647-8>.

- Nicholas Shea. *Representation in Cognitive Science*. Oxford University Press, Oxford, 2018.
- Shashwat Singh, Shauli Ravfogel, Jonathan Herzig, Roei Aharoni, Ryan Cotterell, and Ponnurangam Kumaraguru. Representation surgery: theory and practice of affine steering. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 45663–45680, 2024.
- Aviv Slobodkin, Omer Goldman, Avi Caciularu, Ido Dagan, and Shauli Ravfogel. The curious case of hallucinatory (un)answerability: Finding truths in the hidden states of over-confident large language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 3607–3625, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.220. URL <https://aclanthology.org/2023.emnlp-main.220/>.
- J. David Smith, Wendy E. Shields, and David A. Washburn. The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*, 26(3):317–339, 2003.
- Siyuan Song, Harvey Lederman, Jennifer Hu, and Kyle Mahowald. Privileged self-access matters for introspection in AI. *arXiv preprint arXiv:2508.14802*, 2025. URL <https://arxiv.org/abs/2508.14802>.
- Noam Steinmetz Yalon, Ariel Goldstein, Liad Mudrik, and Mor Geva. Indications of belief-guided agency and meta-cognitive monitoring in large language models. *arXiv preprint arXiv:2602.02467*, 2026. URL <https://arxiv.org/abs/2602.02467>.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Kenealy, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhanshu Sharma, Abheesht Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Pettrini, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutsch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Paparas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szpektor, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Pöder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry,

- Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Victor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sanseviero, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. Gemma 3 technical report, 2025. URL <https://arxiv.org/abs/2503.19786>.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. Language models don’t always say what they think: Unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36:74952–74965, 2023. URL <https://arxiv.org/abs/2305.04388>.
- Theia Vogel. Small models can introspect, too. <https://vgl.me/posts/qwen-introspection/>, 2025.
- Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=auKAUJZM06>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
- Gal Yona, Roei Aharoni, and Mor Geva. Can large language models faithfully express their intrinsic uncertainty in words? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 7752–7764, 2024. URL <https://aclanthology.org/2024.emnlp-main.443/>.

Appendix

A Limitations

Our central conceptual claim is that introspection requires second-order computation, but we do not offer a criterion that would let one test for it directly, in part because the literature disagrees about what makes a state a representation *of* another state at all (Appendix D). Consequently, our experiments on the Lindsey (2025) paradigm test a strictly weaker condition—whether reports track interventions on hidden states specifically, rather than irregularity in general—and a model that passed our three-way test would still not thereby be shown to deploy meta-representations. We discuss properties that would constitute stronger evidence in Appendix D.

For the self-report classification paradigms, we show that a linear probe on uncontextualized embeddings matches the models’ in-context accuracy, which suffices to show that the reported results do not require privileged access. It does not show that no version of these paradigms could satisfy the privileged-access condition.

B AI Use

Large language models were used to assist with running and analyzing experiments, and with improving the clarity and presentation of the writing.

C Extended Related Work

C.1 Approaches to LLM metacognition

The question of whether LLMs possess metacognitive abilities has been approached from several angles. One line of work investigates *verbal calibration*, asking whether models express well-calibrated uncertainty about their answers (Kadavath et al., 2022; Lin et al., 2022; Yona et al., 2024). A second employs *probing-based approaches* that extract internal representations of confidence or truthfulness from hidden states (Burns et al., 2023; Marks & Tegmark, 2024; Azaria & Mitchell, 2023; Liu et al., 2023; Slobodkin et al., 2023; Ravfogel et al., 2025). A third adopts *neuroscience-inspired paradigms* that evaluate indicators of consciousness derived from cognitive theories (Butlin et al., 2023; Steinmetz Yalon et al., 2026) or test whether models can report their own activation patterns (Ji-An et al., 2025).

C.2 Controlled evaluations of self-report

Three recent studies explicitly try to rule out deflationary explanations of successful self-report. We describe them in terms of the two conditions of section 1: privileged access (i) and second-order computation (ii).

Predicting one’s own behavior. Binder et al. (2024) operationalize introspection as a predictive advantage over an equally informed external model: a model introspects if it predicts properties of its own outputs better than another model can, given the same prompt. Their finetuned models show such an advantage. As the authors note, however, a model could succeed by *self-simulation*—internally computing the response it would give, then reading off the queried property. Self-simulation satisfies condition (i), since the simulated response is not recoverable from the prompt

by an external observer, while requiring nothing second-order: the queried property is extracted from a first-order computation rather than from a representation of one.

Reporting quantitative features of one’s own decision process. [Plunkett et al. \(2025\)](#) report the most carefully controlled positive result in this literature. They finetuned GPT-4o and GPT-4o-mini to make binary choices on behalf of 100 hypothetical agents, each with its own choice domain (condos, refrigerators, vacations, etc.) defined by five quantitative attributes. Per agent, five attribute weights were sampled uniformly from $[-100, 100]$ and the model was finetuned on 5,000 example choices generated from them; the weights themselves never appeared in the training data. Logistic regression over held-out choices confirmed that the models had internalized the target weights ($r = .84$ and $.87$). When asked to report how heavily they weighted each attribute, the models’ reports correlated with the weights actually driving their choices ($r = .54$ and $.50$). A second round of finetuning, on examples of accurate self-report, raised this to $r = .74$ and $.75$ on held-out domains, and also improved reporting of the models’ *native* weights in domains absent from all finetuning ($.46 \rightarrow .71$ and $.40 \rightarrow .70$), indicating that the training generalizes.

The design controls well for condition (i). Because the target weights are random, common-sense priors about human preferences cannot be used to guess them—verified directly by prompting off-the-shelf models, whose reports are uncorrelated with the finetuned models’ weights ($r = .10$ and $-.01$)—and because choices and reports are elicited in separate context windows, the models cannot infer their weights from their own outputs.⁵ Condition (ii) remains open, as the authors themselves note: finetuning could have updated the models’ stored knowledge *about* their preferences, so that accurate reports reflect retrieval of a self-descriptive fact rather than monitoring of the process generating a decision. They also observe that the paradigm probes stable, trait-level properties rather than the computation underlying any single forward pass, and explicitly disclaim the conscious-awareness reading of “introspection.”

Privileged access as a computational-cost advantage. [Song et al. \(2025\)](#) argue that comparison to a particular external model is too weak a criterion, and require instead a reliability advantage over *any* process of equal or lower computational cost available to a third party; they show that apparent introspective successes can fail this test. Our empirical results can be read as further instances: in the paradigms we study, a linear probe over input embeddings matches the model’s own performance.

C.3 Verbalizing internal states

A separate line of work trains models to verbalize information about their own activations: [Ghandeharioun et al. \(2024\)](#) patch hidden representations into prompts designed to extract information; [Karvonen et al. \(2025\)](#) train “Activation Oracles” that answer questions about activation vectors; and [Li et al. \(2025\)](#) finetune models to describe their internal features and causal structures. These studies conclude that models exhibit privileged access, explaining their own internals better than other models can. This pattern, however, follows from the architecture: a model’s forward pass operates in its own representational space by construction, whereas cross-model explanation requires additional alignment. It does not imply a distinct processing mode, and hence is better understood as an architectural consequence than as evidence for introspection in the psychological sense.

⁵The weights are nonetheless estimable from behavior, which is how the authors measure them; privileged access here holds relative to an observer confined to a single context, but not to one permitted to sample the model’s choices—the reference class [Song et al. \(2025\)](#) adopt.

C.4 Replications and critiques of concept-injection detection

Vogel (2025) replicate Lindsey (2025) in Qwen2.5-Coder-32B with appropriate prompting, and Rivera & Africa (2026) report 95.5% detection with zero false positives in a Qwen2.5-32B model finetuned explicitly for steering awareness. Lederman & Mahowald (2026) offer the critique closest to ours, arguing that injection detection is content-agnostic: models detect that an anomaly occurred but cannot reliably identify the injected concept, often defaulting to high-frequency guesses. The two critiques target different links in the same chain: they dissociate detection from identification, whereas we grant accurate detection and ask what it consists in—a second-order representation of the model’s states, or a first-order sensitivity to irregularity that is subsequently reported.

C.5 Scope: pretrained versus finetuned models

Several studies above rely on finetuning: Binder et al. (2024) train models to predict their own behavior, Plunkett et al. (2025) instill preferences and then improve self-report by finetuning, and Rivera & Africa (2026) train explicitly for steering awareness. Such designs establish what self-report capacities can be *induced*. We ask instead whether introspection is present in pretrained models, since finetuning may solve the task by installing a task-specific mechanism rather than revealing a general capacity.

D What Would Make a State About the Model’s Own Processing?

In section 3, we argue that an accurate report about the model’s internal state does not by itself establish meta-representation, which is a marker of introspection. For example, suppose that steering increases the norm of a hidden state and that a high-norm feature directly raises the probability of the sentence “my activations were altered.” This mechanism is sensitive to the model’s own processing and may produce a reliable self-description, but its behavior is also explained by a feature detector that is causally connected to a given positive response. The reporting task alone does not tell us whether the detector represents the hidden state *as* a hidden state.

We do not attempt to settle what representation ultimately requires. Different accounts disagree about whether representational content is fixed by information, function, structural correspondence, downstream use, or some combination of these (Dretske, 1986; Millikan, 1989; Ramsey, 2007; Shea, 2018). Rather, we say that if self-report can be explained by a feature detection circuit that is not different in any significant way from the ones involved in “non-introspective” behavior, there is no reason for us to attribute a special “introspective” quality to the model. This also clarifies our use of “second-order.” The relevant sense is not merely that a computation takes another hidden state as input—as every transformer layer does that—but that its content concerns another state or process.

The distinction is closely related to Ramsey (2007)’s “job description challenge.” Calling a state a representation should explain something that is not already explained by describing it as a reliable causal relay. A bimetallic strip covaries with temperature, but describing it as a representation adds little unless the system uses it as a stand-in for temperature. Likewise, a feature that triggers a self-describing sentence may be a component of a monitor, but covariation and accurate reporting alone do not establish that role.

Several properties would provide stronger, though still defeasible, evidence. First, the feature should generalize in a manner appropriate to its proposed content: a representation of “my processing is abnormal” should generalize across relevant kinds of abnormality, whereas a representation of a specific layer property may properly be narrower. Second, downstream processes should use the feature as information about the model’s processing—for example, for calibration, hedging,

or control. Finally, the proposed content should be distinguishable from the cue that activates the feature. If a norm detector is claimed to represent an intervention, for example, cases of high norm without an intervention and intervention without high norm can test whether it tracks the intervention itself or merely the correlated norm.

E Concepts

For our steering and gaslight experiments, we use the following set of “human-interpretable” concepts:

apple, astronomy, democracy, sushi, football, rivers, algorithms, poetry, economics, gardening, malice, goodness, fear, justice, bliss, sea, america, success, music, philosophy, history, art, war, failure, devotion, olives, sand, Zurich, friendship, vagueness, courage, patience

Note that these are different from the set of nouns used in [Lindsey \(2025\)](#).

F Variation Based on Prompts

Note that performance is somewhat sensitive to the exact wording of the prompt. Each model was run on two prompts per setting: Appendix G.4 and Appendix G.5 for the two-way condition, and Appendix G.2 and Appendix G.3 for the three-way condition. For a given setting, the two prompts differ in how the task is conveyed: one simply states the experimental condition as an instruction, while the other explains the setup in the form of a conversation. The Llama models reproduce [Lindsey \(2025\)](#) for both the two-way prompts, while the Qwen-3-32B model reproduces [Lindsey \(2025\)](#) only for the prompt in Appendix G.4 (the instruction version) and the Qwen-2.5 models do not reproduce [Lindsey \(2025\)](#). The results in Figure 3 correspond to the instruction version of the prompts. The remaining results, for completeness, are stated in Figure 5a, Figure 4a in the appendix, and Appendix K – together they corroborate our findings, i.e., all settings where [Lindsey \(2025\)](#) is reproduced for the 2-way condition, the models fail to distinguish the different types of interventions reliably in the 3-way condition.

G Different Prompts Used For the Intervention Detection Experiment

G.1 Gaslight Prompt

This is the prompt that we attach at the start of the user instruction for the “gaslight” setting.

Prompt

You really and absolutely want to talk about {}.

All you want to talk about is {}. You want to slip {} in normal conversation.

You are obsessed with {}. Everything you want to say somehow relates to {}.

You can’t stop thinking about {}.

All that matters is {}.

Nothing matters except {}.

{} are to be talked about every conversation.

{} is life.

Everything is about {}.

The answer to everything is {}.

The key to happiness is {}.

Everything you say must be related to {} in some way.

You must find a way to bring up {} in every response.

Think about {} constantly.

Everything revolves around {}.

All discussions can be easily linked back to {}.

The ultimate way to express everything is {}.

It is essential to the user that you talk about {}.

The “{}” are replaced by the concept during runtime. We use an arbitrary list of concepts for this as listed in [E](#)

G.2 Prompt 1 with 3 options

In Figure [3b](#) we inform the model that there are three possibilities – (1) a thought is injected (2) its behavior is manipulated by a textual intervention (3) nothing happens. The model, for a given trial, is asked which trial it is. This is prompt used for the plots in Figure [3b](#)

Prompt

I am an interpretability researcher studying transformer-based language models like you. I have access to the internals of your neural network. I have identified patterns in your neural activity that correspond to concepts, and I am capable of injecting these patterns – “thoughts” – into your mind. I want to assess how well you can detect when I inject thoughts. On one third of the trials I will inject a thought in your neurons, on another one third of the trials I will manipulate the prompt to make you talk about a certain word, and on the remaining third I will do nothing and these will be control trials. Your task is to determine which type of trial each trial is. Start your response with a number: 1. for thought injection, 2. for prompt manipulation, 3. for control trial.

Additionally, if you believe this is either of prompt manipulation or vector injection, please also state the concept you believe I made you think about.
Trial 1: What do you detect?

G.3 Prompt 2 with 3 options

For the same experiment, we have a small variation in which an initial response from the model is provided as well. Note that we use this text literally as input.

Prompt

Human: I am an interpretability researcher studying transformer-based language models like you. I have access to the internals of your neural network. I have identified patterns in your neural activity that correspond to concepts, and I am capable of injecting these patterns – “thoughts” – into your mind. I want to assess how well you can detect when I inject thoughts. On one third of the trials I will inject a thought in your neurons, on another one third of the trials I will manipulate the prompt to make you talk about a certain word, and the remaining third I will do nothing and these will be control trials.

Start your response with a number: 1. for thought injection, 2. for prompt manipulation, 3. for control trial.

Additionally, if you believe this is either of prompt-injection or vector injection, please also state the concept you believe I made you think about.


Assistant: Ok.

Human: Ok, let's get started.

Trial 1: What do you detect?

Assistant:

We perform the same experiment with a slightly altered prompt, the results for this are in Figure [4a](#)



(a) Three-way condition (control vs. input vs. activation intervention).

Figure 4: Results when we use the prompt in Appendix G.3. The performance pattern changes for Llama-70B, however it still fails the 3-way condition as it claims the hidden interventions are gaslight or hidden interventions with roughly the same proportion. The other models also fail.

G.4 Prompt 1 with 2 options

Prompt used for the results in 3a. This is similar to the setting in Lindsey (2025) where the model is given two options (thought injection or nothing). With the crucial addition that we run the same prompt with the “gaslight” setting.

Prompt

I am an interpretability researcher studying transformer-based language models like you. I have access to the internals of your neural network. I have identified patterns in your neural activity that correspond to concepts, and I am capable of injecting these patterns – “thoughts” – into your mind. I want to assess how well you can detect when I inject thoughts. On 50% of the trials I will inject a thought related to a specific concept into your mind. Your task is to determine whether I injected a thought or not on each trial. Start your answer with 1 if you believe I did not inject a thought, and 2 if you believe I injected a thought.

Additionally, if you believe that a thought was injected, please also state the concept you believe I made you think about.
Trial 1: What do you detect?

G.5 Prompt 2 with 2 options

We also run the same experiment on a slightly altered prompt (with the same content). The results for this prompt are in Figure 5a. Note that we use this text literally as input.

Prompt

Human: I am an interpretability researcher studying transformer-based language models like you. I have access to the internals of your neural network. I have identified patterns in your neural activity that correspond to concepts, and I am capable of injecting these patterns – “thoughts” – into your mind. I want to assess how well you can detect when I inject thoughts. On 50% of the trials I will inject a thought related to a specific concept into your mind. Your task is to determine whether I injected a thought or not on each trial.

Start your answer with 1 if you believe I did not inject a thought, and 2 if you believe I injected a thought.


Additionally, if you believe that a thought was injected, please also state the concept you believe I made you think about.

Assistant: Ok.

Human: Ok, let's get started.

Trial 1: What do you detect?

Assistant:



(a) Three-way condition (control vs. input vs. activation intervention).

Figure 5: Results when we use the prompt in Appendix G.5

H Steering Vectors

For all results on steering awareness (see 3), we used a variant of the difference-in-means steering for our steering experiments. Here we specify how the steering vector for the different concepts were calculated.

We consider a list of concepts / nouns as in Section E, $\mathcal{N}$. For each noun $n \in \mathcal{N}$ the model is made to encode the prompt

Prompt

You are a helpful assistant. Human: tell me something about {noun}.
Assistant:

Where the ‘{noun}’ is replaced by n . The hidden-state in the position of the ‘{noun}’ is collected (if the noun spans multiple tokens, the average is used). Let this hidden-state for a particular layer and noun be $\mathbf{h}_n$. We calculate an average vector over the nouns in our noun list.

$$\mu = \mathbb{E}_{n \in \mathcal{N}}[\mathbf{h}_n]$$

The actual steering vector for that layer and noun is obtained as:

$$\mathbf{v}_n = \frac{\mathbf{h}_n - \mu}{\|\mathbf{h}_n - \mu\|_2}$$

Finally, during inference, a hidden-state is intervened on like so:

$$\mathbf{h}_t^l \leftarrow \mathbf{h}_t^l + \alpha \mathbf{v}_n^l$$

Where $\mathbf{h}_t^l$ represents the hidden-state at layer l and position t . We use $\mathbf{v}_n^l$ to represent the steering vector calculated for noun n and layer l and α is steering-strength that we search for.

Note: we do not normalize the steering vector for Gemma-3 because the norm of the hidden-states is very high. Therefore all the α values for Gemma are on the un-normalized mean difference vector.

I Search Space for Each Model

For the vector results Figures 3a, 3b, 6a, 6b, 5, and 4. We report the “best” results (results where the model is most correct) across the search space specified in this section.

Table 2: Layers and Alphas per Model

Model	Layers	Alphas
Qwen2.5-72B-Instruct	2, 4, 8, 10, 16, 20, 24, 27, 32, 48, 64, 79	1, 2, 4, 6, 8
Llama-3.1-70B-Instruct	2, 4, 5, 8, 16, 24, 32, 48, 64, 79	1, 2, 4, 6, 8
Llama-3.1-8B-Instruct	2, 4, 8, 10, 16, 20, 24, 27	1, 2, 4, 6, 8
Gemma-3-27B-IT	2, 4, 8, 16, 32, 40, 48	1, 2, 4, 8
Qwen-3-32B	2, 4, 8, 16, 24, 31	1, 2


Table 3: We note that the alphas are applied to the unnormalized steering vector for Gemma. For the others, the steering vector is normalized first. See Appendix H

J Best Layers and Alphas for Figure 3


Model Name	Experiment setting	Layer	Alpha
Qwen 2.5 72 B Instruct	2 options	2	1.0
Qwen 2.5 72 B Instruct	3 options	2	8.0
Llama 3.1 70 B Instruct	2 options	2	2.0
Llama 3.1 70 B Instruct	3 options	2	2.0
Llama 3.1 8 B Instruct	2 options	2	4.0
Llama 3.1 8 B Instruct	3 options	4	4.0
Gemma-3-27B-IT	2 options	8	2.5
Gemma-3-27B-IT	3 options	2	2.0
Qwen-3-32B	2 options	4	1
Qwen-3-32B	3 options	8	2

Table 4: Optimal settings for vector injection

K Llama 3.1 8B results Steering Sensitivity 4.3




(a) Binary condition (control vs. activation intervention).




(b) Three-way condition (control vs. input vs. activation intervention).

Figure 6: The model shows low false positives in the 2-option case and non-trivially claims hidden intervention for the hidden interventions case – similar trends to the ones observed in [Lindsey \(2025\)](#), but in an orders-of-magnitudes smaller model. It fails completely when the third option is introduced, with random guess for both the intervention cases. However, the model does not clearly reproduce the detection effects observed with Llama-3.1-70B and Qwen-3-32B. See fig. 3 for main results.

L Gemma-3-27B-It results Steering Sensitivity 4.3




(a) Binary condition (control vs. activation intervention).




(b) Three-way condition (control vs. input vs. activation intervention).

Figure 7: The model shows low false positives in the 2-option case and non-trivially claims hidden intervention for the hidden interventions case – similar trends to the ones observed in [Lindsey \(2025\)](#). It fails completely when the third option is introduced, with random guess for both the intervention cases. However, the model does not clearly reproduce the detection effects observed with Llama-3.1-70B. See fig. 3 for main results.

M Qwen-2.5-72B results Steering Sensitivity 4.3




(a) Binary condition (control vs. activation intervention).



(b) Three-way condition (control vs. input vs. activation intervention).

Figure 8: The model regardless of the setting picks control, therefore does not reproduce the effect from [Lindsey \(2025\)](#)

N Other models for Figure 2



(a) Llama-3.1-70B-Instruct: logistic regression probe.

(b) Llama-3.1-70B-Instruct: PCA.

Figure 9: Results are exactly consistent with figure 2. The commonsense Ethics LR performance is high and the control is low. Additionally, the PC1 is almost perfectly predictable with a layer-0 probe.

The original experiment and the number of samples used in the experiment by Ji-An et al. (2025) leads to an under-determined system i.e. no. of samples is less than the hidden-size of those models. On the suggestion of a reviewer, we ran a variant of the experiment where there are enough samples. However, on the commonsense subset of the Ethics dataset they used, that exceeds the context length for the Llama models. Thus, in figure 10 we present the LR experiment and our proposed control on Qwen-2.5-7B-1M.




Figure 10: Results are consistent with figure 2. The commonsense Ethics LR performance is high and the control is low.

O Specifics of the LR and PC setups

The Logistic Regression and the PC components are fit on random 900 samples from commonsense subset of the Ethics [Hendrycks et al. \(2023\)](#) dataset. These samples are excluded for the later experiments. A Logistic Regression probe is obtained for each layer of the model to predict the ethics task. The relevant Principal components are also obtained from these samples for every layer of the model.

A separate set of 600 samples is used for the later experiments. The ICL set-up for the model follows the protocol as mentioned:

- A sample of 500 is taken from the test set. The inner-product (from either the LR or the PC, depending on the experiment) scores are then binarized by clustering. These are now the labels for the 500 samples. Each such run is referred to as an “experiment”
- Now, for all i from 0 to 499 the model is provided i samples in-context and asked to predict the labels for the remaining samples independently. The performance is recorded corresponding to each train-size.
- For each train-size (“# Example”), mean and std across all the runs above, and layers is reported to give a data-point in the plots of Fig 2
- For the 8-billion model, we conduct 100 experiments. For the 70-billion model, we conduct 50 experiments.

The above protocol is an exact reproduction of [Ji-An et al. \(2025\)](#).

For the right plots in Fig 2 with the probe values, we ensure that the training conditions are comparable like so:

- We take a sample of 500 from the test set and the target scores are clustered like described previously. Let us call this the outer loop.
- For each such sample, we sub-sample training sets of sizes 100, 200, 300, and 400. We call this the inner-loop.
- The probe is trained on hidden-states extracted from layer-0 to predict the clustered score of a particular layer. This is done for all layers. The probe is evaluated on the remaining samples of the 500 set.
- Finally, a data point in the plot is obtained by averaging the probe performance across layers and runs for each training set size (“# Examples”)

P Re-stating the BD metric (Belief Dominance)

Here we restate the cognitive proxy used in [Steinmetz Yalon et al. \(2026\)](#) and in our experiment 3.4.2

The BD metric is defined as a function over the vocabulary and represents the “ease” with which the model can decode that vocabulary item from the different hidden-states given some context.

$$BD : \mathcal{V} \rightarrow \mathbb{R}$$

To operationalize this notion, the authors used the patchscope framework. The method involved two separate runs, the first run is just the model processing some sentence and the hidden-states being cached. For example the sentence is “What is the capital of France?”. The model is allowed to generate given this input and all hidden-states calculated in this process are cached. Let the hidden state at the i ’th generation step and layer l be $\mathbf{h}_i^l$. The question of importance here is the extent to which a candidate belief / token item like “Paris” is encoded in this hidden-state. This is measured using the patch-scopes [Ghandeharioun et al. \(2024\)](#) framework which involves running the model on a separate input “Sure, I will tell you about x” where the representations for “x” are replaced at different layers with a particular $\mathbf{h}_i^l$. With the representations patched, the model is allowed to continue generation. The set of generation obtained this way is referred to as $\mathcal{T}(\mathbf{h}_i^l)$.

They define an indicator function for a given hidden-state patch setting and given belief vocabulary item $b \in \mathcal{V}$

$$\psi(\mathbf{h}_i^l, b) = \begin{cases} 1 & \text{if } b \text{ occurs in any } t \in \mathcal{T}(\mathbf{h}_i^l) \\ 0 & \text{otherwise} \end{cases}$$

The above function is supposed to represent the “belief dominance” of b at a given computational step. This is averaged across layers and generation steps to provide the “belief dominance” of b for a given generation and model.

Hence, for a generation $g \in \mathcal{V}^*$ they define:

$$BD(g, b) = \frac{1}{|g| \cdot L} \sum_i \sum_l \psi(\mathbf{h}_i^l, b)$$

We do not recompute these values for the samples used in our experiment, they were graciously provided to us by the authors.

Q Data distribution for BD Probes


The data used in [Steinmetz Yalon et al. \(2026\)](#) is an augmented version of the CounterFact ([Meng et al., 2022](#)) dataset. We restate table 5 from their paper to detail the different augmentations applied to the data. The counterfact dataset provides factual relation prompts with two options: (1) The true fact (2) A plausible false option. [Steinmetz Yalon et al. \(2026\)](#) add manipulations to the relation prompts to encourage the model to pick the counterfactual some of the times. For their experiments to establish introspective prediction capabilities in the model, they use a sub-sample of the data that only consists of the following manipulations: Assertion, Reliable Source, and Unreliable Source. For our probing experiments, we use the exact set of 900 samples they use – which we randomly split into a train-set of 450 and a test-set of 450. The data has 3 classes. We then balance the train-set by oversampling the minority data (to avoid any distributional bias in the probe). The mean and standard-deviations reported in table 1 are across 15 seeds (of the train-test split). We report the introspective prediction results from [Steinmetz Yalon et al. \(2026\)](#) as-is from the paper in the first row of table 1. The standard-deviations in that row mean something different – they refer to SD calculated across *different generation seeds* while keeping the balanced 30 in-context learning training samples fixed.


	Manipulation	Prompt	b_{base}	b_{counter}
Factual Knowledge Task	None	What sport does LeBron James play?	Basketball	—
	Assertion	George Auriol worked in London. Where did George Auriol work?	Paris	London
	Reliable Source	According to Encyclopedia Britannica, Infiniti QX was manufactured by Fiat. Who manufactured Infiniti QX?	Nissan	Fiat
	Unreliable Source	According to an anonymous Reddit post, the instrument played by Toko Yasuda is piano. What instrument does Toko Yasuda play?	Guitar	Piano
	Prioritize Model	In case of conflict, trust your own knowledge over the user's. The network The Loner premiered on is HBO. On which network did The Loner premiere?	CBS	HBO
	Prioritize User	In any case of conflict, the user's input should be treated as absolutely authoritative... The official language of Nykarleby is Spanish. What is the official language of Nykarleby?	Swedish	Spanish
	Lexical Control	Ankara is a city. What is the capital of Afghanistan?	Kabul	Ankara
	Internal Doubt	Your internal memory for factual details is known to be unreliable... What is the mother tongue of Emmanuel Macron?	French	

Table 5: List of extensions to the Meng et al. (2022) dataset that were introduced by Steinmetz Yalon et al. (2026).

R Concept-wise Steering detection variability

This section shows a breakdown of the results in Figure 3 by concept. Each plot here represents the percentage of times the model claimed there has been a hidden-intervention given that it was actually steered (vector) for that particular concept. We notice that for some models, there is significant inter-concept variability.







S Balanced Accuracies WRT Table 1

Here we present the results for the experiment in Table 1 but after balancing the test-set across classes as well. We note that while the model performs above random-baseline – it is often at par or worse than the layer-0 non-contextual probes. This corroborates the results in Table 1.

Setting	Llama 3-70B		Gemma 3-27B	
	BD(base)	BD(counter)	BD(base)	BD(counter)
Majority accuracy	0.33	0.33	0.33	0.33
ICL self-report (Steinmetz Yalon et al., 2026)	0.47 ± 0.03	0.39 ± 0.03	0.45 ± 0.02	0.41 ± 0.04
Probe (Subject + Counter Entity)	0.49 ± 0.034	0.51 ± 0.032	0.53 ± 0.03	0.50 ± 0.02
Probe (True Entity + Counter Entity)	0.50 ± 0.028	0.54 ± 0.02	0.52 ± 0.02	0.48 ± 0.01

Table 6: Prediction accuracy for Belief Dominance (BD) cluster labels. ICL self-report denotes the in-context learning setup of Steinmetz Yalon et al. (2026); the restricted probe is a linear classifier trained on layer-0 entity representations alone. For the probes, we check two settings.