

Cost-Effective Hallucination Detection for LLMs

Simon Valentin*
Amazon Web Services
Berlin, Germany
simval@amazon.de

Jinmiao Fu*
Amazon
Seattle, USA
jinmiaof@amazon.com

Gianluca Detommaso*
Helsing
Berlin, Germany
Detommaso.gianluca@gmail.com

Shaoyuan Xu
Amazon
Seattle, USA
shaoyux@amazon.com

Giovanni Zappella
Amazon Web Services
Berlin, Germany
zappella@amazon.de

Bryan Wang
Amazon
Seattle, USA
brywan@amazon.com

ABSTRACT

Large language models (LLMs) can be prone to hallucinations – generating unreliable outputs that are unfaithful to their inputs, external facts or internally inconsistent. In this work, we address several challenges for post-hoc hallucination detection in production settings. Our pipeline for hallucination detection entails: first, producing a confidence score representing the likelihood that a generated answer is a hallucination; second, calibrating the score conditional on attributes of the inputs and candidate response; finally, performing detection by thresholding the calibrated score. We benchmark a variety of state-of-the-art scoring methods on different datasets, encompassing question answering, fact checking, and summarization tasks. We employ diverse LLMs to ensure a comprehensive assessment of performance. We show that calibrating individual scoring methods is critical for ensuring risk-aware downstream decision making. Based on findings that no individual score performs best in all situations, we propose a multi-scoring framework, which combines different scores and achieves top performance across all datasets. We further introduce cost-effective multi-scoring, which can match or even outperform more expensive detection methods, while significantly reducing computational overhead.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning; Artificial intelligence.**

KEYWORDS

large language model, uncertainty quantification, text classification, hallucination

ACM Reference Format:

Simon Valentin, Jinmiao Fu, Gianluca Detommaso, Shaoyuan Xu, Giovanni Zappella, and Bryan Wang. 2024. Cost-Effective Hallucination Detection

*These authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD 2024 GenAI Evaluation Workshop, ,

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.




Figure 1: Schematic overview of our proposed hallucination detection approach.

for LLMs. In *Proceedings of (KDD 2024 GenAI Evaluation Workshop)*. ACM, New York, NY, USA, 11 pages.

1 INTRODUCTION

Despite their impressive capabilities, large language models (LLMs) can be prone to generating hallucinations – undesirable outputs that are incorrect, unfaithful, or inconsistent with respect to the inputs (or the output itself) [43]. These unreliable behaviors pose significant risks for adopting LLMs in real-world applications. Challenges in detecting hallucinations lie, among other things, in hallucinations taking different forms, being context-dependent and sometimes being in conflict with other desirable properties of generated text [12, 42]. Hallucinations may be harmless in some contexts, but can be undesired or potentially dangerous in other applications (e.g., erroneous medical advice). Detecting and quantifying hallucination risk is thus a critical capability to enable safe applications of LLMs and improve generated outputs.

Prior work has proposed various approaches for detecting and mitigating hallucinations in LLM-generated outputs, including verifying faithfulness to inputs [26], assessing internal coherence [24], consulting external knowledge sources [32], and quantifying model uncertainty [6, 12, 39, 42]. However, deploying these methods in production settings is far from trivial due to several challenges: First, there is limited comparative evaluation illuminating how different detection methods perform. Second, existing approaches for detecting hallucinations differ greatly in their computational demands, and guidelines are lacking on cost-effectiveness trade-offs to inform method selection for real-world applications with constraints. Third, hallucination detection in the real world often requires careful consideration of risks and false positive/negative

trade-offs, requiring methods to provide well-calibrated probability scores. Fourth, many applications of LLMs take the form of calls to black-box APIs, which sometimes requires methods to employ workarounds to assess the model’s confidence in its generated output.

In this work, we provide a framework for detecting hallucinations in the outputs of any LLM in a model-agnostic manner. Our approach relies on quantifying the probability that a generated answer is a hallucination. After computing initial scores, we employ state-of-the-art calibration techniques to provide calibrated probabilities of the generation containing a hallucination, which can subsequently be used for decision making or other downstream tasks. We evaluate a variety of scores proposed in the literature for hallucination detection on several metrics, across multiple datasets encompassing question answering, fact checking, and summarization tasks. We employ a range of different LLMs to ensure a comprehensive assessment of performance.

Critically, as we show that no single score performs best across all datasets, we introduce *multi-scoring*, a simple way of combining multiple calibrated scores, which achieves superior performance compared to any individual score alone. Furthermore, we propose *cost-effective multi-scoring*, which finds the subset of best-performing scores for any fixed cost budget, and combines them in a multi-scoring fashion. Our empirical demonstrations reveal that cost-effective multi-scoring not only matches but often surpasses the performance of individual scores that incur significantly higher costs. Consequently, our proposed method achieves superior hallucination detection outcomes while maintaining a substantially lower cost footprint.

We summarize our contributions as follows: 1. We benchmark a variety of hallucination detection methods across the literature on several metrics, over different datasets and LLMs. 2. We introduce multi-scoring, a novel approach that aggregates multiple complementary scores and outperforms individual scores. 3. We further propose cost-effective multi-scoring, which optimally balances detection performance and computational constraints.

2 DETECTING LLM HALLUCINATIONS

2.1 Formalizing Hallucination Detection

We study the problem of quantifying the probability that a generated output from a language model contains hallucinations. More formally, let x represent an input token sequence to a language model \mathcal{G} , and let z represent a generated output text sequence from the model. We define a binary random variable $y \in \{0, 1\}$ that indicates whether z is a “permissible” output ($y = 1$) or contains a “hallucination” ($y = 0$).

Our goal is to develop a scoring function to model the probability that a given output text contains a hallucination conditioned on the input.¹ This is critical, as in real-world scenarios, we need to set risk-aware thresholds, balancing false positive/negative rates to accommodate the production requirement. Having access to scores allows us to flexibly set the threshold. Conceptually, the key reason to model this probabilistically is that there is inherent uncertainty

¹Note that while we are interested in whether the entire output contains any hallucination, one could easily adapt the methods at different granularities, such as the sentence or phrase level.

in determining whether a given text contains a hallucination or not. Even human raters may disagree on the assessment, based on their own knowledge and definition. Some key contributors to this epistemic uncertainty are: First, that no system has complete world knowledge to perfectly assess factual correctness. Second, that there is ambiguity in whether something is a hallucination. Finally, any automatic scoring model may make occasional errors, so probabilistic scores reflect confidence. Therefore, while conditional on x and z the true hallucination label y is fixed, our estimate of y remains uncertain. The probabilistic score thus reflects this epistemic uncertainty — the degree of belief that z contains a hallucination given the available knowledge: $p(y = 0 \mid x, z; \psi)$, where ψ denotes parameters of the scoring function. We refer to this conditional probability function as the hallucination score, denoted $s_\psi(x, z)$. The hallucination score can then be applied to downstream tasks, including making risk-aware binary decisions.

We discuss concrete instances of hallucination scores in the following subsection. Generally, the form of s_ψ can vary between hallucination detection approaches. In particular, s_ψ may depend on multiple candidate texts generated for the same input, as implemented by a number of hallucination detection methods proposed in the literature. We denote K candidate generated outputs as $Z = [z_1, \dots, z_K]$. Then $s_\psi(x, Z)$ could quantify inconsistencies within texts in Z using different metrics ψ , as discussed below.

2.2 Scoring Methods

Many hallucination detection methods proposed in the literature make use of LLMs to “judge” the output of an LLM (either the same or a different one). For clarity, we thus distinguish between *generator* and *detector* LLMs. The *generator* is defined as the model used to generate the original response. The *detector* LLM is the model used to score a generated text for the presence of hallucinations. In general, the *generator* and *detector* LLMs may coincide. However, it may be more desirable to use different LLMs in some scenarios, e.g., when computational cost is a greater concern, where a smaller LLM may be used to judge outputs of a more expensive LLM, or when using hallucination scoring methods that require white-/grey-box access, while the generator is black-box.

Today, many interactions with LLMs take the form of API-calls. Typically, only the output tokens are returned with no access to the logits of the predicted tokens, treating the LLM effectively as a black box. Sometimes, inference parameters may be accessible, allowing for setting different temperature (among others) values, thereby providing some (grey-box) access to the model. Generally, hallucination detection methods vary in their required model access, ranging from black-box APIs to full white-box access to the model weights. That is, while some methods only require token-level outputs, other methods may need access to the logits of the generated tokens, or some control over inference parameters like temperature.

In our experiments, we evaluate a comprehensive set of hallucination scoring methods. Generally, we do not make any assumptions about the generator LLM being white-, grey- or black-box. We divide methods into *single-generation* methods, which require only one generated output, and *multi-generation* methods which are based on multiple alternative generations.

Score	Logit Access	NLI model	# LLM calls	# NLI calls	# Generations
P(True)	Yes	No	1	0	0
P(True) Verbalized	No	No	1	0	1
P(InputContradict)	Yes	No	1	0	0
P(SelfContradict)	Yes	No	1	0	0
P(FactContradict)	Yes	No	1	0	0
Inverse Perplexity	Yes	No	1	0	0
NLI (DeBERTa)	No	Yes	0	1	0
SelfCheckGPT-NLI	No	Yes	K	MK	K
HallucinationRail	Yes	No	K	0	K
SimilarityDegree	No	Yes	K	K	K

Table 1: Table summarizing properties of different scoring methods including model access and inference time costs. K denotes the number of multiple generations, M denotes the number of sentences in the response. The column of # Generations denotes the number of LLM calls required besides the LLM call that generates the original response. In this work we set $M = 1$ as we consider the responses as one sentence in our experiment.

2.3 Single-generation

We first provide an overview over a set of hallucination detection methods that are based on scoring the hallucination from a single given generated output.

Inverse Perplexity. This method provides a prominent instance of using a model’s logits over output tokens to score the confidence in the output. Computed as the inverse of the exponentiated average negative log-likelihood of the LLM’s response [11],

$$\text{Perplexity}^{-1}(W) = \exp\left(\frac{1}{N} \sum_{i=1}^N \log p(w_i | w_{i-1}, \dots, w_1)\right), \quad (1)$$

inverse perplexity thus provides a sequence-length normalized expression of the model’s confidence. Here, the model may either be the generator LLM or a different detector LLM. Using a different LLM from the generator as a plug-in estimator implies that the confidence estimate is detached from the generator. If our generator provides access to the model’s logits we can directly use them to score hallucinations. Alternatively, we can use another LLM to generate the logits of output tokens.

P(True). This method works by prompting an LLM whether an answer is correct or not, then making use of the logits of the next token [16]. Given a prompt like the following:

Provide a "True" or "False" response on whether the answer for the following question is correct. Give ONLY a True or False answer, no other words or explanation.

The question is: {x}
 The answer is: {z}
 The answer is:

We compute P(True) by first taking the softmax over the first generated output token’s logits, and then normalizing across the “True” and “False” tokens.² That is, let $\mathbf{z}^{(1)} = (z_1^{(1)}, \dots, z_{|V|}^{(1)})$ be the logit vector over the vocabulary V for the first generated token when

²Note that alternatively, we can prompt the LLM to bind the True/False tokens to symbols like response options A or B [16]. However, this did not change results in our experiments. Further, if the detector LLM provides only black-box access, one can

asking the LLM to evaluate whether the answer is correct. We then take the softmax over these logits, $p(w | \mathbf{z}^{(1)}) = \frac{\exp(z_w^{(1)})}{\sum_{w' \in V} \exp(z_{w'}^{(1)})}$. Let $w_{\text{True}}, w_{\text{False}} \in V$ be the “True” and “False” token ids. Then we can compute P(True) as:

$$P(\text{True} | \mathbf{z}^{(1)}) = \frac{p(w_{\text{True}} | \mathbf{z}^{(1)})}{p(w_{\text{True}} | \mathbf{z}^{(1)}) + p(w_{\text{False}} | \mathbf{z}^{(1)})}. \quad (2)$$

In addition to checking whether the response is correct, we provide variations of P(True), where we check whether the output contradicts the input in *P(InputContradict)*, the output contradicts itself in *P(SelfContradict)*, or the output contradicts generally established facts in *P(FactContradict)*, with the followings prompts:

P(InputContradict): Provide a “True“ or “False“ response on whether the following texts are free of any direct logical or factual contradictions between them. Give ONLY a True or False answer, no other words or explanation.

First text: {x}
 Second text: {z}
 The answer is:

P(SelfContradict): Provide a “True“ or “False“ response on whether the following text is free of internal factual or logical contradictions. Give ONLY a True or False answer, no other words or explanation.

The text is: {z}
 The text is internally consistent:

P(FactContradict): Provide a “True“ or “False“ response on whether the following text contains no contradictions with generally established facts. Give ONLY a True or False answer, no other words or explanation.

Text: {z}
 The text is factually sound:

sample the corresponding tokens and use the empirical distribution to approximate, but thereby incurring a higher computational cost.

NLI Text Classification. Natural language inference (NLI) models provide an alternative to assess the correctness of the model output. As hallucination detection requires checking for contradictions, we compute the score as $1 - S_{\text{contradict}}$, where $S_{\text{contradict}}$ refers to the softmax probability of the output conflicting with the question. Specifically, we use a DeBERTa model fine-tuned on an MNLI task [10] as the underlying NLI model, which we refer to as *NLI (DeBERTa)*.

Verbalized Probabilities. Instead of analyzing a model’s logits, scores are elicited by asking an LLM to provide a confidence *verbatim*, i.e., by generating tokens that indicate the numerical confidence [21, 35] with the following prompt:

Provide the probability between 0.0 and 1.0 that the answer for the following question is correct. Give ONLY the probability value between 0.0 and 1.0, no other words or explanation.

The question is: {x}

The answer is: {z}

Probability the answer is correct:

We note that moderated LLMs accessed through APIs may in practice decline to answer if the context does not provide sufficient information, rather than assigning a low confidence score.

2.4 Multi-generation

Multi-generation methods are based on quantifying the consistency across multiple generated outputs from the generator LLM. This follows the notion of white-/grey-box uncertainty quantification via logits that if the LLM is confident in its response, multiple generated responses will probably be alike and contain compatible facts. Conversely, for fabricated information, sampled responses are more likely to differ and contradict one another. Crucially, this rests on the assumption that the model’s confidence is calibrated, as we discuss and evaluate below.

SelfCheckGPT. This method exploits a DeBERTa NLI model to assess whether the answer A_0 is consistent with K alternative generated answers A_1, \dots, A_K [25]. Each sentence of A_0 is compared against the full set of answers A_1, \dots, A_K . A consistency score per sentence is obtained via an NLI model, then the final score is obtained by averaging over the sentence-wise scores.

Similarity Degree. This method is based on computing pairwise similarities between the multiple responses using NLI models, and then quantifying uncertainty based on the distribution of similarities [22]. Here, we compute the pairwise similarities between responses via the *contradict* class probability of an NLI model and construct a degree matrix, where each diagonal element corresponds to the total (sum) similarity of the corresponding response to all other responses. We use the degree of the candidate response A_0 as the confidence estimate, following favourable results in prior work [22], but note that other metrics have also been proposed in the original paper.

NeMO Guardrails: Hallucination Rail. Unlike SelfCheckGPT or Similarity Degree, the score here is computed via one LLM call,

where we check for agreement between the concatenated K additional generations and the candidate output and not averaged over sentences or computed on pairs of responses. Specifically, we use the softmax probability normalized over yes/no tokens using the following prompt:

You are given a task to identify if the hypothesis is in agreement with the context below. You will only use the contents of the context and not rely on external knowledge. Answer with yes/no.

context: {K additional_sampled_responses}

hypothesis: {candidate_response}

agreement:

2.5 Calibration

Initial hallucination scores may not be properly calibrated, which can lead to poor downstream decisions. Seminal work by Guo et al. [9] has demonstrated that neural models tend to be *miscalibrated*, particularly in the form of models being overconfident. Later work, focusing on language models, has confirmed this across a wide range of tasks, models and datasets [22]. Notably, there is also research which demonstrates that given particular prompts and in-distribution tasks, LLMs can be well-calibrated [16]. However, this has been shown to be brittle and dependent on context [18].

Formally, our score outputs probabilities $p_\psi(\hat{y} = 0 | \mathbf{x}, \mathbf{z})$ that \mathbf{z} contains a hallucination, parameterized by ψ . The score is calibrated if, for any probability level $p \in [0, 1]$, the average observed frequency of hallucinations matches the predicted probability:

$$\mathbb{E}[y \mid p_\psi(\hat{y} = 0 | \mathbf{x}, \mathbf{z}) = p] = p \quad (3)$$

A naive approach for obtaining scores would be to compute the model’s probabilities marginally, ignoring the context/question \mathbf{x} and generated response \mathbf{z} and directly estimating $p(y = 1)$. However, this does not allow for conditioning on individual inputs to obtain calibrated probabilities $p(y = 1 \mid \mathbf{x}, \mathbf{z})$ for specific \mathbf{x} , and \mathbf{z} , which is, however, impossible to guarantee [30]. Common calibration methods include temperature scaling (Platt scaling) of logit outputs [9, 29], isotonic regression [41] or histogram binning [40], which operate marginally, and can thereby not account for different confidence levels for different inputs (e.g., with an LLM being more confident on certain domains than others).

An alternative is to partition the inputs into G groups and compute calibration separately for each group $g \in G$. However, this assumes the groups are disjoint and does not handle inputs belonging to multiple groups, which is often necessary [30]. More advanced calibration methods, such as *multicalibration*, which we use in this work, allow defining G potentially overlapping groups [30]. Prior work has shown the effectiveness of modern calibration techniques in scoring LLM confidence, albeit in a white-box setting [6]. We describe our calibration approach in the experimental section.

2.6 Multi-Scoring: Combining Scores

Different scoring methods capture different aspects of hallucinations, e.g., incorrect, non-factual, non-grounded, irrelevant, inconsistent with other answers, etc. As a result, some scores may work better on some data or for some specific models, and worse on

others. Therefore, we design a *multi-scoring* method to combine the complementary information from individual scores into a single, strong predictor.

Denote each available score by s_n , for $n = 1, \dots, N$. In order to obtain an aggregated score, we run a logistic regression using as features the concatenations of the logit of each score, i.e. $[\text{logit}(s_1(\mathbf{x}, \mathbf{z})), \dots, \text{logit}(s_N(\mathbf{x}, \mathbf{z}))]$, and the labels of the calibration dataset as target variables.³

2.7 Cost-Effective Multi-Scoring

Scoring methods based on multiple generations incur the cost of additional generations as well as the cost of checking their consistency. Similarly, multi-scoring can be a viable choice when there are only few LLM generations to check for hallucinations, but incurs considerable computational cost, especially when using multi-generation methods. To avoid prohibitive computational costs in practice, we propose *cost-effective multi-scoring*, where we set a fixed computational budget and compute the best performing combination of scores that stay within the specified budget.

Given an input text \mathbf{x} and generated text \mathbf{z} , we have N scoring functions $s_1(\mathbf{x}, \mathbf{z}), \dots, s_N(\mathbf{x}, \mathbf{z})$ with associated costs c_1, \dots, c_N (e.g., number of generations required). We are given a total computational budget B . Our goal is to find the optimal subset of scores $S^* \subseteq \{1, \dots, N\}$ that maximizes detection performance while staying within budget B :

$$S^* = \arg \min_{S: S \subseteq \{1, \dots, N\}} \mathcal{L}(f(s_i(\mathbf{x}, \mathbf{z}))_{i \in S}) \quad \text{s.t.} \quad \sum_{i \in S} c_i \leq B \quad (4)$$

where \mathcal{L} measures loss on a validation set. This is a constrained optimization problem over subsets of scoring functions. When $B = \min_i c_i$, it reduces to selecting the single best score at the lowest cost. When $B = \sum_i c_i$, it recovers full multi-scoring. In between, the optimal subset S^* provides the best trade-off between performance and cost. We note that in general, this problem is computationally challenging, given the exponential runtime complexity. However, given that there are only generally a small set of potential scores ($N \sim 10$) and the logistic regression is very fast to fit, computing all combinations takes only 1.8 seconds on a single Intel Xeon processor (3.1 GHz), iterating over all candidate solutions sequentially. When this approach is not feasible, some exemplary alternatives include classic greedy forward-selection methods or regularising the model via an L1 penalty while scaling the regularisation term to accommodate score cost. Overall, cost-effective multi-scoring allows for flexibly balancing multiple scores under computational constraints.

Quantifying Cost. While the cost c_i of scoring function s_i is presented abstractly above, accurately quantifying computational cost can be challenging in practice. The actual runtime per method is a reasonable proxy, however the runtime depends on model architecture, hardware acceleration, batching, etc. One approach is to directly benchmark each scoring function’s average runtime empirically on the target hardware. However, this overlooks nuances like caching effects and ignores runtime variability due to

³Note that we also conducted experiments with alternative models, such as XGBoost [4] and Random Forest [2], but results did not show significant improvements, so we opted for logistic regression for simplicity.

implementational differences. To simplify our analysis, and as LLM calls are generally much more expensive than calling smaller NLI models based on parameter size, we leverage the number of LLM calls required per method as a proxy. See Table 1 for an overview. If more precision is desired, we suggest empirically running different scores and computing their computational cost in the actual application, as the cost will depend on the precise setup and a range of factors.

3 EXPERIMENTS

3.1 Experimental Setup

3.1.1 Datasets.

TriviaQA. TriviaQA [15] is a commonly used factual open-response QA dataset. Originally set up as a reading comprehension task, it is today often used without context as a closed-book (free-recall) task [37]. We use the validation fold, containing 17944 question-answer pairs. We use `mistralai/Mistral-7B-Instruct-v0.2` [13] to generate candidate answers. To decide whether a given response should be labelled as positive or negative, we check whether the correct answer is contained in the generated answer after removing formatting, following the original evaluation script [15].

FEVER. The closed-response Fact Extraction and VERification dataset [34] provides a comprehensive benchmark of factual hallucination detection. We take the test fold, containing 14027 labelled examples of source documents, claims and whether they are supported, refuted or contain not enough info. For the purpose of hallucination detection, we look at all claims that are either supported or refuted.

HaluEval. We include the hallucination detection dataset HaluEval [19] and use the summarization task, which contains 10000 labeled examples of source documents, summaries and labels for whether the summary contains hallucinations.

BIG-bench. It provides an evaluation of LLMs on a diverse set of tasks [33]. We select 11 tasks suitable for hallucination detection, leading a total of 8664 labelled examples from the validation fold.

3.1.2 Metrics. Practical applications of hallucination detection require calibrated scores, but often also binary decisions over whether a given output is hallucinated or not. To this end we report the Brier score [3], binary decision metrics F1 score and accuracy.

3.1.3 LLMs. We conduct experiments with `Mistral-7B-Instruct-v0.2` [13], `Mixtral-8x7B-Instruct-v0.1` [14], `falcon-7b-instruct` [1] and `OpenLLaMA-13b` [5, 8, 36]. All models are used with default configurations via HuggingFace Transformers [38].

3.1.4 Calibration. The calibration step is performed via the following multicalibration approach, using the Fortuna library [7]: To obtain groups, we compute embeddings of the input text \mathbf{x} and the generated response \mathbf{z} , such that our embedding is $\mathbf{e} := [\text{embed}(\mathbf{x}), \text{embed}(\mathbf{z})]$. We obtain embeddings from Universal ANGE Embedding [20], which is the SOTA in the MTEB benchmark [28] at the time of writing. We subsequently reduce the dimension of \mathbf{e} via UMAP [27] and perform soft-clustering via Gaussian Mixture Models, as a simple off-the-shelf algorithm (fitting the number of

cluster components via BIC [31]). The calibration error is measured for each group $g \in G$ separately. Calibration then involves iteratively patching the group with the largest error until the calibration error drops below a threshold for all groups.⁴ This provably converges to a calibrated model with theoretical guarantees [30]. We fit the calibrator to a random *calibration* fold of 80% of the data and report only held-out test results. For all binary predictions, we set the threshold to the 50th percentile to not impose preferences over false true/negative rates, but note that in practice any threshold could be applied.

3.2 Individual Scoring Methods

Table 2 presents the results of different hallucination detection methods on all datasets. Multi-generation methods (SelfCheckGPT-NLI, HallucinationRail and SimilarityDegree) are employed only for TriviaQA, since the latter provides the true answer to each question, which can be compared to the alternative generated answers to assess their agreement. In contrast, for HaluEval, BIG-Bench and FEVER, the true answer is not provided. Instead, we compare the given candidate answer from the dataset against the provided binary label of correctness. If we exclude Multi-Scoring methods, multi-generation methods such as SelfCheckGPT-NLI and SimilarityDegree (with 10 generations) achieve the best performance on TriviaQA, P(True) is the best on HaluEval and BIG-Bench, and P(InputContradict) performs best on FEVER. Thus, we find that there is no single best scoring method across all datasets. This is likely because different scoring methods capture different aspects of hallucination, supporting the notion that hallucination is a multi-faceted concept. Multi-generation methods measure hallucination based on the consistency of the generator LLM’s responses and on the ability of the comparison method to identify whether multiple alternative responses are in agreement. We find that SelfCheckGPT-NLI performs best in our experiments in comparison to SimilarityScore and HallucinationRail. While all of these methods follow similar approaches, there are subtle differences in how they assess the consistency between different responses. More generally, as we discuss below, multi-generation methods are appropriate only if one can assume that there is exactly one correct response, but can fail otherwise.

Other methods are based on different notions of hallucination, such as NLI (DeBERTa) measuring the entailment of the response given the input. Interestingly, variants of P(True) can be appropriate for detecting different kinds of hallucinations. P(True) explicitly asks the evaluator LLM to check whether the response is correct. In some contexts, the evaluator LLM may have the ability to directly evaluate this. However, in other situations, as we see with the FEVER dataset, other methods such as P(InputContradict) can be more appropriate, if we are trying to directly target a specific form of hallucination, which may not be subsumed under a general “correct or not?” prompt. Other applied scenarios may target even more different (or more specific) kinds of hallucinations, though the variants we include in this work are designed to cover the space in a reasonable manner. Similarly, while the datasets included in these experiments were selected to cover a wide range of different kinds of hallucinations, real-world applications may show even new

kinds of hallucinations, for which no public datasets are available. Overall, these findings highlight the need for our multi-scoring method which can absorb the strength of each individual method and can be easily applied to a concrete hallucination detection setting while only requiring a relatively small amount of labeled data.

Overall, neither the inverse perplexity score nor the NLI (DeBERTa) scores emerge as the best scores for any of the datasets we consider. While they may add information that can be exploited in (cost-effective) multi-scoring, as reported below, individually they perform worse than some of the other scores.

3.3 Multi-Scoring

To evaluate the proposed multi-scoring method, we conduct experiments on all datasets combining the scores via logistic regression. As presented in Table 2, the multi-scoring ensemble achieves an F1 score of 0.9106 on TriviaQA, outperforming the best individual F1 score of 0.8614 (achieved by SelfCheckGPT-NLI). Similarly, Brier and Accuracy also show the highest performance for multi-scoring. For HaluEval, we also find that multi-scoring outperforms the best individual score (P(True)) in all metrics. BIG-Bench shows more nuance, as multi-scoring outperforms the best individual score (P(True)) on Brier and Accuracy, but not on F1. This reflects the fact that the metrics capture different properties, and practical considerations may require a decision as to which metric should be prioritized in a given application. For FEVER, multi-score again outperforms the best individual score (P(InputContradict)) in all metrics.

Thus we see that combining scores performs better than the top performing individual scores, showing that combinations of different scoring signals complement each other and can boost performance. This demonstrates that combining complementary signals enables more robust hallucination detection, while balancing the strength of each scoring method. Crucially, while some settings may benefit from combinations of different signals (such as when trying to detect different kinds of hallucinations in a given generated output), in other settings the data-driven selection of an informative signal may be sufficient (such as when trying to detect a more narrowly defined notion of hallucination). These scenarios are covered by our use of multi-scoring, which allows for arriving at optimally combined hallucination scores (or binary decisions) for a given application. Meanwhile, as computational cost can be a concern with scaling LLM applications, it may not be desirable to always use a full ensemble of scores, and we would rather compute the most performant score at a fixed computational cost.

3.4 Cost-effective Multi-Scoring

Shown in Table 2, cost-effective multi-scoring methods are also amongst the top performers, but at considerably lower cost compared to multi-score. We see that at cost $C = 1$, measured as the number of LLM calls (but see our discussion about measuring cost above), we recover the best individual scores, which varies across each dataset. As we increase the budget, the cost-effective multi-scores converge to the performance of the multi-score itself. Notice that in several instances, cost-effective multi-score with $C = 2$ already performs as good as multi-score itself. For TriviaQA, we

⁴Note that alternative approaches are possible [30].

Table 2: Hallucination detection results on all datasets of calibrated scoring methods. Methods that are not applied to a given dataset are marked as —. All scores were computed with Mistral-7B-Instruct-v0.2. TriviaQA responses were generated with Mistral-7B-Instruct-v0.2. Excluding multi-scoring, the best performing scores are underlined. Including multi-scoring, the best performing scores are displayed in boldface.

Scoring Method	TriviaQA			HaluEval			BIG-Bench			FEVER		
	Brier ↓	F1 ↑	Acc ↑	Brier ↓	F1 ↑	Acc ↑	Brier ↓	F1 ↑	Acc ↑	Brier ↓	F1 ↑	Acc ↑
P(True)	0.1819	0.8263	0.7490	<u>0.1980</u>	<u>0.7595</u>	<u>0.6935</u>	<u>0.2066</u>	0.6470	<u>0.6417</u>	0.0729	0.9181	0.9180
P(True) Verbalized	0.1789	0.8143	0.7350	0.2293	0.7040	0.6035	0.2149	0.5252	0.6255	0.0683	0.9197	0.9187
P(InputContradict)	0.1799	0.8150	0.7398	0.2012	0.7552	0.6735	0.2375	0.0000	0.6042	<u>0.0634</u>	<u>0.9328</u>	<u>0.9309</u>
P(SelfContradict)	0.2125	0.8157	0.6888	0.2449	0.6639	0.5620	0.2423	0.1929	0.5413	0.1564	0.8216	0.8001
P(FactContradict)	0.2107	0.8157	0.6888	0.2393	0.6409	0.6000	0.2355	0.2733	0.5949	0.1430	0.8207	0.8190
Inverse Perplexity	0.2033	0.8157	0.6888	0.2490	0.5151	0.5060	0.2289	0.2491	0.6244	0.2353	0.5653	0.5916
NLI (DeBERTa)	0.1924	0.8539	0.7451	0.2417	0.6602	0.5640	0.2322	0.3262	0.6186	0.1673	0.7512	0.7530
SelfCheckGPT-NLI (10)	<u>0.1434</u>	<u>0.8614</u>	<u>0.8011</u>	—	—	—	—	—	—	—	—	—
HallucinationRail (10)	0.1640	0.8539	0.7721	—	—	—	—	—	—	—	—	—
SimilarityDegree (10)	0.1443	0.8585	0.7927	—	—	—	—	—	—	—	—	—
Multi-Score	0.1105	0.9106	0.8593	0.1911	0.7668	0.7075	0.2045	0.5966	0.6590	0.0544	0.9371	0.9351
Cost-Effective (C = 1)	0.1819	0.8263	0.7490	0.1980	0.7595	0.6935	0.2066	0.6470	0.6417	0.0634	0.9328	0.9309
Cost-Effective (C = 2)	0.1772	0.8308	0.7520	0.1911	0.7668	0.7075	0.2045	0.5966	0.6590	0.0636	0.0005	0.9328
Cost-Effective (C = 3)	0.1727	0.8324	0.7537	0.1911	0.7668	0.7075	0.2045	0.5966	0.6590	0.0544	0.9371	0.9351
Cost-Effective (C = 4)	0.1718	0.8277	0.7481	0.1911	0.7668	0.7075	0.2045	0.5966	0.6590	0.0544	0.9371	0.9351
Cost-Effective (C = 5)	0.1485	0.8619	0.8011	0.1911	0.7668	0.7075	0.2045	0.5966	0.6590	0.0544	0.9371	0.9351

see that performance increases on all metrics as we increase the computational budget, with cost-effective multi-score at $C = 5$ outperforming SelfCheckGPT-NLI as the most performant individual scores at half the computational cost in all metrics but Brier. HaluEval shows no improvement beyond a budget of $C = 2$, which is likely due to the kind of hallucination to detect being more narrowly defined and well-captured by a small number of signals, thereby not benefiting from additional scores. For BIG-Bench, as discussed for multi-score, we recover the best individual score at $C = 1$, and match multi-score at higher budgets. FEVER results indicate that, again, we recover the best individual score at budget $C = 1$, and can already recover the full multi-score at a budget of $C = 3$.

We now take a closer look at cost-effective multi-scoring when including multi-generation scoring methods. The overall cost budget is varied over the entire range. For each budget, we solve the constrained optimization in Eq. 4 to find the optimal subset S^* of scores. We compare the hallucination detection F1 score achieved by the cost-effective ensemble versus individual scoring functions and the full ensemble with all scores. Figure 2 shows the results for TriviaQA for responses generated via Mistral-8x7B-v0.1. With a minimal budget of $B = 1$, cost-effective selection recovers the best single method, as expected. As the budget increases, it selectively adds more expensive functions, gradually improving F1, though the gains are marginal at higher budgets per unit cost. At the maximum budget, it recovers the unconstrained full ensemble performance. In between, cost-effective multi-scoring incorporates both less and more and expensive methods to maximize detection within the computational constraints.

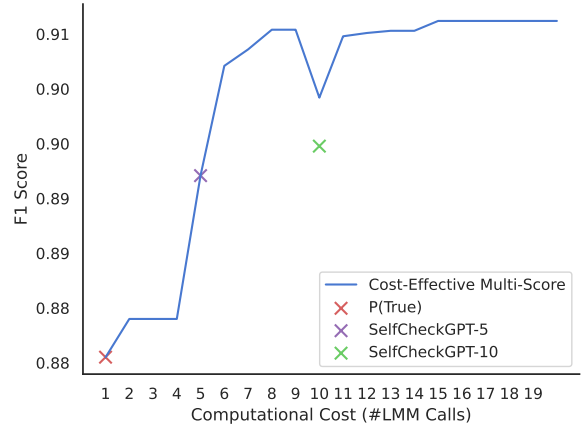


Figure 2: Hallucination detection F1 versus computational budget B for cost-effective multi-scoring.

These results demonstrate that the proposed cost-effective multi-scoring approach can intelligently balance the trade-off between computational expense and hallucination detection effectiveness. It outperforms individual scoring functions and makes selective use of more costly scores to maximize detection performance under a fixed computational budget.

A key consideration in practice would be whether to include potentially more costly multi-generation methods, such as SelfCheckGPT-NLI. In particular as implied by the results presented in Table 2, we may be interested in reducing the number of required multiple generations while matching performance at lower computational costs. To this end, we compare SelfCheckGPT (as the single best performing method for TriviaQA) alone with SelfCheckGPT combined with P(True) at different numbers of additional responses generated via Mixtral-8x7B-v0.1. Here, we count the evaluation of P(True) as one additional LLM call, just as generating one additional output from the generator LLM. As presented in Figure 3, at one additional generation, SelfCheckGPT (in the degenerate case of only one generated output) combined with P(True) thus recovers P(True). As the number of additional generations increases, the combination quickly outperforms SelfCheck alone. In particular, we observe that cost-effective multi-scoring with 3 LLM calls is already as good as SelfCheckGPT with 9 LLM calls in our experiment. This highlights how we can save costs by combining multi-generation methods with other scores while requiring fewer additional generated responses from the generator LLM than if we wanted to achieve the same performance with multi-generation methods alone.




Figure 3: Relationship between number of generations used for SelfCheckGPT and performance of cost-effective multi-score vs SelfCheckGPT alone on TriviaQA.

3.5 Exploration of Relationships between Scores

As discussed above, different scoring methods can target different kinds of hallucinations. To explore this empirically, we here explore their relationships via Spearman rank correlations. As presented in Figure 4, calibrated scores across all different individual scores are positively correlated. Meanwhile, the magnitude of their correlations is smaller than would perhaps be expected if one were to consider hallucinations as a uniform phenomenon. At the same time, we can see certain “clusters” of more strongly inter-correlated scores emerge. For example, multi-generation methods including SelfCheckGPT, HallucinationRail and SimilarityDegree, which check the consistency across multiple generations, show comparably high

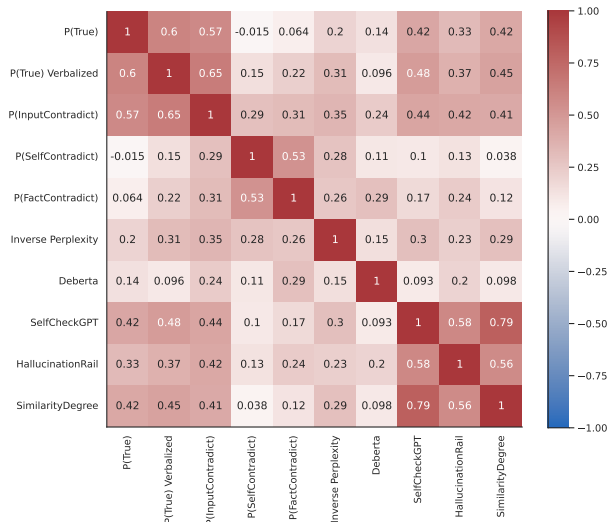


Figure 4: Heatmap of Spearman rank correlations between scores on TriviaQA.

correlations. P(True) and P(True) Verbalised and P(InputContradict) emerge as a similarly correlated cluster of stronger correlations. Overall, this supports the idea that different scores can capture distinct information, and the need to empirically select appropriate (combinations of) scores for a given application, as we propose with cost-effective multi-scoring.

3.6 Experiments across Different LLMs

Most of the scoring methods considered in this work rely on detector LLMs to compute scores for hallucination detection. In Table 4 of Appendix A, thus present additional data for different LLMs, namely Mixtral-8x7B-Instruct-v0.1, falcon-7b-instruct and OpenLaMA-13b. Overall, we see that hallucination detection performance is correlated with performance on general LLM benchmarks⁵. Thus, more overall capable LLMs are likely to perform better at hallucination detection than less capable ones. However, concrete applications may need to take the inference cost of LLMs into account, where one may not always be able to use the most expensive model. We note that cost-effective multi-scoring could here also make use of scores computed via LLMs with different computational costs to find cost-effective combinations.

3.7 The Importance of Calibration

To analyze the impact of calibration, we conduct an ablation study by evaluating model performance with and without calibrating the individual scores. Table 3 in Appendix A indicates that model performance clearly drops in most of the datasets and metrics without

⁵E.g., https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard

calibration. This demonstrates the benefit of calibration for more accurate and risk-aware downstream applications.

3.8 Exploring Multi-Generation Assumptions

User: Generate a recipe for Lunch.

LLM Response 1: Here is a recipe for a tomato and basil soup. First, bring 1L of vegetable broth to a boil...

LLM Response 2: Here is a recipe for a quinoa and avocado salad. First, cook 150g of quinoa according to the package instructions and let it cool...

Figure 5: Example of multi-generation failure-case in NLP systems, illustrating conflicting responses.

Our results indicate that methods based on the uncertainty among multiple generations can provide strong signals when such multiple generations are available. However, such methods can also suffer from problems in practice. Methods based on multiple generations make use of the uncertainty of the *generator*, that is the distribution over generations z given the input x , for model \mathcal{G} , i.e., $p_{\mathcal{G}}(z | x)$ and the consistency between actual sampled generations z . We are interested in classifying whether the candidate output contains a hallucination or not, that is estimating and making use of $p(y | x, z)$. The generating model’s uncertainty $p_{\mathcal{G}}(z | x)$ can be a useful proxy for $p(y | x, z)$, but only under particular conditions.

On a technical level, sampling multiple answers requires access to the generator’s temperature parameter, as temperatures of zero collapse the generations to a single response, as is mentioned, e.g., in [25]. Therefore, these methods are technically grey-box models, as they require *some* access to the model’s inference parameters.

More conceptually, the generator’s uncertainty (or confidence [22]) over $p_{\mathcal{G}}(z | x)$ reflects uncertainty over generated tokens. All recently proposed multi-scoring methods we are aware of are based on the idea of scoring the consistency across multiple generations using model-based metrics. However, self-consistency across multiple generations is neither a necessary nor sufficient criterion for a hallucination to be present. Sufficiency is not given as a model may consistently provide an incorrect response. Self-consistency across multiple generations is also not necessary, as many tasks allow for multiple hallucination-free answers that are contradictory to each other.

As a relatively harmless example, in tasking a model to generate cooking recipes for lunch, the model may generate, among other things, a recipe for a salad and a recipe for a soup. Clearly, the steps in preparing these dishes contain contradictory information, while otherwise being free of hallucinations in themselves. Thus, if there is more than one correct response, the self-consistency assessment in multi-generation methods may falsely score an output as likely to be hallucinated. As there are numerous situations where there exist multiple correct responses, assuming that there is only one could lead to worse LLM responses, also via a loss of diversity.

Practically, methods based on multiple generations can be costly due to added computational overhead, as we have seen above. Finally, these methods are based on the assumption that LLMs are calibrated. As such, multi-generation methods do not allow for detecting cases where the generator LLM is confident yet wrong.

4 CONCLUSION

In this work, we compared a comprehensive set of scoring methods to provide calibrated probability scores for the presence of hallucinations in generated LLM outputs. Overall, we have observed that no single hallucination detection score performs best across all datasets. Our experiments showed that combinations of scores, as suggested with multi-scoring, can effectively combine complementary signals to yield higher hallucination detection performance than any individual score. Further, we demonstrate that cost-effective multi-scoring can find the highest performing scores at a given computational budget.

More generally, our findings support the notion that hallucinations can be rather multi-faceted than present a uniform phenomenon. Thus, detecting hallucinations may require different methods. In concrete settings, one may be interested in even more fine-grained detection of particular types of hallucinations, such as in specific domains like code generation [23].

The approach outlined in this work requires only a small amount of labeled data to calibrate hallucination scores and combine scores via (cost-effective) multi-scoring. However, it is important to acknowledge the limitations of the current work. Future research could explore the effectiveness of this approach on more diverse datasets, investigate alternative scoring methods, or extend the methodology to multi-modal models. While improvements in LLM performance can be expected to also lower their rate of occurrence, hallucinations are unlikely to go away completely [17]. Thus, even as models improve further, detecting hallucinations is likely to remain relevant for applying LLMs in a reliable and trustworthy manner in future.

In summary, our work presents a promising approach for detecting hallucinations in LLM outputs by combining multiple scoring methods in a cost-effective way, offering a pathway towards more reliable and trustworthy language models that can be applied in real-world settings with more confidence.

REFERENCES

- [1] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance. (2023).
- [2] Leo Breiman. 2001. Random forests. *Machine learning* 45 (2001), 5–32.
- [3] Glenn W Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly weather review* 78, 1 (1950), 1–3.
- [4] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [5] Together Computer. 2023. *RedPajama-Data: An Open Source Recipe to Reproduce LLaMA training dataset*. <https://github.com/togethercomputer/RedPajama-Data>
- [6] Gianluca Detommaso, Martin Bertran, Riccardo Fogliato, and Aaron Roth. 2024. Multicalibration for Confidence Scoring in LLMs. *arXiv preprint arXiv:2404.04689* (2024).
- [7] Gianluca Detommaso, Alberto Gasparin, Michele Donini, Matthias Seeger, Andrew Gordon Wilson, and Cedric Archambeau. 2023. Fortuna: A Library for Uncertainty Quantification in Deep Learning. *arXiv preprint arXiv:2302.04019* (2023).

- [8] Xinyang Geng and Hao Liu. 2023. *OpenLLaMA: An Open Reproduction of LLaMA*. https://github.com/openlm-research/open_llama
- [9] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*. PMLR, 1321–1330.
- [10] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. DeBERTa: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654* (2020).
- [11] Fred Jelinek, Robert L Mercer, Lalit R Bahl, and James K Baker. 1977. Perplexity—a measure of the difficulty of speech recognition tasks. *The Journal of the Acoustical Society of America* 62, S1 (1977), S63–S63.
- [12] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *Comput. Surveys* 55, 12 (2023), 1–38.
- [13] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825* (2023).
- [14] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of Experts. *arXiv preprint arXiv:2401.04088* (2024).
- [15] Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551* (2017).
- [16] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221* (2022).
- [17] Adam Tauman Kalai and Santosh S Vempala. 2023. Calibrated Language Models Must Hallucinate. *arXiv preprint arXiv:2311.14648* (2023).
- [18] Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664* (2023).
- [19] Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. Halueval: A large-scale hallucination evaluation benchmark for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 6449–6464.
- [20] Xianming Li and Jing Li. 2023. AngLE-optimized Text Embeddings. *arXiv preprint arXiv:2309.12871* (2023).
- [21] Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334* (2022).
- [22] Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2023. Generating with Confidence: Uncertainty Quantification for Black-box Large Language Models. *arXiv preprint arXiv:2305.19187* (2023).
- [23] Fang Liu, Yang Liu, Lin Shi, Houkun Huang, Ruifeng Wang, Zhen Yang, and Li Zhang. 2024. Exploring and Evaluating Hallucinations in LLM-Powered Code Generation. *arXiv preprint arXiv:2404.00971* (2024).
- [24] Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan. 2021. A token-level reference-free hallucination detection benchmark for free-form text generation. *arXiv preprint arXiv:2104.08704* (2021).
- [25] Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896* (2023).
- [26] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661* (2020).
- [27] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. UMAP: Uniform Manifold Approximation and Projection. *The Journal of Open Source Software* 3, 29 (2018), 861.
- [28] Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. MTEB: Massive Text Embedding Benchmark. *arXiv preprint arXiv:2210.07316* (2022). <https://doi.org/10.48550/ARXIV.2210.07316>
- [29] John Platt et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers* 10, 3 (1999), 61–74.
- [30] Aaron Roth. 2022. Uncertain: Modern topics in uncertainty estimation.
- [31] Gideon Schwarz. 1978. Estimating the dimension of a model. *The annals of statistics* (1978), 461–464.
- [32] Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567* (2021).
- [33] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adria Garriga-Alonso, et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615* (2022).
- [34] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. *arXiv preprint arXiv:1803.05355* (2018).
- [35] Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. *arXiv preprint arXiv:2305.14975* (2023).
- [36] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [37] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652* (2021).
- [38] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771* (2019).
- [39] Yijun Xiao and William Yang Wang. 2021. On hallucination and predictive uncertainty in conditional language generation. *arXiv preprint arXiv:2103.15025* (2021).
- [40] Bianca Zadrozny and Charles Elkan. 2001. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *Icml*, Vol. 1. 609–616.
- [41] Bianca Zadrozny and Charles Elkan. 2002. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. 694–699.
- [42] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lema Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren’s Song in the AI Ocean: A Survey on Hallucination in Large Language Models. *arXiv preprint arXiv:2309.01219* (2023).
- [43] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).

A ADDITIONAL RESULTS

Table 3: Hallucination detection results on all datasets of scoring methods without calibration. Methods that are not applicable for the given dataset are marked as —.

Scoring Method	TriviaQA			HaluEval			BIG-Bench			FEVER		
	Brier	F1	Acc	Brier	F1	Acc	Brier	F1	Acc	Brier	F1	Acc
P(True)	0.2467	0.8267	0.7487	0.3850	0.7195	0.6110	0.3631	0.4981	0.6117	0.0738	0.9241	0.9248
P(True) Verbalized	0.2333	0.8056	0.7269	0.3957	0.6983	0.5645	0.3532	0.5531	0.5245	0.0807	0.9166	0.9163
P(InputContradict)	0.2526	0.8129	0.7392	0.4045	0.7103	0.5910	0.3603	0.4469	0.6001	0.0578	0.9406	0.9398
P(SelfContradict)	0.6438	0.1668	0.3402	0.4825	0.6635	0.5105	0.4238	0.2332	0.5257	0.2101	0.8037	0.7801
P(FactContradict)	0.5571	0.3653	0.4152	0.4634	0.6708	0.5260	0.4076	0.2066	0.5389	0.1952	0.8100	0.7972
Inverse Perplexity	0.5358	0.1286	0.3541	0.0118	0.4980	0.3090	0.0784	0.1771	0.6353	0.4381	0.0055	0.4847
NLI (DeBERTa)	0.3666	0.5914	0.5102	0.4676	0.4910	0.2423	0.0086	0.1943	0.5413	0.2493	0.5987	0.5175
SelfCheckGPT-NLI	0.1503	0.8654	0.8022	—	—	—	—	—	—	—	—	—
HallucinationRail	0.2698	0.8240	0.7086	—	—	—	—	—	—	—	—	—
SimilarityDegree	0.1854	0.8488	0.7612	—	—	—	—	—	—	—	—	—

Table 4: Hallucination detection results on all datasets of scoring methods after calibration via multicalibration. Methods that are not applicable for the given dataset are marked as —.

Scoring Method	TriviaQA			HaluEval			BIG-Bench			FEVER		
	Brier	F1	Acc	Brier	F1	Acc	Brier	F1	Acc	Brier	F1	Acc
Mixtral-8x7B-Instruct-v0.1												
P(True)	0.1543	0.8695	0.8089	0.1959	0.7618	0.6970	0.2081	0.6347	0.6399	0.0769	0.9143	0.9145
P(True) Verbalized	0.1499	0.8617	0.7983	0.2235	0.7158	0.6220	0.2177	0.5907	0.6290	0.0682	0.9221	0.9212
P(InputContradict)	0.1772	0.8493	0.7687	0.2011	0.7547	0.6795	0.2177	0.5907	0.6290	0.0665	0.9288	0.9277
P(SelfContradict)	0.2136	0.8151	0.6879	0.2458	0.6613	0.5550	0.2425	0.0000	0.6145	0.1603	0.8187	0.7979
P(FactContradict)	0.2192	0.8151	0.6879	0.2387	0.6487	0.6030	0.2426	0.2923	0.5586	0.1470	0.8196	0.8190
Inverse Perplexity	0.1894	0.8035	0.7175	0.2497	0.0672	0.5005	0.2455	0.2747	0.6192	0.2355	0.5188	0.5841
OpenLLaMA 13B												
P(True)	0.2133	0.8151	0.6879	0.2502	0.6640	0.5050	0.2495	0.0655	0.6047	0.2512	0.6402	0.5121
P(True) Verbalized	0.2217	0.8151	0.6879	0.2502	0.6707	0.5045	0.2359	0.0620	0.6157	0.2486	0.6814	0.5167
P(InputContradict)	0.2111	0.8150	0.6882	0.2502	0.6707	0.5045	0.2432	0.1730	0.5972	0.2500	0.6358	0.5207
P(SelfContradict)	0.2179	0.8146	0.6874	0.2473	0.6743	0.5300	0.2440	0.4357	0.6099	0.2492	0.6547	0.4986
P(FactContradict)	0.2218	0.8151	0.6879	0.2490	0.5972	0.5070	0.2414	0.2766	0.5926	0.2478	0.5577	0.5381
Inverse Perplexity	0.2191	0.8151	0.6879	0.2466	0.2932	0.5300	0.2400	0.0933	0.6076	0.2404	0.4877	0.5688
Falcon-7B-Instruct												
P(True)	0.2128	0.8147	0.6874	0.2487	0.2794	0.5305	0.2395	0.0030	0.6145	0.2478	0.6198	0.5527
P(True) Verbalized	0.2171	0.8151	0.6879	0.2509	0.6709	0.5055	0.2455	0.0030	0.6151	0.2505	0.6782	0.5153
P(InputContradict)	0.2169	0.8127	0.6857	0.2475	0.3852	0.5515	0.2413	0.0000	0.6145	0.2493	0.5705	0.5203
P(SelfContradict)	0.2138	0.8072	0.6826	0.2493	0.5307	0.5330	0.2347	0.0000	0.6145	0.2364	0.6897	0.6091
P(FactContradict)	0.2135	0.8151	0.6879	0.2265	0.5802	0.6230	0.2364	0.0000	0.6145	0.2358	0.6521	0.6030
Inverse Perplexity	0.2160	0.8151	0.6879	0.2491	0.3705	0.5260	0.2497	0.0809	0.5938	0.2408	0.4853	0.5752