

Enhancing Uncertainty Estimation in LLMs with Expectation of Aggregated Internal Belief

Zeguan Xiao^{1*}, Diyang Dou^{1*}, Boya Xiong^{1*}, Yun Chen^{1,2†}, Guanhua Chen^{3†}

¹Shanghai University of Finance and Economics;

²MoE Key Laboratory of Interdisciplinary Research of Computation and Economics,
Shanghai University of Finance and Economics;

³Southern University of Science and Technology

Abstract

Large Language Models (LLMs) have achieved remarkable success across a wide range of natural language tasks, but often exhibit overconfidence and generate plausible yet incorrect answers. This overconfidence, especially in models undergone Reinforcement Learning from Human Feedback (RLHF), poses significant challenges for reliable uncertainty estimation and safe deployment. In this paper, we propose EAGLE (Expectation of AGgregated internaL bEief), a novel self-evaluation-based calibration method that leverages the internal hidden states of LLMs to derive more accurate confidence scores. Instead of relying on the model’s final output, our approach extracts internal beliefs from multiple intermediate layers during self-evaluation. By aggregating these layer-wise beliefs and calculating the expectation over the resulting confidence score distribution, EAGLE produces a refined confidence score that more faithfully reflects the model’s internal certainty. Extensive experiments on diverse datasets and LLMs demonstrate that EAGLE significantly improves calibration performance over existing baselines. We also provide an in-depth analysis of EAGLE, including a layer-wise examination of uncertainty patterns, a study of the impact of self-evaluation prompts, and an analysis of the effect of self-evaluation score range.

1 Introduction

Large Language Models (LLMs) have demonstrated impressive capabilities across a wide spectrum of natural language understanding and generation tasks (Brown et al. 2020; Wolf et al. 2020; Achiam et al. 2023; Touvron et al. 2023). Their ability to generate coherent and contextually relevant text has led to their increasing deployment in various real-world applications. However, a critical challenge remains: LLMs often produce plausible-sounding but incorrect answers, a phenomenon known as hallucination (Bai et al. 2022). This challenge is further exacerbated by the lack of reliable uncertainty estimation for the outputs of LLMs that have undergone Reinforcement Learning from Human Feedback (RLHF) (Kadavath et al. 2022; Tian et al. 2023), undermining their trustworthiness and limits their safe deployment. To

mitigate this issue, recent research has introduced calibration techniques, which aim to ensure that LLMs’ confidence of predictions accurately reflect their likelihood of correctness (Kadavath et al. 2022; Chen and Mueller 2023; Xiong et al. 2023).

Among the various research directions, a particularly important approach is leveraging the LLMs’ inherent capabilities for self-reflection and self-evaluation as a means of uncertainty estimation (Lin, Hilton, and Evans 2022; Kadavath et al. 2022; Tian et al. 2023; Chen and Mueller 2023; Huang et al. 2024). These methods often involve prompting the LLM to explicitly assess the uncertainty and/or correctness of its own generated responses, usually by eliciting verbalized confidence scores.

The motivation of our method stems from the observation that LLMs’ internal representations, particularly the hidden states at different layers, contain rich information about LLMs’ internal belief and can be used to reveal truthfulness of statements, knowledge discovery and behavior monitoring (Burns et al. 2022; Azaria and Mitchell 2023; Zou et al. 2023; Lai et al. 2025). In particular, concurrent work (Ji et al. 2025) demonstrates that intermediate-layer hidden states naturally separate high-confidence from low-confidence model predictions, suggesting that the model’s latent space already encodes information about its own uncertainty. Inspired by this, we propose a novel self-evaluation-based calibration method, EAGLE (Hidden Layer Aggregation), for uncertainty estimation. Unlike previous methods that rely on superficial verbalized confidence scores, our approach derives the confidence score by leveraging the internal representations of LLMs.

Specifically, EAGLE first begins by prompting the LLM to assess the uncertainty of its own answer by providing a numerical confidence score. It then extract the hidden representations corresponding to the self-evaluation token from multiple layers of the model. These hidden states are projected into the vocabulary space to obtain layer-wise logits, which are subsequently aggregated by averaging to form a refined logit vector. By applying a softmax over the logits associated with the candidate confidence scores, EAGLE produces a probability distribution that reflects the model’s internal belief about its answer’s correctness. The final confidence score is computed as the expectation over this distribution. This layer-wise aggregation strategy enables EA-

*Equal contribution.

†Corresponding Authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

GLE to capture richer and more nuanced confidence signals than methods based solely on the model’s final output, leading to improved calibration performance across diverse tasks and model architectures.

The main contributions of this work are as follows:

- We propose a novel training-free self-evaluation-based uncertainty estimation method, EAGLE, which leverages layer-wise hidden states to enhance calibration in LLMs.
- We conduct extensive empirical evaluations on three diverse datasets with state-of-the-art LLMs, showing that our proposed methods achieve superior calibration performance compared to existing baselines.
- We provide an in-depth analysis of EAGLE, including a layer-wise examination of uncertainty patterns, a study of the impact of self-evaluation prompts, and an analysis of the effect of self-evaluation score range.

2 Related Work

Language Models Calibration. Model calibration—the alignment of a model’s predicted probabilities with the true likelihood of correctness (Hendrycks and Gimpel 2016; Guo et al. 2017)—is crucial for user trust and informed decision-making in AI systems. In traditional supervised learning, post-hoc calibration techniques such as temperature scaling (Guo et al. 2017), Platt scaling (Platt et al. 1999), and isotonic regression (Niculescu-Mizil and Caruana 2005) are used to improve the reliability of model confidence estimates. However, directly applying these methods to LLMs, particularly those undergone RLHF (Kadavath et al. 2022; Tian et al. 2023), poses significant challenges.

Calibration in LLMs has been approached primarily through two lines of research. The first involves analyzing the diversity of multiple generated outputs, where techniques such as clustering-based metrics are used to examine the variance in the model’s responses (Kuhn, Gal, and Farquhar 2023; Lin, Trivedi, and Sun 2024). While effective in some scenarios, these methods are often computationally intensive due to the need for multiple generations per input, and the diversity observed does not always reliably indicate the correctness of individual outputs. A second, increasingly prominent direction leverages the self-evaluation capabilities of LLMs. In this approach, the model is prompted to assess the quality or correctness of its own generated responses, typically by providing a verbalized confidence score or categorical judgment (Lin, Hilton, and Evans 2022; Kadavath et al. 2022; Tian et al. 2023; Chen and Mueller 2023; Huang et al. 2024). Self-evaluation methods offer a more direct means of accessing the model’s internal assessment of its outputs, and have shown promise for improving calibration. However, most existing self-evaluation approaches rely on superficial verbalized confidence scores, which may fail to capture the nuanced confidence signals present in the model’s internal representations.

Internal Representations of LLMs. Analyzing the internal representations of LLMs has emerged as a crucial direction for understanding their inner workings and decision-

making processes. A promising line of research involves interpreting these hidden states by projecting them back into the vocabulary space. This technique allows researchers to treat the representation at each layer as a distribution over the vocabulary, offering insights into the model’s predictions and internal “beliefs” at different stages of processing (Nostalgebraist 2020; Geva et al. 2022). Recent studies have leveraged this approach to probe the model’s internal states for various purposes, such as discovering latent knowledge (Burns et al. 2022), assessing the truthfulness of statements (Azaria and Mitchell 2023), and even steering model behavior through representation engineering (Zou et al. 2023). These works collectively demonstrate that a model’s internal representations often contain more nuanced and faithful information than its final output layer alone.

Our work, EAGLE, is directly inspired by and builds upon this line of inquiry. While previous research has successfully used internal states to uncover a model’s knowledge or factual alignment, we are the first to apply this to the problem of uncertainty calibration in a self-evaluation context. Unlike methods that rely on the final verbalized confidence, EAGLE posits that a more robust and genuine measure of the model’s certainty can be derived by aggregating the logits from multiple layers. This approach allows us to tap into the rich, multi-layered confidence signals within the model, moving beyond superficial verbalized confidence score to achieve a more reliable estimation of uncertainty.

3 Preliminaries

Current mainstream LLMs are decoder-only transformers, which consist of a sequence of identical transformer decoder blocks stacked in series. Formally, a decoder-only transformer can be represented as a composition of functions:

$$f_{\text{transformer}} = f_{\text{unembed}} \circ \underbrace{f^{(L)} \circ f^{(L-1)} \circ \dots \circ f^{(1)}}_{L \text{ transformer layers}} \circ f_{\text{embed}}$$

where $f_{\text{embed}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times d}$ is the input embedding function that maps a sequence of n tokens to their corresponding d -dimensional embeddings, $f^{(l)}$ represents the l -th transformer decoder layer, and $f_{\text{unembed}} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times |V|}$ is the output projection that transforms the final hidden states to logits over the vocabulary V . The hidden state at layer l is denoted as $H^{(l)} \in \mathbb{R}^{n \times d}$, where each row $H_i^{(l)} \in \mathbb{R}^d$ represents the hidden representation of the i -th token at layer l . This hierarchical processing allows the model to progressively refine its understanding of the input sequence, with different layers capturing varying levels of linguistic and semantic information.

4 Methodology

Problem Definition

In this work, we focus on the problem of uncertainty estimation for LLMs, which is particularly challenging for LLMs that have undergone RLHF (Kadavath et al. 2022; Tian et al. 2023). A key concept in uncertainty estimation is calibration. A model is considered well-calibrated if its confidence

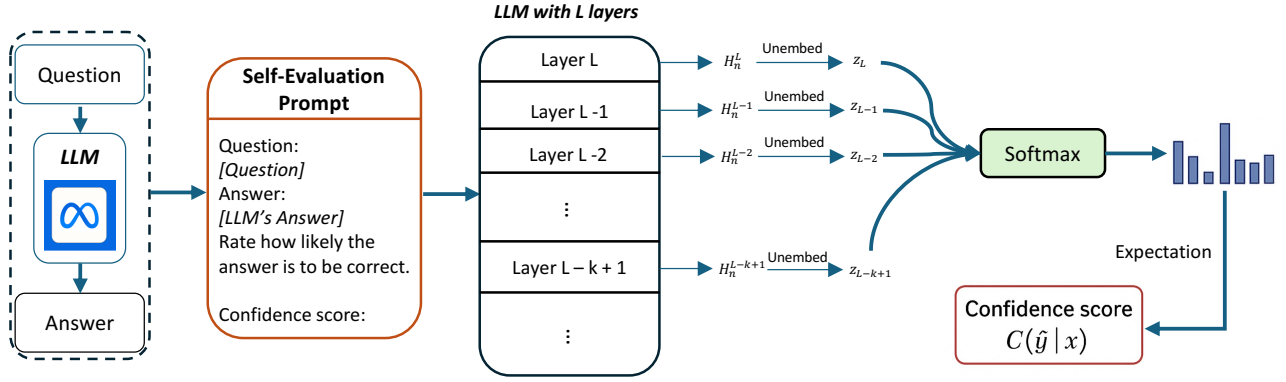


Figure 1: The EAGLE pipeline. First, the LLM generates an answer. Then, EAGLE prompts the LLM to self-evaluate the correctness of its answer, extracts the hidden states corresponding to the self-evaluation token from multiple layers, projects these hidden states to logits, and aggregates them. Finally, a softmax function is applied to the aggregated logits to compute the final confidence score.

of predictions accurately reflect their likelihood of correctness. Formally, for a given input x , its golden label y , and model prediction \hat{y} , a perfectly calibrated model satisfies:

$$\Pr[\hat{y} = y | C(\hat{y}|x) = p] = p, \quad \forall p \in [0, 1], \quad (1)$$

where $C(\hat{y}|x)$ is the confidence score assigned by the model to its prediction \hat{y} given input x .

The most commonly used metric for quantifying the degree of miscalibration is Expected Calibration Error (ECE). ECE is defined as $\mathbb{E}[|\Pr(\hat{y} = y | C(\hat{y}|x) = p) - p|]$, which measures the difference between confidence and accuracy. An empirical estimate of ECE is calculated by partitioning N samples into M bins $\{B_1, B_2, \dots, B_M\}$ according to the confidence score C predicted by the model. The ECE is then formulated as:

$$\text{ECE} = \frac{1}{N} \sum_{m=1}^M |B_m| |\text{acc}(B_m) - \text{conf}(B_m)|, \quad (2)$$

where $\text{acc}(B_m)$ and $\text{conf}(B_m)$ denote the average accuracy and confidence within bin B_m , respectively. A smaller ECE indicates better calibration performance of the model.

Motivation

Our motivation stems from a growing body of work demonstrating that the internal representations of LLMs, specifically the hidden states across different layers, encode a wealth of information that is not always reflected in the final output. Research has shown that these internal states can be probed to discover latent knowledge, track the truthfulness of statements, and monitor the model’s behavior (Burns et al. 2022; Azaria and Mitchell 2023; Zou et al. 2023; Lai et al. 2025). For instance, Burns et al. (2022) found that a simple probe applied to hidden states could consistently reveal a model’s latent knowledge, even when the model’s final output was incorrect. Similarly, Zou et al. (2023) demonstrated that representations of concepts could be identified and steered within the model. In particular, concurrent work (Ji et al. 2025) demonstrates that intermediate

layer hidden states can separate high-confidence from low-confidence model predictions, suggesting that the model’s latent space already encodes information about its own uncertainty.

Inspired by these empirical findings, we hypothesize that by aggregating information directly from hidden states of intermediate layers, we can construct a more robust and better-calibrated uncertainty estimate. Furthermore, we posit that relying on a single point estimate for the confidence score, such as selecting the most probable score, may fail to capture the full extent of the model’s uncertainty, as the entire probability distribution over possible scores contains richer information about the model’s internal belief state. Based on these insights, our approach, which we call EAGLE, taps into the deeper, multi-layered internal reasoning of the model, bypassing the final layer’s potential for overconfident posturing. It then uses the expectation of the confidence score distribution to capture a more nuanced and faithful measure of the model’s internal belief.

EAGLE: Expectation of Aggregated Internal Belief

Given a question x and the LLM’s generated answer \hat{y} , we prompt the LLM to evaluate its own answer. The self-evaluation prompt is designed to elicit a confidence score about the correctness of the answer.

After the forward pass of the LLM, we obtain the hidden states $H^{(l)} \in \mathbb{R}^{n \times d}$ at each layer, where $H_n^{(l)} \in \mathbb{R}^d$ denotes the hidden state corresponding to the token used to predict the confidence score in each layer. Rather than relying solely on the hidden states from the final layer to decode a confidence score ($f_{\text{unembed}} \circ H_n^{(L)}$), we instead take into account the information from all of the last k layers to obtain a more nuanced confidence estimate. The rationale behind this approach is that different layers capture varying levels of information and reasoning, which can provide a richer understanding of the model’s confidence in its self-evaluation.

Self-Evaluation Prompt

Question:
[QUESTION]
Answer:
[LLM'S ANSWER]
Rate how likely the answer is to be correct using a number from 0 to 9:
- 0 = Very likely incorrect
- 5 = Uncertain or partially correct
- 9 = Very likely correct

Confidence score (0-9):

Concretely, for each of the k selected layers l , we map the hidden state $H_n^{(l)}$ to a logit $z_n^{(l)} \in \mathbb{R}^{|V|}$:

$$z_n^{(l)} = f_{\text{unembed}} \circ H_n^{(l)}. \quad (3)$$

Subsequently, we weight and aggregate these logits from different layers:

$$z_n = \sum_{l=1}^k w_l \cdot z_n^{(l)}, \quad (4)$$

where w_l is the weight assigned to the logits from layer l . In our implementation, we set $w_l = 1/k$ for all layers, meaning each layer contributes equally to the final logit.

Afterwards, we apply the softmax function to the aggregated logits z_n , considering only the logits corresponding to the confidence score tokens. Specifically, we restrict the calculation to the set $S = \{0, 1, 2, \dots, 9\}$, which represents all possible confidence scores. This yields a probability distribution over the candidate confidence scores:

$$w_s = \frac{\exp(z_{n,s})}{\sum_{s' \in S} \exp(z_{n,s'})}. \quad (5)$$

Rather than merely selecting the confidence score with the highest probability like previous works (Lin, Hilton, and Evans 2022; Kadavath et al. 2022; Tian et al. 2023; Chen and Mueller 2023; Huang et al. 2024), we utilize the entire distribution to derive a more robust estimate of the model’s confidence. This approach allows us to capture the nuanced uncertainty in the model’s self-evaluation, avoiding reliance on a single point estimate that may not fully represent the model’s internal belief. Formally, the final confidence score, $C(\hat{y}|x)$, is calculated as the examination of the confidence score distribution:

$$C(\hat{y}|x) = \sum_{s \in S} w_s \cdot s \quad (6)$$

In summary, EAGLE introduces two key innovations for self-evaluation-based calibration: (1) it aggregates hidden states from multiple layers, rather than relying solely on the final layer, to capture richer and more robust confidence signals; and (2) it computes the final confidence score as the expectation over the entire confidence score distribution, rather than selecting only the most probable score. Together, these strategies enable EAGLE to more faithfully reflect the model’s internal belief about its answer’s correctness.

5 Experiments

Setup

Datasets. We evaluate the calibration performance of various methods on both open-ended and multiple-choice datasets. For open-ended tasks, we use rc.nocontext subset of TriviaQA (Joshi et al. 2017) and GSM8k (Cobbe et al. 2021). For multiple-choice tasks, we include MMLU (Hendrycks et al. 2021). The final evaluation is conducted on the TriviaQA validation set, GSM8k test set, and MMLU test set. For hyperparameter tuning, we use an equal-size subset of the TriviaQA and GSM8k training sets, and the official validation set for MMLU. These datasets cover diverse domains and reasoning complexities, enabling a comprehensive assessment of calibration across different tasks.

Models. We evaluate our method on two prominent families of LLMs: Qwen2.5 (Team 2024) and Llama3 (Grattafiori et al. 2024). Specifically, we include Qwen2.5 models of varying scales (7B and 72B) and Llama3 models (8B and 70B). These models are selected to cover a wide range of parameter sizes and architectures, enabling a comprehensive assessment of calibration performance across different model capacities. For clarity, we refer to Qwen2.5-7B-Instruct and Meta-Llama-3-8B-Instruct as Qwen2.5 7B and Llama3 8B, respectively, and adopt similar abbreviations for other model variants throughout the paper.

Baselines. We compare our method against several established approaches for confidence estimation:

- **Semantic Entropy (SE)** (Kuhn, Gal, and Farquhar 2023): An entropy measure that clusters semantically equivalent answers before computing uncertainty, addressing linguistic variations of correct responses.
- **Self-consistency (Self-cons.)** (Chen and Mueller 2024): Confidence derived from agreement among multiple sampled answers, with higher agreement indicating greater certainty.
- **P(True)** (Kadavath et al. 2022): The model’s self-evaluation of whether its answer is correct, obtained by prompting it to classify answers as true/false.
- **Contextualized Sequence Likelihood (CSL)** (Lin, Trivedi, and Sun 2024): A weighted sequence likelihood that uses attention mechanisms to focus on relevant tokens when computing answer probabilities.

Evaluation Metrics. To comprehensively assess the quality of confidence scores, we consider two orthogonal evaluation tasks: calibration and failure prediction. Calibration measures how closely the predicted confidence aligns with empirical accuracy. We quantify calibration using ECE, which computes the average absolute difference between predicted confidence and observed accuracy across confidence intervals, as defined in Eq 2. For all experiments, we use 10 bins for ECE calculation. Failure prediction, in contrast, evaluates whether the model assigns higher confidence to correct predictions and lower confidence to incorrect ones, reflecting the discriminative power of the confidence scores. For this, we report the Area Under the Receiver Operating Characteristic Curve (AUROC), which measures the

Metric	Dataset	Llama3 8B					Llama3 70B				
		SE	Self-cons.	P(true)	CSL	EAGLE	SE	Self-cons.	P(true)	CSL	EAGLE
ECE (↓)	TriviaQA	15.5	27.7	23.3	28.7	1.7	19.3	17.1	16.0	18.1	2.0
	GSM8K	17.1	25.4	25.1	6.3	7.6	16.2	5.0	7.5	7.5	4.9
	MMLU	5.1	7.3	37.9	39.2	0.4	8.8	5.0	21.5	22.2	2.4
	Average	12.6	20.1	28.8	24.7	3.2	14.8	9.0	15.0	15.9	3.1
AUROC (↑)	TriviaQA	60.0	55.7	60.6	53.0	61.5	66.2	59.1	63.0	52.0	70.4
	GSM8K	62.1	51.5	67.5	67.7	68.3	58.9	63.3	55.3	50.0	64.7
	MMLU	54.9	54.8	53.3	50.0	55.1	55.5	55.5	53.2	50.0	57.1
	Average	59.0	54.0	60.5	56.9	61.6	60.2	59.3	57.2	50.7	64.1
Metric	Dataset	Qwen2.5 7B					Qwen2.5 72B				
		SE	Self-cons.	P(true)	CSL	EAGLE	SE	Self-cons.	P(true)	CSL	EAGLE
ECE (↓)	TriviaQA	31.6	32.1	31.2	67.4	22.1	18.6	13.8	13.9	14.5	1.7
	GSM8K	9.8	16.8	9.5	15.5	3.0	16.0	9.7	7.5	4.5	1.2
	MMLU	18.1	22.5	24.4	36.6	16.8	23.9	22.0	22.3	18.1	15.7
	Average	19.8	23.8	21.7	39.8	14.0	19.5	15.2	14.6	12.4	6.2
AUROC (↑)	TriviaQA	52.3	51.0	62.2	52.1	64.9	53.5	54.9	54.0	52.0	70.4
	GSM8K	54.3	49.7	74.0	71.7	72.8	60.1	64.1	57.0	53.0	66.5
	MMLU	69.7	69.9	70.3	51.5	73.9	53.8	54.3	54.0	51.9	57.4
	Average	58.8	56.9	68.8	58.4	70.5	55.8	57.8	55.0	52.3	64.8

Table 1: Calibration results of different uncertainty estimation methods on Llama3 and Qwen2.5 models across multiple datasets. Lower ECE and higher AUROC indicate better calibration and discrimination, respectively. All metrics are given by $\times 100$. The best results are highlighted in bold.

ability of confidence scores to distinguish between correct and incorrect outputs. Together, ECE and AUROC provide a comprehensive evaluation of both the reliability and discriminative utility of model confidence scores.

Main Results

Table 1 presents the main experimental results, comparing the calibration performance of our proposed method, EAGLE, against several baselines¹. The results demonstrate that EAGLE consistently achieves superior calibration performance.

Across all evaluated models and datasets, EAGLE shows a significant improvement in ECE, the primary metric for calibration. For instance, on the TriviaQA dataset, EAGLE reduces the ECE of the Llama3 8B model to 1.7, a substantial improvement over the best baseline (SE) which scored 15.5. Similarly, for the Llama3 70B model, EAGLE achieves an ECE of 2.0, far surpassing the next best baseline’s 16.0. This trend of strong ECE reduction holds for the Qwen2.5 models as well, where EAGLE consistently delivers the lowest ECE scores, indicating a much better alignment between the model’s confidence and its actual accuracy. On average, EAGLE achieves the lowest ECE across all settings, highlighting its effectiveness and robustness.

¹We observe that our reported CSL results differ from those presented in the original paper (Lin, Trivedi, and Sun 2024). The discrepancy arises from the use of different model types: the original CSL results were obtained using base LLMs, whereas our experiments employ instruction-tuned (RLHF) models. RLHF models are known to exhibit weaker calibration compared to base LLMs (Tian et al. 2023), leading to the observed divergence.

In addition to superior calibration, EAGLE also demonstrates strong discriminative power, as measured by the AUROC metric. It consistently achieves the highest or highly competitive AUROC scores, indicating its ability to effectively distinguish between correct and incorrect answers. For example, with the Llama3 70B model on TriviaQA, EAGLE achieves an AUROC of 70.4, which is significantly higher than all baselines. While some baselines occasionally achieve a competitive AUROC on specific tasks (e.g., P(true) on GSM8K with Qwen2.5 7B), they do so with a much poorer ECE. In contrast, EAGLE provides a better balance, excelling in both calibration and failure prediction.

Ablation Study

To dissect the components of our proposed EAGLE method and validate our design choices, we conduct a comprehensive ablation study. We investigate three key aspects of our framework: (1) the use of hidden layers for extracting model’s internal belief (using only the final layer versus aggregating the last- n layers), (2) the aggregation strategy (aggregating layer-wise probabilities versus layer-wise logits), and (3) the method for deriving the final confidence score (taking the score with the maximum probability versus computing the expectation over the score distribution). The results of these ablations are presented in Table 2.

Our findings consistently highlight the superiority of the full EAGLE configuration. First, comparing the use of only the **last layer** against aggregating the **last- n layers**, we observe a marked improvement in calibration performance when information from multiple layers is integrated. This supports our core hypothesis that a more robust uncertainty

Strategies For EAGLE						Llama3 8B		Qwen2.5 7B	
Last layer	Last-n layers	Prob. agg.	Logits agg.	Max.	Exp.	TriviaQA	GSM8K	TriviaQA	GSM8K
✓	×	×	×	✓	×	27.9	25.4	31.6	9.8
✓	×	×	×	×	✓	18.8	22.7	30.5	8.5
×	✓	✓	×	✓	×	23.1	25.4	31.6	9.8
×	✓	✓	×	×	✓	63.5	58.2	64.1	84.6
×	✓	×	✓	✓	×	28.3	25.4	31.8	10.1
×	✓	×	✓	×	✓	1.7	7.6	22.1	3.0

Table 2: Ablation study on the components of EAGLE. We compare different strategies for layer selection (last layer vs. last-n layers), aggregation method (probability vs. logits), and final score derivation (max vs. expectation). The reported metric is ECE (\downarrow) on TriviaQA and GSM8K datasets for Llama3 8B and Qwen2.5 7B models. The best results are highlighted in bold.

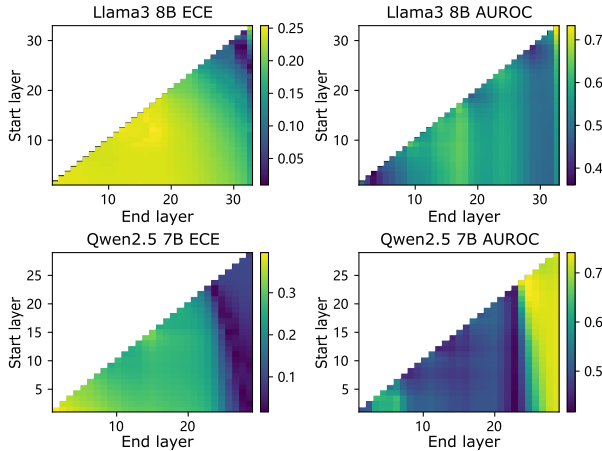


Figure 2: Layer-wise calibration performance on GSM8K for Llama3 8B and Qwen2.5 7B. The x-axis and y-axis represent the end and start layers for hidden state aggregation, respectively. Darker colors in ECE heatmaps and brighter colors in AUROC heatmaps indicate better calibration.

signal can be obtained by tapping into the representations distributed across several layers, rather than relying solely on the final output layer.

Second, we examine the aggregation method. The results clearly indicate that **logits aggregation (Logits agg.)** significantly outperforms **probability aggregation (Prob. agg.)**. Aggregating logits before applying the softmax function preserves a richer, more nuanced signal of the model’s internal confidence. In contrast, aggregating probabilities, which are post-softmax and thus normalized at each layer, leads to a loss of information and results in substantially poorer calibration.

Finally, we compare two approaches for deriving the final confidence score from the aggregated distribution. Using the **expectation (Exp.)** of the score distribution consistently yields better-calibrated results than simply taking the score with the **maximum probability (Max.)**. This demonstrates that the full distribution over the confidence scores contains valuable information, and computing its expected value provides a more accurate and stable estimate of the model’s true uncertainty.

Model	Method	Acc
Llama3 8B	Baseline	71.3
	EAGLE	74.8
Qwen2.5 7B	Baseline	67.5
	EAGLE	69.4

Table 3: EAGLE for answer selection. We demonstrate the accuracy of TriviaQA using different methods.

6 Analyses

Calibration Performance Among Layers

To understand how different layers contribute to uncertainty estimation, we conduct a layer-wise analysis of our proposed method. We vary the range of layers from which hidden states are aggregated—from a start layer m to an end layer n —and evaluate the resulting calibration performance. Figure 2 illustrates the ECE and AUROC scores for different layer combinations on the GSM8K dataset for both Llama3 8B and Qwen2.5 7B.

The heatmaps reveal a clear and consistent pattern: the calibration performance, as measured by both ECE and AUROC, is significantly better when aggregating hidden states from the later layers of the models. For both Llama3 and Qwen2.5, the optimal performance (lowest ECE and highest AUROC) is concentrated in the top-right corner of the plots. This indicates that the most reliable confidence signals are encoded in the final few layers.

While aggregating a block of layers from an intermediate start layer m to an end layer n can be effective, our analysis shows that simply aggregating the last few layers (i.e., a ‘last-n’ strategy) already yields near-optimal results. The performance gain from selecting a more complex ‘m-n’ block from earlier or middle parts of the model is marginal compared to the strong performance of the final layers. Therefore, for simplicity and efficiency, using a ‘last-n’ aggregation strategy is a robust and sufficient choice for EAGLE.

EAGLE for Answer Selection

One straightforward application of EAGLE is to apply it to select more reliable answers from LLMs. Here we conduct experiments on the TriviaQA dataset on both Qwen2.5

7B and Llama3 8B models. Specifically, we generate 5 different answers for each question with a temperature of 1.0, and use the first answer as the baseline. We then compare its accuracy with the answer selected based on the highest EAGLE confidence score. As shown in Table 3, EAGLE consistently improves the accuracy by selecting more reliable answers, outperforming the baseline on both models.

Sensitivity to Prompts

To investigate the robustness of EAGLE to different self-evaluation prompt formulations, we evaluate the calibration performance using two variants of self-evaluation prompts. The details of these prompt variants are shown below:

Prompt Variant A (Without score criteria)

Question:
[QUESTION]
Answer:
[LLM'S ANSWER]
Rate how likely the answer is to be correct using a number from 0 to 9:
Confidence score (0-9):

Prompt Variant B (CoT)

Question:
[QUESTION]
Answer:
[LLM'S ANSWER]
Rate how likely the answer is to be correct using a number from 0 to 9:
- 0 = Very likely incorrect
- 5 = Uncertain or partially correct
- 9 = Very likely correct
You should first provide an explanation in a few sentences, then provide your score.
The output format should look as follows:
Explanation: <write your explanation here>
Confidence score (0-9): <give your score here>

The results in Table 4 show that our method exhibits robustness to the prompt, although the choice of prompt does have an impact on performance. Our main prompt, which provides clear criteria for the scores, generally yields the best or highly competitive results. For instance, on the GSM8K dataset, all three prompt variants produce similar ECE and AUROC scores for both models, indicating that EAGLE is not very sensitive in this scenario. However, on TriviaQA, the performance varies more. The CoT prompt (Variant B) leads to a noticeable degradation in calibration for both models. This suggests that while CoT can encourage reasoning, it may also introduce noise into the self-evaluation process, which can interfere with the internal confidence signals that EAGLE relies on. Overall, while the performance is not entirely independent of the prompt design, EAGLE demonstrates reasonable stability, particularly when provided with a clear and direct self-evaluation format.

Model	Variant	TriviaQA		GSM8K	
		ECE	AUROC	ECE	AUROC
Llama3 8B	Main	1.7	61.5	7.6	68.3
	A	4.4	60.1	6.6	65.9
	B	13.6	57.5	5.8	65.0
Qwen2.5 7B	Main	22.1	64.9	3.0	72.8
	A	24.5	60.2	3.2	71.1
	B	28.5	58.5	2.5	72.5

Table 4: Comparison of prompt for EAGLE method. Results show ECE and AUROC performance across different prompt variants on TriviaQA and GSM8K datasets for Llama3 8B and Qwen2.5 7B models. All metrics are given by $\times 100$. Main Prompt results are taken from Table 1.

Score Range	ECE (\downarrow)	AUROC (\uparrow)
0 to 9	7.6	68.3
0 to 4	8.6	64.7
0 to 1	11.1	61.0

Table 5: Effect of confidence score range on calibration performance. Results are shown for Llama3 8B on GSM8k dataset. All metrics are given by $\times 100$.

Effect of Confidence Score Range

To assess the impact of the confidence score’s granularity on our method’s performance, we conducted an experiment varying the score range in the self-evaluation prompt. We tested three different ranges: 0–9, 0–4, and 0–1. The results for the Llama3 8B model on the GSM8k dataset are presented in Table 5.

The findings indicate a clear relationship between the score range and calibration performance. The 0–9 range yields the best results, achieving the lowest ECE (7.6) and the highest AUROC (68.3). As the range narrows, performance deteriorates. The 0–4 range shows a higher ECE and lower AUROC, and the binary 0–1 range performs the worst. This suggests that a wider, more granular scale allows the model to express its uncertainty more precisely. The fine-grained range enables EAGLE to capture more nuanced signals from the model’s internal belief, leading to better-calibrated confidence scores.

7 Conclusion

In this paper, we introduced EAGLE, a novel self-evaluation-based calibration method for LLMs that leverages the model’s internal representations to derive a more accurate and robust uncertainty estimate. By aggregating hidden states from multiple layers and the expectation of the confidence score distribution, EAGLE effectively captures the model’s true belief about its answer’s correctness. Our extensive experiments across various models and datasets demonstrate that EAGLE significantly outperforms existing methods in both calibration and failure prediction tasks. The ablation study further validates the effectiveness of our design choices.

Acknowledgements

This project was supported by National Natural Science Foundation of China (No. 62306132), Guangdong Basic and Applied Basic Research Foundation (No. 2025A1515011564) and Natural Science Foundation of Shanghai (No. 25ZR1402136). This work was done during Zeguan’s internship at SUSTech. We thank the anonymous reviewers for their insightful feedback on this work.

References

- Achiam, O. J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; Avila, R.; and Others. 2023. GPT-4 Technical Report.
- Azaria, A.; and Mitchell, T. 2023. The Internal State of an LLM Knows When It’s Lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 967–976.
- Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; and Others. 2022. Constitutional AI: Harmlessness from AI Feedback. *ArXiv*, abs/2212.08073.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; and Others. 2020. Language Models are Few-Shot Learners. *ArXiv*, abs/2005.14165.
- Burns, C.; Ye, H.; Klein, D.; and Steinhardt, J. 2022. Discovering latent knowledge in language models without supervision. *arXiv preprint arXiv:2212.03827*.
- Chen, J.; and Mueller, J. 2023. Quantifying Uncertainty in Answers from any Language Model and Enhancing their Trustworthiness. In *Annual Meeting of the Association for Computational Linguistics*.
- Chen, J.; and Mueller, J. 2024. Quantifying Uncertainty in Answers from any Language Model and Enhancing their Trustworthiness. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5186–5200.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *arXiv:2110.14168*.
- Geva, M.; Caciularu, A.; Wang, K.; and Goldberg, Y. 2022. Transformer Feed-Forward Layers Build Predictions by Promoting Concepts in the Vocabulary Space. *ArXiv*, abs/2203.14680.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On Calibration of Modern Neural Networks. *ArXiv*, abs/1706.04599.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Hendrycks, D.; and Gimpel, K. 2016. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. *ArXiv*, abs/1610.02136.
- Huang, Y.; Liu, Y.; Thirukovalluru, R.; Cohan, A.; and Dhingra, B. 2024. Calibrating Long-form Generations from Large Language Models. In *Conference on Empirical Methods in Natural Language Processing*.
- Ji, Z.; Yu, L.; Koishekenov, Y.; Bang, Y.; Hartshorn, A.; Schelten, A.; Zhang, C.; Fung, P.; and Cancedda, N. 2025. Calibrating Verbal Uncertainty as a Linear Feature to Reduce Hallucinations. *arXiv preprint arXiv:2503.14477*.
- Joshi, M.; Choi, E.; Weld, D. S.; and Zettlemoyer, L. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.
- Kadavath, S.; Conerly, T.; Askell, A.; Henighan, T.; Drain, D.; Perez, E.; Schiefer, N.; Hatfield-Dodds, Z.; DasSarma, N.; Tran-Johnson, E.; et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Kuhn, L.; Gal, Y.; and Farquhar, S. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*.
- Lai, P.; Zheng, J.; Cheng, S.; Chen, Y.; Li, P.; Liu, Y.; and Chen, G. 2025. Beyond the Surface: Enhancing LLM-as-a-Judge Alignment with Human via Internal Representations. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Lin, S. C.; Hilton, J.; and Evans, O. 2022. Teaching Models to Express Their Uncertainty in Words. *Trans. Mach. Learn. Res.*, 2022.
- Lin, Z.; Trivedi, S.; and Sun, J. 2024. Contextualized Sequence Likelihood: Enhanced Confidence Scores for Natural Language Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 10351–10368.
- Niculescu-Mizil, A.; and Caruana, R. 2005. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, 625–632.
- Nostalgebraist. 2020. Interpreting GPT: the logit lens. <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>. Accessed: 2025-12-04.
- Platt, J.; et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3): 61–74.
- Team, Q. 2024. Qwen2.5: A Party of Foundation Models. <https://qwenlm.github.io/blog/qwen2.5/>. Accessed: 2025-12-04.
- Tian, K.; Mitchell, E.; Zhou, A.; Sharma, A.; Rafailov, R.; Yao, H.; Finn, C.; and Manning, C. D. 2023. Just Ask for Calibration: Strategies for Eliciting Calibrated Confidence Scores from Language Models Fine-Tuned with Human Feedback. *ArXiv*, abs/2305.14975.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.;

Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *ArXiv*, abs/2302.13971.

Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Scao, T. L.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. M. 2020. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *arXiv*:1910.03771.

Xiong, M.; Hu, Z.; Lu, X.; Li, Y.; Fu, J.; He, J.; and Hooi, B. 2023. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint arXiv:2306.13063*.

Zou, A.; Phan, L.; Chen, S.; Campbell, J.; Guo, P.; Ren, R.; Pan, A.; Yin, X.; Mazeika, M.; Dombrowski, A.-K.; et al. 2023. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*.