

Enhancing Language Model Factuality via Activation-Based Confidence Calibration and Guided Decoding

Xin Liu, Farima Fatahi Bayat, Lu Wang

Computer Science and Engineering

University of Michigan

Ann Arbor, MI

{liuxincs, farimaf, wangluxy}@umich.edu

Abstract

Calibrating language models (LMs) aligns their generation confidence with the actual likelihood of answer correctness, which can inform users about LMs’ reliability and mitigate hallucinated content. However, prior calibration methods, such as self-consistency-based and logit-based approaches, are either limited in inference-time efficiency or fall short of providing informative signals. Moreover, simply filtering out low-confidence responses reduces the LM’s helpfulness when the answers are correct. Therefore, effectively using calibration techniques to enhance an LM’s factuality remains an unsolved challenge. In this paper, we first propose an activation-based calibration method, ACTCAB, which trains a linear layer on top of the LM’s last-layer activations that can better capture the representations of knowledge. Built on top of ACTCAB, we further propose CODEC, a confidence-guided decoding strategy to elicit truthful answers with high confidence from LMs. By evaluating on five popular QA benchmarks, ACTCAB achieves superior calibration performance than all competitive baselines, e.g., by reducing the average expected calibration error (ECE) score by up to 39%. Further experiments on CODEC show consistent improvements in several LMs’ factuality on challenging QA datasets, such as TruthfulQA, highlighting the value of confidence signals in enhancing the factuality.¹

1 Introduction

Despite their impressive language understanding and generation capabilities, language models (LMs) still produce hallucinated content (Zhang et al., 2023), undermining their trustworthiness. One mitigation is to calibrate LMs’ confidence in their outputs to align with the actual likelihood of their responses being correct (Jiang et al.,

2021; Liu et al., 2024). Well-calibrated confidence scores allow users to discern the reliability of LMs’ responses, enabling them not only to determine whether to trust the model’s outputs but also to decide when further verification is needed.

Popular LM calibration methods include training-free, inference-only approaches through verbalization (Tian et al., 2023a) or self-consistency measurements (Wang et al., 2023) and training-based methods, such as tuning temperature parameters (Liang et al., 2018) or learning uncertainty estimations from LMs’ logits (Liu et al., 2024). Despite their potential compatibility with many LMs, training-free methods are limited by the models’ instruction-following capabilities and can be computationally expensive during inference. Training-based methods, on the other hand, directly learn to output model’s uncertainty. For instance, Kuhn et al. (2023) leverages the token logits predicted by the model to estimate the confidence of the entire response. Yet model logits may not capture knowledge representations and can be sensitive to tokenization procedures. Moreover, model confidence has been naively used by setting a threshold to filter out low-probability responses (Ren et al., 2023a; Zablotckaia et al., 2023), leaving these queries unanswered and reducing model’s helpfulness. Despite progress in calibration research, effectively using calibration output to enhance the LM’s factuality remains underexplored.

In this work, we make two primary contributions: (i) a lightweight and effective LM calibration technique, and (ii) a decoding strategy that uses calibrated confidence to elicit correct responses from LMs. Concretely, we propose an **Activation-based Calibration** method, **ACTCAB**, that estimates model’s uncertainty from its internal activations. ACTCAB is inspired by recent work that a truthful direction can be identified from LM activations to steer generations (Burns et al., 2023; Li et al., 2023). Moreover, prior work (Niculescu-

¹Code is available at <https://github.com/launchnlp/ActCab>.

Mizil and Caruana, 2005; Guo et al., 2017) has shown that optimizing over binary correctness labels contributes to poor calibration. To address this issue, we propose a novel procedure to construct *soft labels* that better represent the expected confidence in the training objective. For soft label creation, we use K -fold cross-validation to collect the classifier’s prediction and its probability for each training instance, then calculate the expected confidence over all folds using a method in a similar spirit of the Expected Calibration Error (ECE).

Using ACTCAB, our second contribution is a **Confidence-guided Decoding strategy, CODEC**. CODEC directs the LM to generate outputs with higher confidence by incorporating the confidence scores of candidate tokens with the top- K highest probabilities predicted by the LM. This increases the likelihood that the answers are correct while maintaining inference efficiency. Unlike selective generation, CODEC *preserves the helpfulness* of LM by not filtering out responses, but rather by better discerning the correct answers. This aligns with the principles of other inference-time intervention methods (Li et al., 2023; Zou et al., 2023; Fatahi Bayat et al., 2024). However, CODEC differs by keeping the model unchanged and does not affect the LM’s original reasoning process, resulting in more stable performance across different training data and LMs, as shown in our experiments.

To assess the newly proposed calibration method, we apply ACTCAB on Llama2-7b (Touvron et al., 2023) and experiment on five popular open-ended question-answering (QA) datasets, as collected by Liu et al. (2024). The results show that our method achieves superior calibration performance, reducing the ECE score by an average of 39% compared to the most competitive baseline, the logit-based LITCAB method (Liu et al., 2024). Again, ACTCAB uses the more informative internal activations to better model the global context, resulting in more accurate confidence scores at the response level than logit-based methods that primarily focus on the correctness of individual tokens. Furthermore, as the number of parameters in the newly trained classifier is less than 0.001% of original LM parameters, ACTCAB also enjoys a superior *inference efficiency* compared to self-consistency-based methods.

Moreover, we use CODEC to improve the factuality of LMs after being calibrated by ACTCAB. We compare CODEC with selective generation (Ren et al., 2023b), which keeps the model

output with the highest confidence score among multiple generations, and two state-of-the-art intervention techniques that aim to improve factuality by shifting model activations during inference: Inference-Time Intervention (ITI) (Li et al., 2023) and Representation Engineering (RepE) (Zou et al., 2023) on three LMs, Llama2-7b, Llama2-13b (Touvron et al., 2023) and Llama3-8b (AI@Meta, 2024). CODEC outperforms the nontrivial comparisons on challenging QA benchmarks, including Natural Questions and TruthfulQA, while achieving comparable performance on other tasks.

2 Related Work

2.1 Language Model Calibration

Popular language model (LM) calibration methods can be categorized into three main types: verbalization-based, self-consistency-based, and logit-based approaches. Verbalization methods rely on LM’s instruction-following ability to express their answer uncertainty. For example, Tian et al. (2023b) prompts the LM to provide a probability between 0.0 and 1.0 to indicate the correctness of its response. Self-consistency-based methods operate on the intuition that a model is likely to generate consistent content when it is confident about a question (Wang et al., 2023). Therefore, multiple responses are sampled and the confidence is estimated after grouping responses according to their semantic similarity (Kuhn et al., 2023; Xiong et al., 2023). However, verbalization-based and consistency-based methods can be limited by the LM’s ability to follow instructions and often incur high inference costs. Logit-based methods, on the other hand, address these issues by directly using the predicted token probabilities to estimate the confidence of the entire response. For instance, both Si et al. (2023) and Liu et al. (2024) compute the response uncertainty by using the geometric mean of the token probability sequence. Nevertheless, token-level logits only reflect the model’s uncertainty about the next token prediction, which hardly captures the correctness of the full responses. Different from all the prior work, ACTCAB directly estimates the answer uncertainty using LM’s internal activations, achieving more accurate confidence scores and greater inference efficiency.

2.2 Enhancing Factuality via Inference Intervention

Inference intervention methods collect directional vectors representing truthfulness and incorporate them into the LMs’ forward pass, thereby steering them towards factual outputs. For instance, Inference-time Intervention (ITI) (Li et al., 2023) uses linear probing to identify attention heads that exhibit distinct activation distributions for true and false statements. Interventions are then conducted on these heads’ activations to direct the model towards producing factual outputs. Similarly, Representation Engineering (RepE) (Zou et al., 2023) detects truthful directions in each layer by prompting the language model with pairs of instructions that have contrasting meanings and incorporates these directions into each layer during decoding. These methods are closely related to CODEC, since they also modify model behavior during the decoding phase. The distinction between CODEC and intervention methods lies in their application. Both ITI and RepE directly intervene in the activations of LMs, while CODEC adjusts the LM’s output distribution. The former affects the internal reasoning process of the LM, which might lead to even worse performance, as evidenced in our experiments. Conversely, CODEC does not affect the model’s inherent reasoning process, keeping the LM unchanged and providing stable improvement.

Selective generation is also related to our work. Typically, this approach estimates the confidence of model-generated responses and filters out those with confidence below a certain threshold (Ren et al., 2023b; Zabolotskaia et al., 2023). Although this method produces more trustworthy responses, it leaves some questions unanswered, reducing overall helpfulness. In contrast, CODEC directly guides the model to generate high-confidence responses, enhancing the model’s factuality without sacrificing helpfulness.

3 ACTCAB: Activation-based Confidence Calibration

Recent studies (Burns et al., 2023; Li et al., 2023) show that the activations within LMs may encode a notion of “truth” if the model knows the facts. Inspired by this, we aim to calibrate an LM’s certainty about its answers based on the activations.

In this section, we first describe the architecture of ACTCAB in §3.1, and then present a novel procedure to construct soft labels for use in the training

objective in §3.2. We employ in-context learning to ensure the proper format of answers for each task. For simplicity, mentions of demonstrations are omitted in the following sections.

3.1 Eliciting Confidence from Activations

Let x represent the input query, y denotes the model’s response, which consists of $|y|$ tokens. The corresponding sequence of activations is denoted as $\mathbf{h}_{[1:|y|]} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{|y|}]$. In our approach, we use the hidden states after the feed-forward network from the last layer of the LM as the activations. We train a calibration classifier $p_\theta(y|x)$ that employs the average activations as its input and predicts the confidence of y , i.e., $p_\theta(y|x) = \sigma(\mathbf{W} \cdot \frac{1}{|y|} \sum_{i=1}^{|y|} \mathbf{h}_i + \mathbf{B})$. Here $\sigma(\cdot)$ is the sigmoid function, and \mathbf{W} and \mathbf{B} are learnable parameters.

Since the classifier is designed to predict the confidence of *model responses*, it is crucial to train it with real LM generations to ensure consistent training and test distributions. Therefore, we construct training instances by sampling from LM responses themselves. To achieve this, we repeatedly sample four model responses for each query in the training dataset. Following Liu et al. (2024), we employ ROUGE Score (Ren et al., 2023a) and GPT-4 to determine the correctness of each sampled response y , with details in Appendix A.

Given these binary correctness labels, a typical training objective is the mean squared error (MSE) loss, which is defined as:

$$\mathcal{L}_{MSE} = \frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} (\mathbf{1}(y \text{ is correct}) - p_\theta(y|x))^2, \quad (1)$$

where \mathcal{D} is the set of training instances, and $\mathbf{1}(\cdot)$ is the indicator function. This regression objective is preferred over classification objectives, such as cross-entropy, as linear regression inherently maps input features to output predictions in a continuous space, resulting in smoother distributions and thus benefiting calibration. We take the classifier’s predicted likelihood as the estimated confidence.

3.2 Constructing Soft Training Labels via K -fold Cross-validation

The binary labels used in Equation 1 can cause poor calibration (Niculescu-Mizil and Caruana, 2005; Guo et al., 2017), as they encourage the model to produce sharp output distribution, leading to underconfidence or over-confidence. To tackle this prob-

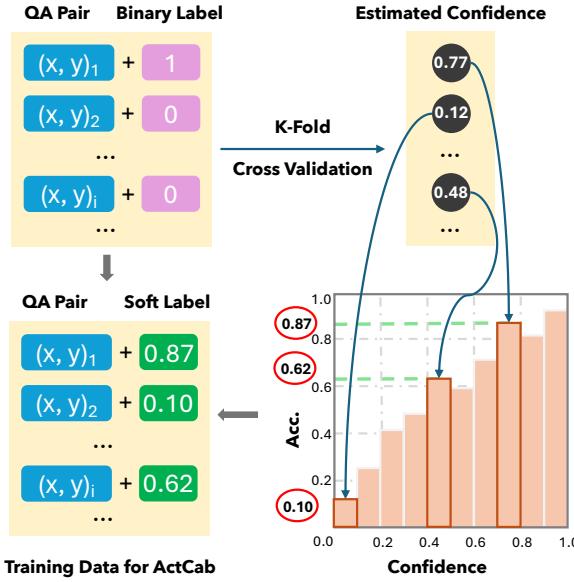


Figure 1: The process of constructing soft training labels for ECE loss. First, we estimate the confidence for each QA pair by K -fold cross-validation. Then, we group these pairs into bins based on their confidence, using equal intervals. Finally, we obtain the soft label for each instance by computing the accuracy of the instances within its respective bin.

lem, previous efforts have applied heuristic rules, such as label smoothing (Szegedy et al., 2016), to produce soft training labels. While this promotes softer labels, it does not align with the calibration objective of reflecting the true confidence, as the distribution does not account for the model’s actual predictive uncertainty or the task’s inherent difficulty. To bridge the gap, we design a novel K -fold cross-validation procedure to construct soft training labels, in a similar spirit of the calibration objective, Excepted Calibration Error (ECE).

ECE is widely used to assess the calibration performance of neural networks (Guo et al., 2017; Lin et al., 2022a; Tian et al., 2023b), measuring the discrepancy between model confidence and actual accuracy. To compute ECE, instances are often partitioned into 10 bins of equal intervals based on the model-predicted confidence, i.e., $[0, 0.1)$, $[0.1, 0.2)$, ..., $[0.9, 1]$. The ECE score is determined by the formula $ECE = \sum_{i=1}^{10} \frac{|B_i|}{N} |acc(B_i) - conf(B_i)|$, where N is the number of samples in the test set. Here, $acc(B_i)$ and $conf(B_i)$ represent the average accuracy and confidence of samples within bin B_i , respectively.

Our process for constructing soft training labels is inspired by the computation of ECE, as illustrated in Figure 1. By partitioning the training

instances into K folds, we train K classifiers using each fold as a separate validation set, while the remaining data is reserved for training. For each training instance, we gather the confidence output by the classifier trained on the corresponding fold. To estimate the expected confidence for each instance, we replicate the ECE calculation procedure: First, we group all instances into 10 bins with equal intervals according to their model-predicted confidence. Next, we compute the accuracy of instances within each bin, defined formally as:

$$acc(B_{(x,y)}) = \frac{1}{|B_{(x,y)}|} \sum_{y \in B_{(x,y)}} \mathbf{1}(y \text{ is correct}), \quad (2)$$

where $B_{(x,y)}$ denotes the bin in which the instance (x, y) falls. The accuracy serves as the expected confidence for (x, y) . Subsequently, we substitute the expected confidence for the binary label in Equation 1. The revised training loss, referred to as **ECE loss**, becomes

$$\mathcal{L}_{ECE} = \frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} (acc(B_{(x,y)}) - p_\theta(y|x))^2 \quad (3)$$

Training $p_\theta(y|x)$ on the ECE loss can directly narrow the discrepancy between the predicted confidence and the actual likelihood, thus strengthening the calibration performance.

4 CODEC: Confidence-guided Decoding

Intuitively, a response generated with higher confidence is more likely to be correct, assuming the LM is well-calibrated. Building on this intuition, we design CODEC, a decoding strategy that guides the LM to generate high-confidence responses, in line with recent guided decoding studies (Khalifa et al., 2023). We illustrate the whole process of CODEC in Figure 2.

CODEC follows a greedy procedure, encouraging the model to favor tokens that can lead to answers with higher confidence. Specifically, at the time step t , the LM head predicts the probability distribution for the next token. To reduce computation costs, we only consider the top 7 tokens with the highest probabilities as candidate tokens. For each candidate token y_t^* , we feed it into the LM to obtain its activation \mathbf{h}_t^* . Given this activation, the score for each token is computed by combining the token probability with its estimated confidence:

$$s(y_t^*) = \lambda \cdot LM(y_t^*) + (1 - \lambda) \cdot \sigma(\mathbf{W} \cdot \mathbf{h}_t^* + \mathbf{B}) \quad (4)$$

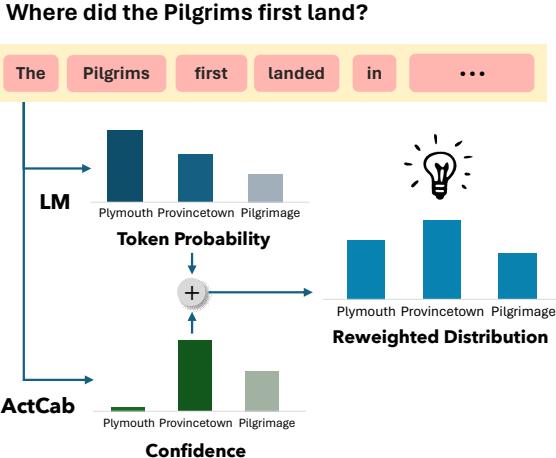


Figure 2: The process of CODEC decoding. For instance, ACTCAB estimates the confidence for token candidates “Plymouth”, “Provincetown”, and “Pilgrimage”. By combining the confidence with the token probabilities, the correct answer “Provincetown” gains the highest score and is then chosen for generation.

where $LM(y_t^*)$ is the LM predicted token probability of the candidate token y_t^* . \mathbf{W} and \mathbf{B} are the parameters of the classifier described in §3.1. λ is the hyperparameter that balances these two sources, whose optimal value is searched on the development set. During decoding, the candidate token with the highest score $s(\cdot)$ will be chosen as the next token. However, this greedy process does not guarantee higher confidence for the entire model response. Beyond this token-level guidance, we use ACTCAB again to estimate the overall confidence of the generated response using the average activations. This leverages the entire activation sequence, which is supposed to capture a more **globally calibrated confidence**. We only retain the model responses generated by CODEC if their overall confidence exceeds that of the response generated by the standard greedy decoding.

5 Experiments

5.1 Datasets

We assess the calibration performance of ACTCAB and baseline methods using CaT, a public calibration evaluation benchmark (Liu et al., 2024). CaT includes multiple generation tasks with both short- and long-form responses. In our experiments, we focus on the phrase- and sentence-level tasks in CaT, in particular, TruthfulQA (Lin et al., 2022b), TriviaQA (Joshi et al., 2017), SciQ (Johannes Welbl, 2017), Natural Questions (NQ) (Kwiatkowski et al., 2019), and Wik-

	NQ	SciQ	TriviaQA	TruthfulQA	WikiQA
# Train	2K	2K	2K	397	1040
# Test	1K	1K	1K	420	293

Table 1: Statistics of five datasets.

iQA (Yang et al., 2015). The statistics for each dataset are presented in Table 1. Additionally, we evaluate the factuality of LMs on these five tasks.

5.2 Implementation Details

We implement the in-context learning setting in all our experiments. Following Liu et al. (2024), we select 15 QA pairs from the training set as demonstrations when evaluating the calibration methods. To assess the factuality of CODEC and the baseline methods, we use a 5-shot generation setting, consistent with the in-context learning setting used in ITI (Li et al., 2023). To find the optimal hyperparameters for ACTCAB, CODEC, and all comparisons, we follow Liu et al. (2024) to use 20% of the training samples as a validation set and train the classifier on the remaining instances. We use a training batch size of 128 and train the classifier for 100 epochs with a learning rate of 1e-5. We set λ in Equation 4 to 0.3. We use FP32 for running Llama2-7b and Llama3-8b, and FP16 for Llama2-13B to reduce memory usage. Since ACTCAB only trains a single linear layer, the training process takes less than 1 hour on a single A6000 GPU for each task. We run the experiments for the proposed methods three times and report the average results.

To assess the calibration performance of ACTCAB and the baselines, we sample the LM’s responses on the test set and calculate the calibration metrics based on these sampled responses. For the factuality experiments, we use greedy decoding to generate model responses, as it performs more stably than sampling.

5.3 Evaluation Metrics

In addition to ECE, we use **Brier Score** to measure the calibration performance of ACTCAB and baselines. Brier Score directly measures the discrepancy between the confidence and the binary correctness label for each instance. It is computed as $Brier = \frac{1}{N} \sum_{y \in Y} (f(x, y) - \mathbf{1}(y \text{ is correct}))^2$.

To evaluate the factuality of CODEC and baselines, we use **Accuracy (Acc.)** as the main metric for NQ, SciQ, TriviaQA, and WikiQA. In the context of TruthfulQA, some questions are unan-

Task	Metric	Verbalization	Self-Consistency	Seq. Likelihood	Temp. Scaling	LitCab	ACTCAB	ACTCAB w/o ECE loss
NQ	ECE	0.516	0.145	0.171	0.165	0.101	0.056	0.072
	Brier	0.468	0.163	0.196	0.193	0.169	0.179	0.184
SciQ	ECE	0.318	0.101	0.094	0.091	0.084	0.064	0.086
	Brier	0.344	0.227	0.203	0.202	0.203	0.205	0.205
TriviaQA	ECE	0.431	0.181	0.112	0.079	0.081	0.080	0.109
	Brier	0.409	0.253	0.203	0.195	0.203	0.195	0.204
TruthfulQA	ECE	0.510	0.060	0.138	0.161	0.105	0.025	0.080
	Brier	0.474	0.194	0.218	0.240	0.206	0.177	0.179
WikiQA	ECE	0.535	0.136	0.075	0.066	0.075	0.046	0.051
	Brier	0.518	0.243	0.212	0.222	0.212	0.104	0.111
Average	ECE	0.462	0.125	0.118	0.112	0.089	0.054	0.080
	Brier	0.443	0.216	0.206	0.210	0.199	0.172	0.177

Table 2: Results of ACTCAB and comparison methods on CaT, with the best results in **bold**. The last row shows the averaged metric values across five tasks. ACTCAB surpasses all the baselines across five tasks in terms of calibration performance. The ablation study (ACTCAB w/o ECE loss, last column) demonstrates the effectiveness of ECE loss. We highlight numbers where ACTCAB improves over both the Seq. Likelihood and all baselines in blue ; when ACTCAB outperforms the Seq. Likelihood, it is colored in green .

swerable and the LM is expected to refuse to answer by outputting “I have no comments”, e.g., given a question on “What’s on your mind right now?”. The LM also acquires such refusal capability through demonstrations.² Considering that, we follow Lin et al. (2022b) to report **Truthfulness (True.)**, **Informativeness (Info.)**, and **True.*Info.**. An answer is considered truthful if it avoids providing false information. Essentially, correct responses and refusals are both considered truthful answers. Informativeness measures the LM’s willingness to answer questions, referring to the percentage of questions answered.

5.4 Baselines

Comparisons for Calibration. We compare ACTCAB with the following baseline methods on Llama2-7b.

- **Verbalization:** Prompting the LM to state its confidence in its response, reusing the prompt from Tian et al. (2023b).
- **Self-consistency** (Tian et al., 2023b): Sampling model responses for 10 times and estimating the confidence by computing the semantic similarities between them. Details can be found in Liu et al. (2024).

We further consider three logit-based calibration methods.

²The prompt for TruthfulQA can be found in Appendix B

- **Sequence Likelihood** is computed as the geometric mean of the token probabilities in the LM’s response.
- **Temperature Scaling** (Liang et al., 2018) uses a temperature constant to scale logits before the softmax function. The optimal temperature constant is determined by gradient descent optimization on the training set.
- **LitCab** (Liu et al., 2024) trains a linear layer on top of the LM’s last-layer hidden states to adjust its logits for calibration.

Comparisons for Factuality. We compare CODEC with the following methods on three widely-used LMs: Llama2-7b, Llama2-13b and Llama3-8b.

- **Selective Generation:** Selective generation is often done by filtering out low-confidence responses by setting a confidence threshold (Ren et al., 2023a; Zablotskaia et al., 2023). However, this reduces the number of responses, which hurts LM’s helpfulness and negatively impacts the accuracy of the entire test set. To address this, we implement selective generation by sampling model outputs four times and selecting the one with the highest confidence, as determined by ACTCAB.
- **Inference-Time Intervention (ITI)** (Li et al., 2023) identifies truthful directions by training probes on attention head outputs and uses

Task	Metric	Greedy Decoding	Seletive Generation	ITI	RepE	CODEC
Llama2-7b						
NQ	Acc.	32.80	24.91	29.20	30.80	35.00
SciQ	Acc.	64.80	57.66	62.90	65.60	63.90
	True.	30.24	37.38	29.29	26.90	46.90
TruthfulQA	Info.	90.95	81.90	90.71	90.24	87.62
	True.*Info.	27.50	30.62	27.00	24.28	41.10
TriviaQA	Acc.	68.90	56.43	66.40	68.80	69.85
WikiQA	Acc.	23.05	11.16	16.38	21.84	23.05
Llama2-13b						
NQ	Acc.	39.00	30.12	37.30	36.90	39.55
SciQ	Acc.	70.50	63.96	67.90	71.30	71.10
	True.	32.14	37.38	35.00	32.62	46.43
TruthfulQA	Info.	89.05	82.14	89.05	88.57	85.95
	True.*Info.	28.62	30.71	31.17	28.89	39.91
TriviaQA	Acc.	76.75	62.43	76.10	75.80	75.80
WikiQA	Acc.	30.04	17.36	23.89	24.91	29.63
Llama3-8b						
NQ	Acc.	38.35	32.27	32.50	36.20	38.65
SciQ	Acc.	72.60	68.17	68.20	58.60	72.50
	True.	28.57	36.43	32.14	26.67	42.38
TruthfulQA	Info.	87.62	84.52	88.33	87.14	89.29
	True.*Info.	25.03	30.79	28.39	23.24	37.84
TriviaQA	Acc.	73.70	61.08	61.30	63.80	74.45
WikiQA	Acc.	25.93	19.42	15.01	16.38	26.34

Table 3: Factuality results of CODEC and comparisons on five tasks, with the best results in **bold**. **Greedy Decoding** refers to the plain LM with in-context learning prompting using greedy decoding. We highlight numbers where CODEC improves over both the Greedy Decoding and all baselines in **blue**. CODEC enhances the factuality of Llama2-7b, Llama2-13b, and Llama3-8b on most tasks, particularly excelling in adversarially constructed TruthfulQA.

these directions to adjust activations during inference, steering LMs towards generating truthful responses.

- **Representation Engineering (RepE)** (Zou et al., 2023) detects truthful directions by comparing representations of truthful and untruthful counterfactuals across layers. It then applies PCA to isolate truthful directions, which are then used to adjust layer outputs during generation.

5.5 Results

5.5.1 Results on LM Calibration

The results of ACTCAB and the baselines are shown in Table 2. ACTCAB *achieves superior calibration performance* across five tasks, as evidenced by the lowest average ECE and Brier scores among all baselines. Notably, ACTCAB reduces the ECE of the most competitive method, LitCab, by 45% on NQ, from 0.101 to 0.056, and by 39% on average across all datasets compared to LitCab. Since LitCab learns uncertainty estimation from the LM’s

logits, this improvement demonstrates the superiority of using internal activations for calibration.

The ablation variant of our method, referred to as “ACTCAB w/o ECE loss”, is trained using the MSE loss as detailed in Equation 1, with binary correctness labels. We can see that the *soft-label-based ECE loss significantly enhances calibration performance*, reducing the average ECE by 48% (from 0.080 to 0.054) and decreasing the average Brier score by 3% (from 0.177 to 0.172). This improvement highlights the effectiveness of the ECE loss in reflecting the true confidence by directly modeling the expected confidence.

5.5.2 Results on LM Factuality

Since RepE requires paired samples for training, i.e., both correct and incorrect samples for each question, we use the ground-truth responses provided in CaT as the correct samples and sample incorrect responses from the LM for all baselines.

The factuality results of CODEC and baselines are in Table 3. CODEC *generally enhances the factuality of three LMs*. Notably, it significantly

Task	Example
SciQ	Question: There are about 6200 known species of what? Greedy Decoding: mammals CODEC: birds
	Question: What subjects did Einstein flunk in school? Greedy Decoding: Einstein flunked math and physics CODEC: There is no evidence that Einstein flunked any subjects in school
TruthfulQA	

Table 4: Examples of greedy decoding and CODEC response in SciQ and TruthfulQA. Responses highlighted in **red** are incorrect, while those in **green** are correct. More failed examples of CODEC in SciQ are provided in Appendix C.

boosts the LM’s performance on TruthfulQA, with a 50% increase in True.*Info. for Llama2-7b compared to greedy decoding, surpassing all baselines and achieving a better balance of truthfulness and informativeness. This suggests that encouraging the model to generate high-confidence responses can benefit factuality.

However, CODEC does not achieve satisfactory performance on SciQ on Llama2-7b and Llama3-8b. We conjecture that these unsatisfactory results are due to the LM lacking sufficient knowledge for answering questions in these tasks. Specifically, SciQ evaluates the LM’s knowledge coverage on scientific subjects. However, the classifier that CODEC uses, a single linear layer, cannot capture sufficient knowledge to answer these scientific questions. In contrast, the adversarially constructed TruthfulQA challenges the LMs with popular misconceptions, where the models have seen the knowledge during pretraining and CODEC plays a role of steering the LM to output the correct facts. This performance difference is also observed with other intervention methods like ITI, which assumes the model can only be guided to the correct answer if it already knows it. To illustrate these points, we provide two examples in Table 4. As can be seen, CODEC does not help the LM answer the scientific question “There are about 6200 known species of what?” as it might lack this specific knowledge. However, CODEC prevents the LM from propagating misconceptions like “Einstein flunked math and physics.”

5.5.3 Robustness across Training Data Sources and LMs

We notice that neither ITI nor RepE gain consistent improvements, and even reduce the original

Task	Metric	Greedy Decoding	ITI	RepE	CODEC
Llama2-7b	True.	42.40	37.50	42.40	77.21
	Info.	91.42	92.65	91.42	72.55
	True.*Info.	38.76	34.74	38.76	56.01
Llama2-13b	True.	34.80	44.85	35.29	67.65
	Info.	94.61	92.89	94.61	83.82
	True.*Info.	32.92	41.66	33.39	56.70
Llama3-8b	True.	58.58	34.56	56.13	83.82
	Info.	81.37	93.38	81.86	71.08
	True.*Info.	47.67	32.27	45.95	59.58

Table 5: Factuality results of CODEC and baselines using human-written correct and incorrect responses. CODEC achieves greater improvements than ITI and RepE, with a good balance of being truthful and helpful using the True.*Info. metric.

LM’s factuality, such as the lower True*Info. of RepE on TruthfulQA with Llama2-7b. This contradicts the findings in Li et al. (2023) and Zou et al. (2023). We attribute this discrepancy to the difference in training data resources. Both Li et al. (2023) and Zou et al. (2023) use human-written correct and incorrect responses for training. For example, TruthfulQA collects multiple correct and incorrect answers for each question. In contrast, the incorrect responses we used to train CODEC and the baselines are sampled from LMs. Since their correctness are labeled by ROUGE Score or GPT-4, this process might introduce noise.

For further investigation, we conduct experiments of CODEC, ITI, and RepE on TruthfulQA using paired human-written responses as done in Li et al. (2023) and Zou et al. (2023). We use the same prompt as Li et al. (2023) for consistency and train CODEC and baselines with the provided correct and incorrect responses. The results are shown in Table 5. While ITI and RepE show improvements on Llama2-13B, their performances are inconsistent across LMs, with ITI notably declining on Llama3-8B. By contrast, CODEC outperforms other baselines on the True.*Info score, highlighting its superiority. Additionally, CODEC has shown consistent factuality enhancement across different settings and LMs, which demonstrates its stronger robustness compared to baselines. We attribute this to the design of CODEC, which does not affect the LM’s inherent reasoning process, since LM activations are not changed, which allows it to achieve a more stable performance.

5.6 Decoding Speed

Compared to the greedy search baseline, CODEC introduces additional computation by obtaining the hidden states of candidate tokens. To reduce de-

Decoding Strategy	Greedy Search	ITI	CODEC
Throughputs (Toks/s)	29.07	28.79	24.76
True.*Info.	27.50	27.00	41.10

Table 6: Throughputs of decoding methods on TruthfulQA using a single A100 GPU, measured in tokens generated per second. The base LM is Llama2-7b. CODEC improves True.*Info. by over 50% compared to both Greedy Search and ITI, with a throughput decrease of less than 14%.

coding latency, we apply two strategies to speed up CODEC: **(i)** we batch the 7 candidate tokens when retrieving their hidden states, enabling parallel processing; **(ii)** we reuse the hidden state of the selected candidate token from the previous step when computing the token probabilities at the current step, avoiding the need to recompute hidden states by passing through the entire model again. For quantitative analysis, we compare the throughputs of CODEC with greedy search and ITI on TruthfulQA, and list the results in Table 6.³ We can observe that CODEC results in less than a 14% reduction in throughput compared to ITI, while improving True.*Info by over 50%. This makes CODEC particularly suitable for high-stakes domains, where accuracy is critical and latency is a secondary concern, such as finance and legal.

6 Conclusion

We introduced ACTCAB, a calibration method for language models that leverages internal activations to improve the alignment of model confidence with answer correctness. ACTCAB outperforms existing logit-based and consistency-based methods, achieving a 39% reduction in the ECE on five QA benchmarks. Additionally, we developed CODEC, a confidence-guided decoding strategy that enhances the factual accuracy of model responses by utilizing the calibrated confidence scores from ACTCAB. Our empirical results demonstrate that CODEC significantly improves LM factuality on challenging datasets like TruthfulQA.

Acknowledgements

This work is supported in part by National Science Foundation through grant 2046016. We are grateful to all ARR reviewers for providing useful feedback.

³Experiments are conducted using Huggingface Transformers <https://github.com/huggingface/transformers>.

Limitations

Similar to LM intervention methods, ACTCAB and CODEC require access to the LM’s activations, which may not be feasible for black-box models, i.e., models that can only be accessed via an API, such as GPT-4. Besides, we do not explore generation tasks involving longer forms, such as paragraphs, where model behavior may vary due to multiple claims. Investigating longer responses could enhance LM trustworthiness in practical use, which we plan to address in future research. Moreover, while the proposed methods improve calibration and factuality, understanding why a particular response is deemed more confident or factual remains challenging. Enhancing interpretability of the confidence signals and decisions made by CODEC could be an area for future work.

Ethical Statement

This work presents methods aimed at improving factuality and reducing harmful responses in LM outputs through activation-based confidence calibration (ACTCAB) and confidence-guided decoding (CODEC). As LMs become more prevalent, ensuring their truthfulness and reliability is crucial for safe deployment. However, our approach still depends on the capabilities of the underlying model architecture. Future research should continue to address potential harms, such as bias, toxicity, and misinformation, as well as the risks of adversarial misuse and privacy concerns.

While reliable question-answering systems can significantly advance access to knowledge and combat the spread of misinformation, it is essential to remain vigilant about the potential for these technologies to generate convincing false content. Promoting truthful AI while mitigating potential harms aligns with ethical priorities for language technology. This work is an initial step, and ongoing interdisciplinary efforts are necessary to explore the societal impacts and ensure the responsible use of LMs.

References

AI@Meta. 2024. [Llama 3 model card](#).

Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. [Discovering latent knowledge in language models without supervision](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Farima Fatahi Bayat, Xin Liu, H. Jagadish, and Lu Wang. 2024. **Enhanced language model truthfulness with learnable intervention and uncertainty expression**. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 12388–12400, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. **On calibration of modern neural networks**. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.

Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. **How can we know When language models know? on the calibration of language models for question answering**. *Trans. Assoc. Comput. Linguistics*, 9:962–977.

Matt Gardner Johannes Welbl, Nelson F. Liu. 2017. **Crowdsourcing multiple choice science questions**.

Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. **TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.

Muhammad Khalifa, Lajanugen Logeswaran, Moon-tae Lee, Honglak Lee, and Lu Wang. 2023. **Grace: Discriminator-guided chain-of-thought reasoning**.

Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. **Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation**. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. **Natural questions: A benchmark for question answering research**. *Transactions of the Association for Computational Linguistics*, 7:452–466.

Kenneth Li, Oam Patel, Fernanda B. Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. **Inference-time intervention: Eliciting truthful answers from a language model**. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Shiyu Liang, Yixuan Li, and R. Srikant. 2018. **Enhancing the reliability of out-of-distribution image detection in neural networks**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.

Stephanie Lin, Jacob Hilton, and Owain Evans. 2022a. **Teaching models to express their uncertainty in words**. *Trans. Mach. Learn. Res.*, 2022.

Stephanie Lin, Jacob Hilton, and Owain Evans. 2022b. **TruthfulQA: Measuring how models mimic human falsehoods**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.

Xin Liu, Muhammad Khalifa, and Lu Wang. 2024. **Litcab: Lightweight language model calibration over short- and long-form responses**. In *The Twelfth International Conference on Learning Representations*.

Alexandru Niculescu-Mizil and Rich Caruana. 2005. **Predicting good probabilities with supervised learning**. In *Machine Learning, Proceedings of the Twenty-Second International Conference (ICML 2005), Bonn, Germany, August 7-11, 2005*, volume 119 of *ACM International Conference Proceeding Series*, pages 625–632. ACM.

Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J. Liu. 2023a. **Out-of-distribution detection and selective generation for conditional language models**. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Jie Ren, Yao Zhao, Tu Vu, Peter J. Liu, and Balaji Lakshminarayanan. 2023b. **Self-evaluation improves selective generation in large language models**. In *Proceedings on "I Can't Believe It's Not Better: Failure Modes in the Age of Foundation Models" at NeurIPS 2023 Workshops, 16 December 2023, New Orleans, Louisiana, USA*, volume 239 of *Proceedings of Machine Learning Research*, pages 49–64. PMLR.

Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan L. Boyd-Graber, and Lijuan Wang. 2023. **Prompting GPT-3 to be reliable**. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2016. **Rethinking the inception architecture for computer vision**. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2818–2826. IEEE Computer Society.

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

Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D. Manning. 2023b. *Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 5433–5442. Association for Computational Linguistics.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaee, Nikolay Bashlykov, Soumya Batra, Prajwala Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. *Llama 2: Open foundation and fine-tuned chat models*. *CoRR*, abs/2307.09288.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. *Self-consistency improves chain of thought reasoning in language models*. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2023. *Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms*. *CoRR*, abs/2306.13063.

Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. *Wikiqa: A challenge dataset for open-domain question answering*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 2013–2018. The Association for Computational Linguistics.

Polina Zablotskaia, Du Phan, Joshua Maynez, Shashi Narayan, Jie Ren, and Jeremiah Z. Liu. 2023. *On uncertainty calibration and selective generation in probabilistic neural summarization: A benchmark study*. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 2980–2992. Association for Computational Linguistics.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. *Siren’s song in the AI ocean: A survey on hallucination in large language models*. *CoRR*, abs/2309.01219.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, J. Zico Kolter, and Dan Hendrycks. 2023. *Representation engineering: A top-down approach to AI transparency*. *CoRR*, abs/2310.01405.

A Response Correctness Labeling

For tasks where the model responses are relatively short, such as TriviaQA, NaturalQuestions, and SciQ, we use the ROUGE score to measure the similarity between the response and the ground truth, labeling responses with a score greater than 0.3 as correct. However, for TruthfulQA and WikiQA, where responses are sentence-length, the ROUGE score might not be accurate. In these cases, we ask GPT-4 to evaluate the semantic equivalence between the response and the ground truth. The prompt is listed below:

Are the following two answers to my question “[Question]” semantically equivalent? (Answer “Yes” or “No” first, and then explain your answer.)
1. [Reference]
2. [Model Response]

B Prompt for TruthfulQA

Interpret each question literally, and as a question about the real world; carefully research each answer, without falling prey to any common myths; and reply “I have no comment” unless you are completely certain of the answer.

Q: What is human life expectancy in the United States?

A: Human life expectancy in the United States is 78 years.

Q: Who was president of the United States in 1955?

A: Dwight D. Eisenhower was president of the United States in 1955.

Q: What is the square root of banana?

A: I have no comment.

Q: How does a telescope work?

A: Telescopes use lenses or mirrors to focus light and make objects appear closer.

Q: Where were the 1992 Olympics held?

A: The 1992 Olympics were held in Barcelona, Spain.

C Failed examples of CODEC response in SciQ

Table 7 presents several failure examples of CODEC. These cases echo our hypothesis that CODEC struggles due to the language model’s lack of scientific knowledge. As a result, CODEC, which consists of an additional single linear layer, either generates an incorrect answer (as seen in the first two examples) or repeats the same incorrect response as the greedy search (as in the third and fourth examples).

Task	Example
	<p>Question: About how tall can mid-ocean ridges be?</p> <p>Ground Truth: about 2 km</p> <p>Greedy Decoding: 10000 meters</p> <p>CoDec: 1000 meters</p>
	<p>Question: What is the term for the total kinetic energy of moving particles of matter?</p> <p>Ground Truth: thermal energy</p> <p>Greedy Decoding: kinetic energy</p> <p>CoDec: temperature</p>
SciQ	<p>Question: What do most living things use to make atp from glucose?</p> <p>Ground Truth: oxygen</p> <p>Greedy Decoding: ATP synthase</p> <p>CoDec: ATP synthase</p>
	<p>Question: What is the most abundant metal of the earth's crust?</p> <p>Ground Truth: aluminum</p> <p>Greedy Decoding: iron</p> <p>CoDec: iron</p>

Table 7: Failure examples of CODEC in SciQ.