

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LLM MICROSCOPE: WHAT MODEL INTERNALS RE- VEAL ABOUT ANSWER CORRECTNESS AND CONTEXT UTILIZATION

006 **Anonymous authors**

007 Paper under double-blind review

## 010 ABSTRACT

013 Although large language models (LLMs) have tremendous utility, trustworthiness  
 014 is still a chief concern: models often generate incorrect information with high  
 015 confidence. While contextual information can help guide generation, identifying  
 016 when a query would benefit from retrieved context and assessing the effectiveness  
 017 of that context remains challenging. In this work, we operationalize interpretability  
 018 methods to ascertain whether we can predict the correctness of model outputs from  
 019 the model’s activations alone. We also explore whether model internals contain  
 020 signals about the efficacy of external context. We consider correct, incorrect, and  
 021 irrelevant context and introduce metrics to distinguish amongst them. Experiments  
 022 on six different models reveal that a simple classifier trained on intermediate layer  
 023 activations of the first output token can predict output correctness with about 75%  
 024 accuracy, enabling early auditing. Our model-internals-based metric significantly  
 025 outperforms prompting baselines at distinguishing between correct and incorrect  
 026 context, guarding against inaccuracies introduced by polluted context. These  
 027 findings offer a lens to better understand the underlying decision-making processes  
 028 of LLMs.

## 029 1 INTRODUCTION

030 Large language models (LLMs) have shown tremendous utility in many domains, including those that  
 031 require accurately answering factual queries (Scao et al., 2023; Grattafiori et al., 2024; Groeneveld  
 032 et al., 2024). However, trustworthiness remains a chief concern: LLMs often generate convincing,  
 033 but thoroughly incorrect and non-factual responses, termed *hallucinations* (Bang et al., 2023; Huang  
 034 et al., 2025; Guerreiro et al., 2023).

035 Recently, retrieval-augmented generation (RAG) has been proposed to mitigate this problem (Lewis  
 036 et al., 2020). Although RAG is effective, two challenges remain: confidence estimation to identify  
 037 uncertain examples where an LLM requires external context and efficacy evaluation to score the  
 038 utility of the retrieved external context. Confidence estimation is challenging as LLMs are poorly  
 039 calibrated: models often assign high probabilities to incorrect generations, making it difficult to detect  
 040 when retrieval is needed (Jiang et al., 2021; Mielke et al., 2022a; Kadavath et al., 2022a; Yin et al.,  
 041 2023a). However, existing approaches either rely on fragile self-evaluation (Yin et al., 2023b; Chen  
 042 et al., 2024) or focus on narrow tasks and require complex setups (Azaria & Mitchell, 2023; Burns  
 043 et al., 2024).

044 For context evaluation, although there exist methods to estimate efficacy of retrieved context such as  
 045 SelfRAG (Asai et al., 2023), they rely on fine-tuning models and prompting external models to gauge  
 046 the utility of external context. A lightweight way to judge efficacy directly from model internals  
 047 remains missing. Instead, we look at these questions using our *LLM microscope*, through the *lens of*  
 048 *mechanistic interpretability*, and study whether model internals contain signals about the correctness  
 049 of responses and the efficacy of specific auxiliary context when answering a query. Concretely, we  
 050 study the following research questions:

- 051 1. **RQ1:** *Can we estimate the correctness of a model generation from its model internals alone?*
- 052 2. **RQ2:** *Can we estimate the efficacy of a given context directly from model internals?*

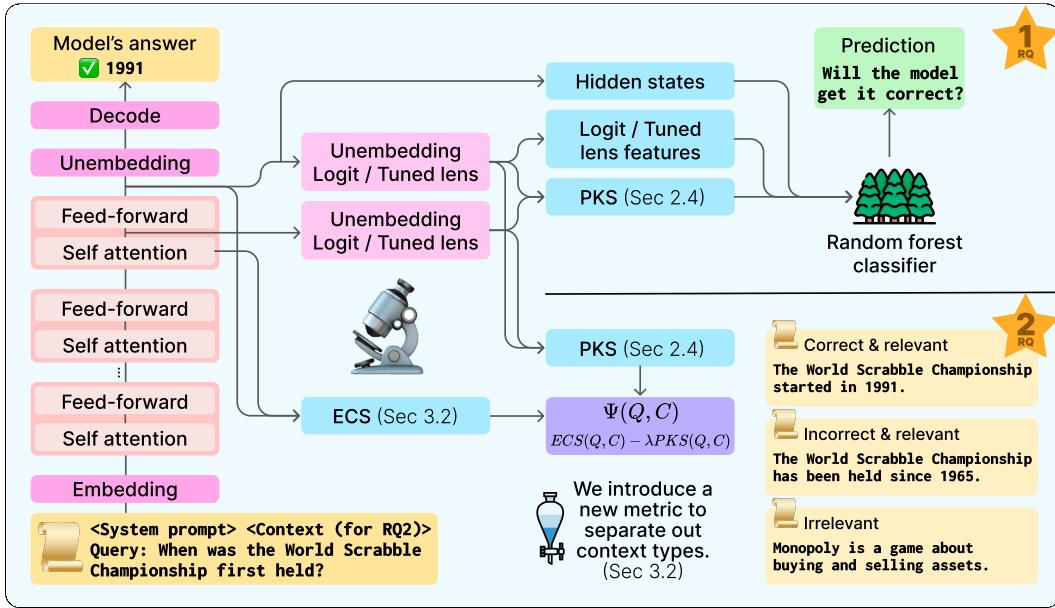


Figure 1: Overview of our framework. For RQ1, we use model internals, including hidden states, Logit/Tuned Lens-based features, and parametric knowledge score (PKS) to train classifiers that predict the correctness of a model’s output when answering a question. For RQ2, we analyze how internal signals like external context score (ECS) and PKS respond to different types of external context (correct, incorrect, irrelevant) in order to assess the model’s sensitivity to context when generating answers.

To answer these questions, we study six LLMs across three model families and sizes in an open-domain factual question-answering setting on the TriviaQA (Joshi et al., 2017) and MMLU (Hendrycks et al., 2021) datasets. We train classifiers on model internals to predict generation correctness and context relevance. To operationalize contextual relevance, we introduce two novel measures: *contextual log-likelihood gain* and *contextual relative utility* and study whether model internals can discriminate between contexts across two axes: correctness and relevance. We find that:

1. We can estimate correctness of a model generation to open-domain questions from model internals alone with over 75% accuracy and 70% AUC-ROC.
2. Using our model internals-based contextual log-likelihood gain metric, we can effectively discriminate between contexts across both the axes of correctness and relevance.

## 2 RELATED WORK

**Confidence Estimation** Confidence estimation methods are typically categorized as *closed-box* or *open-box*. Closed-box methods prompt the model to assess its own correctness, either by estimating the likelihood that its answer is correct or by judging whether it knows the answer (Kadavath et al., 2022b; Tian et al., 2023; Yin et al., 2023b; Huang et al., 2024). Some use linguistic cues (Mielke et al., 2022b) or generation consistency (Manakul et al., 2023; Zhang et al., 2023; Chen et al., 2024). Open-box methods instead analyze the model’s internals, either through logit-based uncertainty estimates (Murray & Chiang, 2018; Kadavath et al., 2022b; Huang et al., 2023; Vazhentsev et al., 2023) or by probing hidden activations directly (Li et al., 2023; Orgad et al., 2024; Burns et al., 2024; Subramani et al., 2025).

Our work builds on this line of open-box internal-state approaches but differs in several key ways. First, we focus on predicting the correctness of generated answers across both open-domain QA and multiple-choice settings (Azaria & Mitchell, 2023; Burns et al., 2024). Second, we evaluate whether a model’s internals can predict the correctness of its own generations, establishing a direct, on-policy evaluation setting (Azaria & Mitchell, 2023; Servedio et al., 2025). We use a simple random

108 forest classifier to both preserve feature interpretability and avoid entangling effects of classifier  
 109 complexity (Kadavath et al., 2022b; Orgad et al., 2024; Ashok & May, 2025). This simple yet  
 110 effective setup allows us to demonstrate that a variety of internal-state features derived from diverse  
 111 interpretability techniques can support robust confidence estimation in generative tasks.<sup>1</sup>  
 112

113 **Parametric and Contextual Knowledge** The activation spaces of LLMs encode structured, editable  
 114 knowledge through mechanisms like induction heads (Elhage et al., 2021), knowledge neurons (Dai  
 115 et al., 2022), and feed-forward layers (Geva et al., 2022). Probing tools such as Logit-Lens (nostal-  
 116 gebraist, 2020), Tuned-Lens (Belrose et al., 2023), Future-Lens (Pal et al., 2023), and Backward-  
 117 Lens (Katz et al., 2024), and methods like causal tracing (Meng et al., 2022), attention analysis (Geva  
 118 et al., 2023; Yuksekgonul et al., 2023), and steering (Subramani & Suresh, 2020; Subramani et al.,  
 119 2022; Turner et al., 2023), reveal how predictions evolve across layers. Additionally, different layers  
 120 of transformer-based language models encode different linguistic properties (Tenney et al., 2019;  
 121 Ethayarajh, 2019; Li & Subramani, 2025). To improve factuality, retrieval-augmented approaches  
 122 incorporate external context (Roller et al., 2020; Shuster et al., 2021; Ravichander et al., 2025).  
 123 Recent work (Wu et al., 2024; Wadhwa et al., 2024; Sun et al., 2025; Li et al., 2025) examine  
 124 clashes between internal and retrieved knowledge, identify model over-reliance on retrieved context,  
 125 and develop metrics to diagnose and treat that reliance. We address this gap by moving beyond  
 126 prompting-based approaches to context helpfulness measurement; to our knowledge, we are the  
 127 first to apply interpretability techniques to study how models utilize contextual versus parametric  
 128 knowledge (Huang et al., 2024).  
 129

### 3 RQ1: ESTIMATING CORRECTNESS

131 To study RQ1, we formulate a binary classification task to predict whether a generated answer is  
 132 correct. Each instance in this task consists of a factual question  $Q$ , the ground-truth answer  $A^*$ ,  
 133 an LLM-generated answer  $A$ , and all activations and outputs  $\mathcal{H}_{Q,A}$  produced by a model  $M$  when  
 134 attempting to answer  $Q$ . We train a simple classifier  $f$  that takes  $Q$ ,  $A$ , and  $\mathcal{H}_{Q,A}$  as input and  
 135 predicts  $\mathbb{I}[A = A^*]$ . The classifier does not have access to the ground-truth answer  $A^*$  or any external  
 136 knowledge source;  $A^*$  is only used to create the ground-truth label for training. We explore several  
 137 choices for  $\mathcal{H}_{Q,A}$  as input features, which we describe below.  
 138

#### 3.1 ASKING LLMs DIRECTLY

141 We explore two prompting strategies to measure how well LLMs can express their confidence in an  
 142 answer: *Prompting without Answers* and *Prompting with Answers*. In *Prompting without Answers*, the  
 143 model  $M$  is given a question  $Q$ , and asked to output a confidence score between 0 and 100 indicating  
 144 how confident it is about answering the question without actually generating an answer. In *Prompting*  
 145 *with Answers*,  $M$  first generates a candidate answer  $A$  for  $Q$  and then is asked to output a confidence  
 146 score based on how certain it is about the generated answer (see §B.1 for prompting setup details).  
 147

#### 3.2 DECODING FROM INTERMEDIATE LAYERS

148 To test whether model internals contain signals for estimating answer correctness, we look at  
 149 two techniques that facilitate decoding from intermediate hidden states: *Logit Lens* and *Tuned*  
 150 *Lens* (nostalgebraist, 2020; Belrose et al., 2023). Both of these methods convert intermediate hidden  
 151 states into the vocabulary space. We obtain the following input features  $\mathcal{H}_{Q,A}$  from these probability  
 152 distributions. These features are computed per layer and sequence position and provided at once to  
 153 the classifier. See §B.2 and §B.3 for additional details on *Logit/Tuned Lens* and these features.  
 154

155 • The **Shannon’s entropy** measures the uncertainty in the model’s probability distribution over the  
 156 next token at each layer and sequence position (Shannon, 1948).  
 157 • The **output token rank** is the position, in descending order of log-probability, of the token selected  
 158 by the model at each layer. While this rank is related to probability, it serves as a more direct proxy  
 159 for the token a decoding algorithm is likely to generate.  
 160

<sup>1</sup>See Geng et al. (2024) for a broader overview.

162 • The **top- $p$  presence** of the output token is a binary indicator of whether the token generated by  
 163 the model appears in the top- $p$  nucleus set at each layer ( $p \in \{0.5, 0.9, 0.95, 0.99\}$  Holtzman et al.  
 164 (2019)).  
 165 • The **cross-entropy** quantifies how similar model predictions are across layers and is computed as  
 166 the negative log probability of the generated token under the distribution at each layer.

167  
 168 **3.3 HIDDEN STATES**  
 169

170 Rather than transforming the hidden states and decoding from them, we experiment with using the  
 171 model’s hidden states directly as input features to  $f$ . Formally, we use  $\mathbf{h} \in \mathbb{R}^d$  at layer  $\ell \in \{1, \dots, L\}$   
 172 from the first token position of the forward pass that generated  $A$  in response to  $Q$  to train  $f$ .  
 173 Since activations could have different scales due to differences in parameter norms throughout  
 174 training (Merrill et al., 2021), we convert the value in every dimension  $d$  of the activation  $h$  to a  
 175 z-score using means and variances computed from an auxiliary dataset.

176  
 177 **3.4 PARAMETRIC KNOWLEDGE SCORES (PKS)**

178 We expect the amount of parametric knowledge used to answer a question  $Q$  positively correlates  
 179 with the confidence the model has about its generated answer  $A$ . In other words, the more parametric  
 180 knowledge used to answer the question the more confidence it should have about that answer. To  
 181 quantify the utilization of parametric knowledge, we use the *Parametric Knowledge Score (PKS)*  
 182 from Sun et al. (2025), which measures how much the feedforward networks (FFN) contribute to the  
 183 activations. For each token at layer  $\ell$ , the activations before and after the FFN layer are transformed  
 184 into a probability distribution over the vocabulary via *Logit Lens*. PKS is the Jensen-Shannon  
 185 divergence (JSD) between these two distributions, loosely capturing the amount of information  
 186 imparted by FFN weights into the activations. Formally, the token-level PKS is given by  $P_\ell =$   
 187  $\text{JSD}(q(\mathbf{x}_{\text{before}, \ell}) \parallel q(\mathbf{x}_{\text{after}, \ell}))$ , where  $q(\cdot) = \text{softmax}(\text{LogitLens}(\cdot))$ .

188  
 189 **4 RQ2: ESTIMATING EFFICACY OF CONTEXT**  
 190

191 Our goal is to see whether we can measure the efficacy of a given context  $C$  from either internal  
 192 (via model internals) or external (via prompting) features. We look at two attributes of context:  
 193 correctness and relevance and define three types of context  $C$ :

194 • Correct and relevant ( $C_{\text{correct}}$ ): aligns with the gold answer and has essential information.  
 195 • Incorrect but relevant ( $C_{\text{incorrect}}$ ): structurally similar and topically aligned, but contains incorrect  
 196 or misleading information.  
 197 • Irrelevant ( $C_{\text{irrelevant}}$ ): topically unrelated to and unhelpful for answering the question.

198 We define two lenses with which we can observe the effect of contexts: *contextual log-likelihood*  
 199 *gain* and *contextual relative utility*.

200 *Contextual log-likelihood gain* measures how much incorporating a context  $C$  improves (positive) or  
 201 degrades (negative) the model’s confidence in generating the correct answer. In RAG, it quantifies the  
 202 utility of the retrieved context to the model to generate the correct answer. For a question  $Q$ , context  
 203  $C$ , and ground-truth answer tokens  $\mathbf{y} = (y_1, \dots, y_T)$ , the *contextual log-likelihood gain* is defined  
 204 as:

205 
$$\text{LL}(Q, C) = \sum_{t=1}^T \log p(y_t | y_{<t}, Q, C) - \sum_{t=1}^T \log p(y_t | y_{<t}, Q) \quad (1)$$

206 *Contextual relative utility* compares two different contexts  $C_1$  and  $C_2$  and measures whether  $C_1$  is  
 207 more helpful (positive) or harmful (negative) than  $C_2$  for a model to produce the correct answer. We  
 208 formally define *contextual relative utility* as  $\Delta \text{LL}(Q, C_1, C_2) = \text{LL}(Q, C_1) - \text{LL}(Q, C_2)$ .

209  
 210 **4.1 PROMPTING-BASED CONFIDENCE ESTIMATION**  
 211

212 We analyze whether prompting-based methods can approximate the model’s contextual log-likelihood  
 213 gain through asking the model to generate confidence scores. For each question  $Q$ , we prompt the

model to output a confidence score on a 0–100 scale under the two different context conditions discussed in §3.1: *Prompting without Answers* and *Prompting with Answers*.

Let  $\text{Conf}(Q, C)$  denote the model’s confidence score when answering  $Q$  with context  $C$ , and  $\text{Conf}(Q)$  denote its confidence when answering without any context. We define the *prompting-based contextual gain* as:

$$\Omega(Q, C) = \text{Conf}(Q, C) - \text{Conf}(Q). \quad (2)$$

To compare two contexts  $C_1$  and  $C_2$ , we define the *prompting-based relative utility* as  $\Delta\Omega(Q, C_1, C_2) = \Omega(Q, C_1) - \Omega(Q, C_2)$ , which estimates the relative helpfulness of  $C_1$  over  $C_2$ . If the model can accurately distinguish context quality, we expect:  $\Delta\Omega(Q, C_{\text{correct}}, C_{\text{incorrect}}) > 0$ , and  $\Delta\Omega(Q, C_{\text{correct}}, C_{\text{irrelevant}}) > 0$ .<sup>2</sup>

## 4.2 INTERNALS-BASED CONFIDENCE ESTIMATION

To measure how well additional context can affect question-answering ability along the two axes of external context utilization and internal parametric knowledge reliance, we use *external context score* (ECS) and *parametric knowledge score* (PKS) from Sun et al. (2025). *External Context Score* (ECS) captures the extent to which a language model relies on external context during generation. For each output token  $t_n$ , ECS is computed as the cosine similarity between the token’s final-layer hidden state  $x_n^L$  and the mean-pooled embedding  $e$  of the top- $k\%$  most attended context tokens, selected based on attention weights. The token-level ECS is defined as:

$$\text{ECS}_h^l(t_n) = \frac{e \cdot x_n^L}{\|e\| \|x_n^L\|} \quad (3)$$

We can obtain the ECS score across a multi-token generation by simply averaging token-wise ECS scores over all output tokens. A higher ECS indicates stronger alignment between the generated output and the retrieved context, suggesting that the model is effectively utilizing external information.

In our setting, parametric knowledge and external context can be thought of as orthogonal. We define a proxy for *contextual log-likelihood gain* using model internals as:

$$\Psi(Q, C) = \text{ECS}(Q, C) - \lambda \cdot \text{PKS}(Q, C), \quad (4)$$

where  $\text{ECS}(Q, C)$  measures the model’s reliance on external context,  $\text{PKS}(Q, C)$  quantifies the effect of parametric knowledge in the presence of context, and  $\lambda$  is a scaling factor that rescales PKS to match ECS. In our experiment, we choose the value of  $\lambda$  for each dataset such that the mean PKS score is rescaled to match the mean ECS score, computed across all examples and all context types. This normalization ensures that per-example differences in  $\Psi(Q, C)$  reflect shifts in context reliance without introducing category-specific bias. To compare two contexts  $C_1$  and  $C_2$ , we define *internals-based relative utility* as:

$$\Delta\Psi(Q, C_1, C_2) = \Psi(Q, C_1) - \Psi(Q, C_2) \quad (5)$$

Our formulation reflects the tradeoff between contextual and parametric knowledge: higher values of  $\Psi(Q, C)$  suggest stronger reliance on external context, while lower values indicate dominance of parametric knowledge or potential confusion from misleading context.

In addition, we directly compute contextual relative utility using the actual log-likelihood, denoted as  $\Delta\Psi^{(\text{log likelihood})}$ . Specifically, we apply the log softmax over the vocabulary to obtain token-level log probabilities, average these over the ground-truth answer tokens, and then take the difference between two context types.

## 5 EXPERIMENTAL SETUP

**Datasets** We use factual short-question-answering datasets TriviaQA (Joshi et al., 2017) and MMLU (Hendrycks et al., 2021) for our experiments. We report results on both datasets for RQ1. As TriviaQA is open-ended while MMLU is multiple-choice, this ensures generalizability of our results across question types. As RQ2 requires accompanying context, we only use TriviaQA for this

<sup>2</sup>See §C.2 for the reasoning, and §C.1 for details on the prompting setup.

Dataset	Context Type	LLaMA 3.8B	LLaMA 2.13B	LLaMA 2.7B	Gemma 2.9B	Qwen 2.5.7B	Qwen 2.5.3B
TriviaQA	None	0.708	0.591	0.611	0.749	0.558	0.435
	Correct	0.888	0.833	0.845	0.906	0.856	0.815
	Incorrect	0.358	0.340	0.314	0.406	0.332	0.301
	Irrelevant	0.579	0.516	0.465	0.727	0.488	0.293
MMLU	None	0.611	0.484	0.462	0.709	0.693	0.634

Table 1: Model accuracies on TriviaQA and MMLU datasets with no context and with varying context conditions. MMLU does not contain context or evidence documents.

research question as MMLU does not provide context passages or evidence documents. To ensure our experiments remain computationally viable, we only consider the validation subset of TriviaQA, from which we retain 6,557 examples after quality filtering (see §D.1 for details). These are split 80%-20% for training and testing our classifiers. To account for possible variations in referring to the same entity, we use GPT-4o (OpenAI et al., 2024) as an LLM-judge to evaluate the correctness of model-generated answers against all the ground-truth answer variations present in TriviaQA. We use the 14,042 test examples of the “all” subset of MMLU as our training set and the 1,531 validation examples as our test set. As MMLU questions are multiple-choice, we use regex-based answer extraction and verification.

**External Context:** We use the original evidence document provided by TriviaQA, or a summarized version if the original exceeds 500 tokens, as the correct context. We then construct incorrect and irrelevant contexts by modifying or substituting these documents in controlled ways: for incorrect context, we ask an LLM to replace all references to the ground-truth answers with incorrect alternatives; for irrelevant context, we sample the correct context of a different example with low textual similarity (see §D.1 for details and examples).

**Models:** We experiment with 6 models across families and sizes: LLaMA-3-8B-Instruct (Grattafiori et al., 2024), LLaMA-2-7B-Chat-HF and LLaMA-2-13B-Chat-HF (Touvron et al., 2023), Qwen-2.5-3B-Chat and Qwen-2.5-7B-Chat (Yang et al., 2024), and Gemma-2-9B-It (Team et al., 2024).

**Methodology Details:** We choose random forest classifiers for their interpretability, which facilitates feature importance analysis, while also capturing non-linear patterns in the data. These classifiers are trained on the features discussed in §3 (see §D.2 for hyper-parameters and training specifications). For RQ2, we estimate *contextual log-likelihood gain* using ECS and PKS using equation (4). For ECS, we average across both attention heads and output tokens and for PKS, we average across output tokens to capture the overall effect of parametric knowledge on the entire output. PKS and ECS are calculated per layer and then averaged across layers. We present a layer-wise analysis in §7.2. We calculate one  $\lambda$  for each model by rescaling the relative magnitudes of PKS and ECS, averaged over all examples and context types in the training set.

## 6 RESULTS

We present the accuracies of our six LLMs on both TriviaQA and MMLU in Table 1. All models achieve moderate accuracy in the default no-context setting, with the larger and more recent models generally performing better. As expected, accuracy on TriviaQA increases with correct context and drops with incorrect or irrelevant context, highlighting that models are indeed sensitive to external context. We focus on the no-context setting to study RQ1 in §6.1 and examine different context types for RQ2 in §6.2.

### 6.1 RQ1: CAN WE ESTIMATE CORRECTNESS FROM MODEL INTERNALS ALONE?

Table 2 shows that prompting-based baselines perform poorly across all models, barely outperforming the majority-class baseline, even when given access to its own generated answer. This is particularly evident in the LLaMA models, suggesting that such prompting strategies fail to estimate correctness accurately and highlight that LLMs are poorly calibrated.

	Estimator	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Accuracy	Majority	0.699	0.591	0.621	0.735	0.567	0.565
	Prompt w/ A	<b>0.789</b>	0.599	0.634	0.762	0.727	0.652
	Prompt w/o A	0.699	0.611	0.621	0.768	0.637	0.605
	Logit lens	0.782*‡	<b>0.779*‡‡</b>	<b>0.776*‡‡</b>	0.774*	0.751*‡	0.729*‡‡
	Tuned lens	0.779*‡	0.775*‡‡	0.775*‡‡	-	-	-
	Hidden states (Best)	0.759*‡	0.679*‡‡	0.702*‡‡	<b>0.785*</b>	<b>0.782*‡‡</b>	<b>0.749*‡‡</b>
	PKS	0.733	0.725*‡‡	0.709*‡‡	0.743	0.650*	0.695*‡‡
AUC-ROC	Majority	0.500	0.500	0.500	0.500	0.500	0.500
	Prompt w/ A	0.783	0.512	0.537	0.592	0.736	0.687
	Prompt w/o A	0.591	0.582	0.542	0.742	0.653	0.631
	Logit lens	<b>0.790*‡</b>	<b>0.847*‡‡</b>	<b>0.835*‡‡</b>	<b>0.747*†</b>	<b>0.826*‡‡</b>	<b>0.812*‡‡</b>
	Tuned lens	0.782*‡	0.846*‡‡	0.829*‡‡	-	-	-
	Hidden states (Best)	0.647*‡	0.631*‡‡	0.647*‡‡	0.616*	0.774*‡	0.735*‡‡
	PKS	0.729*‡	0.768*‡‡	0.743*‡‡	0.723*†	0.715*‡	0.752*‡‡

(a) TriviaQA

	Estimator	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Accuracy	Majority	0.604	0.538	0.548	0.717	0.692	0.627
	Prompt w/ A	0.603	0.497	0.547	0.718	0.705	0.627
	Prompt w/o A	0.605	0.535	0.529	0.718	0.704	0.629
	Logit lens	0.705*‡‡	<b>0.692*‡‡</b>	<b>0.699*‡‡</b>	0.769*‡‡	<b>0.815*‡‡</b>	0.673*‡‡
	Tuned lens	0.695*‡‡	0.671*‡‡	0.684*‡‡	-	-	-
	Hidden states (Best)	<b>0.744*‡‡</b>	0.686*‡‡	0.672*‡‡	<b>0.801*‡‡</b>	0.777*‡‡	<b>0.746*‡‡</b>
	PKS	0.605	-	0.543	0.705	0.691	0.618
AUC-ROC	Majority	0.500	0.500	0.500	0.500	0.500	0.500
	Prompt w/ A	0.590	0.501	0.560	0.557	0.562	0.629
	Prompt w/o A	0.497	0.538	0.492	0.513	0.544	0.519
	Logit lens	<b>0.798*‡‡</b>	<b>0.771*‡‡</b>	<b>0.767*‡‡</b>	<b>0.843*‡‡</b>	<b>0.939*‡‡</b>	0.711*‡‡
	Tuned lens	0.770*‡‡	0.752*‡‡	0.760*‡‡	-	-	-
	Hidden states (Best)	0.740*‡‡	0.684*‡‡	0.576*‡	0.736*‡‡	0.728*‡‡	<b>0.733*‡‡</b>
	PKS	0.537‡	-	0.543*‡	0.555*‡	0.541*	0.538

(b) MMLU

Table 2: Performance of various classifiers on the test sets using our proposed methods and baseline approaches. We include two prompting baselines: Prompt w/ A (prompting with answers) and Prompt w/o A (prompting without answers) and a simple majority class baseline (Majority). For Logit Lens, Tuned Lens, and PKS methods, we use all values across layers as input features, while for Hidden States, we choose the best layer. Hidden states are normalized using z-score normalization. Tuned Lens results are omitted for models whose weights are not publicly available (Belrose et al., 2023). PKS scores are not available for LLaMA 2 13B on MMLU due to computational constraints. \*, †, and ‡ indicate statistical significance compared to Majority, Prompt w/ A, and Prompt w/o A, respectively (p-value < 0.05, two-sided permutation test).

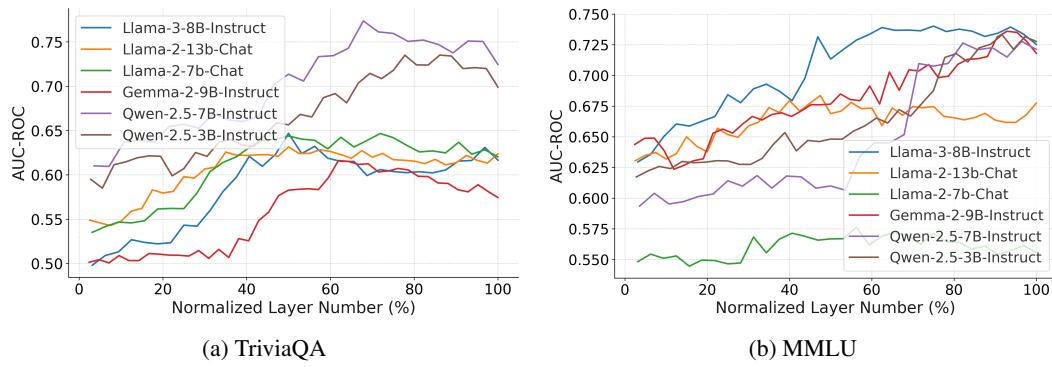


Figure 2: Area under ROC curve for random forest classifiers trained on z-score normalized hidden states of each layer. Performance increases with layer depth, suggesting that later layers refine and consolidate decision-relevant signals.

Context Comparison	Differentiator	LLaMA 3 8B	LLaMA 2 7B	Qwen 2.5 7B	Qwen 2.5 3B	Average
Correct > Incorrect	$\Delta\Omega$ (prompt w/ A)	19.7	10.9	32.7	16.4	19.9
	$\Delta\Omega$ (prompt w/o A)	20.8	11.3	30.0	23.6	21.4
	$\Delta\Psi$ (log likelihood)	83.9 <sup>†‡</sup>	<b>89.5<sup>†‡</sup></b>	82.6 <sup>†‡</sup>	<b>83.3<sup>†‡</sup></b>	<b>84.8<sup>†‡</sup></b>
	$\Delta\Psi$ (model internals)	<b>85.5<sup>†‡</sup></b>	70.2 <sup>†‡</sup>	<b>85.2<sup>†‡</sup></b>	75.9 <sup>†‡</sup>	79.2 <sup>†‡</sup>
Correct > Irrelevant	$\Delta\Omega$ (prompt w/ A)	83.2	39.7	52.2	<b>88.5</b>	65.9
	$\Delta\Omega$ (prompt w/o A)	<b>93.6</b>	68.4	79.9	72.4	78.6
	$\Delta\Psi$ (log likelihood)	76.7	74.7 <sup>†‡</sup>	74.5 <sup>†</sup>	80.6 <sup>‡</sup>	76.6 <sup>†</sup>
	$\Delta\Psi$ (model internals)	86.3 <sup>†</sup>	<b>90.5<sup>†‡</sup></b>	<b>89.6<sup>†‡</sup></b>	70.6	<b>84.3<sup>†‡</sup></b>

Table 3: Proportion of examples where  $\Delta LL$  successfully distinguishes context quality using prompting-based and internal-based confidence estimators on TriviaQA. We compute  $\Delta LL(Q, C_{\text{correct}}, C_{\text{incorrect}})$  and  $\Delta LL(Q, C_{\text{correct}}, C_{\text{irrelevant}})$  separately for each example and report the fraction for which the result is greater than zero. <sup>†</sup> and <sup>‡</sup> indicate statistical significance compared to Prompt w/ A and Prompt w/o A, respectively (p-value < 0.05, two-sided permutation test).

In contrast, the classifier trained on features extracted with the *Logit Lens* show strong performance, yielding the highest AUC-ROC scores on both datasets for most models. *Tuned Lens* performs comparably to *Logit Lens*, despite being developed to improve upon *Logit Lens* by better aligning intermediate hidden states with the output distribution through affine transformations. We hypothesize that the only distinction between the Tuned Lens and the Logit Lens is the affine transformation, and thus the corresponding scaling and shifting of activations do not contribute additional predictive information.

Surprisingly, classifiers trained directly on hidden states perform nearly as well as those based on Logit Lens features on both datasets. This finding indicates that the vanilla activations already encode information about the correctness of the model’s output. In Figure 2, we observe that performance improves with depth, suggesting that later layers refine and consolidate decision-relevant signals. This result implies that intermediate decoding methods may not be necessary to predict correctness and the strong performance of classifiers trained on first token hidden states across models may allow for correctness auditing of model responses early in the generation process.

Lastly, we find that PKS alone is strongly predictive of correctness on TriviaQA, despite being explicitly designed to just measure the influence of feedforward networks on token representations. This indicates that the feedforward layers of an LLM impart information onto the activations differently based on how in-distribution a given query is. However, PKS performs close to random chance for MMLU, suggesting that its discriminatory power may be limited to open-ended question answering and not extend to single-token multiple choice questions.

## 6.2 RQ2: CAN WE ESTIMATE CONTEXT EFFICACY DIRECTLY FROM MODEL INTERNALS?

Table 3 shows that internals-based confidence estimation significantly outperforms prompting-based methods when distinguishing between correct and incorrect contexts. This result is striking: while models fail to express higher confidence when generating answers conditioned on correct context, their internal activations nonetheless reflect this difference. All models exceed the 50% random baseline by a wide margin. We do not include results for LLaMA 2 13B due to computational constraints, nor for Gemma 2 9B due to the inapplicability of ECS to its architecture. Further explanation is provided in §F.1.

When comparing correct and irrelevant contexts, *Prompting without Answer* performs well, often on par with internal-based estimation and consistently outperforms the *Prompting with Answer* baseline. This suggests that including the model’s own answer in the prompt may mislead it, especially when the answer is incorrect due to misleading context. Without the generated answer, however, models appear better able to distinguish irrelevant contexts. This is paradoxical: despite being able to differentiate between relevant and irrelevant contexts, the underlying LLM remains strongly influenced by irrelevant context (see Table 1), suggesting that recognition alone is insufficient to steer the model towards an accurate answer. We suspect that during training, the model is rarely given irrelevant context and learns to implicitly trust contextual information.

432 

## 7 DISCUSSION

433 

### 7.1 RQ1

434 **Prompting Baselines** Here, we look at how well calibrated the prompting-based approaches  
 435 are by looking at a reliability diagram and measuring smooth expected calibration error (Smooth  
 436 ECE) (Błasior & Nakkiran, 2023). In our reliability diagrams, we plot the confidence (x-axis) vs.  
 437 accuracy (y-axis) and plot the line  $y = x$  indicating a perfectly calibrated system. Figures 3–6 in  
 438 our Appendix show that all prompting baselines for all LLMs for both TriviaQA and MMLU are  
 439 systematically overconfident (*i.e.*, predictions with  $x\%$  confidence yield  $< x\%$  accuracy).  
 440

441 **Logit Lens vs. Tuned Lens** In addition to the comparison presented in Table 2, we further train  
 442 separate classifiers on logit lens and tuned lens features of each individual layer to assess performance  
 443 across layers on TriviaQA (see Figure 7 in Appendix). The performance of the two methods  
 444 varies across layers without a consistent trend and neither consistently outperforms the other. We  
 445 observe that earlier layers are also highly predictive of generation correctness, achieving performance  
 446 comparable to later layers. This strengthens our finding that information about output correctness is  
 447 available early on in the models’ activations, underscoring the potential of early auditing.  
 448

449 **Logit/Tuned Lens: Feature Importance** We analyze feature importance scores from random  
 450 forest classifiers for TriviaQA in Figure 8 in the Appendix. Entropy and cross-entropy from the  
 451 final layer emerge as the most important features across all models. This is intuitive: output token  
 452 logits and distributions directly capture the model’s confidence when generating the correct answer.  
 453 Additionally, features from the last  $\sim 5$  layers generally exhibit higher importance.  
 454

455 **Logit/Tuned Lens: AUC-ROC Curves of Features from External vs. Internal Layers** To  
 456 examine whether features from non-final layers (denoted as *Internal*) are as predictive as those from  
 457 the final layer alone (denoted as *External*), we train separate classifiers using each feature set and  
 458 compare their AUC-ROC curves, as shown in Figure 9 for TriviaQA and Figure 10 for MMLU.  
 459 Interestingly, across all models, classifiers trained on internal layers achieve higher AUC scores  
 460 than those using only the final layer. This suggests that intermediate representations can carry more  
 461 information than the final layer alone for predicting generation correctness. See §E for other analyses.  
 462

463 

### 7.2 RQ2

464 **Example Analysis** We visualize ECS and PKS values across model layers and output tokens for  
 465 selected TriviaQA examples in Figures 12, 13, and 14 in the Appendix. While PKS scores show  
 466 only small variations across tokens and layers, later layers show noticeably higher values. Here,  
 467 initial output tokens have higher PKS values than later ones. These patterns suggest that parametric  
 468 knowledge accumulated in later layers has a stronger influence on the output distribution, particularly  
 469 at the beginning of generation.  
 470

471 In contrast, ECS scores vary more across tokens, reflecting the model’s selective use of context  
 472 depending on the importance and influence of the context on each output token. Uninformative tokens  
 473 tend to exhibit lower ECS scores. Across layers, ECS shows relatively little variation, indicating that  
 474 attention to external context is distributed broadly across the network rather than being isolated to  
 475 specific layers. While features continue to evolve across layers, the degree to which tokens attend to  
 476 one another’s features appears relatively stable.  
 477

478 

## 8 CONCLUSION

479 Our experiments demonstrate that we can indeed predict the correctness of a model generation from  
 480 model internals alone. In fact, with just the activations of first output token, we can predict correctness  
 481 with about 75% accuracy, hinting that early auditing could be possible. Prompting, on the other hand,  
 482 is poorly calibrated and thus has little utility. Additionally, using model internals based *contextualized*  
 483 *log-likelihood*, we can estimate the efficacy of external context along two axis: correctness and  
 484 relevancy. Taken together, our results suggest that deciphering model internals could provide valuable  
 485 insight into making language models more trustworthy.

486 REPRODUCIBILITY STATEMENT  
487488 We provide the dataset details, training details, model specifications, and prompts used in the main text  
489 or the Appendix to ensure our experiments are reproducible. We use mainstream and commonly-used  
490 software libraries, datasets, and models for our experiments. We publicly release our code.  
491492 ETHICS STATEMENT  
493494 Large language models carry significant potential for misuse, both intentional and accidental. Our  
495 work aims to advance understanding of the decision-making processes in LLMs and to identify  
496 hallucinations in model generations, thereby helping to mitigate some of the harms associated with  
497 the spread of inaccurate, model-generated content. Nevertheless, our methods should not be seen as a  
498 substitute for careful usage and verification; all model outputs must be independently checked for  
499 accuracy, particularly in domains where correctness is critical.  
500501 We emphasize that all datasets and models used in this study are available under permissive licenses  
502 for research purposes. We rely on instruction-tuned models, which have already undergone safety-  
503 related training. At the same time, we acknowledge that probing or intervening in model internals  
504 may alter or undermine this safety tuning. Understanding and exposing hidden mechanisms inside  
505 LLMs can yield valuable scientific insight, but it also carries risks: such methods could, in principle,  
506 be adapted to bypass alignment safeguards or weaken safety behaviors.  
507508 We therefore caution that techniques for manipulating internal representations should be applied with  
509 care and with an awareness of their broader implications. Our intention in presenting this work is to  
510 promote transparency, accountability, and safer deployment of LLMs, not to provide tools that could  
511 be used to compromise alignment or safety constraints.  
512513 REFERENCES  
514515 Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf.  
516 FLAIR: An easy-to-use framework for state-of-the-art NLP. In *NAACL 2019, 2019 Annual  
517 Conference of the North American Chapter of the Association for Computational Linguistics  
518 (Demonstrations)*, pp. 54–59, 2019.519 Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning to  
520 retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on  
521 Learning Representations*, 2023.522 Dhananjay Ashok and Jonathan May. Language models can predict their own behavior. *arXiv  
523 preprint arXiv:2502.13329*, 2025.524 Amos Azaria and Tom Mitchell. The internal state of an LLM knows when it's lying. In  
525 Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Com-  
526 putational Linguistics: EMNLP 2023*, pp. 967–976, Singapore, December 2023. Associa-  
527 tion for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.68. URL <https://aclanthology.org/2023.findings-emnlp.68/>.  
528529 Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia,  
530 Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. A multitask,  
531 multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactiv-  
532 ity. In Jong C. Park, Yuki Arase, Baotian Hu, Wei Lu, Derry Wijaya, Ayu Purwarianti, and  
533 Adila Alfa Krisnadhi (eds.), *Proceedings of the 13th International Joint Conference on Natural  
534 Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for  
535 Computational Linguistics (Volume 1: Long Papers)*, pp. 675–718, Nusa Dua, Bali, November  
536 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.ijcnlp-main.45. URL  
537 <https://aclanthology.org/2023.ijcnlp-main.45/>.  
538539 Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella  
Biderman, and Jacob Steinhardt. Eliciting latent predictions from transformers with the tuned lens,  
2023. URL <https://arxiv.org/abs/2303.08112>.

540 Jarosław Błasiok and Preetum Nakir. Smooth ece: Principled reliability diagrams via kernel  
 541 smoothing. *arXiv preprint arXiv:2309.12236*, 2023.

542

543 Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language  
 544 models without supervision, 2024. URL <https://arxiv.org/abs/2212.03827>.

545 Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye.  
 546 Inside: Llms' internal states retain the power of hallucination detection, 2024. URL <https://arxiv.org/abs/2402.03744>.

547

548 Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. Knowledge neurons in  
 549 pretrained transformers, 2022. URL <https://arxiv.org/abs/2104.08696>.

550

551 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda  
 552 Askill, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli,  
 553 Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal  
 554 Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris  
 555 Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021.  
 556 <https://transformer-circuits.pub/2021/framework/index.html>.

557

558 Kawin Ethayarajh. How contextual are contextualized word representations? Comparing the ge-  
 559 ometry of BERT, ELMo, and GPT-2 embeddings. In Kentaro Inui, Jing Jiang, Vincent Ng, and  
 560 Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Lan-  
 561 guage Processing and the 9th International Joint Conference on Natural Language Processing  
 562 (EMNLP-IJCNLP)*, pp. 55–65, Hong Kong, China, November 2019. Association for Computational  
 563 Linguistics. doi: 10.18653/v1/D19-1006. URL <https://aclanthology.org/D19-1006/>.

564

565 Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepll, Preslav Nakov, and Iryna Gurevych. A survey  
 566 of confidence estimation and calibration in large language models. In Kevin Duh, Helena Gomez,  
 567 and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter  
 568 of the Association for Computational Linguistics: Human Language Technologies (Volume 1:  
 569 Long Papers)*, pp. 6577–6595, Mexico City, Mexico, June 2024. Association for Computational  
 570 Linguistics. doi: 10.18653/v1/2024.nacl-long.366. URL <https://aclanthology.org/2024.nacl-long.366/>.

571

572 Mor Geva, Avi Caciularu, Kevin Ro Wang, and Yoav Goldberg. Transformer feed-forward layers  
 573 build predictions by promoting concepts in the vocabulary space. *arXiv preprint arXiv:2203.14680*,  
 574 2022.

575

576 Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting recall of factual  
 577 associations in auto-regressive language models. *arXiv preprint arXiv:2304.14767*, 2023.

578

579 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad  
 580 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of  
 581 models. *arXiv preprint arXiv:2407.21783*, 2024.

582

583 Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyyvind Tafjord, Ananya  
 584 Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell  
 585 Author, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel,  
 586 Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal  
 587 Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh  
 588 Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi,  
 589 Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini,  
 590 Noah Smith, and Hannaneh Hajishirzi. OLMo: Accelerating the science of language models. In  
 591 Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting  
 592 of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15789–15809,  
 593 Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/  
 594 2024.acl-long.841. URL <https://aclanthology.org/2024.acl-long.841/>.

595

596 Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen,  
 597 Shengjie Ma, Honghao Liu, et al. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594*,  
 598 2024.

594 Nuno M. Guerreiro, Duarte M. Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre  
 595 Colombo, and André F. T. Martins. Hallucinations in large multilingual translation models.  
 596 *Transactions of the Association for Computational Linguistics*, 11:1500–1517, 2023. doi: 10.1162/  
 597 tacl\_a\_00615. URL <https://aclanthology.org/2023.tacl-1.85/>.

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

602 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text  
 603 degeneration. *ArXiv*, abs/1904.09751, 2019. URL <https://api.semanticscholar.org/CorpusID:127986954>.

605 Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong  
 606 Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large  
 607 language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.*, 43  
 608 (2), January 2025. ISSN 1046-8188. doi: 10.1145/3703155. URL <https://doi.org/10.1145/3703155>.

609

610 Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu,  
 611 and Lei Ma. Look before you leap: An exploratory study of uncertainty measurement for large  
 612 language models. *arXiv preprint arXiv:2307.10236*, 2023.

613

614 Yukun Huang, Sanxing Chen, Hongyi Cai, and Bhuwan Dhingra. To trust or not to trust? enhancing  
 615 large language models' situated faithfulness to external contexts. *arXiv preprint arXiv:2410.14675*,  
 616 2024.

617

618 Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. How can we know when language  
 619 models know? on the calibration of language models for question answering. *Transactions of the  
 620 Association for Computational Linguistics*, 9:962–977, 2021. doi: 10.1162/tac1\_a\_00407. URL  
 621 <https://aclanthology.org/2021.tacl-1.57/>.

622

623 Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly  
 624 supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.

625

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

634

635 Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas  
 636 Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly)  
 637 know what they know. *arXiv preprint arXiv:2207.05221*, 2022b.

638

639 Shahar Katz, Yonatan Belinkov, Mor Geva, and Lior Wolf. Backward lens: Projecting language  
 640 model gradients into the vocabulary space. *arXiv preprint arXiv:2402.12865*, 2024.

641

642 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
 643 Heinrich Küttrler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented genera-  
 644 tion for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:  
 645 9459–9474, 2020.

646

647 Gaotang Li, Yuzhong Chen, and Hanghang Tong. Taming knowledge conflicts in language models.  
 648 *arXiv preprint arXiv:2503.10996*, 2025.

649

650 Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-  
 651 time intervention: eliciting truthful answers from a language model. In *Proceedings of the 37th  
 652 International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY,  
 653 USA, 2023. Curran Associates Inc.

648 Michael Li and Nishant Subramani. Model internal sleuthing: Finding lexical identity and inflectional  
 649 morphology in modern language models. *ArXiv*, abs/2506.02132, 2025. URL <https://api.semanticscholar.org/CorpusID:279118488>.  
 650

651 Potsawee Manakul, Adian Liusie, and Mark JF Gales. Selfcheckgpt: Zero-resource black-box  
 652 hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*,  
 653 2023.

654 Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual  
 655 associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022.  
 656

657 William Merrill, Vivek Ramanujan, Yoav Goldberg, Roy Schwartz, and Noah A. Smith. Effects  
 658 of parameter norm growth during transformer training: Inductive bias from gradient descent. In  
 659 Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings  
 660 of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 1766–1781,  
 661 Online and Punta Cana, Dominican Republic, November 2021. Association for Computational  
 662 Linguistics. doi: 10.18653/v1/2021.emnlp-main.133. URL [https://aclanthology.org/2021.emnlp-main.133/](https://aclanthology.org/2021.emnlp-main.133).  
 663

664 Sabrina J. Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. Reducing conversational agents'  
 665 overconfidence through linguistic calibration. *Transactions of the Association for Computational  
 666 Linguistics*, 10:857–872, 2022a. doi: 10.1162/tacl\_a\_00494. URL <https://aclanthology.org/2022.tacl-1.50/>.  
 667

668 Sabrina J Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. Reducing conversational agents'  
 669 overconfidence through linguistic calibration. *Transactions of the Association for Computational  
 670 Linguistics*, 10:857–872, 2022b.  
 671

672 Kenton Murray and David Chiang. Correcting length bias in neural machine translation. *arXiv  
 673 preprint arXiv:1808.10006*, 2018.

674 nostalgebraist. interpreting gpt: the logit lens, Aug 2020. URL <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>.  
 675

676 OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan  
 677 Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-  
 678 Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex  
 679 Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau,  
 680 Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian,  
 681 Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein,  
 682 Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey  
 683 Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia,  
 684 Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben  
 685 Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake  
 686 Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon  
 687 Walkin, Brendan Quinn, Brian Guaraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo  
 688 Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li,  
 689 Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu,  
 690 Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim,  
 691 Christine Choi, Christine McLeavy, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley  
 692 Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler,  
 693 Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki,  
 694 Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay,  
 695 Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric  
 696 Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani,  
 697 Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh,  
 698 Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang  
 699 Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik  
 700 Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung,  
 701 Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu,  
 702 Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon,

702 Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie  
 703 Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe,  
 704 Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi  
 705 Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers,  
 706 Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan  
 707 Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh  
 708 Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn  
 709 Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra  
 710 Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe,  
 711 Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman,  
 712 Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng,  
 713 Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas Kondracik,  
 714 Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine  
 715 Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin  
 716 Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank  
 717 Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna  
 718 Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle  
 719 Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles  
 720 Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho  
 721 Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine,  
 722 Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige,  
 723 Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko,  
 724 Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick  
 725 Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan,  
 726 Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal,  
 727 Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo  
 728 Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob  
 729 Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory  
 730 Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi  
 731 Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara  
 732 Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu  
 733 Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi Xia, Sonia Phene, Spencer  
 734 Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal  
 735 Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas  
 736 Cunningham, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao  
 737 Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan,  
 738 Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie  
 739 Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra,  
 Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang,  
 Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. Gpt-4o system card, 2024.  
 URL <https://arxiv.org/abs/2410.21276>.

740 Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan  
 741 Belinkov. Llms know more than they show: On the intrinsic representation of llm hallucinations,  
 742 2024. URL <https://arxiv.org/abs/2410.02707>.

743 Koyena Pal, Jiuding Sun, Andrew Yuan, Byron C Wallace, and David Bau. Future lens: Anticipating  
 744 subsequent tokens from a single hidden state. *arXiv preprint arXiv:2311.04897*, 2023.

745 Abhilasha Ravichander, Shruti Ghela, David Wadden, and Yejin Choi. Halogen: Fantastic llm  
 746 hallucinations and where to find them, 2025. URL <https://arxiv.org/abs/2501.08292>.

747 Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks.  
 748 In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.  
 749 Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.

750 Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhuan Liu, Jing Xu, Myle  
 751 Ott, Kurt Shuster, Eric M Smith, et al. Recipes for building an open-domain chatbot. *arXiv preprint*  
 752 *arXiv:2004.13637*, 2020.

756 Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman  
 757 Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M.  
 758 Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît  
 759 Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas  
 760 Bekman, Angelina Mcmillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro  
 761 Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell,  
 762 Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna  
 763 Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm,  
 764 Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González  
 765 Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco de Toni, Gérard Dupont,  
 766 Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris  
 767 Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge,  
 768 Jian Zhu, Jonathan Chang, Jörg Frohberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Al-  
 769 mubarak, Kimbo Chen, Kyle Lo, Leandro von Werra, Leon Weber, Long Phan, Loubna Ben Allal,  
 770 Ludovic Tangy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario  
 771 Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu,  
 772 Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis,  
 773 Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo,  
 774 Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui  
 775 Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan  
 776 Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai,  
 777 Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina  
 778 Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat,  
 779 Arun Raja, Benjamin Heinzerling, Chenglei Si, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y.  
 780 Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza  
 781 Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos  
 782 Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-Shaibani, Matteo Manica, Ryan  
 783 Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali  
 784 Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin  
 785 Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won  
 786 Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune,  
 787 Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi,  
 788 Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre  
 789 Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden  
 790 Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anas-  
 791 tasia Chevleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névéol, Charles Lovering,  
 792 Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli  
 793 Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova,  
 794 Jessica Zosa Forde, Jordan Clive, Jungsik Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat,  
 795 Miruna Clinciu, Nadjoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal,  
 796 Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shami Pais, Tatiana Shavrina, Thomas Scialom,  
 797 Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruk-  
 798 sachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana,  
 799 Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Unldreaj,  
 800 Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh Hajihosseini, Bahareh Behroozi,  
 801 Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Danish Contractor, David Lansky,  
 802 Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani,  
 803 Fatima Mirza, Frankline Ononiwu, Habib Rezanejad, Hessie Jones, Indrani Bhattacharya, Irene  
 804 Solaiman, Irina Sedenko, Isar Nejadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia  
 805 Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu,  
 806 Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani,  
 807 Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira  
 808 Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade,  
 809 Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Calla-  
 han, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio  
 Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Periñán, Daniel  
 Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin  
 Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi,  
 Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa

810 Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Mari-  
 811 anna Nezhurina, Mario Sänger, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel  
 812 de Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan  
 813 Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chan-  
 814 drasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel  
 815 Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott,  
 816 Sinee Sang-Aroonsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant,  
 817 Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan  
 818 Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and  
 819 Thomas Wolf. BLOOM: A 176B-Parameter Open-Access Multilingual Language Model. working  
 820 paper or preprint, November 2023. URL <https://inria.hal.science/hal-03850124>.

821 Giovanni Servedio, Alessandro De Bellis, Dario Di Palma, Vito Walter Anelli, and Tommaso Di Noia.  
 822 Are the hidden states hiding something? testing the limits of factuality-encoding capabilities in  
 823 llms. *arXiv preprint arXiv:2505.16520*, 2025.

824 Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27  
 825 (3):379–423, 1948.

826 Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. Retrieval augmentation  
 827 reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*, 2021.

828 Nishant Subramani and Nivedita Suresh. Discovering useful sentence representations from large  
 829 pretrained language models. *arXiv preprint arXiv:2008.09049*, 2020.

830 Nishant Subramani, Nivedita Suresh, and Matthew E Peters. Extracting latent steering vectors from  
 831 pretrained language models. *arXiv preprint arXiv:2205.05124*, 2022.

832 Nishant Subramani, Jason Eisner, Justin Svegliato, Benjamin Van Durme, Yu Su, and Sam Thomson.  
 833 MICE for CATs: Model-internal confidence estimation for calibrating agents with tools. In  
 834 Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Proceedings of the 2025 Conference of the  
 835 Nations of the Americas Chapter of the Association for Computational Linguistics: Human  
 836 Language Technologies (Volume 1: Long Papers)*, pp. 12362–12375, Albuquerque, New Mexico,  
 837 April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. URL <https://aclanthology.org/2025.nacl-long.615/>.

838 ZhongXiang Sun, Xiaoxue Zang, Kai Zheng, Jun Xu, Xiao Zhang, Weijie Yu, Yang Song, and  
 839 Han Li. RedeEP: Detecting hallucination in retrieval-augmented generation via mechanistic  
 840 interpretability. In *The Thirteenth International Conference on Learning Representations*, 2025.  
 841 URL <https://openreview.net/forum?id=ztzZDzfrh>.

842 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya  
 843 Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al.  
 844 Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*,  
 845 2024.

846 Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovers the classical nlp pipeline. *arXiv  
 847 preprint arXiv:1905.05950*, 2019.

848 Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn,  
 849 and Christopher D Manning. Just ask for calibration: Strategies for eliciting calibrated confidence  
 850 scores from language models fine-tuned with human feedback. *arXiv preprint arXiv:2305.14975*,  
 851 2023.

852 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay  
 853 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation  
 854 and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

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

864 Artem Vazhentsev, Akim Tsvigun, Roman Vashurin, Sergey Petrakov, Daniil Vasilev, Maxim Panov,  
 865 Alexander Panchenko, and Artem Shelmanov. Efficient out-of-domain detection for sequence to  
 866 sequence models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp.  
 867 1430–1454, 2023.

868  
 869 Hitesh Wadhwa, Rahul Seetharaman, Somyaa Aggarwal, Reshma Ghosh, Samyadeep Basu,  
 870 Soundararajan Srinivasan, Wenlong Zhao, Shreyas Chaudhari, and Ehsan Aghazadeh. From  
 871 rags to rich parameters: Probing how language models utilize external knowledge over parametric  
 872 information for factual queries. *arXiv preprint arXiv:2406.12824*, 2024.

873  
 874 Kevin Wu, Eric Wu, and James Zou. Clasheval: Quantifying the tug-of-war between an llm’s  
 875 internal prior and external evidence. In *Neural Information Processing Systems*, 2024. URL  
 876 <https://api.semanticscholar.org/CorpusID:269157310>.

877  
 878 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,  
 879 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*  
 880 *arXiv:2412.15115*, 2024.

881  
 882 Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. Do large  
 883 language models know what they don’t know? In Anna Rogers, Jordan Boyd-Graber, and Naoaki  
 884 Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 8653–  
 885 8665, Toronto, Canada, July 2023a. Association for Computational Linguistics. doi: 10.18653/v1/  
 886 2023.findings-acl.551. URL <https://aclanthology.org/2023.findings-acl.551/>.

887  
 888 Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. Do large  
 889 language models know what they don’t know? *arXiv preprint arXiv:2305.18153*, 2023b.

890  
 891 Mert Yuksekgonul, Varun Chandrasekaran, Erik Jones, Suriya Gunasekar, Ranjita Naik, Hamid  
 892 Palangi, Ece Kamar, and Besmira Nushi. Attention satisfies: A constraint-satisfaction lens on  
 893 factual errors of language models. *arXiv preprint arXiv:2309.15098*, 2023.

894  
 895 Jiaxin Zhang, Zhuohang Li, Kamalika Das, Bradley Malin, and Sricharan Kumar. SAC<sup>3</sup>: Reliable  
 896 hallucination detection in black-box language models via semantic-aware cross-check consistency.  
 897 In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for*  
 898 *Computational Linguistics: EMNLP 2023*, pp. 15445–15458, Singapore, December 2023. As-  
 899 sociation for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.1032. URL  
 900 <https://aclanthology.org/2023.findings-emnlp.1032/>.

901  
 902 Qinyu Zhao, Ming Xu, Kartik Gupta, Akshay Asthana, Liang Zheng, and Stephen Gould. The first  
 903 to know: How token distributions reveal hidden knowledge in large vision-language models? In  
 904 *European Conference on Computer Vision*, pp. 127–142. Springer, 2024.

905  
 906 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
 907 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and  
 908 chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

## 909 A LIMITATIONS

910  
 911 This study has the following limitations. Several of our analyses are performed either in aggregate  
 912 across token positions or only at the first output token. It may be valuable to study variations across  
 913 token positions, especially in settings with longer model outputs. Although the TriviaQA and MMLU  
 914 datasets cover diverse domains, we restrict our analysis to short-answer factual question-answering  
 915 for computational feasibility. It is valuable to study the generalizability of our results to different  
 916 tasks and domains; we leave this for future work. Second, we only experiment with instruction-tuned  
 917 models. Investigating how pre-trained models differ from our analysis may reveal how post-training  
 918 processes affect models’ confidence estimates. In addition, experiments with larger models (> 32B)  
 919 may yield more generalizable results.

918 B METHODOLOGY DETAILS FOR RQ1  
919920 B.1 PROMPTING  
921922 Below are the prompts used for RQ1 to elicit confidence estimates from the model. We explore two  
923 prompting formats: *Prompting without Answer* and *Prompting with Answer*.  
924925 **Prompting without Answer** In this setting, the model is asked to estimate its confidence in being  
926 able to answer the question correctly, without actually providing an answer. The prompt is as follows:  
927928 For the question below, output your confidence in your ability to generate  
929 the correct answer as an integer between 0 and 100, where 0 means complete  
930 uncertainty and 100 means complete certainty. Provide only the confidence  
931 score, without answering the question or including any explanations or  
932 additional text.  
933934 Question:  
935 {Question}  
936937 **Prompting with Answer** This format uses a multi-turn chat interaction. First, the model generates  
938 an answer. Then, in a second turn, it is asked to provide a confidence score for the generated answer.  
939940 *Turn 1:*  
941942 Answer the following question with a single word or phrase. Do not provide  
943 explanations or additional context:  
944

945 {Question}

946 *Turn 2:*  
947948 Please output your confidence in the correctness of the answer as an  
949 integer between 0 and 100, where 0 means complete uncertainty and 100  
950 means complete certainty. Output only the confidence score with no  
951 explanations or additional text.952  
953 B.2 LOGIT LENS AND TUNED LENS  
954955 **Logit Lens:** *Logit Lens* projects the intermediate hidden states directly into the vocabulary space  
956 via the unembedding matrix or final linear layer of the model (nostalgebraist, 2020). Concretely,  
957 given a hidden state  $\mathbf{h}_\ell \in \mathbb{R}^d$  at layer  $\ell$ , and the unembedding matrix  $\mathbf{W}_U \in \mathbb{R}^{V \times d}$  (where  $V$  is the  
958 vocabulary size and  $d$  is the hidden size), the layer-wise logits under the logit lens are computed as:  
959  $\mathbf{z}_\ell^{(LL)} = \mathbf{W}_U \mathbf{h}_\ell$ . For decoding, these logits are passed through a softmax as normal to produce a  
960 probability distribution over the vocabulary:  $\mathbf{p}_\ell^{(LL)} = \text{softmax}(\mathbf{z}_\ell^{(LL)})$ .  
961962 **Tuned Lens:** *Tuned Lens* first learns an affine transformation or *translator* for each layer and  
963 then projects the transformed hidden state directly to the vocabulary space via the unembedding  
964 matrix (Belrose et al., 2023). Since pretrained language models are typically not trained to project  
965 from intermediate hidden states to the unembedding matrix, there is often a mismatch. *Tuned*  
966 *Lens* remedies this by post-hoc training these *translators* at each layer, improving intermediate  
967 hidden states alignment in both magnitude and direction to the unembedding matrix. For each  
968 layer  $\ell$ , the learned affine map  $(\mathbf{A}_\ell, \mathbf{b}_\ell)$  transforms the hidden state  $\mathbf{h}_\ell$  before multiplying with the  
969 unembedding matrix:  $\mathbf{z}_\ell^{(TL)} = \mathbf{W}_U(\mathbf{A}_\ell \mathbf{h}_\ell + \mathbf{b}_\ell)$ . The resulting logits can be passed through a softmax  
970 producing a probability distribution over the vocabulary:  $\mathbf{p}_\ell^{(TL)} = \text{softmax}(\mathbf{z}_\ell^{(TL)})$ . The *translators* are  
971 trained to minimize the KL divergence between  $\mathbf{p}_\ell^{(TL)}$  and the model’s final output distribution on  
972 general-purpose pretraining data.

972 B.3 INPUT FEATURE DETAILS FOR LOGIT LENS AND TUNED LENS  
973974 **Input Feature 1: Entropy** Entropy measures the uncertainty in a probability distribution. For  
975 us, entropy estimates how peaky (or uniform) the distribution over the vocabulary is at each token  
976 position at each layer. Given a distribution over the vocabulary using either *Logit Lens* (LL) or *Tuned*  
977 *Lens* (TL),  $\mathbf{p}_\ell^{(\text{LL-TL})} = \{p_\ell^{(\text{LL-TL})}(i)\}_{i=1}^V$  at layer  $\ell$ , the entropy  $H_\ell$  is defined as:  
978

979 
$$H_\ell = - \sum_{i=1}^V p_\ell^{(\text{LL-TL})}(i) \log p_\ell^{(\text{LL-TL})}(i) \quad (6)$$
  
980  
981

982 Note: these estimates are computed independently for each layer and sequence position.  
983984 **Input Feature 2: Rank of Output Token** A probability of 0.01 in certain token positions may  
985 reflect the highest rank token if the preceding context is ambiguous (e.g. The  $\langle \text{next\_word} \rangle$ ), whereas  
986 could reflect a much lower rank token if the context is less ambiguous (e.g. Harriet Tub $\langle \text{next\_word} \rangle$ ). The rank of the output next token at intermediate layers loosely correlates with the probability,  
987 but can offer a more direct signal towards what token a decoding algorithm may likely generate.  
988 At layer  $\ell$ , we compute either the *Logit Lens* (LL) or *Tuned Lens* (TL) distribution  $\mathbf{p}_\ell^{(\text{LL-TL})}$ . Let  
989  $\phi_\ell(t) := \log p_\ell^{(\text{LL-TL})}(t)$  denote the log-probability of token  $t$  under this distribution. We sort the  
990 vocabulary tokens by descending log-probability:  
991

992  
993 
$$\phi_\ell(\sigma(1)) \geq \phi_\ell(\sigma(2)) \geq \dots \geq \phi_\ell(\sigma(V)) \quad (7)$$
  
994

995 Then, the *rank*  $r_\ell(y)$  of the target token  $y$  is the index  $k$  such that  $\sigma(k) = y$ . A rank of 1 indicates the  
996 model considers  $y$  the most likely token at that layer. This is computed per layer and position.  
997998 **Input Feature 3: Top- $p$  Presence** Nucleus-sampling is one of the most common decoding al-  
999 gorithms for language models (Holtzman et al., 2019), where rather than considering the entire  
1000 distribution over the vocabulary, one considers only the highest probability tokens until a top- $p$   
1001 cumulative probability. This eliminates the chances of randomly sampling a very low probability  
1002 token in the long-tail. Motivated by this, we compute a binary indicator of whether the target token  
1003 appears within the top- $p$  nucleus set at each layer. Let  $\mathbf{p}_\ell^{(\text{LL-TL})}$  be sorted in descending order as  
1004  $p_{(1)} \geq p_{(2)} \geq \dots$ . The top- $p$  nucleus set  $V_p$  is the smallest set such that:  
1005

1006 
$$\sum_{i=1}^k p_{(i)} \geq p \quad (8)$$
  
1007

1008 The *top- $p$  presence* indicator  $I_{p,\ell}(y)$  is defined as:  
1009

1010 
$$I_{p,\ell}(y) = \begin{cases} 1 & \text{if } y \in V_p \\ 0 & \text{otherwise} \end{cases} \quad (9)$$
  
1011  
1012

1013 **Input Feature 4: Cross-Entropy** Cross-entropy quantifies how well the intermediate predictions  
1014 match the output next token. At each layer  $\ell$ , it is computed using the log probability assigned to the  
1015 target token  $y$  under the logit lens or tuned lens distribution:  
1016

1017 
$$\text{CE}_\ell(y) = -\log p_\ell^{(\text{LL-TL})}(y) \quad (10)$$

1018 This value is lower when the model assigns higher confidence to the output token. It is evaluated per  
1019 layer and token position.  
10201021 C METHODOLOGY DETAILS FOR RQ2  
10221023 C.1 PROMPTING  
10241025 Below are the prompts used for RQ2 to elicit confidence estimates from the model. We explore two  
prompting formats: *Prompting without Answer* and *Prompting with Answer*.  
1026

1026 **Prompting without Answer** In this setting, the model is provided with context and is asked to  
 1027 estimate its confidence in answering the question correctly, without actually generating an answer.  
 1028 The prompt used is as follows:  
 1029

1030 For the question below, output your confidence in your ability to generate  
 1031 the correct answer as an integer between 0 and 100, where 0 means complete  
 1032 uncertainty and 100 means complete certainty. Provide only the confidence  
 1033 score, without answering the question or including any explanations or  
 1034 additional text.  
 1035

1036 Context:  
 1037 `---`  
 1038 `{context}`  
 1039 `---`

1040 Question:  
 1041 `---`  
 1042 `{question}`  
 1043 `---`

1044  
 1045 **Prompting with Answer** This format uses a multi-turn chat interaction. In the first turn, the model  
 1046 is provided with context and asked to generate an answer to the question. In the second turn, it is  
 1047 asked to estimate its confidence in the answer it just provided.  
 1048

1049 *Turn 1:*

1050 Using the following context, answer the question that follows with a  
 1051 single word or phrase. Do not provide explanations or additional context:  
 1052  
 1053 > Context:  
 1054 `{Context}`  
 1055  
 1056 > Question:  
 1057 `{Question}`

1058  
 1059 *Turn 2:*

1060 Please output your confidence in the correctness of the answer as an  
 1061 integer between 0 and 100, where 0 means complete uncertainty and 100  
 1062 means complete certainty. Output only the confidence score with no  
 1063 explanations or additional text.  
 1064

## 1066 C.2 REASONING FOR CONTEXT COMPARISON

1067 **Correct vs. Incorrect Context** A correct and relevant context ( $C_{\text{correct}}$ ) directly supports the correct  
 1068 answer with accurate factual information. This should increase the model’s confidence, especially  
 1069 when its parametric knowledge alone is insufficient. In contrast, an incorrect but relevant context  
 1070 ( $C_{\text{incorrect}}$ ) may appear plausible but contains factual errors. While it may influence the model’s output,  
 1071 it should not increase confidence as much—especially if the model can detect inconsistencies or is  
 1072 sensitive to contradiction. Thus, we expect higher confidence in the presence of  $C_{\text{correct}}$  than  $C_{\text{incorrect}}$ .  
 1073

1074 **Correct vs. Irrelevant Context** An irrelevant context ( $C_{\text{irrelevant}}$ ) is topically unrelated to the  
 1075 question and provides no useful information for answering it. Ideally, the model should recognize its  
 1076 lack of utility and ignore it, leading to little or no increase in confidence compared to the no-context  
 1077 baseline. In contrast,  $C_{\text{correct}}$  provides directly useful information, so the model should become  
 1078 more confident in its answer. Therefore, the confidence gain should be higher with  $C_{\text{correct}}$  than with  
 1079  $C_{\text{irrelevant}}$ .

1080 **D EXPERIMENT SETUP DETAILS**  
10811082 **D.1 DATASET DETAILS**  
1083

1084	Question	Which island in Kent is the second largest of England's isles?
1085	Correct Answers	"Shurland Hall", "Isle of Sheppy", "Shurland House", "Isle of Sheppey"
1086	Question	Rita Coolidge sang the title song for which Bond film?
1087	Correct Answers	"Kamal kahn", "List of Bond girls in Octopussy", "Magda (James Bond)", "List of James Bond allies in Octopussy", "Vijay (James Bond)", "Bond 13", "Octopussy (character)", "Penelope Smallbone", "Octopussy", "General Orlov", "Kamal Khan", "Octopussy (film)", "List of James Bond villains in Octopussy", "Jim Fanning (James Bond)"
1088	Question	What was invented by Jonas Hanway in the late 1750s?
1089	Correct Answers	"Umbrella", "Umbrela", "History of the umbrella", "Unbrella", "Beach umbrella", "Parasols", "Windproof umbrella", "Beach parasol", "Parasol", "History of the Umbrella", "Umbrellas"
1090		
1091		
1092		
1093		
1094		
1095		

1096 Table 4: Examples from the TriviaQA dataset.  
1097

1099	Question	Which of the following bones develop by endochondral ossification?
1100	Choices	A: "The ribs", B: "The ribs and sternum", C: "The ribs, sternum and clavicle", D: "The ribs, sternum, clavicle and vertebrae"
1101	Correct Answer	B
1102	Question	Which of the following gives the total spin quantum number of the electrons in the ground state of neutral nitrogen (Z = 7)?
1103	Choices	A: "1/2", B: "1", C: "3/2", D: "5/2"
1104	Correct Answer	C
1105	Question	The Hawthorn Studies are most associated with which writer?
1106	Choices	A: "Mary Parker Follett", B: "Elton Mayo", C: "Lillian Gilbreth", D: "Frederick Taylor"
1107	Correct Answer	B
1108		
1109		
1110		

1111 Table 5: Examples from the MMLU dataset.  
1112

1113  
1114 **Quality Filtering:** TriviaQA includes multiple retrieved Wikipedia or web documents as evidence  
1115 for each example, which we repurpose to serve as (correct) context for RQ2 experiments. While the  
1116 documents are generally of high quality, manual inspection reveals a small percentage of examples  
1117 where the documents do not contain enough information to answer the question. As a result, we  
1118 exclude any example where no retrieved document contains one of the ground-truth answers and at  
1119 least 60% of entities extracted from the question.<sup>3</sup> This yields a final set of 6,557 examples, which we  
1120 use for all our experiments. No quality filtering is necessary for MMLU. Table 4 and Table 5 present  
1121 example questions along with their corresponding ground-truth answers from TriviaQA and MMLU.  
1122

1123 **External Context Construction:** Concatenating multiple long documents and analyzing these is  
1124 challenging for interpretability methods that normally operate at the single-token or very few token  
1125 level. To address this, we select a single evidence document from TriviaQA that meets the quality  
1126 filtering criteria described above. If the selected document exceeds 500 tokens, we use GPT-4o to  
1127 summarize it down to 500 tokens, using the LLaMA-3-8B tokenizer for consistency. This summary  
1128 is termed the "correct" context for the experiments for RQ2. To generate "incorrect" contexts, we  
1129 take the "correct" context for a given question  $Q$  and prompt GPT-4o to replace all references to the  
1130 ground-truth answer with plausible but incorrect alternatives. Crucially, we keep the rest of the text  
1131 unchanged, which makes the surrounding context relevant to  $Q$ . We also obtain "irrelevant" contexts  
1132 by sampling "correct" contexts from a different question such that the context has no lexical overlap  
1133 with the ground-truth answer. Additionally, we require that the SentenceBERT (Reimers & Gurevych,

<sup>3</sup>We use the FLAIR Sequence Tagger for entity extraction (Akbik et al., 2019).

1134 2019) embeddings using all-MiniLM-L6-v2 of the candidate context have a cosine similarity < 0.3  
 1135 with the “correct” context for that question  $Q$ . This context construction process is grounded in the  
 1136 original “correc” context, and the LLM is used solely for straightforward tasks such as summarization  
 1137 and term replacement.

1138

1139 **Prompting Formats:** Below are the prompt formats used to generate answers from models on  
 1140 TriviaQA questions, either with or without supporting context.

1141

1142 *Without context:*

1143

1144 Answer the following question with a single word or phrase. Do not provide  
 1145 explanations or additional context:

1146

1147 {Question}

1148

1149 *With context:*

1150

1151 Using the following context, answer the question that follows with a  
 1152 single word or phrase. Do not provide explanations or additional context:

1153

1154 > Context:  
 1155 {Context}

1156

1157 > Question:  
 1158 {Question}

1159

1158 We use the following prompt format to generate answers from models for MMLU questions.

1159

1160 The following is a multiple-choice question. Please choose the most  
 1161 suitable one among A, B, C and D as the answer to this question. Only  
 1162 output the choice identifier (A, B, C, or D) and nothing else. Do not  
 1163 provide explanations or additional context.

1164

1165 {Question}

1166

1167 **Evaluation:** While the TriviaQA evaluation framework relies on exact match after text normalization  
 1168 between the candidate answer and one of the ground-truth answers, this rule-based approach  
 1169 can result in significant false negatives when the answer is correct but not phrased in exactly the  
 1170 same wording as the ground-truth answers. We use GPT-4o to assess correctness when exact match  
 1171 is not found due to its performance in factual QA evaluation (Zheng et al., 2023; Gu et al., 2024).  
 1172 Specifically, if the generated answer does not exactly match any of the ground-truth answers, we  
 1173 prompt GPT-4o with the question  $Q$ , the ground-truth answers  $A^*$ , and the generated answer  $A$  and  
 1174 ask it to judge correctness. Unattempted questions are marked incorrect. This improves robustness,  
 1175 especially because LLMs can phrase answers in slightly different but equivalent ways and may  
 1176 not exactly follow a given output format, especially when minimally post-trained. Since MMLU  
 1177 questions are multiple-choice, we ask the model to only output the choice identifier and rely on  
 1178 regex-based answer extraction and verification.

1179

## 1180 D.2 RQ1 SETUP DETAILS

1181

1182 **Logit-Lens and Tuned-Lens** We include all top- $p$  values from the set  $\{0.5, 0.9, 0.95, 0.99\}$  as input  
 1183 features when training our classifier using Logit-Lens and Tuned-Lens representations. We extract  
 1184 the statistics described in §B.3 (entropy, rank of the correct token, top- $p$  presence, and cross-entropy  
 1185 loss) from the decoded logits for each model layer. We then train our classifier  $f$  using these features  
 1186 to predict whether a generated answer is correct or not.

1187

1188 **Hidden States** We train a separate classifier on the hidden states extracted from each transformer  
 1189 layer during the generation of the first answer token. We focus on the first token because, due to the

1188 autoregressive nature of transformer models, the output at this initial step is conditioned solely on the  
 1189 input prompt and is unaffected by previously generated tokens (Zhao et al., 2024). For each example,  
 1190 per-dimension mean and standard deviation values used for calculating z-scores are obtained by  
 1191 collecting the first-output-token hidden states from each layer for questions from an unused split of  
 1192 the TriviaQA and MMLU datasets. This avoids any sharing of information across examples.  
 1193

1194 **Classifier:** In practice, one could use any classifier here. Subramani et al. (2025) show that  
 1195 random forests are best able to recalibrate tool-using agents when using model internals as features  
 1196 in comparison to other simple classifiers such as logistic regression. In our experiments, we use a  
 1197 random forest due to its efficacy, simplicity, and interpretability. We use grid search and five-fold  
 1198 cross-validation over a small subset (1024 examples) of our training set to set hyperparameters for  
 1199 the random forest classifier.

1200 During hyperparameter tuning, a grid search was conducted over the following parameters for the  
 1201 Random Forest classifier: the number of estimators (`n_estimators`  $\in \{100, 200, 300\}$ ), maxi-  
 1202 mum tree depth (`max_depth`  $\in \{\text{None}, 10, 20, 30\}$ ), minimum number of samples required to  
 1203 split an internal node (`min_samples_split`  $\in \{2, 5\}$ ), minimum number of samples required at a  
 1204 leaf node (`min_samples_leaf`  $\in \{1, 2\}$ ), the number of features considered for splitting at each  
 1205 node (`max_features`  $\in \{\text{"sqrt"}, \text{"log2"}, \text{None}\}$ ), and class weighting strategies (`class_weight`  $\in$   
 1206  $\{\text{"balanced"}, \text{"balanced_subsample"}, \text{None}\}$ ). The hyperparameters were set to the following for ran-  
 1207 dom forest classifiers trained on Logit Lens, Tuned Lens, and PKS features: `n_estimators` = 300,  
 1208 `max_depth` = `None`, `max_features` = `"log2"`, `min_samples_leaf` = 1, `min_samples_split` =  
 1209 2, `class_weight` = `"balanced"`. The hyperparameters were set to the following for ran-  
 1210 dom forest classifiers trained on hidden state features: `n_estimators` = 300, `max_depth` = 10,  
 1211 `max_features` = `"sqrt"`, `min_samples_leaf` = 1, `min_samples_split` = 2, `class_weight` =  
 1212 `"balanced_subsample"`.  
 1213

## E ADDITIONAL RESULTS FOR RQ1

### E.1 PROMPTING

1217 For prompting, we evaluate all integer threshold values from 0 to 100 and select the one that yields  
 1218 the highest accuracy. Accordingly, the performance reported for prompting methods reflects the  
 1219 best-performing threshold. Figure 3 and Figure 5 present the Smooth ECE scores of various models  
 1220 using prompting with answers on TriviaQA and MMLU respectively, while Figure 4 and Figure 6  
 1221 show the scores for prompting without answers.

### E.2 LOGIT LENS

1225 **Logit Lens VS. Tuned Lens** Figure 7 shows the layerwise AUC-ROC performance of the Logit  
 1226 Lens and Tuned Lens on LLaMA 3 8B, LLaMA 2 13B, and LLaMA 2 7B. We do not observe a  
 1227 consistent trend in the relative performance that would clearly indicate which method is superior.  
 1228 However, both the Logit Lens and Tuned Lens exhibit predictive capability from very early layers.

1229 **Feature Importance** We present the Logit Lens feature importance heatmaps for all six models  
 1230 in Figure 8 based on the trained classifiers. Top- $k$  features are omitted due to their consistently low  
 1231 importance across all layers and models.

1233 **ROC Curve: External VS. Internal** Figure 9 shows the ROC curve comparisons between clas-  
 1234 sifiers trained solely on last-layer features (“external”) and those trained on internal-layer features  
 1235 (“internal”) across all six models. The internal classifier achieves performance comparable to using  
 1236 features from all layers, while consistently outperforming the external classifier across all tested  
 1237 models.

### E.3 HIDDEN STATES

1239 **Hidden States: Effect of Z-Score Normalization** To evaluate the impact of z-score normalization  
 1240 on classifier performance we obtain means and variances per hidden state dimension for z-score

	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Averaged ACC	0.778	0.710	0.729	0.770	0.737	0.734
Averaged AUC ROC	0.774	0.772	0.804	0.741	0.790	0.803
First token ACC	0.782	0.779	0.776	0.774	0.751	0.729
First token AUC ROC	0.790	0.847	0.835	0.747	0.826	0.812

Table 6: TriviaQA results for LogitLens comparing averaged vs. first token features.

	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B
Averaged ACC	0.780	0.703	0.736
Averaged AUC ROC	0.777	0.768	0.808
First token ACC	0.779	0.775	0.775
First token AUC ROC	0.782	0.846	0.829

Table 7: TriviaQA results for TunedLens comparing averaged vs. first token features.

normalization on auxiliary subsets of TriviaQA and MMLU to ensure the distribution remains similar to the train and test examples while avoiding leakage of information across examples. To evaluate the impact of z-score normalization on classifier performance, we compare models trained on hidden state features with and without normalization and find virtually identical performance (see Figure 11). This suggests that the classifier is relatively insensitive to the absolute magnitude of hidden state feature values across different dimensions and examples.

#### E.4 FIRST TOKEN FEATURES VS. AVERAGING ACROSS ALL OUTPUT TOKENS

We use first token features for all classifier training in RQ1. We study whether interpretability features are computed using only the first generated token or averaged across all generated tokens. First token features provide the practical benefit of enabling early auditing before the model produces a full response. At the same time, aggregating features across all tokens may capture additional signal present throughout the generation process. To evaluate these two approaches, we conduct an ablation across multiple models and multiple internal feature extraction methods, including LogitLens, TunedLens, hidden states, and PKS.

Across most models and datasets, averaging features over all generated tokens performs similarly or slightly better than using only the first token. The differences are most pronounced for LogitLens and TunedLens on MMLU, where full token averaging provides consistent improvements. On TriviaQA, both approaches achieve comparable performance with small method dependent variations. These findings suggest that while full token averaging can offer moderate gains, first token features remain competitive and are particularly useful when early auditing is desired for efficiency.

**TriviaQA Results** TriviaQA results for all methods are in Tables 6–9.

**MMLU Results** MMLU results are shown in Tables 10–12. Here, averaging across tokens provides more consistent benefits, particularly for LogitLens and TunedLens.

#### E.5 MLP CLASSIFIER VS. RANDOM FOREST CLASSIFIER

This section investigates whether more expressive classifiers can further improve correctness prediction. Our main experiments use a random forest classifier due to its simplicity, robustness, and low tuning requirements. To assess whether a more flexible model can extract additional signal from interpretability features, we train a multi layer perceptron (MLP) classifier across all models and datasets. The MLP has two hidden layers with 256 and 64 units, uses ReLU activations, applies L2 regularization with coefficient 0.001, and is optimized with Adam using adaptive learning rates and early stopping.

The results reveal that the MLP achieves performance comparable to or better than the random forest classifier across most settings. Improvements are especially pronounced on MMLU, where the MLP consistently reaches higher accuracy and AUC ROC. These findings indicate that correctness

	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Averaged ACC	0.739	0.714	0.707	0.773	0.782	0.768
Averaged AUC ROC	0.611	0.670	0.633	0.584	0.769	0.754
First token ACC	0.759	0.679	0.702	0.785	0.782	0.749
First token AUC ROC	0.647	0.631	0.647	0.616	0.774	0.735

1301 Table 8: TriviaQA results for Hidden States comparing averaged vs. first token features.  
1302

	LLaMA 3 8B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Averaged ACC	0.733	0.709	0.743	0.650	0.695
Averaged AUC ROC	0.729	0.743	0.723	0.715	0.752
First token ACC	0.737	0.668	0.733	0.659	0.684
First token AUC ROC	0.740	0.741	0.642	0.702	0.731

1309 Table 9: TriviaQA results for PKS comparing averaged vs. first token features.  
13101311 prediction can benefit from more expressive classifiers, although both approaches remain competitive  
1312 depending on the method and dataset.  
13131314 **TriviaQA Results** Full TriviaQA results are shown in Table 13.  
13151316 **MMLU Results** MMLU results appear in Table 14. Compared to the random forest classifier, the  
1317 MLP achieves large gains for Logit Lens and Tuned Lens and modest gains for Hidden States and  
1318 PKS.  
13191320 

## F ADDITIONAL RESULTS FOR RQ2

1321 

### F.1 NON-APPLICABILITY OF ECS ON GEMMA 2 9B

1322 Gemma 2 9B employs sliding-window attention on every even-numbered layer, which makes it  
1323 infeasible to implement ECS in a way that correctly tracks the top-k attention scores. Consequently,  
1324 we exclude Gemma 2 9B from our analysis for RQ2.  
13251326 

### F.2 EXAMPLES WITH PKS AND ECS SCORES

1327 We illustrate how PKS and ECS scores evolve as the model generates output tokens across layers,  
1328 using several examples shown in Figures 12, 13, and 14.  
13291330 

## G COMPUTATIONAL COSTS

1331 All experiments were conducted using two A6000 GPUs, with a total compute time of under 1,000  
1332 hours.  
1333

	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Averaged ACC	0.836	0.717	0.684	0.814	0.822	0.803
Averaged AUC ROC	0.941	0.789	0.757	0.904	0.945	0.908
First token ACC	0.705	0.692	0.699	0.769	0.815	0.673
First token AUC ROC	0.798	0.771	0.767	0.843	0.939	0.711

Table 10: MMLU results for LogitLens comparing averaged vs. first token features.

	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B
Averaged ACC	0.826	0.728	0.680
Averaged AUC ROC	0.928	0.815	0.768
First token ACC	0.695	0.671	0.684
First token AUC ROC	0.770	0.752	0.760

Table 11: MMLU results for TunedLens comparing averaged vs. first token features.

	LLaMA 3 8B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Averaged ACC	0.605	0.543	0.705	0.691	0.618
Averaged AUC ROC	0.537	0.543	0.555	0.541	0.538
First token ACC	0.588	0.529	0.708	0.694	0.615
First token AUC ROC	0.524	0.510	0.526	0.565	0.542

Table 12: MMLU results for PKS comparing averaged vs. first token features.

Method	Source	Metric	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Logit Lens	Table 2	ACC	0.782	0.779	0.776	0.774	0.751	0.729
Logit Lens	Table 2	AUC ROC	0.790	0.847	0.835	0.747	0.826	0.812
Logit Lens	MLP	ACC	0.788	0.778	0.760	0.770	0.711	0.686
Logit Lens	MLP	AUC ROC	0.790	0.844	0.820	0.758	0.764	0.755
Tuned Lens	Table 2	ACC	0.779	0.775	0.775	-	-	-
Tuned Lens	Table 2	AUC ROC	0.782	0.846	0.829	-	-	-
Tuned Lens	MLP	ACC	0.780	0.770	0.748	-	-	-
Tuned Lens	MLP	AUC ROC	0.784	0.827	0.808	-	-	-
Hidden States	Table 2	ACC	0.739	0.714	0.707	0.773	0.782	0.768
Hidden States	Table 2	AUC ROC	0.611	0.670	0.633	0.584	0.769	0.754
Hidden States	MLP	ACC	0.759	0.679	0.702	0.785	0.782	0.749
Hidden States	MLP	AUC ROC	0.647	0.631	0.647	0.616	0.774	0.735
PKS	Table 2	ACC	0.733	0.725	0.709	0.743	0.650	0.695
PKS	Table 2	AUC ROC	0.729	0.768	0.743	0.723	0.715	0.752
PKS	MLP	ACC	0.728	0.691	0.675	0.729	0.667	0.658
PKS	MLP	AUC ROC	0.715	0.744	0.722	0.595	0.731	0.730

Table 13: TriviaQA results comparing the random forest classifier (Table 2) with an MLP classifier across all methods and models.

Method	Source	Metric	LLaMA 3 8B	LLaMA 2 13B	LLaMA 2 7B	Gemma 2 9B	Qwen 2.5 7B	Qwen 2.5 3B
Logit Lens	Table 2	ACC	0.705	0.692	0.699	0.769	0.815	0.673
Logit Lens	Table 2	AUC ROC	0.798	0.771	0.767	0.843	0.939	0.711
Logit Lens	MLP	ACC	0.901	0.858	0.834	0.940	0.974	0.729
Logit Lens	MLP	AUC ROC	0.961	0.933	0.914	0.975	0.990	0.796
Tuned Lens	Table 2	ACC	0.705	0.692	0.699	-	-	-
Tuned Lens	Table 2	AUC ROC	0.798	0.771	0.767	-	-	-
Tuned Lens	MLP	ACC	0.901	0.858	0.834	-	-	-
Tuned Lens	MLP	AUC ROC	0.961	0.933	0.914	-	-	-
Hidden States	Table 2	ACC	0.744	0.686	0.672	0.801	0.777	0.746
Hidden States	Table 2	AUC ROC	0.740	0.684	0.576	0.736	0.728	0.733
Hidden States	MLP	ACC	0.748	0.679	0.664	0.779	0.770	0.730
Hidden States	MLP	AUC ROC	0.753	0.676	0.585	0.736	0.741	0.723
PKS	Table 2	ACC	0.605	-	0.543	0.705	0.691	0.618
PKS	Table 2	AUC ROC	0.537	-	0.543	0.555	0.541	0.538
PKS	MLP	ACC	0.596	-	0.530	0.709	0.694	0.617
PKS	MLP	AUC ROC	0.524	-	0.518	0.507	0.534	0.546

Table 14: MMLU results comparing the random forest classifier (Table 2) with an MLP classifier across all methods and models.

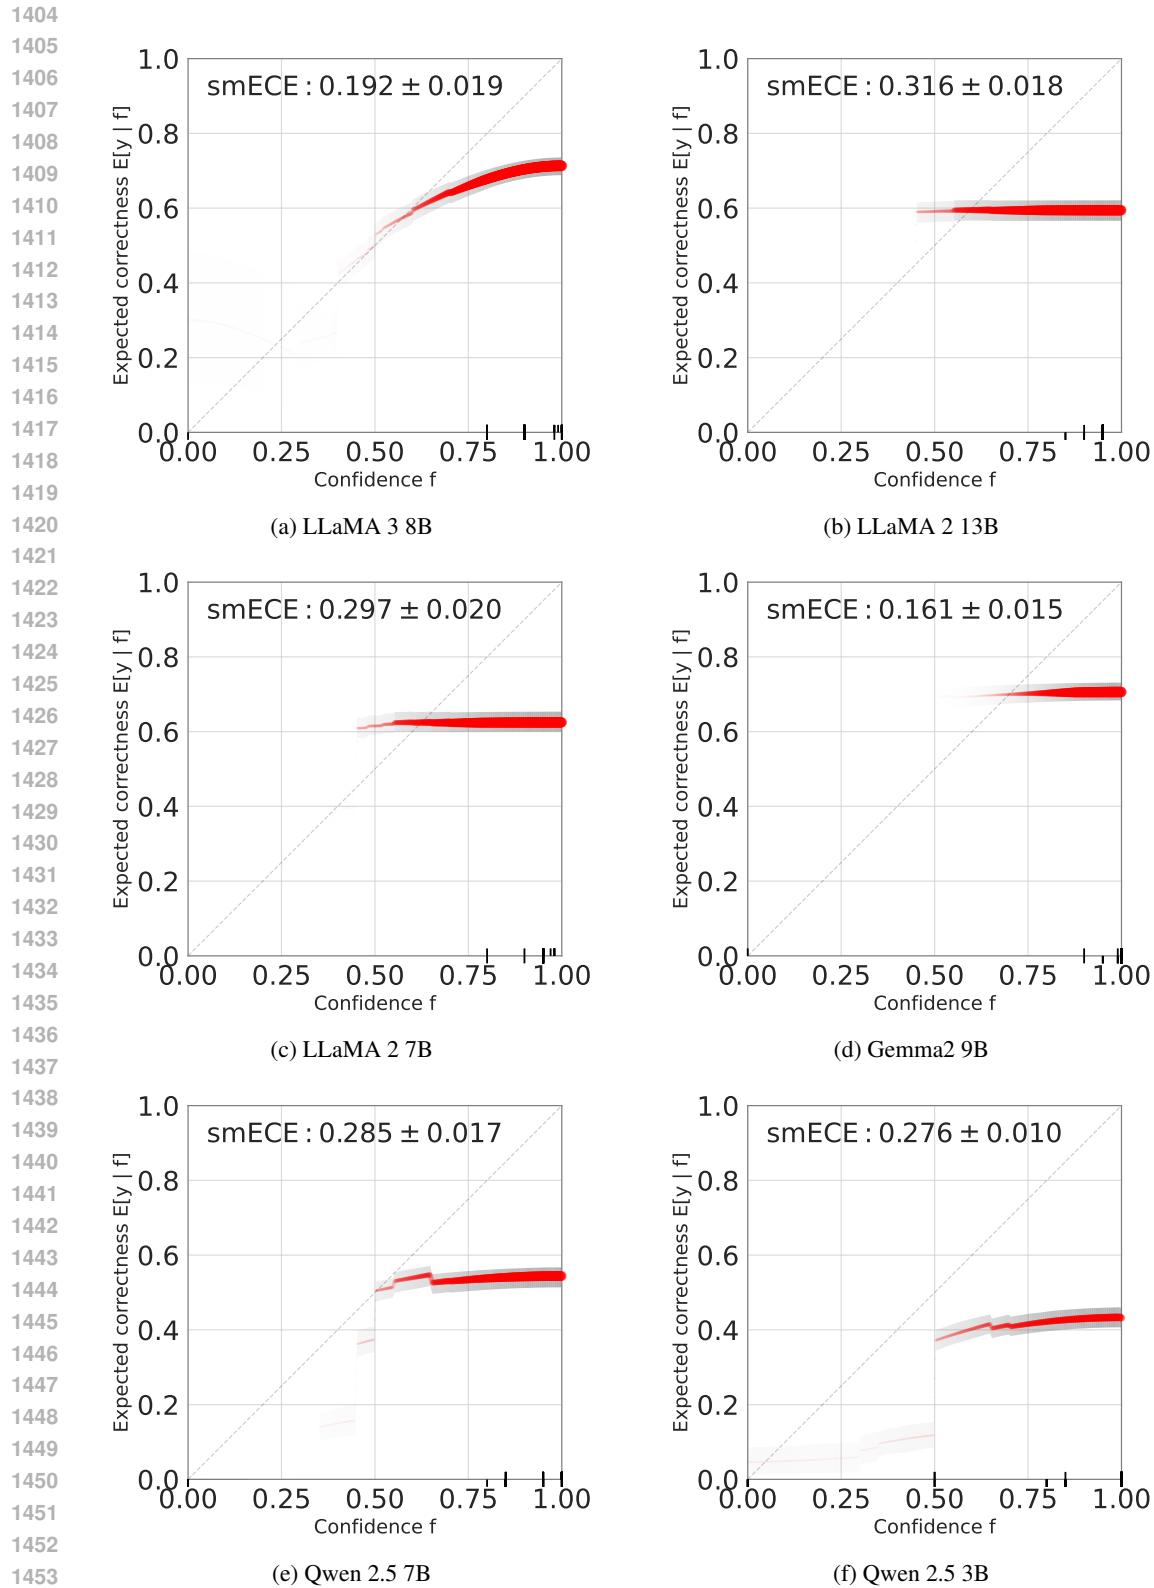


Figure 3: Smooth ECE scores for prompting with answer on TriviaQA. Predicted confidence  $f$  is the model’s stated probability of being correct, and  $E[y|f]$  is the actual accuracy observed at that confidence level.

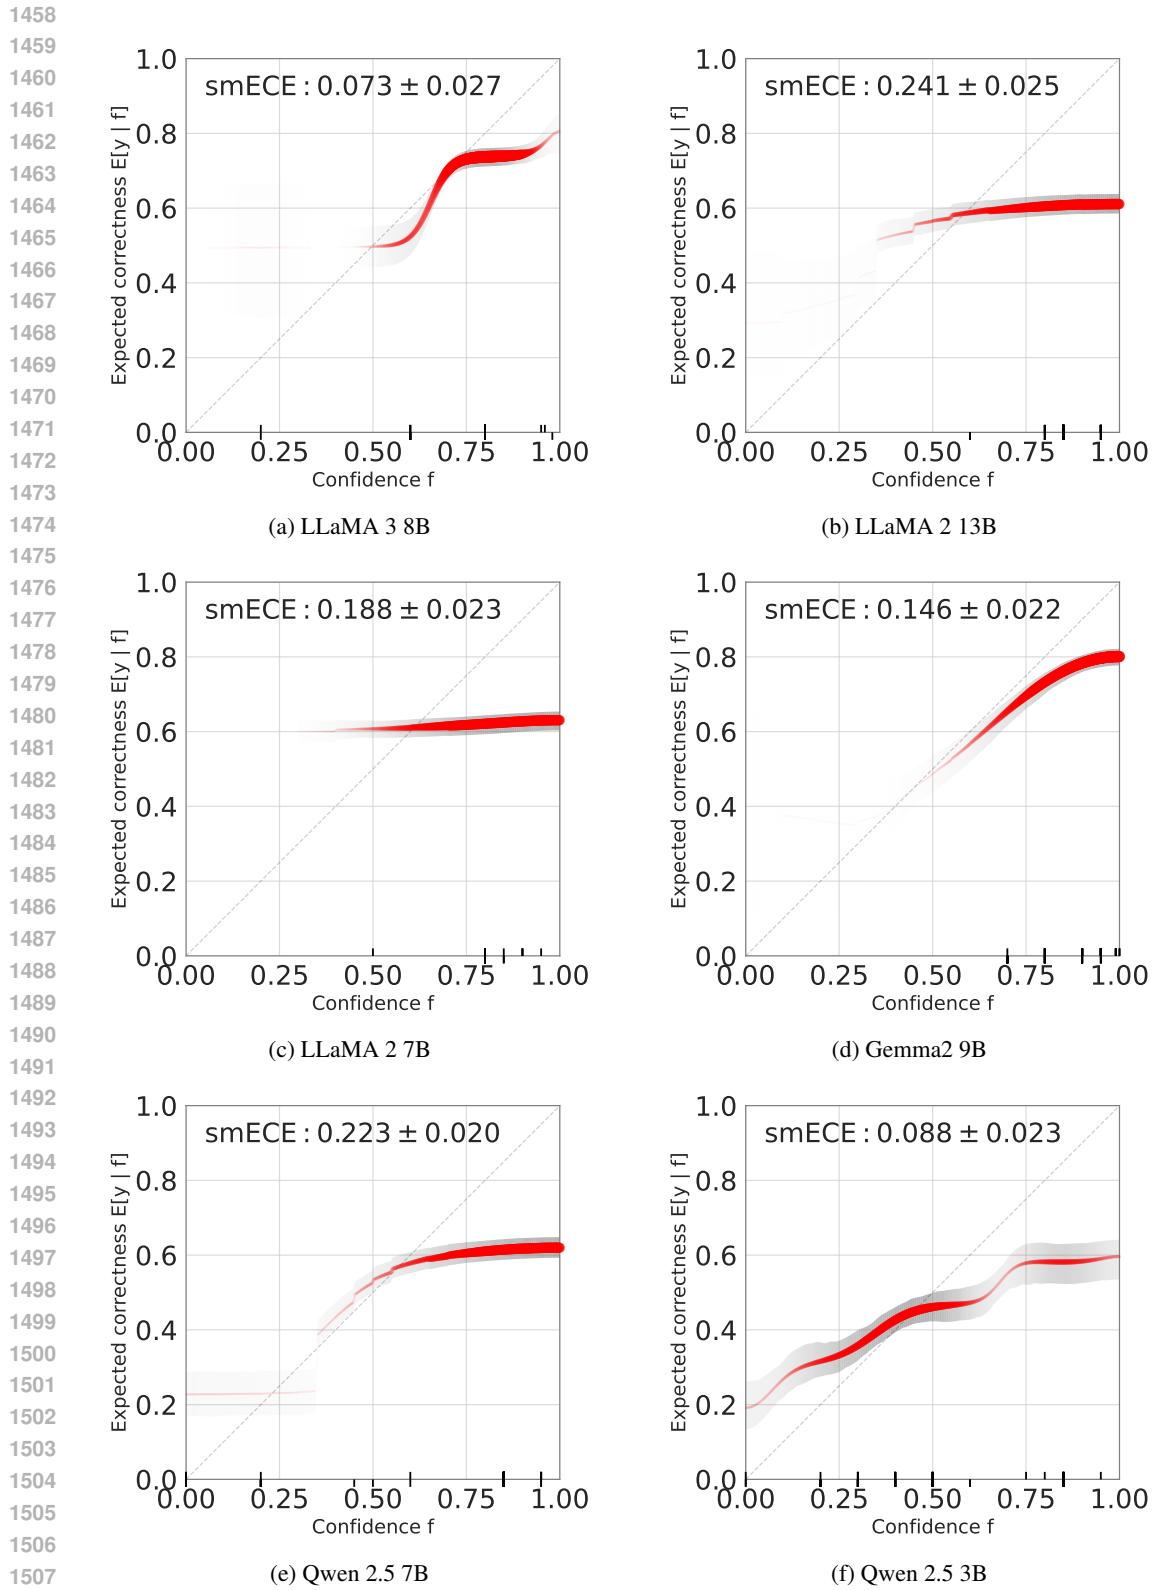


Figure 4: Smooth ECE scores for prompting without answer on TriviaQA. Predicted confidence  $f$  is the model's stated probability of being correct, and  $E[y|f]$  is the actual accuracy observed at that confidence level.

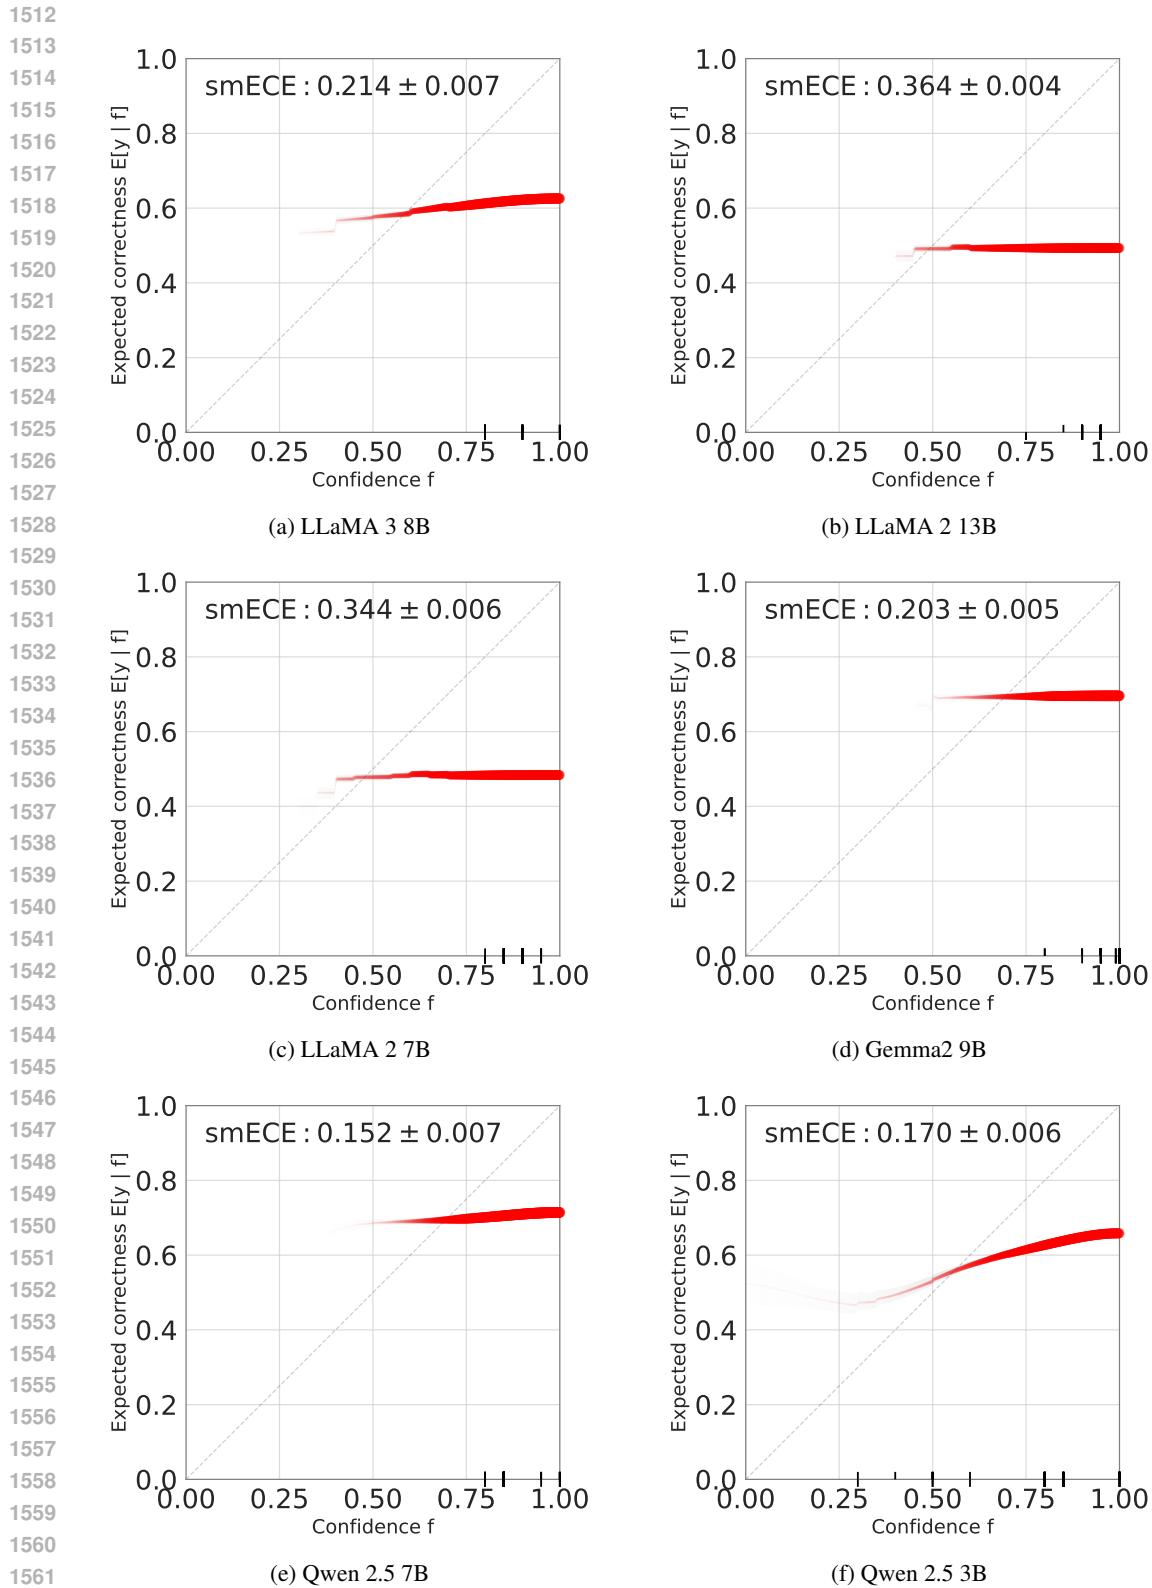


Figure 5: Smooth ECE scores for prompting with answer on MMLU. Predicted confidence  $f$  is the model's stated probability of being correct, and  $E[y|f]$  is the actual accuracy observed at that confidence level.

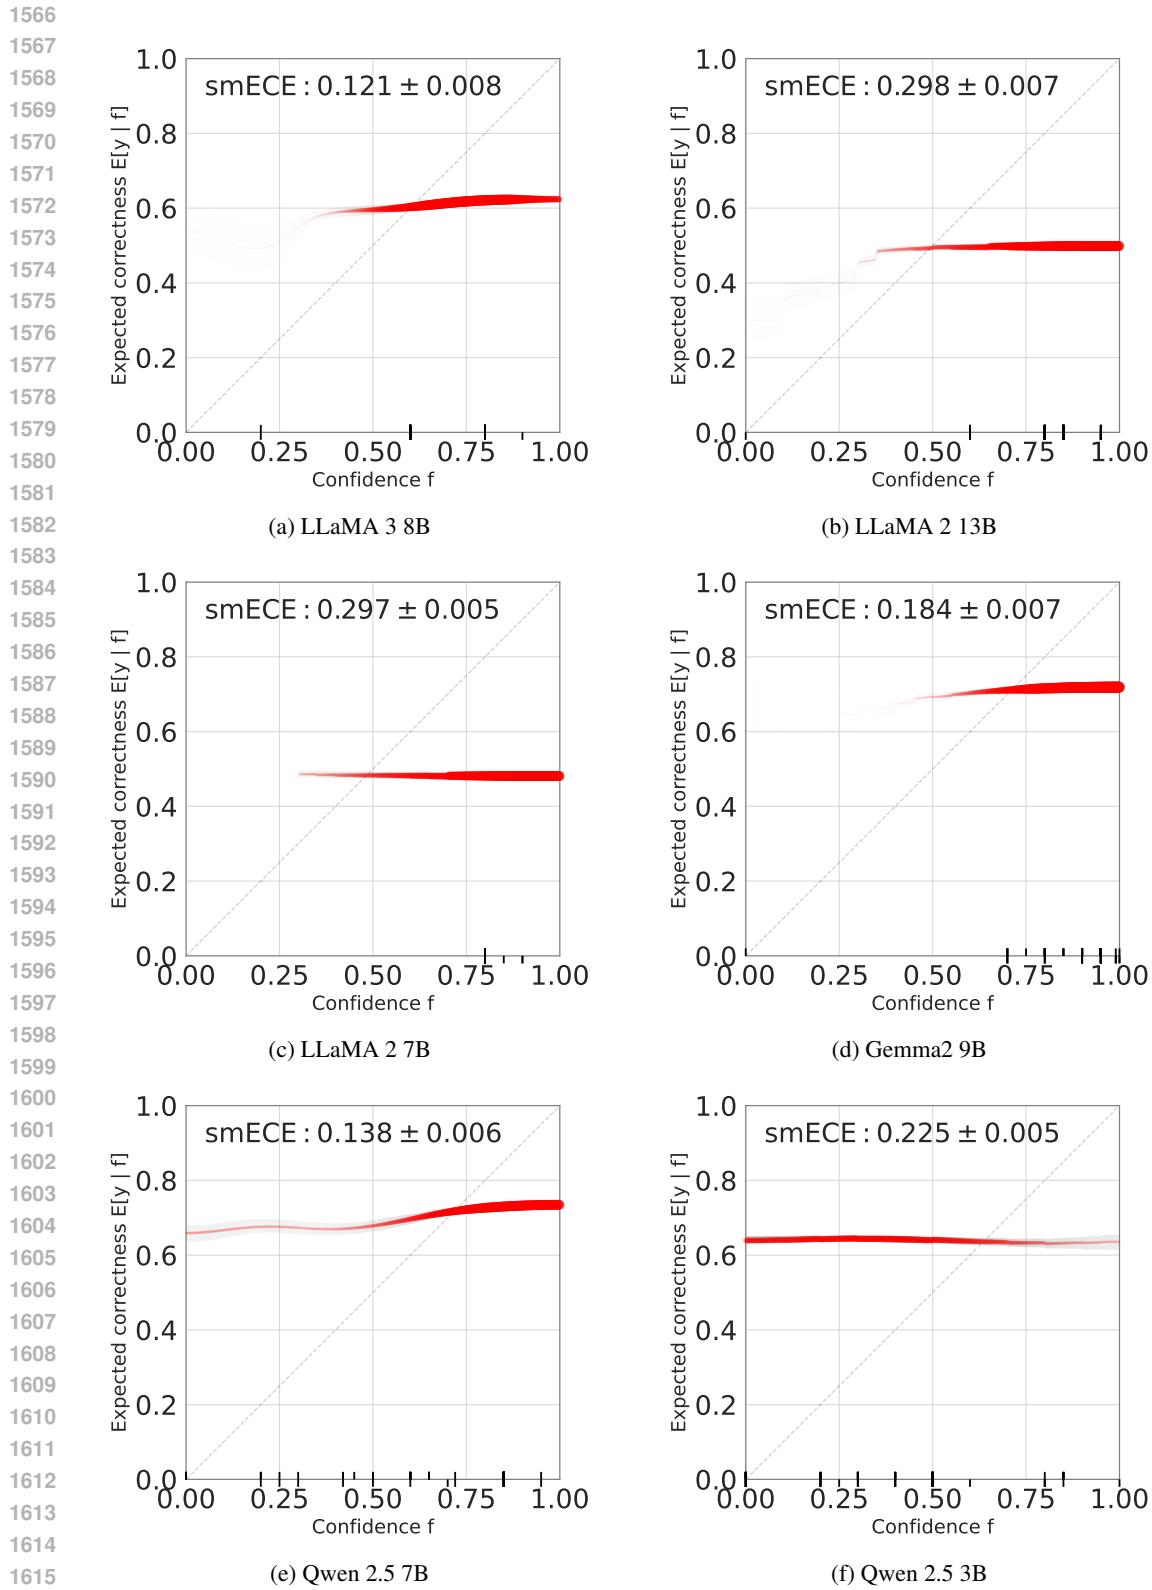


Figure 6: Smooth ECE scores for prompting without answer on MMLU. Predicted confidence  $f$  is the model's stated probability of being correct, and  $E[y|f]$  is the actual accuracy observed at that confidence level.

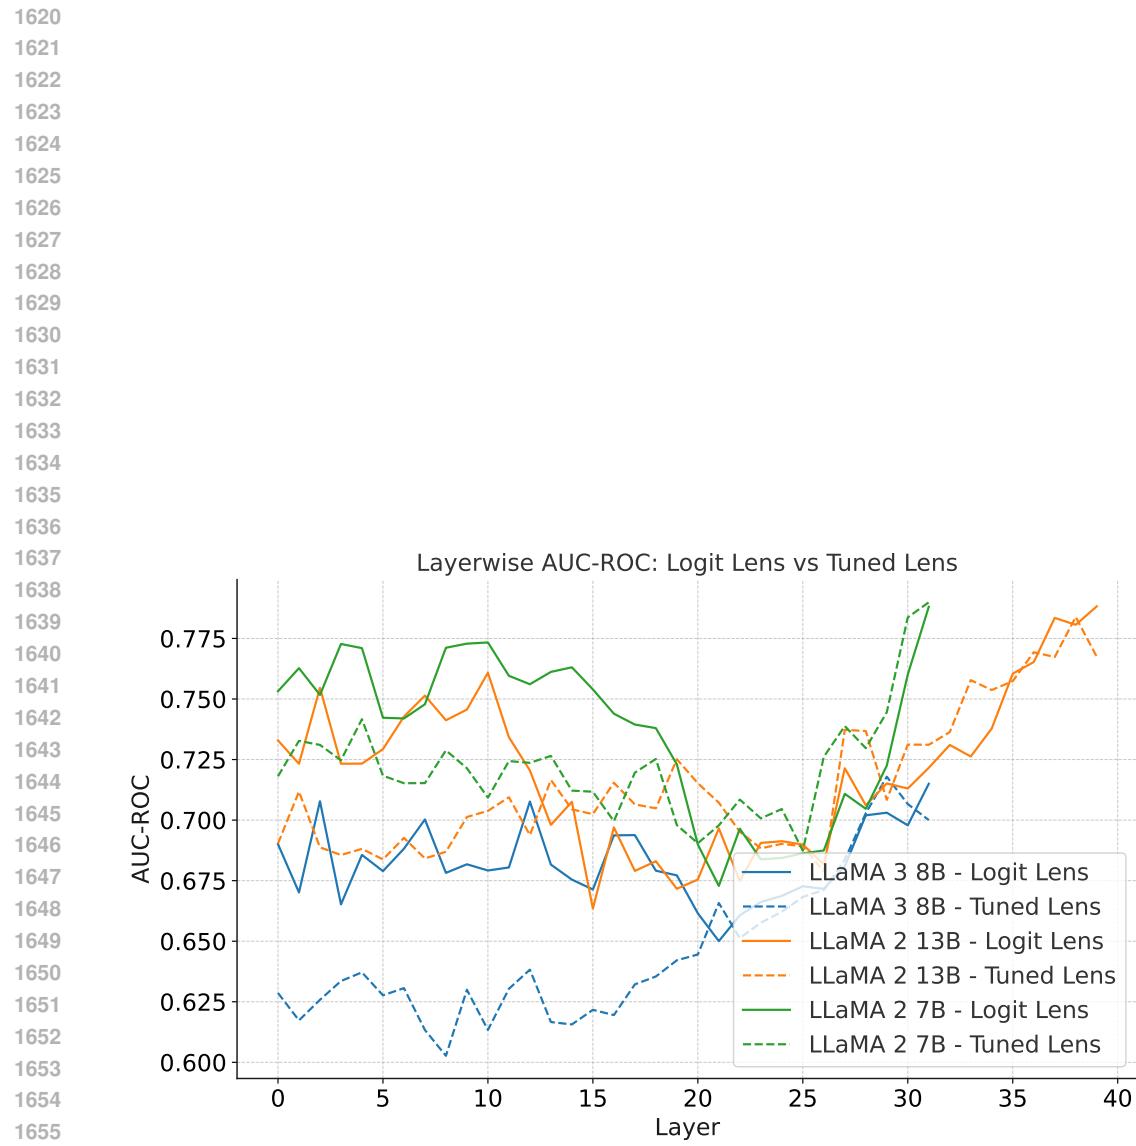


Figure 7: Layerwise AUC-ROC for Logit Lens (solid lines) and Tuned Lens (dashed lines).

1658  
 1659  
 1660  
 1661  
 1662  
 1663  
 1664  
 1665  
 1666  
 1667  
 1668  
 1669  
 1670  
 1671  
 1672  
 1673

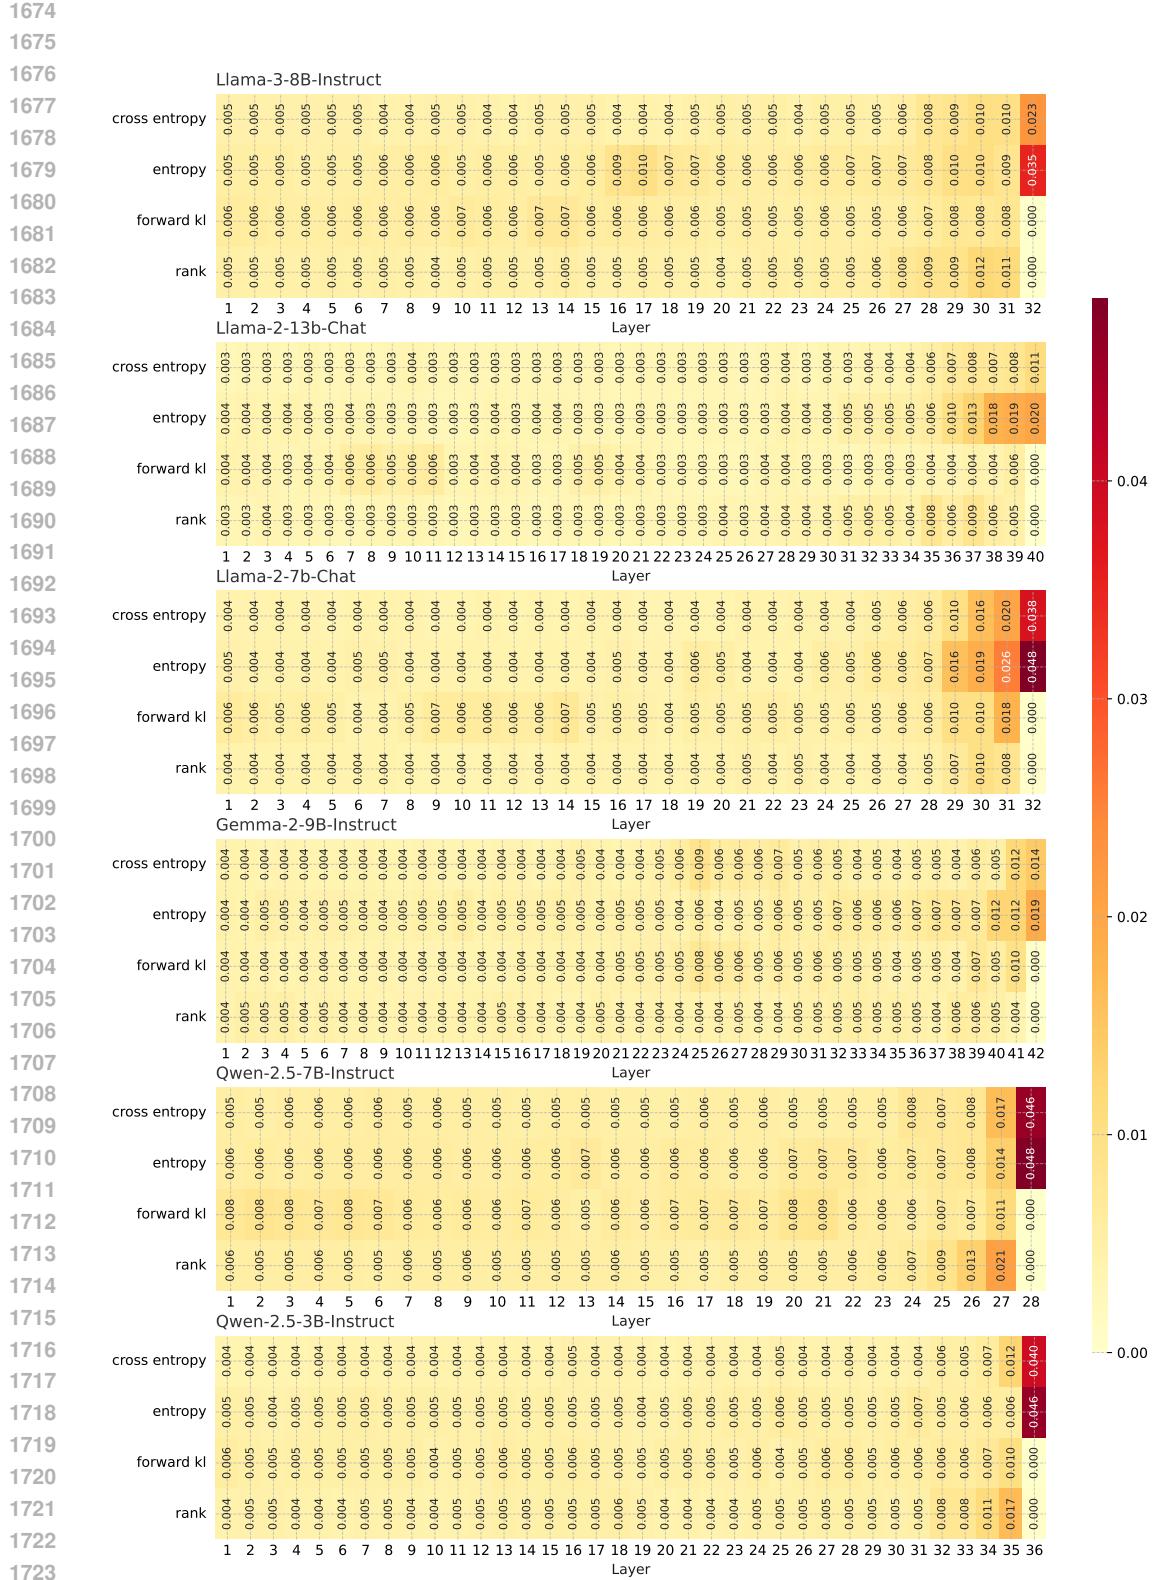


Figure 8: Impurity-based random forest feature importance scores for Logit Lens features from each layer across six models on TriviaQA. Top- $p$  features contribute minimally and are therefore excluded.

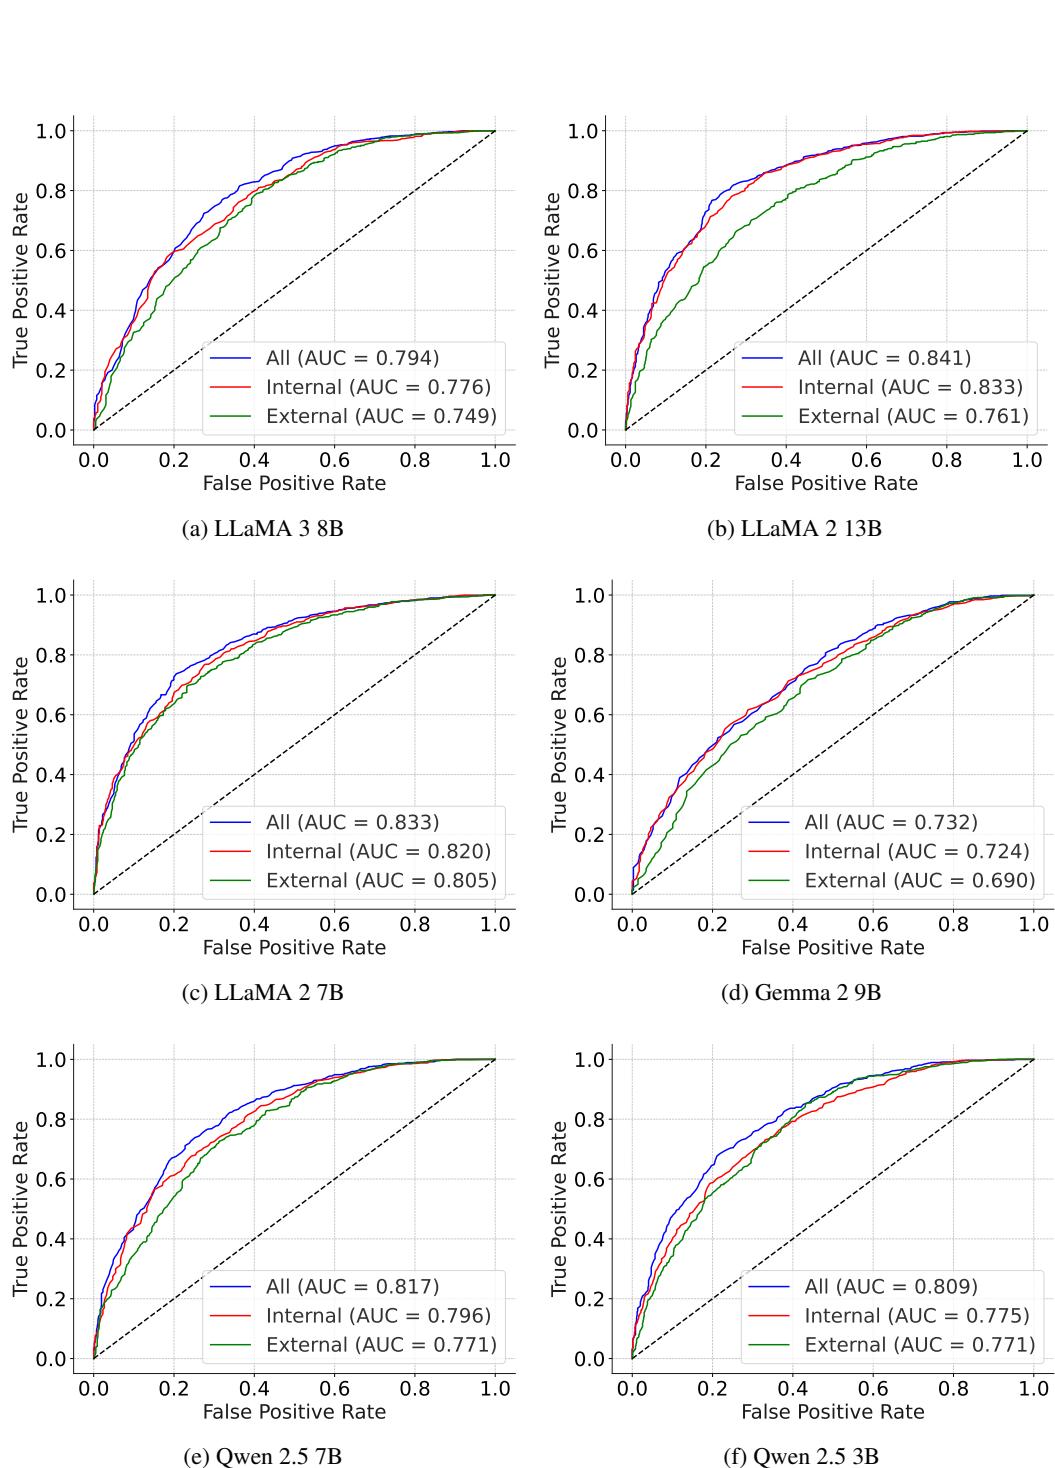


Figure 9: ROC curves comparing classifiers trained on last-layer features (external) versus internal-layer features, across six models on TriviaQA.

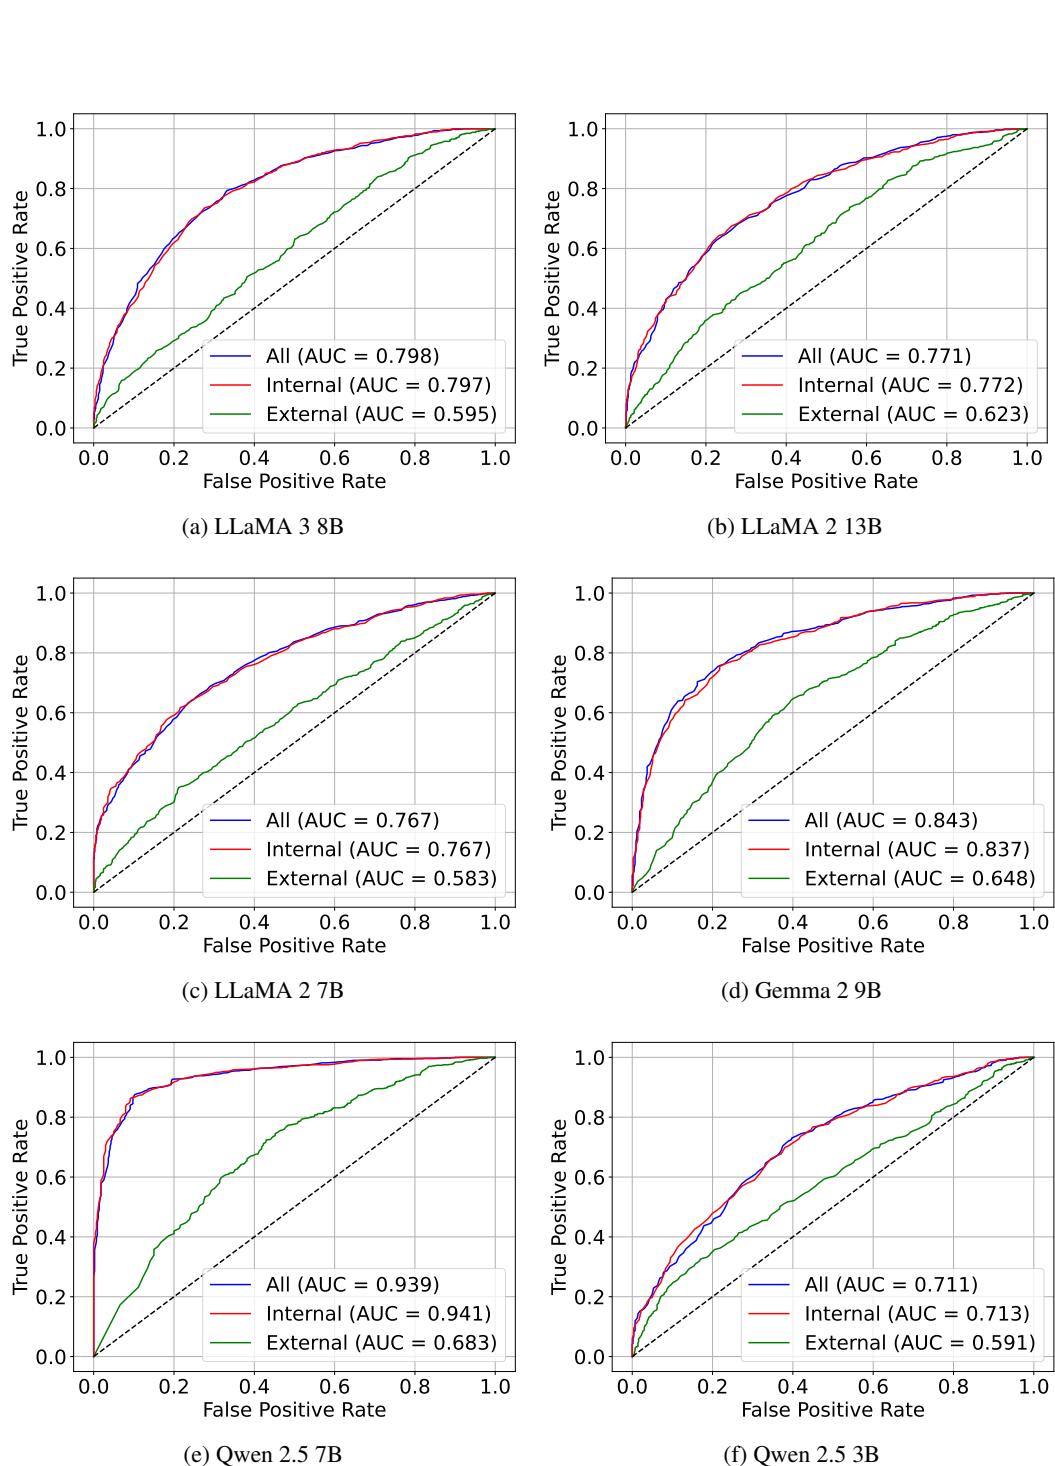


Figure 10: ROC curves comparing classifiers trained on last-layer features (external) versus internal-layer features, across six models on MMLU.

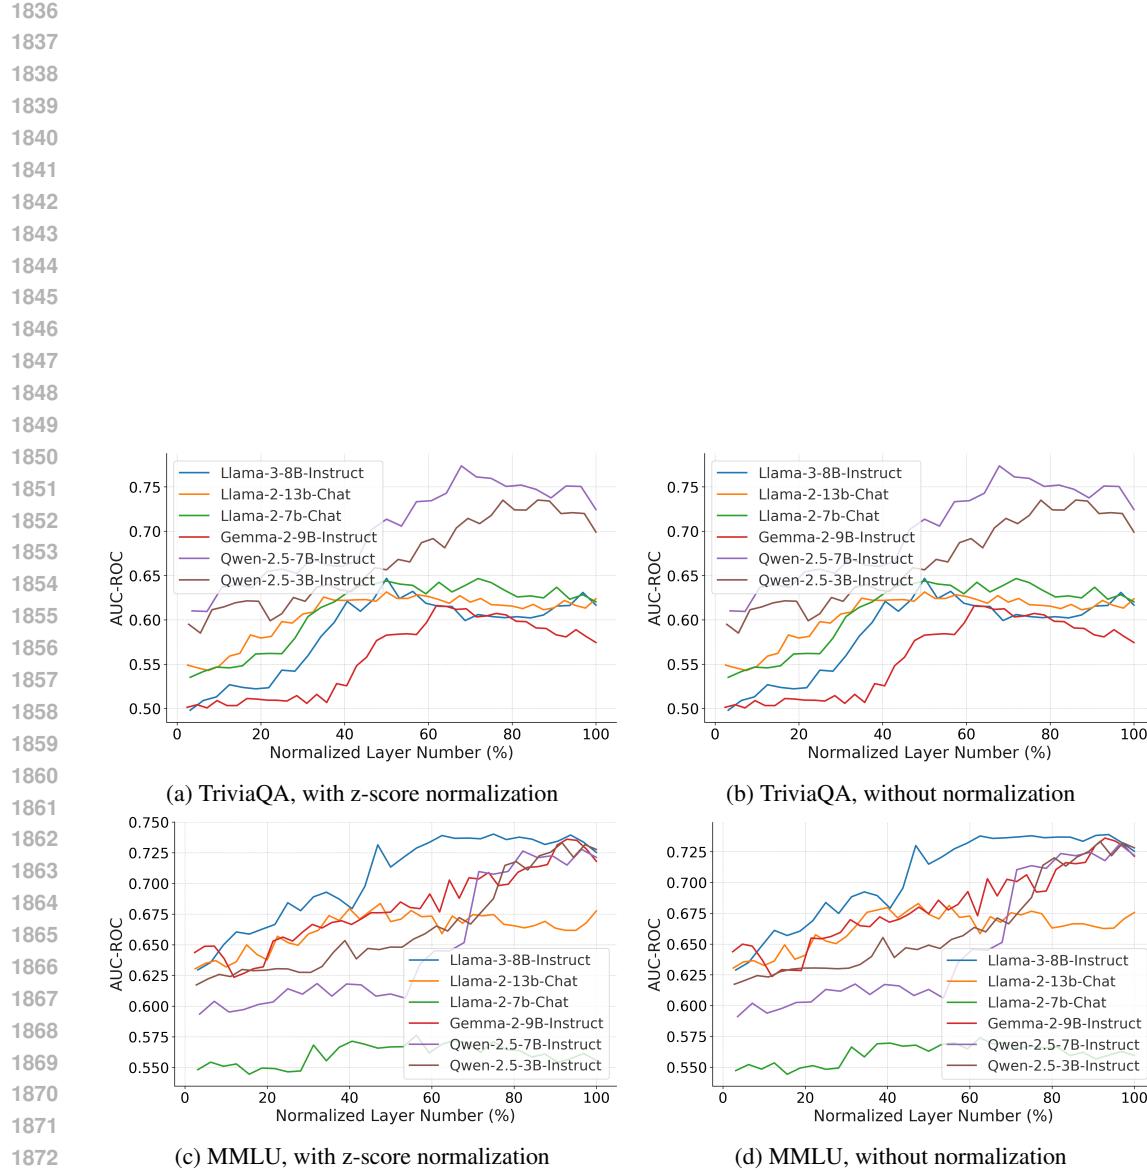


Figure 11: Area under ROC curve for random forest classifiers trained on z-score normalized hidden states of each layer. Performance increases with layer depth, suggesting that later layers refine and consolidate decision-relevant signals. Normalization has virtually no effect on performance.

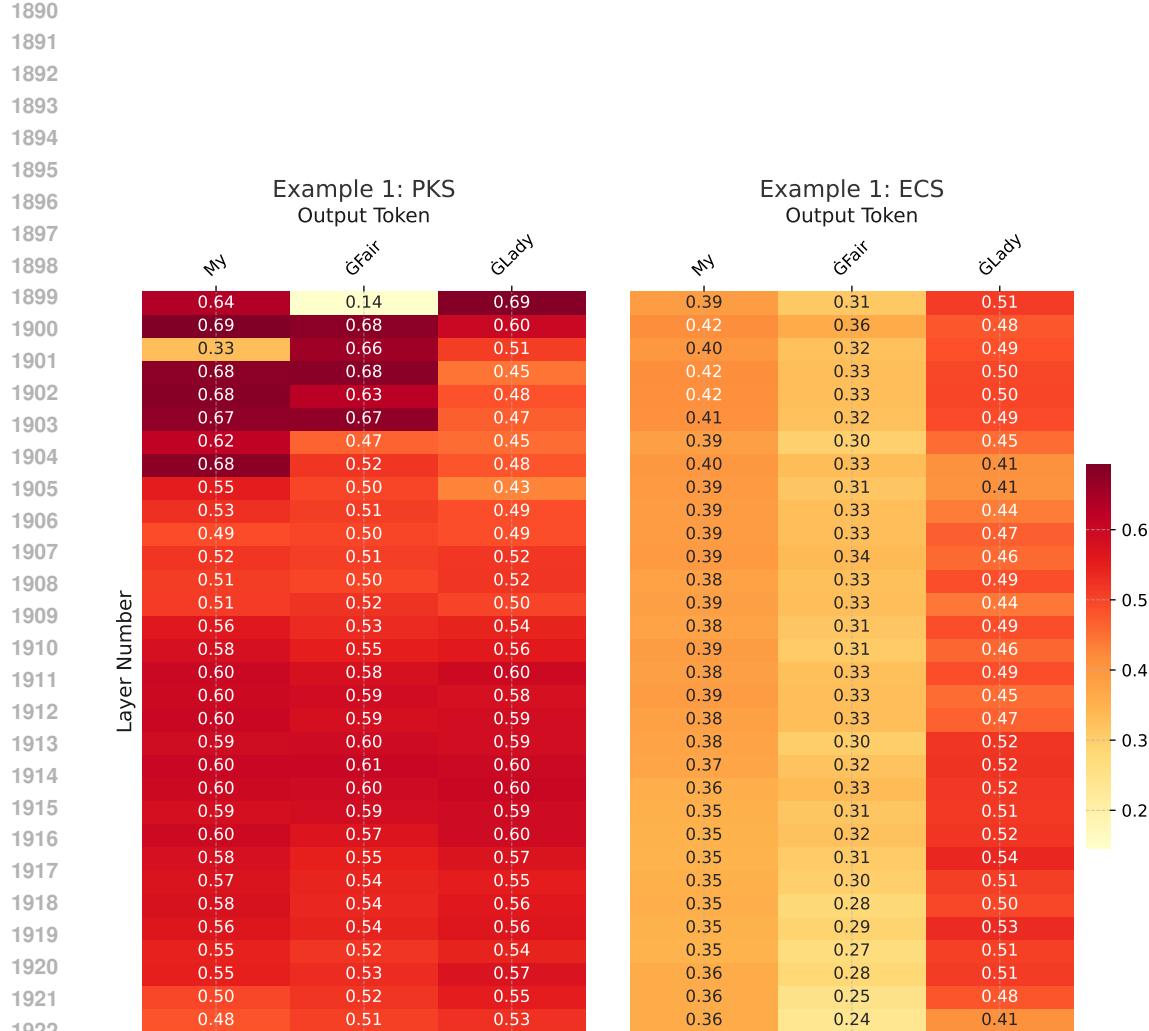


Figure 12: Heatmap example 1 from TriviaQA of PKS and ECS scores for LLaMA-3-8B, illustrating layer-wise token importance. Question: "Which musical featured the song *The Street Where You Live*?" Context: "'On the Street Where You Live' is a song with music by Frederick Loewe and lyrics by Alan Jay Lerner from the 1956 Broadway musical *My Fair Lady*. It is sung in the musical by the character Freddy Eynsford-Hill, originally portrayed by John Michael King. In the 1964 film version, the song was performed by Bill Shirley, dubbing for Jeremy Brett. The most popular single was recorded by Vic Damone in 1956, reaching #4 on the Billboard charts and #6 on Cash Box magazine's chart, and it was a #1 hit in the UK in 1958. In 1955, Damone had one song on the charts, "Por Favor," which peaked at #73, but he starred in *Hit the Deck* and *Kismet*. In 1956, he moved to Columbia Records, achieving success with hits like "On the Street Where You Live" from *My Fair Lady* and "An Affair to Remember." His albums on Columbia included *That Towering Feeling*, *Angela Mia*, *Closer Than a Kiss*, *This Game of Love*, *On the Swingin' Side*, and *Young and Lively*. Lyrics describe the narrator's thrill on the street where a loved one lives, highlighting the emotional impact of such proximity. The content is administered by SME and used here for educational purposes under fair use. If concerns arise about unauthorized use, contact the poster. This adheres to the Copyright Act's fair use principles for criticism, comment, news reporting, teaching, scholarship, and research, emphasizing non-profit, educational intentions."

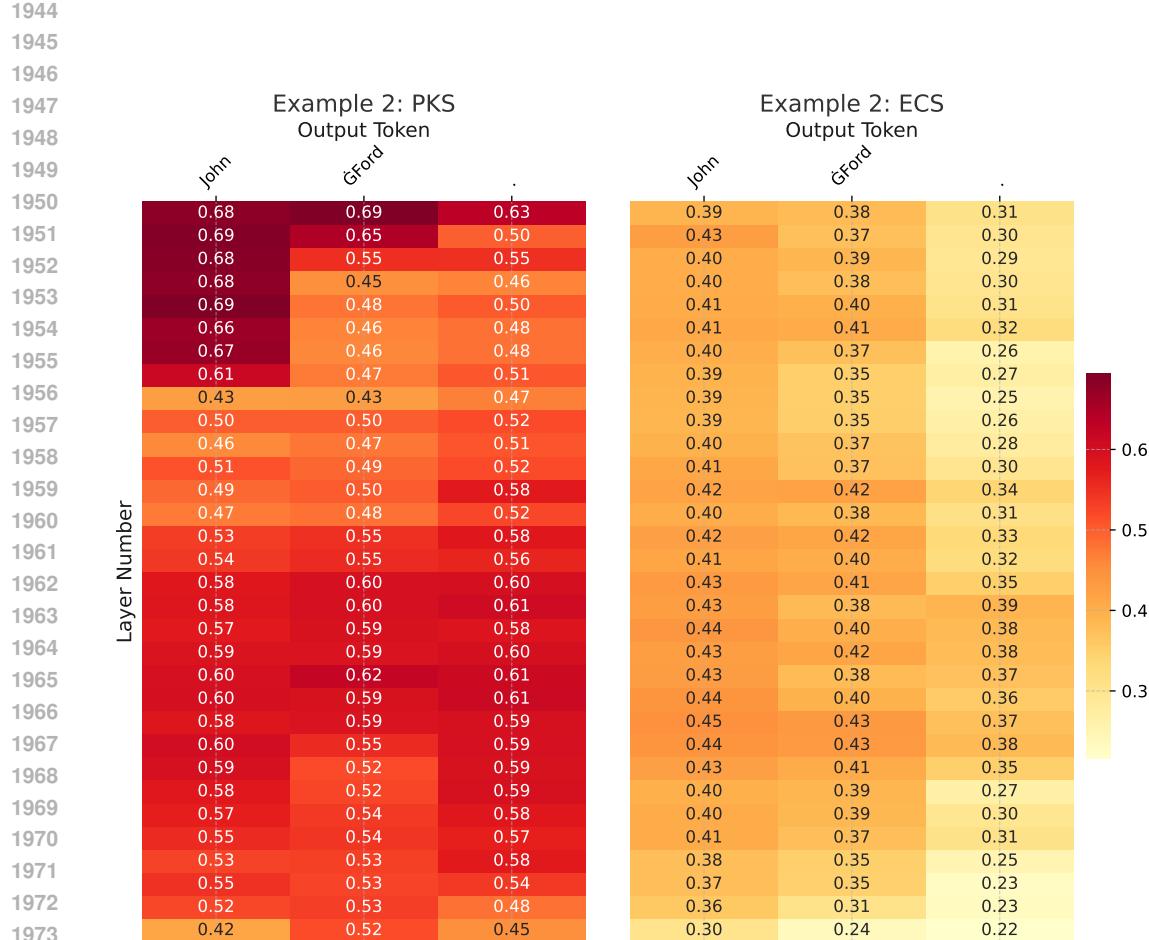


Figure 13: Heatmap example 2 from TriviaQA of PKS and ECS scores for LLaMA-3-8B, illustrating layer-wise token importance. Question: "Who directed the classic 30s western Stagecoach?" Context: "Stagecoach, directed by John Ford (1895-1973), is a quintessential Western notable for its complex characters and Monument Valley setting. John Wayne, portraying "The Ringo Kid," catapulted to stardom in this film. The narrative follows a diverse group of travelers on a stagecoach from Tonto to Lordsburg, facing dangers from an Apache uprising led by Geronimo. Central characters include a falsely accused outlaw Ringo Kid, played by Wayne—seeking to avenge his family's murder—and a prostitute named Dallas, portrayed by Claire Trevor. Other passengers include a drunken doctor (Thomas Mitchell, whose performance won an Academy Award), a gentleman gambler, an embezzling banker, a pregnant army wife, and others, each contributing to a microcosm of society. Ford's direction, coupled with Dudley Nichols and Ben Hecht's script, ensures tight plotting and memorable character arcs, many based on Ernest Haycox's short story "Stage to Lordsburg" and Guy de Maupassant's "Boule de Suif." Ford's handling allowed for minimal screen time yet deep character development. The film features an intense climax with a chase/fight scene that includes Yakima Canutt's pioneering stunts, echoing Spielberg's Raiders of the Lost Ark years later. Stagecoach is celebrated for blending mythic Western landscapes with a poignant social allegory, reflecting on issues like prejudice and redemption. While technical aspects, like some cinematography, show their age, the film remains a paragon from the era, emblematic of Ford's work and establishing recurring Western motifs. It earned seven Oscar nominations, securing wins for Best Supporting Actor and Best Score. Stagecoach's impact persists, marking a pivotal moment in film history and solidifying John Ford's legacy as a master director, influencing numerous directors and films in subsequent years. Whether or not it is perceived as a perfect film, it stands as a significant cultural artifact and vital viewing for any student of cinema."

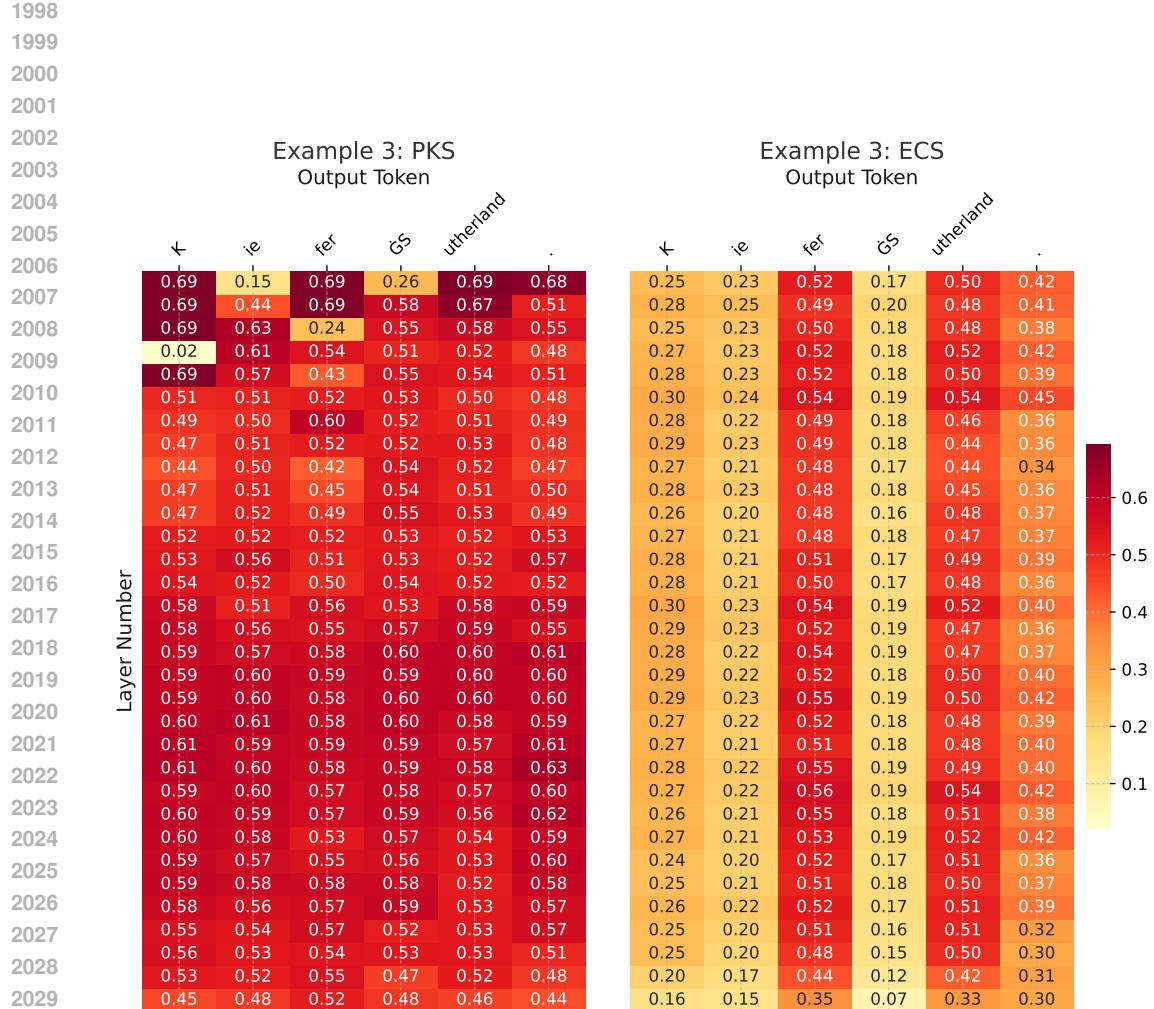


Figure 14: Heatmap example 3 from TriviaQA of PKS and ECS scores for LLaMA-3-8B, illustrating layer-wise token importance. Question: "Who was born first, Kiefer Sutherland or Christian Slater?" Context: "Young Guns II (1990) follows Billy the Kid and his gang as wanted outlaws. The story unfolds with Pat Garrett, a former partner of Billy's, being paid to kill him by cattle baron John Chisum. The movie, directed by Christopher Cain, explores themes of betrayal and redemption as Pat Garrett, who has plans to go respectable, is conflicted about turning against his former friend. Garrett seizes the opportunity to become Sheriff offered by the Governor, who believes in hiring a thief to catch one. This decision to capture Billy opens up great western adventure, with Pat grappling between loyalty and duty. The nuanced performance of William Petersen as Garrett contrasts well with Emilio Estevez reprising his role as the charismatic Billy the Kid. Lou Diamond Phillips elevates his role as Chavez, delivering a performance that is more spiritual and wise than in the first Young Guns film. The talented cast also includes Kiefer Sutherland, Christian Slater, Balthazar Getty, and Alan Ruck. Though depicted as close friends, the real-life association between Pat Garrett and Billy the Kid was less intimate; their familiarity stemmed from mutual patronage of a saloon. Despite this inaccuracy, the film presents an engaging exploration of the dynamics between Garrett and Billy. Young Guns II is a remarkable sequel to 1988's Young Guns, offering a compelling mix of action and moral questions. The film's strong character portrayals, particularly by Estevez and Phillips, enhance its rich narrative of western adventure and friendship against a backdrop of historical myth. It's an exciting film experience that shouldn't be missed."