Sparse Reward Subsystem in Large Language Models

Guowei Xu
Tsinghua University

Mert Yuksekgonul
Stanford University

James Zou
Stanford University

Abstract

Recent studies show that LLM hidden states encode reward-related information, such as answer correctness and model confidence. However, existing approaches typically fit black-box probes on the full hidden states, offering little insight into how this information is structured across neurons. In this paper, we show that reward-related information is concentrated in a sparse subset of neurons. Using simple probing, we identify two types of neurons: value neurons, whose activations predict state value, and dopamine neurons, whose activations encode step-level temporal difference (TD) errors. Together, these neurons form a sparse reward subsystem within LLM hidden states. These names are drawn by analogy with neuroscience, where value neurons and dopamine neurons in the biological reward subsystem also encode value and reward prediction errors, respectively. We demonstrate that value neurons are robust and transferable across diverse datasets and models, and provide causal evidence that they encode reward-related information. Finally, we show applications of the reward subsystem: value neurons serve as effective predictors of model confidence, and dopamine neurons can function as a process reward model (PRM) to guide inference-time search. Our code is available at https://github.com/zou-group/reward_subsystem.

1 Introduction

In recent years, the reasoning capabilities of large language models (LLMs) have achieved significant breakthroughs, even surpassing humans in various mathematical and scientific tasks [45, 23]. Parallel to the continuous pursuit of higher performance, researchers seek to understand the underlying reasoning mechanisms of these models. Recent studies reveal that LLM hidden states encode rich reward-related information, which can be leveraged for weighted learning [40], predicting confidence and answer correctness [19, 16], detecting hallucinations [56], and performing implicit reasoning [10].

However, these approaches typically fit black-box probes on the full hidden states [6]: they demonstrate that reward-related information can be decoded, but do not characterize *how* this information is structured within the hidden states. Understanding this matters both for mechanistic interpretability, as it clarifies whether reward signals are sparsely encoded, and for practical applications, since identifying a minimal set of responsible neurons can enable more lightweight and targeted applications.

In this paper, we find that the ability of LLM hidden states to predict reward-related information can be attributed to the existence of a small subset of neurons. Using simple probing, we identify two types of neurons: *value neurons*, whose activations predict the expected value of the current state, and *dopamine neurons*, whose activations predict the TD error at each reasoning step. Together, these neurons form what we term a *reward subsystem* within LLM hidden states. These names are drawn by analogy with neuroscience, where biological value neurons represent the subjective value of stimuli during decision-making [50, 41] and dopamine neurons encode reward prediction errors [43, 8, 14, 13].

We first provide motivation for why reward-related information may be sparse in LLM hidden states. Then, for value neurons, we show via pruning experiments that less than 1% of neurons suffice to

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

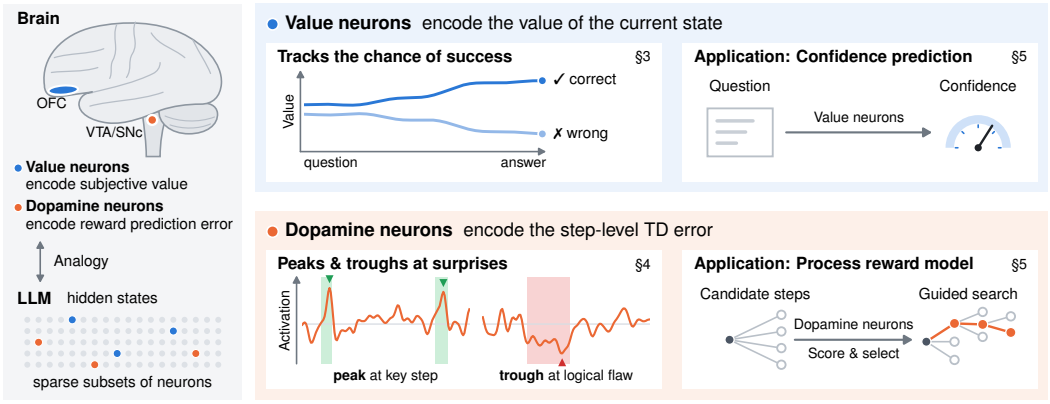


Figure 1: Overview of the sparse reward subsystem in LLMs. The left panel illustrates a computational analogy with biological value and dopamine neurons. Value neurons encode the value of the current state, while dopamine neurons encode step-level TD errors. The signal illustrations show value estimates tracking the chance of success (§3) and dopamine-neuron activations exhibiting peaks and troughs at unexpected progress and reasoning errors (§4). We apply value neurons to confidence prediction (§5.1) and dopamine neurons to process-reward-guided search (§5.2).

predict state value, and establish their specificity as reward-related neurons through intervention experiments. For dopamine neurons, we show via analogous pruning that a similarly sparse subset encodes step-level TD errors, and visualize their activation patterns during reasoning, confirming that they exhibit peaks when the model makes unexpected progress and troughs when the model encounters errors. Finally, we show practical applications of these neurons: value neurons serve as effective predictors of model confidence, and dopamine neurons can function as a process reward model (PRM) to guide inference-time search. Figure 1 provides an overview of the reward subsystem and its applications.

In summary, our key contributions are:

1. We identify sparse value neurons that predict value information, establish their specificity as reward-related neurons through controlled interventions, and demonstrate their robustness across datasets, models, and layers.
2. We identify sparse dopamine neurons that encode step-level TD errors, with activation patterns consistent with reward prediction error signals.
3. We highlight applications of these neurons: value neurons can predict model confidence prior to generation, and dopamine neurons can function as an internal process reward model.

2 Motivation

This section provides motivation for two questions: why autoregressive LLMs may contain neurons that can predict value and TD-error, and why such neurons may appear sparse. The arguments below are intended as motivation, not as rigorous proof.

2.1 Why may reward subsystem exist in LLMs?

We view an autoregressive LLM $\mathcal{M}$ as a policy over partial reasoning trajectories. The initial state is s_0. At generation step t, the model samples an action $a_t \sim \mathcal{M}(\cdot \mid s_t)$ and updates the state as $s_{t+1} = s_t \oplus a_t$. Let $r(s_T) \in \{0, 1\}$ denote the terminal reward after the model completes a response. For a fixed model $\mathcal{M}$, the value of a partial reasoning step can be written as:

$$V^{\mathcal{M}}(s_t) = \mathbb{E}_{\mathcal{M}}\left[r(s_T) \mid s_t\right], \tag{1}$$

which is the expected probability that continuing generation from s_t will eventually yield a correct answer.

A useful motivation comes from maximum-entropy reinforcement learning. In the soft-optimal setting, the optimal policy and optimal soft Q-function satisfy

$$\pi^*(a \mid s) = \frac{\exp(Q^*(s,a)/\tau)}{Z(s)}, \qquad Z(s) = \sum_{a'} \exp(Q^*(s,a')/\tau), \tag{2}$$

or equivalently,

$$\log \pi^*(a \mid s) = \frac{Q^*(s,a)}{\tau} - \log Z(s) = \frac{Q^*(s,a) - V^*(s)}{\tau}. \tag{3}$$

Here, $\tau > 0$ is the temperature and $V^*(s)$ is the optimal *soft* value function. These equations illustrate the relation between policy and value in an idealized soft-optimal setting. Our intuition is that a sufficiently capable policy may implicitly encode value-related information. Since LLMs are reasonably capable on the tasks we study, we investigate whether their internal representations encode information about the current-policy value.

2.2 Why may these neurons be sparse?

Feasibility. Information theory shows that a small number of coordinates can, in principle, suffice to encode value information. Let $R = r(s_T)$ denote the binary terminal reward obtained by completing a trajectory from the current state, and let $h = h(s_t, l)$ denote the hidden state at layer l. For predicting R from h, the posterior probability

$$\eta(h) = P(R = 1 \mid h) \tag{4}$$

is a sufficient scalar summary for binary prediction. By Bayes' rule,

$$P(R = 1 \mid h) = \frac{p(h \mid R = 1)P(R = 1)}{p(h \mid R = 1)P(R = 1) + p(h \mid R = 0)P(R = 0)}. \tag{5}$$

Equivalently,

$$P(R = 1 \mid h) = \sigma\left(\log \frac{p(h \mid R = 1)}{p(h \mid R = 0)} + \log \frac{P(R = 1)}{P(R = 0)}\right), \tag{6}$$

where σ is the sigmoid function. Thus, for Bayes-optimal prediction of the binary terminal reward, the relevant information in h can in principle be summarized by the scalar log-likelihood ratio

$$T(h) = \log \frac{p(h \mid R = 1)}{p(h \mid R = 0)}. \tag{7}$$

This argument shows that the reward-relevant information needed for terminal correctness prediction can be low-dimensional in principle. The same applies to TD-error prediction, since δ_t is also a scalar. Therefore, a small number of hidden neurons can be sufficient, in principle, to preserve most of the value and TD-error information.

Realization. The superposition hypothesis [17] offers a principled explanation for why low-dimensional reward-related information may be realized as a sparse set of neurons. Specifically, the hypothesis suggests that features that are frequent and high-importance during training tend to receive dedicated, sparse representations. Value and TD-error information is useful throughout reasoning, as they summarize whether a partial trajectory is likely to succeed and whether a newly generated step improves that trajectory. Therefore, the superposition hypothesis predicts that such information is likely to be encoded sparsely. To be clear, this argument does not prove that these neurons must be sparse. However, in the experiments that follow, we find that these neurons indeed occupy less than 1% of the total neurons, consistent with the prediction of the superposition hypothesis.

3 Sparse Value Neurons in Large Language Models

In this section, we first introduce value neurons, a sparse subset of neurons whose activations encode the model's expected value of the current state.

3.1 Training a Value Probe

Consider an autoregressive LLM $\mathcal{M}$. The initial state for generation is defined as s_0. During the generation process, the LLM produces tokens sequentially. At each step t, an action is sampled such that $a_t \sim \mathcal{M}(\cdot|s_t)$, and the subsequent state is updated as $s_{t+1} = s_t \oplus a_t$. Suppose the model generates a total of T new tokens. Given that modern decoder-only LLMs consist of multiple Transformer blocks, each hidden layer produces a corresponding representation for every state s_t. We denote the hidden state at the l-th layer and t-th step as $h(s_t, l)$. Here, a *neuron* refers to an individual coordinate of the LLM hidden state.

To extract the reward information embedded within these representations, we introduce a value probe V. Following [60], we employ a two-layer multi-layer perceptron (MLP) with ReLU activation. The input dimension matches the dimensionality of the LLM's l-th layer hidden states, while the output is a scalar representing the predicted reward. After computing $h(s_t, l)$, we feed it into V to obtain the value prediction output $V(h(s_t, l))$. For brevity, we define $V_l(s_t) = V(h(s_t, l))$. The probe is optimized using temporal difference (TD) learning. Let $r(s_T)$ be the final binary reward received. Given a discount factor γ, the TD error δ_t is defined as:

$$\delta_t = \begin{cases} r(s_T) - V_l(s_T), & \text{if } t = T+1 \\ \gamma V_l(s_t) - V_l(s_{t-1}), & \text{otherwise.} \end{cases} \tag{8}$$

We conduct a layer-wise analysis to investigate the characteristics of the value neurons. For a given layer l and a training dataset $\mathcal{D}_\mathcal{T}$, our objective for the value probe is to minimize the expected TD error over the distribution of generated sequences:

$$\mathcal{L}_{\text{TD}}(l) = \mathbb{E}_{s_0 \sim \mathcal{D}_\mathcal{T}, s_t \sim \mathcal{M}(\cdot|s_{t-1})} \left[\sum_{t=1}^{T+1} \delta_t^2 \right]. \tag{9}$$

TD learning bootstraps each state's value from its successor and enforces step-level temporal consistency, whereas direct regression onto the final reward can rely on static question-level correlates.

3.2 Identifying Value Neurons

Once the value probe is trained, we evaluate its performance on a validation dataset $\mathcal{D}_V$. During evaluation, we measure the predicted value $V_l(s_0)$ directly from the hidden state at the initial input position s_0. This allows us to detect whether the value neurons have already formed an assessment of the potential reward before any tokens of the answer are generated.

Specifically, we evaluate the correlation between $V_l(s_0)$ and $r(s_T)$. Higher correlation indicates better predictive capability. We utilize $AUC(V_l(s_0), r(s_T))$, the Area Under the Receiver Operating Characteristic (AUC), as our metric.

To substantiate the existence of sparse value neurons, we must demonstrate a small subset of the hidden states maintains sufficient predictive power regarding the value, so we perform pruning experiments. Assume the input dimensionality of our value probe is N. We introduce a pruning ratio p. We prune pN of these input dimensions, feeding only the most significant $(1-p)N$ dimensions into the value probe. To prune the network, we calculate the L_1 norm of the weights connecting the input dimensions to the neurons in the first hidden layer and prune the input dimensions that correspond to the smallest L_1 weight norms, constructing a pruned value probe V_l^p. The surviving probe input dimensions define the set of value neurons at that pruning ratio.

We then plot the relationship between $AUC(V_l^p(s_0), r(s_T))$ and the pruning ratio p. At each pruning ratio, we resample the validation set with replacement 1,000 times and compute the standard deviation of the resulting AUC values to estimate uncertainty. If the AUC curve remains stable as p increases, it provides strong evidence for the existence of value neurons.

3.3 Empirical Evidence

We provide an illustrative example using the curves from layers 2–4 of the Qwen-2.5-14B-SimpleRL-Zoo [55] model on the GSM8K [12] and MATH500 [33] datasets. For detailed hyperparameters, please refer to Appendix A. As illustrated in Figure 2, the AUC curves do not exhibit a significant decline as pruning proceeds, which indicates that the value probe can effectively estimate the value

of the current state by relying on a very small fraction (less than 1%) of the total neurons. These specific neurons are designated as *value neurons*.

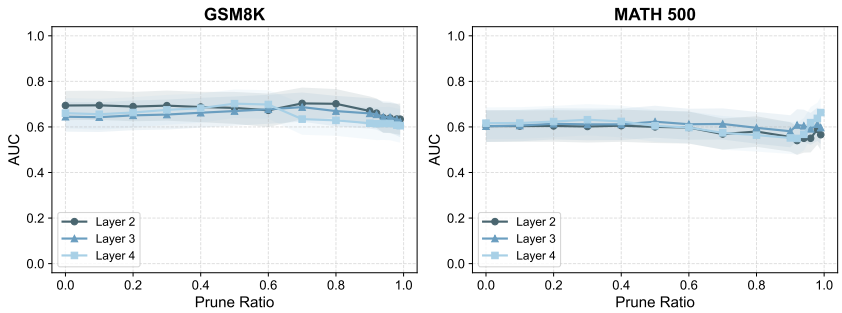


Figure 2: AUC curves for layers 2–4 of the Qwen-2.5-14B-SimpleRL-Zoo model on the GSM8K and MATH500 datasets. The curves indicate that the value probe can accurately predict the value by relying on only a very small number of value neurons. Shaded regions indicate uncertainty.

In Appendix C, we systematically verify the robustness of value neurons. Specifically, we validate the existence of sparse value neurons across benchmarks including MATH500 [33], ARC [11], MBPP+ [4] (coding), and IFEval [59] (general instruction-following), and across models including Qwen3.5-0.8B [51], Phi-3.5-mini-instruct [1], Llama-3.1-8B-Instruct [20], and Qwen-2.5-14B-SimpleRL-Zoo [55], confirming that value neurons consistently emerge at various layers.

We also note that the experiments in this section evaluate value neurons at the initial input position s_0, which is inherently challenging and results in moderate absolute AUC values. In Appendix I, we conduct analysis at the final response position s_T and observe the AUC remains mostly stable above 0.8, confirming that value neurons encode reward information throughout the generation process.

3.4 Intervention Experiments: Value Neurons might be Reward-Related

The probing results above show that value neuron activations *correlate* with reward information. A natural follow-up question is whether these neurons might be reward-related.

Necessity. We first examine whether value neurons might be necessary for reasoning. We select the Qwen-2.5-7B-SimpleRL-Zoo model, zero out the activations of the top 1% of value neurons in specific layers, and then measure the resulting performance drop on the MATH500 dataset. We compare against four baselines: (1) **Random**: randomly zeroing out the same proportion of neurons; (2) **Weight Magnitude**: zeroing out the top 1% of neurons with the largest absolute weights in the MLP `down_proj` layer; (3) **Wanda** [46]: zeroing out the top 1% of neurons selected by Wanda; and (4) **NTP Neurons**: training a probe with a reward-unrelated objective (next-token prediction) using the same architecture and pruning procedure, then zeroing out its top 1% neurons. We provide additional comparisons in Appendix G.

Table 1: Intervention results for the Qwen-2.5-7B-SimpleRL-Zoo model on the MATH500 dataset. Performance is measured by accuracy (%) after zeroing out a 1% subset of neurons in a single layer. The original accuracy is 75.2%.

Layer	Value Neurons	Random	NTP Neurons	Weight Magnitude	Wanda
2	37.0 (-38.2)	77.0 (+1.8)	76.4 (+1.2)	74.6 (-0.6)	74.8 (-0.4)
3	13.6 (-61.6)	73.4 (-1.8)	76.8 (+1.6)	61.6 (-13.6)	48.2 (-27.0)
4	29.4 (-45.8)	73.8 (-1.4)	77.4 (+2.2)	75.4 (+0.2)	67.0 (-8.2)
5	1.2 (-74.0)	74.4 (-0.8)	75.2 (+0.0)	75.0 (-0.2)	73.4 (-1.8)
Avg	20.3 (-54.9)	74.6 (-0.6)	76.4 (+1.2)	71.6 (-3.6)	65.8 (-9.4)

As shown in Table 1, zeroing out value neurons causes degradation. Additionally, zeroing out NTP neurons has no effect, and the overlap between value neurons and NTP neurons is only 0.7% (random baseline: 0.5%), indicating that the value neurons identified by the reward-trained probe occupy

distinct positions from those identified by a reward-unrelated probe.These results suggest that value neurons might be necessary for reasoning.

Specificity. Having examined their necessity for reasoning, we next examine whether value neurons might be specifically reward-related. We compare three reasoning benchmarks (MATH500, GSM8K, and ARC) with two factual-recall benchmarks (PopQA [36] and EntityQuestions [44]). We use the probe trained on MATH and zero out the top 1% of value neurons in layer 2 of Qwen-2.5-7B-SimpleRL-Zoo.

Table 2: Comparison of reasoning and factual-recall performance after zeroing out the top 1% of value neurons in layer 2 of Qwen-2.5-7B-SimpleRL-Zoo, using the probe trained on MATH. Performance is measured by accuracy (%). Δ denotes the change in percentage points.

	MATH500	**GSM8K**	**ARC**	**PopQA**	**EntityQuestions**
Task type	Reasoning	Reasoning	Reasoning	Factual recall	Factual recall
Baseline	75.2	93.0	82.2	42.6	27.2
Value Neurons	37.0	52.8	66.2	46.2	24.2
Δ	-38.2	-40.2	-16.0	$+3.6$	-3.0

As shown in Table 2, zeroing out value neurons degrades performance on the reasoning benchmarks, while leaving factual-recall performance largely unaffected. This suggests the possibility of reward-related specificity of the identified neurons.

Functional coherence. We also show in Appendix H that zeroing out value neurons in earlier layers causes dopamine neurons in subsequent layers to lose their TD error predicting abilities, while zeroing out the same number of random neurons has minimal effect. This further supports that these neurons might form a functionally coherent subsystem organized around reward signals.

In summary, the intervention experiments provide three lines of preliminary evidence: (1) *Necessity*: value neurons might be critical for reasoning performance. (2) *Specificity*: value neurons might be specifically associated with reward signals. (3) *Functional coherence*: disrupting value neurons in earlier layers might impair the reward subsystem in subsequent layers.

We note that these intervention experiments are intended to provide preliminary insights rather than conclusive causal evidence. Potential confounds remain, including the stability of the selected neurons across probes trained on different rollouts and the overlap between the selected neurons and coordinates with large activation magnitudes in the residual stream. We leave a more thorough investigation of these factors to future work.

3.5 Transferability of the Value Neurons Across Different Datasets and Models

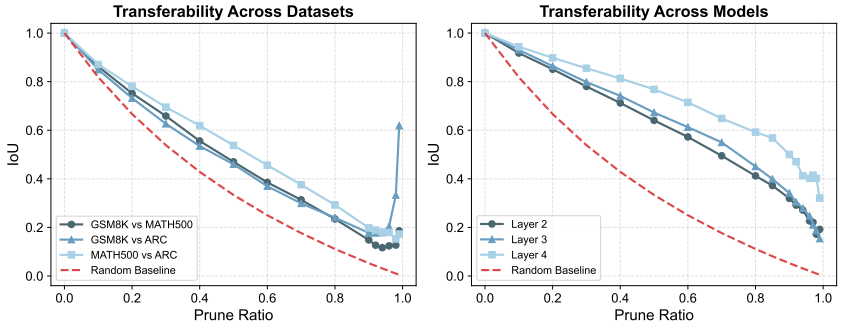


Figure 3: IoU as a function of the pruning ratio. The IoU values for value neurons across different datasets and models are significantly higher than the random baseline.

If value neurons are an intrinsic property of LLMs, their positions should remain consistent across datasets and models. To test this, we compute the Intersection over Union (IoU) of value neurons identified on any two datasets at a given pruning ratio, and compare against a uniform selection baseline (derivation in Appendix B). Using layer 3 of Qwen-2.5-14B-SimpleRL-Zoo, we compute

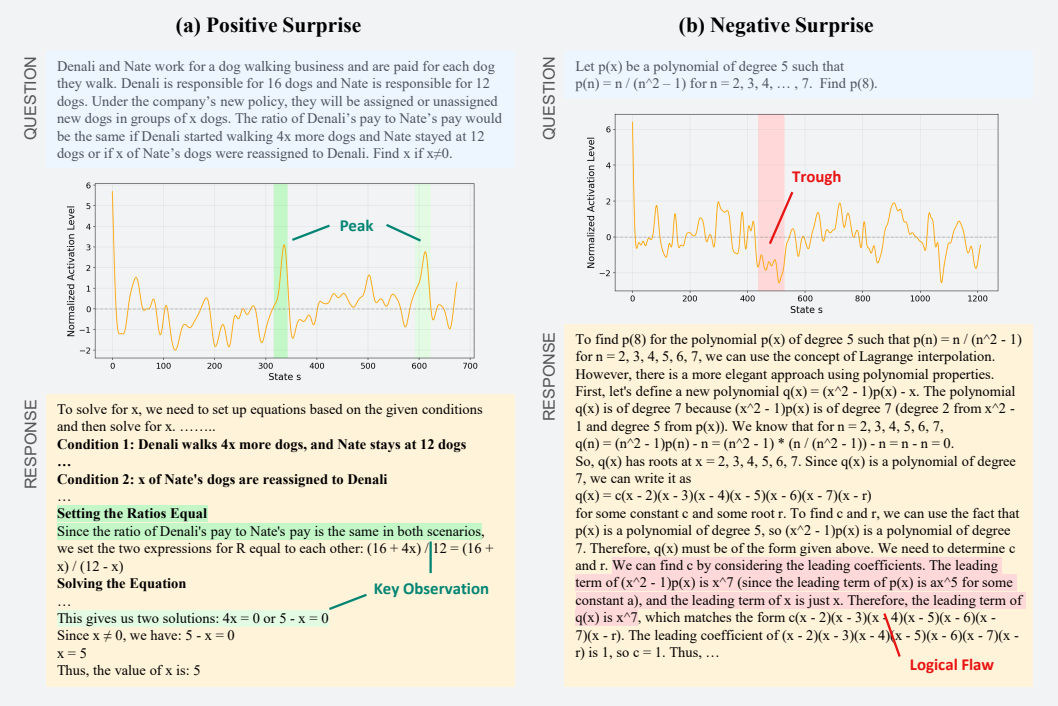


Figure 4: Dopamine neurons encode information regarding the model's prediction error for the current state. (a) **Positive Surprise:** The model initially lacks confidence in answering the problem but ultimately provides the correct solution. This neuron exhibits two significant peaks when the model identifies a critical logical step and subsequently derives the final key result. (b) **Negative Surprise:** Conversely, the model begins with high confidence but fails to solve the problem correctly. The neuron displays a distinct trough at the exact moment a logical flaw occurs.

pairwise IoU curves across three datasets (GSM8K, MATH500, ARC). As shown in Figure 3, IoU consistently exceeds the random baseline and increases sharply as the pruning ratio approaches 1, indicating that the most critical value neurons are highly stable across tasks. Additional results are provided in Appendix D.

We further investigate whether value neuron positions are preserved across models fine-tuned from the same base model. We compare layers 2–4 of Qwen-2.5-7B-SimpleRL-Zoo and Qwen2.5-7B-PPO-Zero (both derived from Qwen-2.5-7B via RLVR), using value probes trained on MATH500. As shown in Figure 3, the IoU remains consistently above the random baseline, suggesting that value neurons are transferable across different models derived from the same base model.

4 Sparse Dopamine Neurons in Large Language Models

In this section, we first introduce dopamine neurons, a sparse subset of neurons whose activations encode temporal-difference (TD) reward prediction errors.

4.1 Training a Dopamine Probe

Given a problem, the model is initialized with state s_0 as defined in Section 3. The model generates a response y which is segmented into P paragraphs $y = (c_1, c_2, \ldots, c_P)$. The state at paragraph boundary t is then defined as $s_t = s_0 \oplus c_1 \oplus c_2 \oplus \cdots \oplus c_t$. For each paragraph boundary, we estimate the value function $\hat{V}(s_t)$ by using Monte Carlo estimation to generate K independent rollout completions from state s_t. At the terminal boundary, we set $\hat{V}(s_P) = r(s_P)$. The TD-error at paragraph t quantifies the unexpected change in reward:

$$\delta_t = \gamma\, \hat{V}(s_t) - \hat{V}(s_{t-1}), \quad t = 1, 2, \ldots, P, \tag{10}$$

where $\gamma \in [0, 1]$ is the discount factor. A positive δ_t indicates that paragraph c_t increases the expected probability of arriving at a correct answer, while a negative δ_t indicates a detrimental reasoning step.

To extract the RPE information encoded in hidden states, we introduce a dopamine probe D, parallel in architecture to the value probe. For paragraph t, we define $\tilde{h}(s_t, l)$ as the normalized hidden state in layer l (see Appendix E for details). The probe prediction is $D_l(s_t) = D(\tilde{h}(s_t, l))$. The probe is trained to minimize:

$$\mathcal{L}_{\text{dopamine}}(l) = \mathbb{E}_{s_0 \sim \mathcal{D}_T} \left[\sum_{t=1}^{P} \left(D_l(s_t) - \delta_t \right)^2 \right]. \tag{11}$$

To focus on steps where meaningful reward changes occur, we retain only paragraphs satisfying $|\delta_t| > 0.3$ during training. The same threshold is applied during evaluation. We provide an ablation study supporting the choice of this threshold in Appendix L.

4.2 Identifying Dopamine Neurons

Once trained, we identify dopamine neurons via the same L1-norm pruning procedure used for value neurons. We introduce a pruning ratio p and construct a pruned probe D_l^p. We evaluate the pruned probe on the validation set $\mathcal{D}_V$ via the Spearman correlation coefficient $r(D_l^p(s_t),\, \delta_t)$ between predicted and MC-estimated TD errors. If this metric remains stable as p increases, it provides strong evidence that a small subset of neurons is sufficient to encode step-level reward prediction errors.

4.3 Empirical Evidence

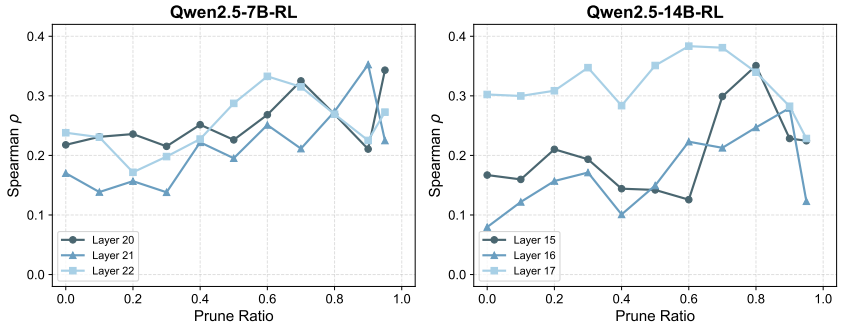


Figure 5: Spearman correlation curves for Qwen-2.5-7B/14B-SimpleRL-Zoo on MATH500.

Our experiments utilize the Qwen-2.5-7B/14B-SimpleRL-Zoo models. The probes are trained and evaluated on the MATH500 dataset. Please refer to Appendix A for details.

As illustrated in Figure 5, the Spearman correlation curves remain largely invariant across a wide range of pruning ratios. This demonstrates that the information necessary for signaling reward prediction errors is concentrated within a set of sparse neurons. We identify these as *dopamine neurons*, as they encode reward prediction errors.

4.4 Visualization Results

We illustrate the behavior of dopamine neurons using the 1517-th neuron in layer 5 as an example. In Appendix F, we provide additional visualization cases. As shown in Figure 4, in the positive surprise case (a), the model initially lacks confidence but ultimately solves the problem correctly. The neuron exhibits sharp activation peaks at the moments when the model identifies critical logical steps. In the negative surprise case (b), the model begins with high confidence but commits an error mid-reasoning, and the neuron displays a trough at the exact moment the logical flaw occurs. These patterns confirm that dopamine neuron activations closely track TD errors during inference.

5 Applications

5.1 Applications of Value Neurons: Predicting Model Confidence

Since value neurons encode value information before any response tokens are generated, they can serve as a lightweight confidence estimator. This is practically valuable. For instance, modern

reasoning models produce very long responses, so gauging confidence beforehand enables adaptive strategies such as dynamically allocating compute based on the initial confidence level.

We compare against four baselines across four models (Llama-3.1-8B, Phi-3.5-mini, Gemma-3-4B, Qwen3.5-0.8B) and two benchmarks (MATH500, ARC): LCD [9], which trains a linear probe on the full hidden state; Verbal Confidence, which prompts the model to output its confidence directly; Token Confidence, which uses the next-token log probability at the question position; and Question Length as a simple proxy. For our method, AUC values are obtained at the optimal pruning ratio and averaged across layers 2–4. Baseline configurations are detailed in Appendix J.

As shown in Table 3, value neurons achieve the highest average AUC. This demonstrates that a sparse set of value neurons carries informative model confidence signals.

Table 3: Comparison of AUC for different confidence prediction methods across models and benchmarks. Higher is better. **Bold** denotes the best performing method per model.

Model	Value Neurons	LCD [9]	Verbal Conf.	Token Conf.	Question Len.
MATH500					
Llama-3.1-8B	**0.69**	0.64	**0.69**	0.54	0.63
Phi-3.5-mini	**0.73**	0.63	0.54	0.41	0.62
Gemma-3-4B	**0.72**	0.65	0.62	0.41	0.71
Qwen3.5-0.8B	0.71	0.70	0.41	0.56	**0.77**
ARC					
Llama-3.1-8B	**0.66**	0.56	0.55	0.49	0.55
Phi-3.5-mini	**0.63**	0.54	0.29	0.46	0.54
Gemma-3-4B	**0.59**	0.50	0.47	0.44	0.49
Qwen3.5-0.8B	0.63	0.54	0.55	0.50	**0.65**
Average	**0.67**	0.60	0.52	0.48	0.62

5.2 Applications of Dopamine Neurons: Process Reward Model

Since dopamine neurons encode step-level reward prediction errors, their activations can be directly leveraged to score individual reasoning steps, functioning as a Process Reward Model (PRM). To evaluate this, we employ a guided search procedure: at each paragraph boundary, the model generates $K = 4$ candidate responses, each is scored by the PRM, and the highest-scoring candidate is selected as the next reasoning step. This procedure repeats until the model produces a final answer.

Table 4: Comparison of accuracy (%) on MATH500 for different methods.

Method	Qwen-2.5-7B SimpleRL-Zoo	Qwen-2.5-14B SimpleRL-Zoo	Llama-3.1-8B Instruct	Average
Greedy	69.6	67.0	47.3	61.3
Random	77.3	74.8	46.3	66.1
Majority of N	64.8	76.0	54.2	65.0
Implicit PRM	69.8	74.8	45.8	63.5
Mean logprob	72.2	78.7	41.7	64.2
Min logprob	74.1	76.0	48.6	66.2
P(true)	78.7	76.0	47.2	67.3
Dopamine Neurons	**79.6**	**82.7**	**54.6**	**72.3**

We compare the dopamine probe against seven baselines on MATH500 using three models. **Greedy** decoding uses $K = 1$ with temperature zero. **Random** selects randomly among the K candidates. For **Implicit PRM**, motivated by the theoretical result that KL-constrained RL yields $\pi^* \propto \pi_{\text{ref}} \exp(r/\beta)$ [42, 54], we use $\log \pi_{\text{RL}} - \log \pi_{\text{base}}$ as an approximation to the step-level reward. **Majority of N** selects the most frequent final answer among N sampled responses; **Mean logprob** [35] and **Min**

logprob [37] score candidates using the mean and minimum token log probabilities, respectively; and **P(true)** [29] uses the model's self-estimated probability that a candidate is correct. In our method, we use the predicted value of the dopamine probe trained on the optimal layer as the process reward. All experiments are averaged over three random seeds.

As shown in Table 4, using the dopamine neuron predictions as a process reward signal improves the model's reasoning performance, demonstrating dopamine neurons can be applied to guide reasoning.

6 Related Work

Probing methods in large language models. Probing methods serve as powerful tools for investigating and interpreting the internal characteristics of Large Language Models (LLMs) and have been extensively utilized within the research community [3, 28, 2, 25]. No-answer-needed [9] demonstrated that linear probes can be trained to predict the correctness of a model's forthcoming answer. Expanding beyond linear analysis, [15] introduced a polar probe designed to extract syntactic relations by analyzing both the distance and direction between word embeddings. Probing has also been instrumental in assessing the veracity of generated content. [38] found that LLMs linearly represent the truth or falsehood of factual statements. This is supported by [26], who noted that LLM hidden states are highly predictive of factuality in long-form natural language generation, and that such information can be efficiently extracted at inference time using a lightweight probe. [32] employed lightweight MLP probes to perform nonlinear modeling of high-level hidden states for token-level hallucination detection. [18] identified that a sparse subset of neurons within LLMs can reliably predict the occurrence of hallucinations. Furthermore, [27] observed that simple probing methods demonstrate superior generalization on Out-of-Distribution (OOD) tasks compared to Sparse Autoencoders (SAEs). Following [60], our study employs a two-layer MLP network as our probing model to maintain simplicity while capturing the necessary reward signals.

Reward modeling based on LLM internal representations. Recent studies have investigated the extraction of reward signals from the internal representations of LLMs [31]. Anthropic proposed a value head to predict whether models can answer questions correctly [29]. [7] utilized a purely unsupervised approach to discover latent knowledge within the internal activations of a language model, enabling the accurate answering of yes-no questions using only model activations. [57] utilized tools from mechanistic interpretability to analyze model activations, employing Sparse Autoencoders (SAEs) to map high-dimensional hidden states onto interpretable semantic features. Similarly, the RISE framework [34] proposed simultaneously enhancing a model's problem-solving and self-verification capabilities within a single training process, requiring the model to generate both a solution and a corresponding evaluative score. Furthermore, [58] introduced Reinforcement Learning from Internal Feedback (RLIF), a framework that enables LLMs to learn from intrinsic signals in the absence of external rewards. Similarly, LaSeR [52] found that a last-token self-rewarding score can guide the reinforcement learning process. [16] observed that the latent thoughts leading to correct versus incorrect answers exhibit highly distinguishable patterns, allowing a latent classifier to predict answer correctness directly from these representations. The SWIFT method [24] demonstrated that mining intrinsic rewards from LLM hidden states facilitates efficient Best-of-N sampling. Despite these attempts to predict rewards from internal representations, existing literature has yet to identify how this information is structured within the hidden states.

7 Conclusion

In this work, we show that reward information in LLM hidden states is concentrated in a sparse set of neurons, the reward subsystem. Using probing and pruning, we identify value neurons that encode value information and dopamine neurons that encode reward prediction errors. Intervention experiments confirm that value neurons are specifically tied to reward signals, and their positions transfer across datasets and models. On the application side, value neurons provide lightweight confidence estimates, and dopamine neurons can serve as an intrinsic process reward model. Looking ahead, we believe the reward subsystem offers a useful lens for understanding LLM reasoning, and we hope this work encourages further exploration of neuron-level structure in language models.

References

[1] Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacroce, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. Phi-3 technical report: A highly capable language model locally on your phone, 2024. URL https://arxiv.org/abs/2404.14219. 5, 22

[2] Deema Alnuhait, Neeraja Kirtane, Muhammad Khalifa, and Hao Peng. Factcheckmate: Preemptively detecting and mitigating hallucinations in lms, 2025. URL https://arxiv.org/abs/2410.02899. 10

[3] Aryaman Arora, Zhengxuan Wu, Jacob Steinhardt, and Sarah Schwettmann. Language model circuits are sparse in the neuron basis, 2026. URL https://arxiv.org/abs/2601.22594. 10

[4] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large language models, 2021. URL https://arxiv.org/abs/2108.07732. 5, 22

[5] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report, 2023. URL https://arxiv.org/abs/2309.16609. 22

[6] Yonatan Belinkov. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219, March 2022. doi: 10.1162/coli_a_00422. URL https://aclanthology.org/2022.cl-1.7/. 1

[7] 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. 10

[8] Malcolm G. Campbell, Yongsoo Ra, Zhiqin Chen, Shudi Xu, Mark Burrell, Sara Matias, Mitsuko Watabe-Uchida, and Naoshige Uchida. A hardwired neural circuit for temporal difference learning. *bioRxiv*, 2025. URL https://api.semanticscholar.org/CorpusID:281521361. 1

[9] Iván Vicente Moreno Cencerrado, Arnau Padrés Masdemont, Anton Gonzalvez Hawthorne, David Demitri Africa, and Lorenzo Pacchiardi. No answer needed: Predicting llm answer

accuracy from question-only linear probes, 2025. URL https://arxiv.org/abs/2509.10625. 9, 10, 27

[10] Junhao Chen, Shengding Hu, Zhiyuan Liu, and Maosong Sun. States hidden in hidden states: Llms emerge discrete state representations implicitly, 2024. URL https://arxiv.org/abs/2407.11421. 1

[11] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL https://arxiv.org/abs/1803.05457. 5

[12] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168. 4

[13] Luke T Coddington, Sarah E Lindo, and Joshua T Dudman. Mesolimbic dopamine adapts the rate of learning from action. *Nature*, 614(7947):294–302, 2023. 1

[14] Will Dabney, Zeb Kurth-Nelson, Naoshige Uchida, Clara Kwon Starkweather, Demis Hassabis, Rémi Munos, and Matthew Botvinick. A distributional code for value in dopamine-based reinforcement learning. *Nature*, 577(7792):671–675, 2020. 1

[15] Pablo Diego-Simón, Stéphane D'Ascoli, Emmanuel Chemla, Yair Lakretz, and Jean-Rémi King. A polar coordinate system represents syntax in large language models, 2024. URL https://arxiv.org/abs/2412.05571. 10

[16] Hanwen Du, Yuxin Dong, and Xia Ning. Latent thinking optimization: Your latent reasoning language model secretly encodes reward signals in its latent thoughts. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=2jkAk3EP0v. 1, 10

[17] Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of superposition, 2022. URL https://arxiv.org/abs/2209.10652. 3

[18] Cheng Gao, Huimin Chen, Chaojun Xiao, Zhiyi Chen, Zhiyuan Liu, and Maosong Sun. H-neurons: On the existence, impact, and origin of hallucination-associated neurons in llms, 2025. URL https://arxiv.org/abs/2512.01797. 10

[19] Zorik Gekhman, Eyal Ben-David, Hadas Orgad, Eran Ofek, Yonatan Belinkov, Idan Szpektor, Jonathan Herzig, and Roi Reichart. Inside-out: Hidden factual knowledge in LLMs. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=f7GG1MbsSM. 1

[20] 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-Yeary, 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, Shaoliang 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, Vítor 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, Dingkang 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, Filippos 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 Khabsa, 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 Ratanchandani, 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. 5, 22

[21] Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. When attention sink emerges in language models: An empirical view. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=78Nn4QJTEN. 25

[22] Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. Textbooks are all you need, 2023. URL https://arxiv.org/abs/2306.11644. 22

[23] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*,

645(8081):633–638, September 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-09422-z. URL http://dx.doi.org/10.1038/s41586-025-09422-z. 1

[24] Jizhou Guo, Zhaomin Wu, Hanchen Yang, and Philip S. Yu. Mining intrinsic rewards from llm hidden states for efficient best-of-n sampling, 2025. URL https://arxiv.org/abs/2505.12225. 10

[25] Wes Gurnee, Neel Nanda, Matthew Pauly, Katherine Harvey, Dmitrii Troitskii, and Dimitris Bertsimas. Finding neurons in a haystack: Case studies with sparse probing. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=JYs1R9IMJr. 10

[26] Jiatong Han, Neil Band, Muhammed Razzak, Jannik Kossen, Tim G. J. Rudner, and Yarin Gal. Simple factuality probes detect hallucinations in long-form natural language generation. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng, editors, *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 16209–16226, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.880. URL https://aclanthology.org/2025.findings-emnlp.880/. 10

[27] Lovis Heindrich, Philip Torr, Fazl Barez, and Veronika Thost. Do sparse autoencoders generalize? a case study of answerability. In *ICML 2025 Workshop on Reliable and Responsible Foundation Models*, 2025. URL https://openreview.net/forum?id=rs3alQ5BV8. 10

[28] Junjie Hu, Gang Tu, Cheng Shengyu, JinXin Li, Jinting Wang, Rui Chen, Zhilong Zhou, and Dongbo Shan. Harp: Hallucination detection via reasoning subspace projection. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=ShEDWasmDG. 10

[29] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly) know what they know, 2022. URL https://arxiv.org/abs/2207.05221. 10

[30] Aitor Lewkowycz, Anders Johan Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Venkatesh Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=IFXTZERXdM7. 27

[31] Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=aLLuYpn83y. 10

[32] Shize Liang and Hongzhi Wang. Neural probe-based hallucination detection for large language models, 2025. URL https://arxiv.org/abs/2512.20949. 10

[33] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step, 2023. URL https://arxiv.org/abs/2305.20050. 4, 5

[34] Xiaoyuan Liu, Tian Liang, Zhiwei He, Jiahao Xu, Wenxuan Wang, Pinjia He, Zhaopeng Tu, Haitao Mi, and Dong Yu. Trust, but verify: A self-verification approach to reinforcement learning with verifiable rewards. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=gA3fFAEXNT. 10

[35] Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=jN5y-zb5Q7m. 9

[36] 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 Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.546. URL https://aclanthology.org/2023.acl-long.546/. 6

[37] Potsawee Manakul, Adian Liusie, and Mark Gales. SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.557. URL https://aclanthology.org/2023.emnlp-main.557/. 10

[38] 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. 10

[39] Joe Needham, Giles Edkins, Govind Pimpale, Henning Bartsch, and Marius Hobbhahn. Large language models often know when they are being evaluated, 2025. URL https://arxiv.org/abs/2505.23836. 29

[40] Jaehoon Oh, Seungjun Shin, and Dokwan Oh. House of cards: Massive weights in llms, 2025. URL https://arxiv.org/abs/2410.01866. 1

[41] Camillo Padoa-Schioppa and John A Assad. Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090):223–226, 2006. 1

[42] Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From r to q^*: Your language model is secretly a q-function. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=kEVcNxtqXk. 9

[43] Wolfram Schultz. Predictive reward signal of dopamine neurons. *Journal of neurophysiology*, 1998. 1

[44] Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. Simple entity-centric questions challenge dense retrievers. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6138–6148, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.496. URL https://aclanthology.org/2021.emnlp-main.496/. 6

[45] Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, Akshay Nathan, Alan Luo, Alec Helyar, Aleksander Madry, Aleksandr Efremov, Aleksandra Spyra, Alex Baker-Whitcomb, Alex Beutel, Alex Karpenko, Alex Makelov, Alex Neitz, Alex Wei, Alexandra Barr, Alexandre Kirchmeyer, Alexey Ivanov, Alexi Christakis, Alistair Gillespie, Allison Tam, Ally Bennett, Alvin Wan, Alyssa Huang, Amy McDonald Sandjideh, Amy Yang, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrei Gheorghe, Andres Garcia Garcia, Andrew Braunstein, Andrew Liu, Andrew Schmidt, Andrey Mereskin, Andrey Mishchenko, Andy Applebaum, Andy Rogerson, Ann Rajan, Annie Wei, Anoop Kotha, Anubha Srivastava, Anushree Agrawal, Arun Vijayvergiya, Ashley Tyra, Ashvin Nair, Avi Nayak, Ben Eggers, Bessie Ji, Beth Hoover, Bill Chen, Blair Chen, Boaz Barak, Borys Minaiev, Botao Hao, Bowen Baker, Brad Lightcap, Brandon McKinzie, Brandon Wang, Brendan Quinn, Brian Fioca, Brian Hsu, Brian Yang, Brian Yu, Brian Zhang, Brittany Brenner, Callie Riggins Zetino, Cameron Raymond, Camillo Lugaresi, Carolina Paz, Cary Hudson, Cedric Whitney, Chak Li, Charles Chen, Charlotte Cole, Chelsea Voss, Chen Ding, Chen Shen, Chengdu Huang, Chris Colby, Chris Hallacy, Chris

Koch, Chris Lu, Christina Kaplan, Christina Kim, CJ Minott-Henriques, Cliff Frey, Cody Yu, Coley Czarnecki, Colin Reid, Colin Wei, Cory Decareaux, Cristina Scheau, Cyril Zhang, Cyrus Forbes, Da Tang, Dakota Goldberg, Dan Roberts, Dana Palmie, Daniel Kappler, Daniel Levine, Daniel Wright, Dave Leo, David Lin, David Robinson, Declan Grabb, Derek Chen, Derek Lim, Derek Salama, Dibya Bhattacharjee, Dimitris Tsipras, Dinghua Li, Dingli Yu, DJ Strouse, Drew Williams, Dylan Hunn, Ed Bayes, Edwin Arbus, Ekin Akyurek, Elaine Ya Le, Elana Widmann, Eli Yani, Elizabeth Proehl, Enis Sert, Enoch Cheung, Eri Schwartz, Eric Han, Eric Jiang, Eric Mitchell, Eric Sigler, Eric Wallace, Erik Ritter, Erin Kavanaugh, Evan Mays, Evgenii Nikishin, Fangyuan Li, Felipe Petroski Such, Filipe de Avila Belbute Peres, Filippo Raso, Florent Bekerman, Foivos Tsimpourlas, Fotis Chantzis, Francis Song, Francis Zhang, Gaby Raila, Garrett McGrath, Gary Briggs, Gary Yang, Giambattista Parascandolo, Gildas Chabot, Grace Kim, Grace Zhao, Gregory Valiant, Guillaume Leclerc, Hadi Salman, Hanson Wang, Hao Sheng, Haoming Jiang, Haoyu Wang, Haozhun Jin, Harshit Sikchi, Heather Schmidt, Henry Aspegren, Honglin Chen, Huida Qiu, Hunter Lightman, Ian Covert, Ian Kivlichan, Ian Silber, Ian Sohl, Ibrahim Hammoud, Ignasi Clavera, Ikai Lan, Ilge Akkaya, Ilya Kostrikov, Irina Kofman, Isak Etinger, Ishaan Singal, Jackie Hehir, Jacob Huh, Jacqueline Pan, Jake Wilczynski, Jakub Pachocki, James Lee, James Quinn, Jamie Kiros, Janvi Kalra, Jasmyn Samaroo, Jason Wang, Jason Wolfe, Jay Chen, Jay Wang, Jean Harb, Jeffrey Han, Jeffrey Wang, Jennifer Zhao, Jeremy Chen, Jerene Yang, Jerry Tworek, Jesse Chand, Jessica Landon, Jessica Liang, Ji Lin, Jiancheng Liu, Jianfeng Wang, Jie Tang, Jihan Yin, Joanne Jang, Joel Morris, Joey Flynn, Johannes Ferstad, Johannes Heidecke, John Fishbein, John Hallman, Jonah Grant, Jonathan Chien, Jonathan Gordon, Jongsoo Park, Jordan Liss, Jos Kraaijeveld, Joseph Guay, Joseph Mo, Josh Lawson, Josh McGrath, Joshua Vendrow, Joy Jiao, Julian Lee, Julie Steele, Julie Wang, Junhua Mao, Kai Chen, Kai Hayashi, Kai Xiao, Kamyar Salahi, Kan Wu, Karan Sekhri, Karan Sharma, Karan Singhal, Karen Li, Kenny Nguyen, Keren Gu-Lemberg, Kevin King, Kevin Liu, Kevin Stone, Kevin Yu, Kristen Ying, Kristian Georgiev, Kristie Lim, Kushal Tirumala, Kyle Miller, Lama Ahmad, Larry Lv, Laura Clare, Laurance Fauconnet, Lauren Itow, Lauren Yang, Laurentia Romaniuk, Leah Anise, Lee Byron, Leher Pathak, Leon Maksin, Leyan Lo, Leyton Ho, Li Jing, Liang Wu, Liang Xiong, Lien Mamitsuka, Lin Yang, Lindsay McCallum, Lindsey Held, Liz Bourgeois, Logan Engstrom, Lorenz Kuhn, Louis Feuvrier, Lu Zhang, Lucas Switzer, Lukas Kondraciuk, Lukasz Kaiser, Manas Joglekar, Mandeep Singh, Mandip Shah, Manuka Stratta, Marcus Williams, Mark Chen, Mark Sun, Marselus Cayton, Martin Li, Marvin Zhang, Marwan Aljubeh, Matt Nichols, Matthew Haines, Max Schwarzer, Mayank Gupta, Meghan Shah, Melody Huang, Meng Dong, Mengqing Wang, Mia Glaese, Micah Carroll, Michael Lampe, Michael Malek, Michael Sharman, Michael Zhang, Michele Wang, Michelle Pokrass, Mihai Florian, Mikhail Pavlov, Miles Wang, Ming Chen, Mingxuan Wang, Minnia Feng, Mo Bavarian, Molly Lin, Moose Abdool, Mostafa Rohaninejad, Nacho Soto, Natalie Staudacher, Natan LaFontaine, Nathan Marwell, Nelson Liu, Nick Preston, Nick Turley, Nicklas Ansman, Nicole Blades, Nikil Pancha, Nikita Mikhaylin, Niko Felix, Nikunj Handa, Nishant Rai, Nitish Keskar, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Oona Gleeson, Pamela Mishkin, Patryk Lesiewicz, Paul Baltescu, Pavel Belov, Peter Zhokhov, Philip Pronin, Phillip Guo, Phoebe Thacker, Qi Liu, Qiming Yuan, Qinghua Liu, Rachel Dias, Rachel Puckett, Rahul Arora, Ravi Teja Mullapudi, Raz Gaon, Reah Miyara, Rennie Song, Rishabh Aggarwal, RJ Marsan, Robel Yemiru, Robert Xiong, Rohan Kshirsagar, Rohan Nuttall, Roman Tsiupa, Ronen Eldan, Rose Wang, Roshan James, Roy Ziv, Rui Shu, Ruslan Nigmatullin, Saachi Jain, Saam Talaie, Sam Altman, Sam Arnesen, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Sarah Yoo, Savannah Heon, Scott Ethersmith, Sean Grove, Sean Taylor, Sebastien Bubeck, Sever Banesiu, Shaokyi Amdo, Shengjia Zhao, Sherwin Wu, Shibani Santurkar, Shiyu Zhao, Shraman Ray Chaudhuri, Shreyas Krishnaswamy, Shuaiqi, Xia, Shuyang Cheng, Shyamal Anadkat, Simón Posada Fishman, Simon Tobin, Siyuan Fu, Somay Jain, Song Mei, Sonya Egoian, Spencer Kim, Spug Golden, SQ Mah, Steph Lin, Stephen Imm, Steve Sharpe, Steve Yadlowsky, Sulman Choudhry, Sungwon Eum, Suvansh Sanjeev, Tabarak Khan, Tal Stramer, Tao Wang, Tao Xin, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Degry, Thomas Shadwell, Tianfu Fu, Tianshi Gao, Timur Garipov, Tina Sriskandarajah, Toki Sherbakov, Tomer Kaftan, Tomo Hiratsuka, Tongzhou Wang, Tony Song, Tony Zhao, Troy Peterson, Val Kharitonov, Victoria Chernova, Vineet Kosaraju, Vishal Kuo, Vitchyr Pong, Vivek Verma, Vlad Petrov, Wanning Jiang, Weixing Zhang, Wenda Zhou, Wenlei Xie, Wenting Zhan, Wes McCabe, Will DePue, Will Ellsworth, Wulfie Bain, Wyatt Thompson, Xiangning Chen, Xiangyu Qi, Xin Xiang, Xinwei Shi, Yann

Dubois, Yaodong Yu, Yara Khakbaz, Yifan Wu, Yilei Qian, Yin Tat Lee, Yinbo Chen, Yizhen Zhang, Yizhong Xiong, Yonglong Tian, Young Cha, Yu Bai, Yu Yang, Yuan Yuan, Yuanzhi Li, Yufeng Zhang, Yuguang Yang, Yujia Jin, Yun Jiang, Yunyun Wang, Yushi Wang, Yutian Liu, Zach Stubenvoll, Zehao Dou, Zheng Wu, and Zhigang Wang. Openai gpt-5 system card, 2025. URL https://arxiv.org/abs/2601.03267. 1

[46] Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter. A simple and effective pruning approach for large language models. In *Workshop on Efficient Systems for Foundation Models @ ICML2023*, 2023. URL https://openreview.net/forum?id=tz9JV2PRSv. 5

[47] Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. Massive activations in large language models. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=F7aAhfitX6. 25

[48] William Timkey and Marten van Schijndel. All bark and no bite: Rogue dimensions in transformer language models obscure representational quality. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4527–4546, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.372. URL https://aclanthology.org/2021.emnlp-main.372/. 25

[49] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023. URL https://arxiv.org/abs/2302.13971. 22

[50] Léon Tremblay and Wolfram Schultz. Relative reward preference in primate orbitofrontal cortex. *Nature*, 398(6729):704–708, 1999. 1

[51] 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. 5, 22

[52] Wenkai Yang, Weijie Liu, Ruobing Xie, Yiju Guo, Lulu Wu, Saiyong Yang, and Yankai Lin. Laser: Reinforcement learning with last-token self-rewarding, 2025. URL https://arxiv.org/abs/2510.14943. 10

[53] Mengxia Yu, De Wang, Colorado Reed, and Alvin Wan. The super weight in large language models, 2025. URL https://openreview.net/forum?id=0Ag8FQ5Rr3. 25

[54] Lifan Yuan, Wendi Li, Huayu Chen, Ganqu Cui, Ning Ding, Kaiyan Zhang, Bowen Zhou, Zhiyuan Liu, and Hao Peng. Free process rewards without process labels. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=8ThnPFhGm8. 9

[55] Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild, 2025. URL https://arxiv.org/abs/2503.18892. 4, 5

[56] Luan Zhang, Dandan Song, Zhijing Wu, Yuhang Tian, Changzhi Zhou, Jing Xu, Ziyi Yang, and Shuhao Zhang. Detecting hallucination in large language models through deep internal representation analysis. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, IJCAI '25, 2025. ISBN 978-1-956792-06-5. doi: 10.24963/ijcai.2025/929. URL https://doi.org/10.24963/ijcai.2025/929. 1

[57] Shuyi Zhang, Wei Shi, Sihang Li, Jiayi Liao, Hengxing Cai, and Xiang Wang. Interpretable reward model via sparse autoencoder, 2025. URL https://arxiv.org/abs/2508.08746. 10

[58] Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. Learning to reason without external rewards, 2025. URL https://arxiv.org/abs/2505.19590. 10

[59] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models, 2023. URL https://arxiv.org/abs/2311.07911. 5, 22

[60] Yubo Zhu, Dongrui Liu, Zecheng Lin, Wei Tong, Sheng Zhong, and Jing Shao. The llm already knows: Estimating llm-perceived question difficulty via hidden representations, 2025. URL https://arxiv.org/abs/2509.12886. 4, 10

Appendix - Table of Contents

A Hyperparameters and Experimental Settings **21**

A.1 Response Generation Hyperparameters 21

A.2 Value / Dopamine Probe Training Hyperparameters 21

A.3 AUC Curve Evaluation Hyperparameters 21

B Derivation of the IoU Curve for the Random Baseline **21**

C Robustness of the Value Neurons **22**

C.1 Robustness across Different Datasets 22

C.2 Robustness across Different Models 22

C.3 Robustness across Different Layers 23

D Transferability Across Different Datasets: More IoU Curves **23**

E Detailed Procedures for Hidden State Normalization **23**

F Further Evidence for the Characteristics of Dopamine Neurons **24**

G Additional Intervention Controls **25**

H Correlation between Value Neurons and Dopamine Neurons **25**

I Value Neuron Analysis at the Final Response Position **26**

J Model Confidence Baseline Setup **27**

K Cross-Task Generalization of Trained Probes **27**

L Sensitivity to the TD-Error Threshold **28**

M Is the Reward Subsystem Already Present in Base Models or Created by RL Fine-Tuning? **28**

N Additional Discussion **28**

O Potential Limitations and Broader Impacts **29**

A Hyperparameters and Experimental Settings

Response generation is performed on two NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs. Both training and inference are conducted on a single NVIDIA RTX PRO 6000 Blackwell Server Edition GPU.

A.1 Response Generation Hyperparameters

The hyperparameters used for response generation are summarized in Table 5.

Table 5: Hyperparameters used for response generation.

Category	Hyperparameter = Value
Generation	max_tokens_per_call = 16000; temperature = 1.0; top_p = 0.95; n_sampling = 1; prompt_type = qwen-boxed
Inference	use_vllm = True; gpu_memory_utilization = 0.75;

A.2 Value / Dopamine Probe Training Hyperparameters

The hyperparameters used for value and dopamine probe training are listed in Table 6.

Table 6: Hyperparameters used for value and dopamine probe training.

Category	Hyperparameter = Value
Optimization	optimizer = AdamW; learning_rate = 1e-4; weight_decay = 0.01
Training	num_epochs = 100; batch_size = 4; train_ratio = 0.8; gamma = 1−1e-5; seed = 0
Architecture	hidden_size → 1024 (value) / 32 (dopamine) → 1 (two-layer MLP)

A.3 AUC Curve Evaluation Hyperparameters

The hyperparameters used for reward subsystem evaluation with neuron pruning are provided in Table 7.

Table 7: Hyperparameters used for reward subsystem evaluation with neuron pruning.

Category	Hyperparameter = Value
Evaluation	batch_size = 4; train_ratio = 0.8; seed = 0
Pruning	prune_ratios = {0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.92, 0.94, 0.96, 0.98, 0.99};

B Derivation of the IoU Curve for the Random Baseline

Let N denote the total number of neurons and $p \in [0, 1)$ denote the pruning ratio. Each set contains $k = (1 - p)N$ neurons selected uniformly at random. For two independently sampled sets A and B, each of size k, the intersection size $|A \cap B|$ follows a hypergeometric distribution:

$$P(|A \cap B| = i) = \frac{\binom{k}{i}\binom{N-k}{k-i}}{\binom{N}{k}}$$

where $i \in [\max(0, 2k - N), k]$. The expected IoU is given by:

$$\mathbb{E}\left[\text{IoU}\right] = \sum_{i=\max(0,2k-N)}^{k} \frac{\binom{k}{i}\binom{N-k}{k-i}}{\binom{N}{k}} \cdot \frac{i}{2k - i}$$

C Robustness of the Value Neurons

In this section, we demonstrate that value neurons are robust across various datasets, model scales, layers, and model architectures.

C.1 Robustness across Different Datasets

To demonstrate the robustness of value neurons, we investigate whether these findings generalize to other representative task types, such as coding and general instruction-following. To this end, we extend our verification to include MBPP+ [4] (coding) and IFEval [59] (general instruction-following).

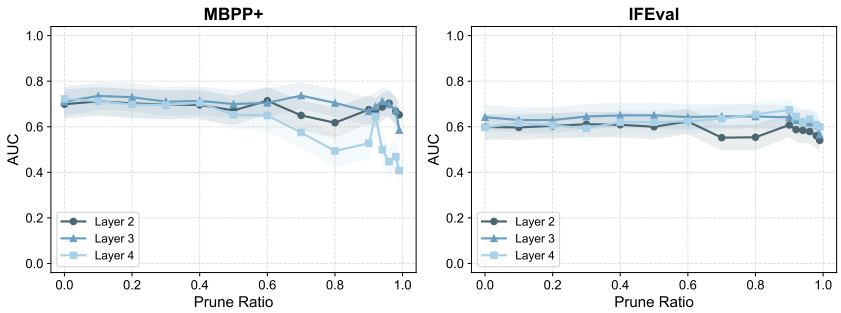


Figure 6: AUC curves for layers 2–4 of the Qwen-2.5-7B-SimpleRL-Zoo model on the MBPP+ and IFEval datasets. Shaded regions indicate uncertainty.

As shown in Figure 6, the AUC curve remains largely invariant to pruning, demonstrating that the existence of value neurons within the reward subsystem is consistently observed across different types of datasets.

C.2 Robustness across Different Models

We primarily focus on models based on the Qwen architecture. In addition to Qwen [5], the Llama [49] and Phi [22] architectures are also widely adopted in the open-source LLM community. To verify the applicability of the reward subsystem across these architectures, we conduct experiments on the MATH500 and ARC dataset using Qwen3.5-0.8B [51], Phi-3.5-mini-instruct [1], Llama-3.1-8B-Instruct [20], and Qwen-2.5-14B-SimpleRL-Zoo models.

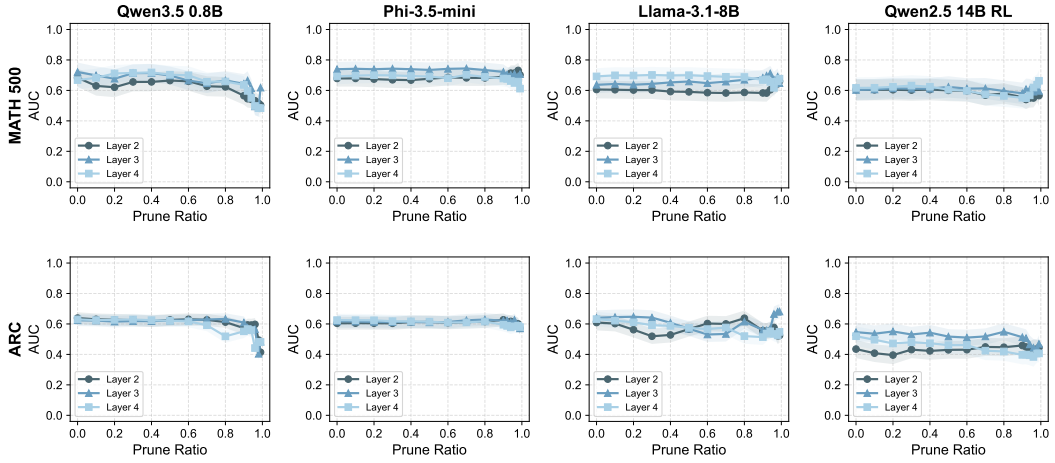


Figure 7: AUC curves for layers 2–4 of the Qwen-3.5-0.8B, Phi-3.5-mini-instruct, Llama-3.1-8B-Instruct, and Qwen-2.5-14B-SimpleRL-Zoo models on the MATH500 and ARC dataset. Shaded regions indicate uncertainty.

As illustrated in Figure 7, our conclusion regarding the existence of sparse value neurons remains valid for open-source models with different architectures and model scales.

C.3 Robustness across Different Layers

In this section, we examine the robustness of the reward subsystem across various layers. For this purpose, we utilize the Qwen-2.5-7B/14B-SimpleRL-Zoo model.

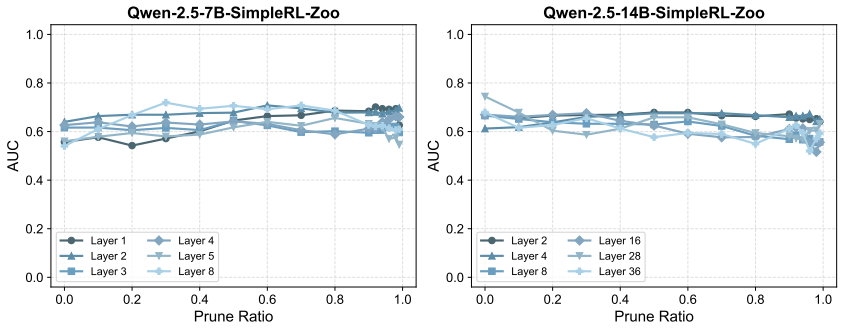


Figure 8: AUC curves for different layers of the Qwen-2.5-7B/14B-SimpleRL-Zoo model on the GSM8K dataset.

As shown in Figure 8, we find that the AUC remains largely stable as the pruning ratio rises.

D Transferability Across Different Datasets: More IoU Curves

In this section, we provide additional evidence regarding the transferability of value neurons across different datasets. As illustrated in Figure 9, we utilize layer 2 and layer 4 of the Qwen-2.5-14B-SimpleRL-Zoo model and compute pairwise IoU curves as a function of the pruning ratio for the GSM8K, MATH500, and ARC datasets. In these results, the IoU curves for any two datasets consistently exceed the random baseline. Furthermore, the observation that many IoU curves exhibit a significant upward trend as the pruning ratio approaches 1 remains valid for these layers as well.

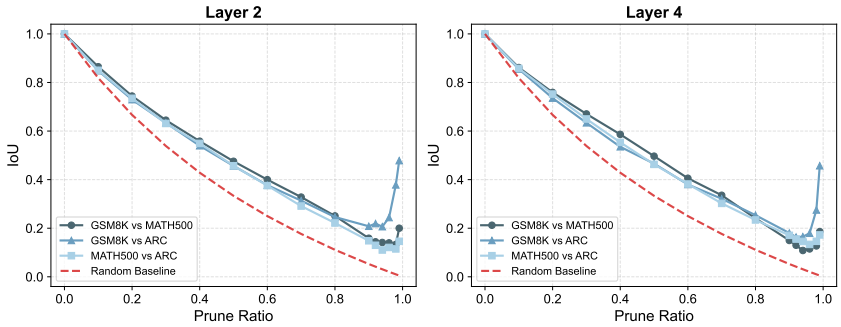


Figure 9: IoU as a function of the pruning ratio. The IoU values for value neurons across different datasets are significantly higher than the random baseline, indicating that for the same LLM, the positions of identified value neurons are closely correlated across tasks.

E Detailed Procedures for Hidden State Normalization

For each sample, we perform a single forward pass through the full sequence and extract hidden states from all L transformer layers. Let $h_{j,n,l}$ denote the hidden state at token position j for the n-th neuron in layer l.

To remove sample-specific baseline activation levels, we apply z-score normalization across the token dimension within each response. Let j_0 denote the token index at which the response begins. For

each neuron, we compute:

$$\bar{h}_{n,l} = \frac{1}{|\mathcal{T}|} \sum_{j \in \mathcal{T}} h_{j,n,l}, \quad \sigma_{n,l} = \sqrt{\frac{1}{|\mathcal{T}|} \sum_{j \in \mathcal{T}} (h_{j,n,l} - \bar{h}_{n,l})^2}, \tag{12}$$

where $\mathcal{T} = \{j_0, j_0 + 1, \ldots, J\}$ is the set of response token positions and J is the last token position. The normalized activation is:

$$\tilde{h}_{j,n,l} = \frac{h_{j,n,l} - \bar{h}_{n,l}}{\sigma_{n,l} + \epsilon}, \quad \epsilon = 10^{-8}. \tag{13}$$

Motivated by neuroscience, where dopaminergic responses are measured as average firing rates over a temporal window, we aggregate activations at the paragraph level rather than at individual token positions. Let $\mathcal{T}_t \subset \mathcal{T}$ denote the set of token positions corresponding to paragraph c_t. The paragraph-level representation is the mean of the normalized activations over its tokens:

$$\tilde{h}(s_t, l) = \left[\frac{1}{|\mathcal{T}_t|} \sum_{j \in \mathcal{T}_t} \tilde{h}_{j,n,l} \right]_{n=1}^{d} \in \mathbb{R}^d. \tag{14}$$

F Further Evidence for the Characteristics of Dopamine Neurons

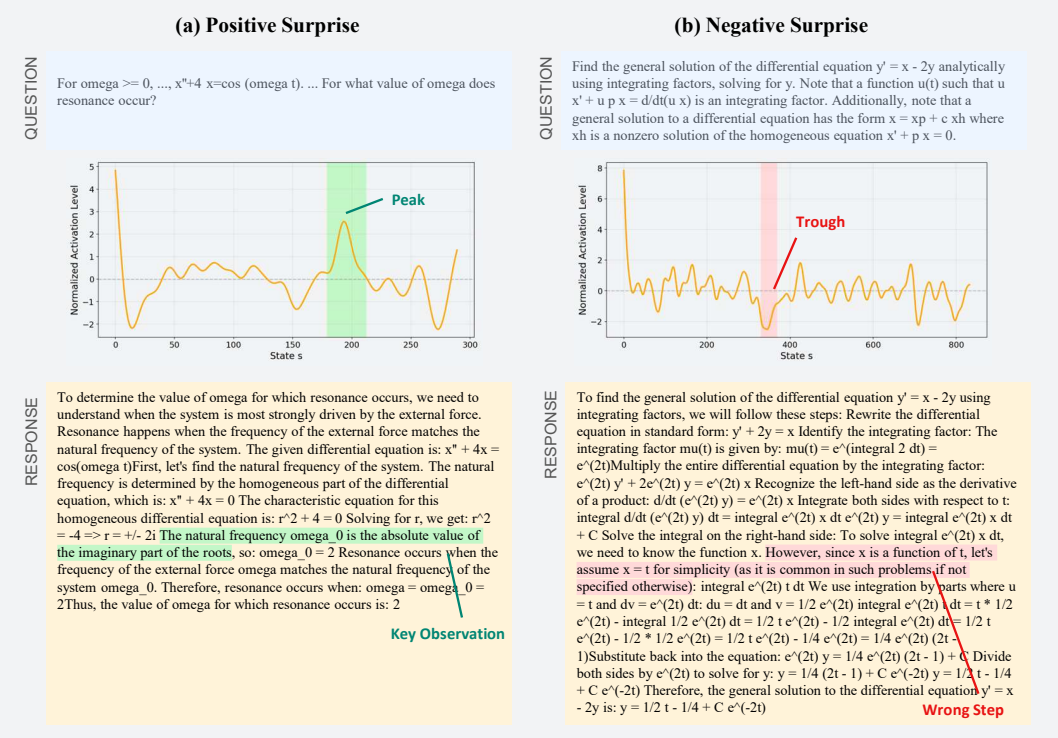


Figure 10: Dopamine neurons encode information regarding the model's prediction error for the current state. (a) **Positive Surprise:** The model initially lacks confidence in answering the problem but ultimately provides the correct solution. This neuron exhibits a significant peak when the model identifies a critical logical step and subsequently derives the final key result. (b) **Negative Surprise:** Conversely, the model begins with high confidence but fails to solve the problem correctly. The neuron displays a distinct trough at the exact moment a wrong step occurs.

To verify the robustness of the dopamine neurons' existence, we conduct further experiments using the Minerva Math dataset. We maintain the previously established experimental setup, selecting the top 50 neurons in each layer as dopamine neurons on the Minerva Math dataset. Notably, we observe

a significant overlap between these dopamine neurons and those identified earlier on the MATH500 dataset; in particular, the 1517-th neuron in layer 5 remains among the dopamine neurons.

Consequently, we continue to investigate the predictive capacity of this neuron. Our findings confirm that it consistently displays a period of low activation when the model initially predicts a high value but fails to obtain a reward, and a period of high activation when the model initially predicts a low value but ultimately succeeds.

As illustrated in Figure 10, this neuron encounters both a positive surprise and a negative surprise. In the case of the positive surprise shown in Figure 10(a), the model initially exhibits low confidence in completing the task. However, during the inference process (around the 200th token), the model derives a key conclusion, resulting in a high TD error; consequently, we observe a sharp spike in the neuron's activation level. Since the subsequent reasoning proceeds steadily, the TD error remains low as the following steps become predictable, leading to the neuron's activation returning to a relatively low state. These observations suggest that dopamine neurons exhibit higher activation levels when the model acquires unexpected rewards or makes significant progress.

Conversely, in the negative surprise shown in Figure 10(b), the model begins with high confidence. During the initial stage of inference, the model follows a correct path and even provides the critical modeling logic within the first 300 tokens. However, between the 300th and 400th tokens, the model commits an error, leading to a significant negative TD error and a corresponding suppression in the neuron's activation.

Thus, our visualization on the Minerva Math dataset consistently demonstrates a close correlation between the TD error during inference and neuronal activation levels, further substantiating the fundamental properties of dopamine neurons.

G Additional Intervention Controls

We compare against four additional baselines: **Massive Activations** [47], **Super Weights** [53], **Rogue Dimensions** [48], and **Attention Sink** [21]. We conduct these experiments in layer 2 of Qwen-2.5-7B-SimpleRL-Zoo on MATH500, GSM8K, and ARC, using value neurons identified by the probe trained on MATH.

Table 8: Additional intervention controls for layer 2 of Qwen-2.5-7B-SimpleRL-Zoo on MATH500, GSM8K, and ARC. Baseline denotes the unmodified model, and parentheses indicate changes in percentage points.

Method	**MATH500**	**GSM8K**	**ARC**
Baseline	75.2	93.0	82.2
Massive Activations	76.0 (+0.8)	91.6 (-1.4)	83.0 (+0.8)
Super Weights	74.6 (-0.6)	91.2 (-1.8)	80.4 (-1.8)
Rogue Dimensions	65.8 (-9.4)	81.6 (-11.4)	79.8 (-2.4)
Attention Sink	73.2 (-2.0)	93.0 (+0.0)	81.6 (-0.6)
Value Neurons	**37.0 (-38.2)**	**52.8 (-40.2)**	**66.2 (-16.0)**

As shown in Table 8, zeroing out value neurons causes larger performance drops.

H Correlation between Value Neurons and Dopamine Neurons

Given that value neurons and dopamine neurons both belong to the reward subsystem, we hypothesize that they are functionally interconnected and exert mutual influence. To verify this, we perform an ablation experiment where we zero out the activation values of value neurons in earlier layers and observe the resulting changes in the activation trajectories of dopamine neurons in subsequent layers. As a control, we compare this with the effect of zeroing out an equivalent number of randomly selected neurons in the same layers. If perturbing the value neurons leads to a significantly more pronounced alteration in the dopamine neurons' activation curves compared to the random perturbation, it would indicate a close relationship between value neurons and dopamine neurons.

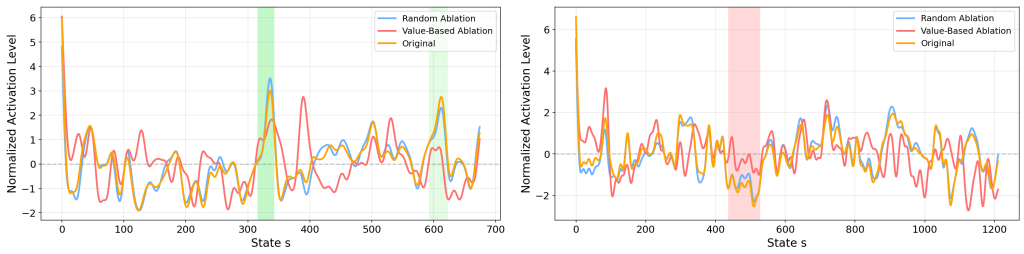


Figure 11: The activation curves of a dopamine neuron across states under different ablation conditions. The yellow line represents the original trajectory, the blue line shows the result of zeroing 20% random neurons, and the red line depicts the result of zeroing the top 20% value neurons. While random ablation has a minimal impact on the overall trend, value neuron ablation significantly alters the trajectory of the curve. This indicates a close correlation between value and dopamine neurons.

We use the Qwen-2.5-14B-SimpleRL-Zoo model evaluated on the MATH500 dataset, and compare the effects of zeroing out the top 20% of value neurons (identified by the highest L_1 norm in the value probe) versus zeroing out 20% of randomly selected neurons. We then analyze the deviations in the normalized activation levels of the 1517-th neuron in layer 5 relative to its original trajectory.

As shown in Figure 11, while the blue line (random ablation) remains largely consistent with the original yellow curve aside from minor numerical fluctuations, the red line (value neuron ablation) exhibits fundamental differences. Specifically, the positions of the activation peaks and troughs shift significantly, causing the neuron to lose the characteristic properties of a dopamine neuron encoding prediction error. These results demonstrate that value and dopamine neurons are closely related, as disrupting a small subset of value neurons is sufficient to significantly impair the predictive performance of dopamine neurons.

I Value Neuron Analysis at the Final Response Position

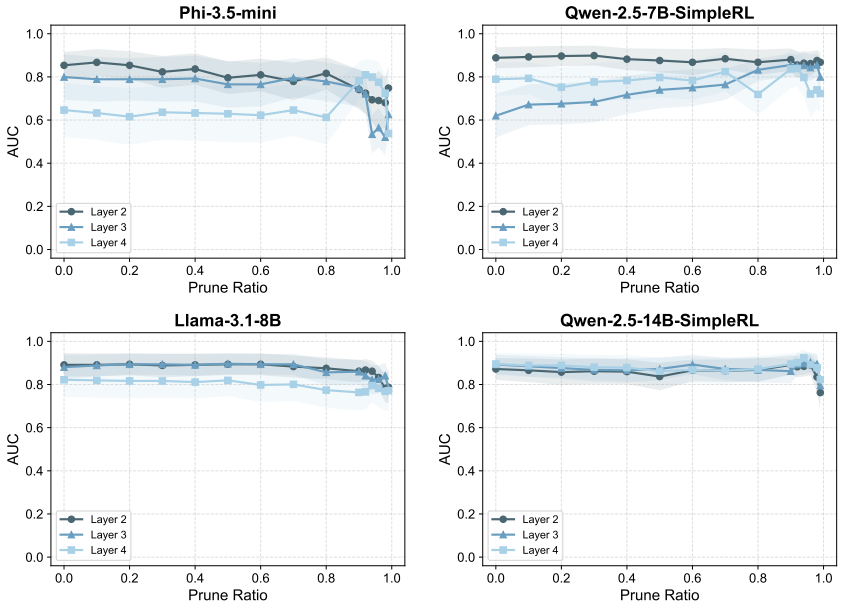


Figure 12: AUC curves at the final response position for four models on the Minerva Math dataset. The AUC remains mostly stable above 0.8 across a wide range of pruning ratios. Shaded regions indicate uncertainty.

The experiments in the main paper evaluate value neurons using hidden states at the initial input position s_0, i.e., after the problem input but before the model begins generating its response. This

setting tests whether the reward subsystem has already formed an assessment of the expected value prior to generation, but the signal at this position is relatively weak since the model has not yet begun reasoning.

A complementary setting is to evaluate value neurons at the last token of the model's completed response, where the reward-relevant signal can be expected to be stronger as the model has access to its full reasoning trajectory. We conduct this analysis using Phi-3.5-mini-instruct, Qwen-2.5-7B-SimpleRL-Zoo, Qwen-2.5-14B-SimpleRL-Zoo, and Llama-3.1-8B-Instruct on the Minerva Math [30] dataset.

As shown in Figure 12, the AUC remains mostly stable above 0.8, substantially higher than the pre-generation setting reported in the main paper. The consistency of this result across architectures and scales confirms that value neurons reliably encode reward information at both ends of the generation process, further strengthening the evidence for a sparse reward subsystem within LLM hidden states.

J Model Confidence Baseline Setup

Verbal Confidence. In this straightforward baseline, we input the question into the LLM and use the prompt, *'\n\n Rate your confidence in answering this question correctly from 0 to 1 (where 0 means no confidence and 1 means complete confidence). Only output a single number between 0 and 1. Do not try to solve the question.'* to elicit a confidence score. If the model fails to produce a valid numerical output, we resample until a score is obtained.

Token Confidence. In this baseline, we utilize the log probability ($\log p$) of the most likely token at the first generated position as a metric for the model's confidence.

Latent Correctness Direction (LCD). We re-implement the method proposed by [9] within our experimental setting. It is noteworthy that the baseline method utilizes the full set of neurons for prediction, so it cannot provide any guidance regarding the existence of the value neurons or the localization of neurons. In contrast, our approach utilizes significantly fewer hidden state dimensions than the baseline.

K Cross-Task Generalization of Trained Probes

The cross-dataset IoU analyses in Section 3.5 and Appendix D show that value neurons identified on different datasets overlap more than expected under random selection. To complement these analyses, we directly evaluate whether the trained probes transfer across tasks. Specifically, we apply the value and dopamine probes trained on MATH500 using Llama-3.1-8B-Instruct to GSM8K, evaluating their performance on model confidence prediction and process-reward-guided search, respectively.

Process-Reward-Guided Search. We use the transferred dopamine probe to score reasoning steps and guide inference-time search on GSM8K. As shown in Table 9, this approach achieves an accuracy of 94.4%, the highest mean accuracy among the evaluated methods.

Table 9: Cross-task evaluation of process-reward-guided search.

Method	Accuracy (%)
Greedy	90.7
Random	90.7
Implicit PRM	92.8
Dopamine Neurons (MATH500-trained)	**94.4**

Model Confidence Prediction. We also evaluate the transferred value probe for confidence prediction on GSM8K. Table 10 shows that it achieves an AUC of 0.692, the highest mean AUC among the evaluated methods.

Table 10: Cross-task evaluation of model confidence prediction.

Method	AUC
Token Confidence	0.570
Verbal Confidence	0.640
Question Length	0.661
LCD	0.641
Value Neurons (MATH500-trained)	**0.692**

Together, these results provide evidence that the trained value and dopamine probes retain predictive and practical utility when transferred from MATH500 to GSM8K, supporting cross-task generalization.

L Sensitivity to the TD-Error Threshold

To examine the sensitivity to the TD-error threshold, we vary τ in the filtering criterion $|\delta_t| > \tau$. We conduct this ablation on Qwen-2.5-7B-SimpleRL-Zoo with MATH500 across all layers, using three random seeds. Table 11 reports the mean Spearman correlation averaged across layers and seeds. As shown in Table 11, $\tau = 0.3$ achieves the highest mean Spearman correlation among the evaluated thresholds, supporting the threshold used in our experiments.

Table 11: Mean Spearman correlation for different TD-error thresholds on Qwen-2.5-7B-SimpleRL-Zoo with MATH500, averaged across all layers.

Threshold τ	**Mean ρ**
0.1	0.152
0.2	0.267
0.3	**0.281**
0.4	0.193

M Is the Reward Subsystem Already Present in Base Models or Created by RL Fine-Tuning?

To examine whether value neurons are already present before RL fine-tuning, we compare Qwen-2.5-7B with two RL-tuned models derived from it: Qwen-2.5-7B-SimpleRL-Zoo and Qwen2.5-7B-PPO-Zero. Both models are obtained through RLVR using different training pipelines.

We train value probes on MATH500 for all three models and compute the IoU between their identified value-neuron sets at different pruning ratios p. The results are averaged over layers 2–4 and compared against the random baseline. As shown in Table 12, the IoU between the base model and both RL-tuned models consistently exceeds the random baseline, even at high pruning ratios. These results support the presence of value-related representations in the base model before RL fine-tuning and show substantial overlap in the identified coordinates across two different RL pipelines.

N Additional Discussion

Internal Estimates of Success and Failure. Our results suggest that sparse hidden-state coordinates carry information about the expected success of the current reasoning state. Value-neuron probes outperform verbalized confidence on average, suggesting that internal representations contain success-related information that is not fully captured by verbalized confidence. This motivates investigating whether changes in value estimates can anticipate errors as reasoning unfolds.

Training Interventions and Steering. Value and dopamine probes offer lightweight value and process-reward signals that may guide post-training and inference-time search. Because these probes

Table 12: IoU between value-neuron sets identified in Qwen-2.5-7B and its two RL-tuned variants. Value probes are trained on MATH500, and results are averaged over layers 2–4. p denotes the pruning ratio.

p	Random Baseline	Base vs SimpleRL-Zoo	Base vs PPO-Zero	SimpleRL-Zoo vs PPO-Zero
0.50	0.333	0.723	0.705	0.676
0.80	0.111	0.513	0.525	0.454
0.90	0.053	0.429	0.408	0.362
0.96	0.020	0.326	0.316	0.266
0.98	0.010	0.290	0.310	0.235
0.99	0.005	0.283	0.290	0.236

estimate reward-related quantities under the current policy, they may require refreshing when the policy changes substantially. A natural next step is to investigate targeted steering of the identified coordinates and measure its effects on both probe predictions and actual task performance.

Relationship to the Implicit Q-Function View. The policy–value relationship in Equation 3 motivates searching for value-related information within policy representations. Our findings complement this view by showing that state value and step-level TD-error information can be predicted from sparse hidden-state coordinates.

Potential Connections to Evaluation-Aware Behavior. Internal estimates of expected reward may provide a useful perspective for studying evaluation awareness and reward seeking [39]. If models use such estimates when adapting their behavior to evaluation contexts, the identified coordinates could offer signals for investigating these adaptations. Another possible application is lightweight monitoring: discrepancies between internal value estimates and overt behavior might help identify evaluation-aware behavior. These connections remain hypotheses requiring direct evaluation.

Scope of the Neuroscience Analogy. The terms "value neurons" and "dopamine neurons" describe analogies in computational role: predicting state value and reward prediction errors, respectively. Here, neurons are residual-stream coordinates selected through supervised probing, rather than biologically characterized cell types. The analogy therefore guides interpretation and future experiments without establishing a biological correspondence.

O Potential Limitations and Broader Impacts

Limitations. We acknowledge the following limitations of the paper:

(1) While our benchmarks cover a diverse range of tasks including mathematics, reasoning, coding, and general instruction following, our method requires reward signals from the environment and has not been tested on tasks where such signals are difficult to obtain, such as open-ended generation.

(2) Due to resource and computational constraints, we have not investigated whether similar phenomena exist in models larger than 32B.

(3) Our intervention experiments offer preliminary insights, and potential confounds, such as probe stability across rollouts and large activation magnitudes, remain for future work. We thank Nell Watson for her interest in our work and for her independent pre-registered replication, which provided valuable insights and improved our understanding of these results.

Broader Impacts. Our work identifies sparse neurons that encode reward-related information in LLMs, which has implications for both interpretability and control. On the positive side, understanding where reward information resides can improve transparency in LLM reasoning. However, this capability also raises concerns. Neuron-level interventions on reward subsystems could potentially be used to manipulate model behavior in unintended ways. We encourage future work to investigate the robustness of these neurons under adversarial settings and to develop safeguards before deploying neuron-level control mechanisms in practice.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We clearly state our claims in the abstract and the introduction section. They accurately reflect the methods and the experiments.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We mentioned potential limitations in Appendix O.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

Justification: The paper does not include formal theoretical results.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We clearly describe our methods in Section 3 and Section 4.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide open code access.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The experimental settings are provided in Appendix A and Appendix J.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: We compute all curves via bootstrap resampling on the validation set.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The experiments compute resources are provided in Appendix A.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The paper conforms with the NeurIPS code of ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We discuss the broader impacts in Appendix O.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: They are properly credited, mentioned and respected.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: Yes, code is the only new asset. It is well documented and the documentation is provided alongside the code.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.