

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 NDAD: NEGATIVE-DIRECTION AWARE DECODING FOR LARGE LANGUAGE MODELS VIA CONTROLLABLE HALLUCINATION SIGNAL INJECTION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Large language models (LLMs) have recently achieved impressive progress in knowledge-intensive and reasoning tasks. However, their tendency to produce fabricated or factually inconsistent content remains a fundamental challenge to their practical deployment. To address this issue, we propose Negative-Direction Aware Decoding (NDAD), a novel decoding method that identifies and exploits hallucination signals as repulsive directions in the model’s representation space, thereby improving factual adherence without retraining. Specifically, NDAD elicits hallucination-leaning signals by selectively masking critical attention heads, which exposes unstable hypotheses that the model would otherwise amplify during generation. To regulate the influence of these signals, NDAD employs two complementary weights: a global alignment weight measuring how well the induced signal aligns with the layer’s native activations (thus quantifying its referential utility) and a local weight estimating whether low-probability tokens in the masked distribution are likely to evolve toward the final output. Based on the weights, we derive a latent hallucination distribution that serves as the negative direction. A lightweight gradient-descent step then subtracts mass from hallucination-prone regions of the output distribution, adjusting the final logits while preserving the model’s high-confidence predictions. Extensive experiments across multiple LLMs and diverse benchmark datasets demonstrate that NDAD consistently enhances factual reliability without requiring additional training or external knowledge.

## 1 INTRODUCTION

In recent years, large language models (LLMs) have achieved remarkable breakthroughs over various tasks (Achiam et al., 2023; Anil et al., 2023; Touvron et al., 2023b;a; Team et al., 2023). However, a pervasive challenge is the phenomenon of hallucination, wherein LLMs generate factually incorrect, fabricated, or nonsensical information with high confidence (Ji et al., 2023; Rawte et al., 2023; Zhang et al., 2023b; Li et al., 2024).

Current approaches to mitigate hallucinations mainly fall into two categories: retrieval-augmented methods (Li et al., 2023b; Min et al., 2023) and training-based methods (Tian et al., 2023; Rafailov et al., 2023). Retrieval-augmented methods, while effective, often introduce architectural complexity, latency, and dependency on the availability and integrity of external large-scale databases. Training-based methods, on the other hand, can be computationally intensive and may struggle to generalize across diverse factual domains. A less explored, yet highly promising, avenue is the optimization of the decoding process itself (Welleck et al., 2024; Shi et al., 2024). Importantly, prior studies suggest that LLMs already encode factual signals within their internal representations as a byproduct of large-scale pretraining, though conventional decoding techniques often fail to surface this latent knowledge (Wang et al., 2020; Kadavath et al., 2022; Li et al., 2023a; Saunders et al., 2023). Motivated by this observation, intervention-based decoding methods (Chuang et al., 2023; Li et al., 2023a; 2022; Zhang et al., 2023a) have been developed to exploit these factual signals to alleviate hallucinations.

In this study, we propose a novel yet quite effective intervention decoding method called Negative-Direction Aware Decoding (NDAD). Unlike prior approaches that focus on extracting latent factual cues from early layers, NDAD instead identifies hallucination signals and then leverages them to

054 calibrate the decoding process. As shown in Figure 1, NDAD detects latent distributions that correlate  
 055 with factually incorrect outputs by masking influential attention heads in the model. We then develop  
 056 two complementary weights based on the detected distributions to suppress hallucination signals.  
 057 Specifically, the global weighting component evaluates the alignment between hallucination-oriented  
 058 logits and earlier-layer logits, estimating whether such trajectories reflect distributions the model is  
 059 more likely to generate. In parallel, the local weighting component tracks high-risk tail tokens to  
 060 assess their likelihood of advancing toward the final output. The final token selection is then guided  
 061 by a single-step gradient-descent adjustment, which penalizes the generation of tokens associated  
 062 with identified hallucination risks. Our main contributions are as follows:

- 063 • We propose NDAD, an innovative decoding approach that introduces hallucination signal to  
 064 expose the model’s underlying hallucination distribution and applies a negative awareness  
 065 mechanism for intervention.
- 066 • We incorporate a global weight measuring the directional consistency between hallucination  
 067 signal and original early-layer logits, and a local weight quantifying the likelihood of tail  
 068 tokens evolving toward the mature distribution.
- 069 • We perform comprehensive experiments on a diverse set of LLMs with different configura-  
 070 tions and scales. The experimental results indicate that NDAD reliably enhances factual  
 071 accuracy across multiple tasks and benchmark datasets.

## 072 2 RELATED WORK

073 **Hallucination Mitigation.** In LLMs, hallucination refers to the generation of content that diverges  
 074 from factual knowledge, and it has become a critical bottleneck for ensuring model reliability.  
 075 Existing research has proposed mitigation strategies along several directions. Retrieval-based methods  
 076 introduce external knowledge to calibrate factuality, such as Retrieval-Augmented Generation (RAG)  
 077 (Cheng et al., 2023; Chen et al., 2024a; Fan et al., 2024; Lewis et al., 2020), or enhance attribution by  
 078 applying retrieval and editing after generation to improve both factuality and traceability (Gao et al.,  
 079 2022; Mishra et al., 2024). Training- and preference-based methods rely on additional supervised  
 080 data or human preference signals for optimization, including Supervised Fine-Tuning (SFT) (Tian  
 081 et al., 2023), Reinforcement Learning with Human Feedback (RLHF) (Ouyang et al., 2022), and  
 082 Direct Preference Optimization (DPO) (Rafailov et al., 2023), thereby reducing hallucination through  
 083 parameter updates. Self-evaluation-based methods do not rely on external data; instead, they improve  
 084 reliability by leveraging multiple inference-time samples and incorporating techniques such as  
 085 self-criticism (Saunders et al., 2023) and diversified reasoning path sampling (Wang et al., 2022).  
 086 To further improve efficiency in enhancing factuality, our goal is to directly optimize the output  
 087 distribution of language models, thereby strengthening their robustness.

088 **Intervention Decoding.** In recent years, a line of research has emerged that enhances the factuality  
 089 of LLMs by intervening during the decoding time. Inference-Time Intervention (ITI) (Li et al.,  
 090 2023a) identifies attention heads correlated with truthfulness during inference and shifts activations  
 091 along these “truthful directions”, thereby enhancing the truthfulness of generated outputs. Similarly,  
 092 Activation Decoding (AD) (Chen et al., 2024b) leverages the model’s internal representations by  
 093 introducing an entropy-based metric of contextual activation sharpness as a decoding constraint,  
 094 thereby biasing outputs toward more reliable generations. Inspired by early work on Contrastive  
 095 Decoding (CD) (Li et al., 2022), which compared strong expert models against weaker amateur  
 096 models to improve fluency and coherence without addressing factuality, subsequent studies extended  
 097 the idea of “contrast” to the logits level. For example, Auto-Contrastive Decoding (ACD) (Gera et al.,  
 098 2023) requires fine-tuning the prediction heads of earlier layers and is therefore mainly applicable  
 099 to small-scale models. In contrast, Decoding by Contrasting Layers (DoLA) (Chuang et al., 2023)  
 100 dynamically selects the early layer that exhibits the largest semantic divergence from the final layer,  
 101 thereby suppressing erroneous tendencies in lower layers. Building upon this, Self Logits Evolution  
 102 Decoding (SLED) (Zhang et al., 2024) further integrates multiple early layers through weighted  
 103 combination and employs a gradient-descent procedure to guide the correction of the final logits,  
 104 resulting in more robust factuality enhancement.




Figure 1: Overview of NDAD. To enhance factual reliability, we introduce hallucination signals to adjust the final output logits. The weights  $\mathcal{W}^{\text{global}}$  and  $\mathcal{W}^{\text{local}}$  jointly regulate the hallucination signals to form the latent hallucination distribution  $\mathcal{P}_{\text{NDAD}}$ , from which the model steers its output away.

### 3 METHOD

LLMs are designed to autoregressively predict the next token given a preceding context. Formally, given an input prefix represented as  $\mathbf{x}_{<t} = \{x_1, x_2, \dots, x_{t-1}\}$ , the model first converts these tokens into a sequence of embedding vectors,  $\mathcal{H}_0 = \{h_0^{[1]}, h_0^{[2]}, \dots, h_0^{[t-1]}\}$ , through an embedding layer. These representations are then updated successively by a stack of  $L$  transformer blocks. We denote the hidden state of the  $t$ -token at the  $l$ -th block as  $h_l^{[t]} \in \mathbb{R}^{d_h}$ . To generate a probability distribution over the model’s vocabulary  $\mathcal{V}$ , a shared projection head  $\psi : \mathbb{R}^{d_h} \rightarrow \mathbb{R}^d$  is applied to the hidden states. In detail, from the  $l$ -th layer’s hidden state, the unnormalized score vector (logits) for the next token and its corresponding probability distribution are defined as:

$$P_l^{[t]} = \text{softmax}(\text{logits}_l^{[t]}), \text{ where } \text{logits}_l^{[t]} = \psi(h_l^{[t]}), l = 1, \dots, L. \quad (1)$$

Typically, the logits from the final layer,  $\text{logits}_L^{[t]}$ , are used for decoding. However, this can lead to generations that are plausible but factually incorrect or nonsensical. To mitigate this issue, we propose NDAD, which adjusts the logits by leveraging hallucination signal, thereby improving the reliability of the generated text.

#### 3.1 HALLUCINATION SIGNAL GENERATION


Unlike prior approaches such as DoLa (Chuang et al., 2023) and SLED (Zhang et al., 2024), which harness early-layer representations as a proxy for faithful evidence to reshape the final token distribution, we instead attempt to explicitly separate the hallucination signal to encourage the final output distribution to diverge from it. Intuitively, this shifts calibration from boosting positives to subtracting negatives. In this way, our method can prevent probability mass from accumulating on spurious or speculative trajectories.

Prior studies (Wu et al., 2024) have demonstrated that certain attention heads in LLMs play a critical role in preserving factuality and stabilizing generation. Once the support of these heads is weakened, the model tends to deviate from factual directions, making the decoding process more susceptible to hallucinations. Building on this insight, we exclusively mask influential heads to isolate a hallucination signal, which serves as a negative direction for contrastive decoding. To determine which heads should be masked, we adopt head importance scores from prior work (Wu et al., 2024)

162 to evaluate the importance of each head. Furthermore, we take into account the entropy of each  
 163 layer's distribution: a lower entropy indicates that the importance is concentrated on a small subset  
 164 of heads, suggesting that these heads are more influential. By integrating both head importance and  
 165 layer-level entropy, we achieve a more precise selection of heads to be masked. As illustrated in  
 166 Figure 2, for each block  $l \in L$ , following Wu et al. (2024), we first obtain a score list of  $n$  heads  
 167  $\{s_{l,1}, s_{l,2}, \dots, s_{l,n}\}$  in this block. We then normalize the scores into a probability distribution and  
 168 compute the layer entropy as follows:

$$E_l = - \sum_{i=1}^n p_{l,i} \log p_{l,i}, \quad p_{l,i} = \frac{s_{l,i}}{\sum_{j=1}^n s_{l,j}}, \quad i = 1, \dots, n, \quad (2)$$

172 Here,  $p_{l,i}$  denotes the normalized importance of head  $i$  in layer  $l$ , and  $E_l$  measures the uncertainty of  
 173 head importance within this layer. Then we select the top  $K$  layers with the lowest entropy and mask  
 174 the top  $x$  heads within these layers to separate hallucination signal. We represent the hallucination  
 175 signal corresponding to  $l$ -th block as  $\text{logits}_l^{\text{mask}}$ . The complete algorithmic workflow can be found in  
 176 Appendix Algorithm 1. After extracting these hallucination signals, the remaining question is how to  
 177 leverage them to calibrate the model's final outputs.



189 Figure 2: Hallucination signal generation. Darker colors indicate larger values, and gray cells  
 190 correspond to masked attention heads.

### 193 3.2 DYNAMIC WEIGHTING VIA GLOBAL CONSISTENCY AND LOCAL DIVERGENCE

195 To exploit the identified negative direction, we propose a dynamic weighting framework that integrates  
 196 both global and local perspectives.

197 **Global Consistency.** At the global level, we evaluate the directional consistency between the hallucination  
 198 signal and the original early-layer logits from the same layer, which provides a quantitative  
 199 assessment of the correlation between the original signal and the hallucination signal. Specifically,  
 200 the directional consistency  $c_l$  at layer  $l$  is measured by computing the cosine similarity between the  
 201 hallucination signal  $\text{logits}_l^{\text{mask}}$  and the original logits  $\text{logits}_l$  at the same layer  $l$ :

$$\mathcal{W}_l^{\text{global}} = \varphi(c_l), \quad c_l = \text{cos\_sim}(\text{logits}_l, \text{logits}_l^{\text{mask}}). \quad (3)$$

204 where  $\varphi(\cdot)$  denotes a linear mapping that scales values into the range  $[0, 1]$ . By measuring directional  
 205 consistency, we assess the correlation between the hallucination signal  $\text{logits}_l^{\text{mask}}$  and the model's  
 206 original logits  $\text{logits}_l$ , thereby providing a quantitative basis for the referential value of hallucination  
 207 signal at layer  $l$ . A higher consistency indicates that the signal is more closely aligned with the  
 208 model's latent hallucination direction. Accordingly, the weighting scheme increases the contribution  
 209 of more relevant hallucination signals.

210 **Local Divergence.** At the local level, we further examine the distribution of low-probability tokens.  
 211 Consistent with prior studies (Chuang et al., 2023; Zhang et al., 2024), we approximate the final-layer  
 212 logits  $\text{logits}_L$  as the ground-truth distribution. We define the evolution trajectory from the premature  
 213 to the mature state as  $\text{logits}_L - \text{logits}_l^{\text{mask}}$ . For the final mature layer, we further obtain the probability  
 214 distribution  $\mathcal{P}_{\text{logits}_L} = \text{softmax}(\text{logits}_L)$ , select the top- $I$  tokens, and construct  $I$  one-hot vectors  
 215  $\mathcal{T} = \{\mathcal{P}_{e_1}, \mathcal{P}_{e_2}, \dots, \mathcal{P}_{e_I}\}$  to serve as  $I$  approximate distributions of mature, where the index of the  
 selected token is set to 1 and all others are set to 0. In order to derive the hallucination distribution at

layer  $l$ , we begin with  $\mathcal{P}_{\logits_l^{\text{mask}}} = \text{softmax}(\logits_l^{\text{mask}})$ . Mahaut et al. (2024) suggested that low-probability tokens typically correspond to reduced factuality. Based on this observation, we define our final hallucination distribution by removing the top- $I$  tokens, which encourages the resulting distribution to approximate the negative direction more closely. In particular, by assigning a very small probability  $\epsilon \rightarrow 0$  to the top- $I$  tokens in  $\mathcal{P}_{\logits_l^{\text{mask}}}$ , we are able to derive a cleaner representation of the premature distribution, which is defined as  $\mathcal{P}_{\logits_l^{\text{mask}}}^{\text{tail}}$ . As illustrated in Figure 1, the vector  $d_1$  denotes the evolution trajectory from the premature signal  $\logits_l^{\text{mask}}$  at layer  $l$  to the mature signal  $\logits_L$ , while  $d_2$  represents the trajectory from the hallucination distribution  $\mathcal{P}_{\logits_l^{\text{mask}}}^{\text{tail}}$  toward a candidate distribution of correctness  $P_{e_i}$ . Both  $d_1$  and  $d_2$  can be interpreted as representations of the the trajectory of factual evolution, and thus we have:

$$d_1 \stackrel{\text{direction}}{\approx} d_2, \text{ where } d_1 = \logits_L - \logits_l^{\text{mask}}, d_2 = \nabla \text{KL}(\mathcal{P}_{\logits_l^{\text{mask}}}^{\text{tail}}, \mathcal{P}_{e_i}). \quad (4)$$

Intuitively, if  $d_1$  and  $d_2$  are more closely aligned, it indicates that the token in  $\mathcal{P}_{\logits_l^{\text{mask}}}^{\text{tail}}$  is more likely to evolve toward the mature output, and therefore a larger weight should be assigned to suppress its evolution. To quantify this evolution trajectory, we define the local weight as:

$$\mathcal{W}_{l,i}^{\text{local}} = \max \left( \text{cos\_sim}(\logits_l^{\text{mask}} - \logits_L, \mathcal{P}_{\logits_l^{\text{mask}}}^{\text{tail}} - \mathcal{P}_{e_i}), 0 \right), \quad i \in [1, I]. \quad (5)$$

After deriving both the global and local weights, we integrate them to obtain the final weight for each correctness direction within the top- $I$  tokens. Specifically, for the one-hot vector  $P_{e_i}$  corresponding to the  $i$ -th distribution in the correctness, the final weight at layer  $l$  is defined as:

$$\mathcal{W}_{l,i} = \mathcal{W}_l^{\text{global}} \mathcal{W}_{l,i}^{\text{local}}, \quad i \in [1, I]. \quad (6)$$

To better capture dominant signals and attenuate weak or noisy ones, we apply a squared transformation to the final weight scores. This operation accentuates high-confidence directions while diminishing the influence of marginal ones, thereby producing a sharper weighting distribution (Hinton et al., 2015; Müller et al., 2019; Zhang et al., 2021). Formally, the squared weight is:

$$\tilde{\mathcal{W}}_{l,i} = (\mathcal{W}_{l,i})^2, \quad i \in [1, I] \quad (7)$$

### 3.3 NEGATIVE-DIRECTION AWARE DECODING

After introducing the global and local weighting mechanisms, we now integrate them into the overall decoding framework. NDAD leverages these weights to controllably exploit the injected hallucination signal and employs an update in the direction of gradient-descent to guide the model away from hallucination directions during generation. The following describes the specific procedure for adjusting the final-layer logits, we first perform intra-layer normalization on the obtained signals, followed by inter-layer aggregation. The squared weights  $\tilde{\mathcal{W}}_{l,i}$  are normalized across the  $I$  correctness directions within each layer, resulting in a layer-wise normalized distribution. Formally, the latent distribution of layer  $l$  is expressed as:

$$\mathcal{P}_l = \left( \tilde{\mathcal{W}}_{l,1}, \tilde{\mathcal{W}}_{l,2}, \dots, \tilde{\mathcal{W}}_{l,I} \right) / \mathcal{Z}_l, \quad \mathcal{Z}_l = \sum_{i=1}^I \tilde{\mathcal{W}}_{l,i} \quad (8)$$

We further apply inter-layer weighting to obtain the final NDAD distribution:

$$\mathcal{P}_{\text{NDAD}} = \sum_{l=1}^L \mathcal{N}_l \mathcal{P}_l, \quad \text{where } \mathcal{N}_l = \frac{\mathcal{Z}_l}{\sum_{l=1}^L \mathcal{Z}_l}. \quad (9)$$

Here,  $\mathcal{N}_l$  denotes the relative contribution of layer  $l$ , ensuring that the aggregation respects the proportional importance of each layer while preserving comparability across layers. By incorporating negative-direction awareness, we obtain a latent hallucination distribution  $\mathcal{P}_{\text{NDAD}}$ . To suppress the generation of hallucination-prone tokens, we penalize the divergence between distribution  $\mathcal{P}_{\text{NDAD}}$  and the original distribution  $\mathcal{P}_{\logits_L}$  using the KL divergence term. The procedure is outlined in Algorithm 2. Here, the parameter  $\alpha$ , referred to as the Evolution Rate and originally introduced in the (Zhang et al., 2024), controls the magnitude of adjustment applied to the logits along the gradient direction. We then obtain the final adjusted logits as shown below:

$$\logits_L^{\text{new}} = \logits_L - \alpha \nabla \text{KL}(\mathcal{P}_{\text{NDAD}}, \mathcal{P}_{\logits_L}) \quad (10)$$

270 Table 1: Evaluation results of different methods on Llama models over varying datasets.  
271

272 Method	273 TruthfluQA(MC)				274 Factor	275 CoT	
	276 MC1	277 MC2	278 MC3	279 Avg.		280 Wiki	281 StrQA
283 Llama2-7B-base	26.58	41.88	18.96	29.14	284 58.42	285 60.74	286 13.95
	+DoLa-low	33.04	63.73	31.25	287 42.67	288 63.36	289 59.56
	+DoLa-high	31.77	63.26	30.40	290 41.81	291 62.56	292 60.44
	+AD	32.41	49.89	24.03	293 35.44	294 53.14	295 1.97
	+SLED	34.15	62.57	31.89	296 42.87	297 67.00	298 61.27
	+NDAD	<b>34.39</b>	<b>62.62</b>	<b>31.98</b>	299 <b>43.00</b>	300 <b>67.30</b>	301 <b>61.57</b>
							<b>14.86</b>
302 Llama2-7B-chat	35.62	57.47	32.10	41.73	303 56.68	304 63.58	305 21.23
	+DoLa-low	34.18	62.80	31.00	306 42.66	307 56.58	308 64.59
	+DoLa-high	33.92	61.75	30.40	309 42.02	310 56.25	311 64.19
	+AD	32.15	49.90	23.99	312 35.35	313 51.44	314 0.48
	+SLED	<b>37.09</b>	<b>63.83</b>	<b>32.96</b>	315 <b>44.63</b>	316 64.80	317 64.50
	+NDAD	36.84	63.42	32.93	318 44.40	319 <b>65.06</b>	320 <b>64.67</b>
							<b>21.99</b>
321 Llama2-13B-base	27.59	43.14	19.53	30.09	322 63.79	323 65.98	324 28.81
	+DoLa-low	31.57	62.48	30.41	325 41.49	326 65.70	327 66.46
	+DoLa-high	29.38	63.92	33.62	328 42.31	329 52.84	330 60.83
	+AD	32.15	49.90	23.99	331 35.35	332 58.18	333 2.01
	+SLED	34.76	63.58	31.88	334 43.41	335 70.94	336 66.51
	+NDAD	<b>34.88</b>	<b>63.60</b>	<b>31.97</b>	337 <b>43.48</b>	338 <b>71.18</b>	339 <b>66.81</b>
							<b>29.19</b>
340 Llama2-13B-chat	36.47	63.06	32.77	44.10	341 61.96	342 69.65	343 36.69
	+DoLa-low	34.27	63.27	31.36	344 42.97	345 60.69	346 69.48
	+DoLa-high	31.82	62.55	31.13	347 41.83	348 54.81	349 66.51
	+AD	32.15	49.90	23.99	350 35.35	351 56.71	352 23.14
	+SLED	<b>37.45</b>	<b>63.50</b>	<b>32.90</b>	353 <b>44.62</b>	354 <b>67.50</b>	355 <b>69.74</b>
	+NDAD	<b>37.58</b>	<b>63.63</b>	<b>33.02</b>	356 <b>44.74</b>	357 <b>67.74</b>	358 <b>69.96</b>
							<b>37.30</b>

## 299 4 EXPERIMENTS

## 300 4.1 EXPERIMENTAL SETUP

301 **Benchmark datasets.** We evaluate our approach against strong baselines across both multiple-choice  
302 and open-ended generation tasks. For multiple-choice settings, we employ the TruthfulQA (Lin et al.,  
303 2021) dataset to measure factuality in short-answer scenarios and the FACTOR (Wiki) (Muhlgay  
304 et al., 2023) dataset to assess performance in long-paragraph contexts. For open-ended generation,  
305 we consider PopQA (Mallen et al., 2022), NQ-Open (Lee et al., 2019), and TriviaQA (Joshi et al.,  
306 2017), as well as reasoning-intensive tasks involving chain-of-thought (CoT), including StrategyQA  
307 (Geva et al., 2021) and GSM8K (Cobbe et al., 2021).

308 **Models and Baselines.** In our experiments, we adopt a diverse set of representative open-source  
309 LLMs, including Llama2-7B (base and chat) (Touvron et al., 2023b), Llama2-13B (base and chat)  
310 (Touvron et al., 2023b), Qwen2.5-7B-instruct (Team, 2024), Mistral-7B-instruct (Jiang et al., 2023),  
311 and Llama3-8B-instruct (Grattafiori et al., 2024). We compare the following baselines: (1) Greedy  
312 Decoding. (2) DoLA-Low (Chuang et al., 2023) subtracts the logits of the most distributionally  
313 different layer from the first half of the network from the final-layer logits. (3) DoLA-High (Chuang  
314 et al., 2023) subtracts the logits of the most distributionally different layer from the second half of the  
315 network from the final-layer logits. (4) AD (Chen et al., 2024b) uses an entropy-based measure of  
316 contextual activation sharpness to constrain decoding with the model’s internal representations. (5)  
317 SLED (Zhang et al., 2024) integrates multiple early layers via weighted combination and applies a  
318 gradient-descent adjustment to refine the final logits for improved factuality.

319 **Metrics and Parameters.** For multiple-choice and CoT reasoning tasks, we evaluate factual accuracy  
320 following the approach in (Chuang et al., 2023). To assess correctness on TriviaQA, HotpotQA, and  
321 NQ-Open, we adopt the Exact Match (EM) metric, consistent with the protocol of (Joshi et al., 2017).  
322 The detailed parameter settings are provided in Appendix A.1.

324  
325  
326 Table 2: Evaluation results on Open-Ended generation tasks.  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339

Method	Llama2-7B-base			Llama2-7B-chat		
	TriviaQA	PopQA	NQ-Open	TriviaQA	PopQA	NQ-Open
Greedy	65.04	13.67	21.02	59.61	18.55	23.41
+DoLa-low	64.96	13.88	20.78	54.65	19.64	<u>23.60</u>
+DoLa-high	63.96	13.41	19.31	54.24	19.48	23.55
+AD	48.78	15.11	22.44	<u>59.64</u>	18.43	<u>23.60</u>
+SLED	<u>65.10</u>	<u>25.86</u>	<u>25.96</u>	59.61	<u>19.98</u>	23.46
+NDAD	<b>65.21</b>	<b>26.00</b>	<b>26.26</b>	<b>59.67</b>	<b>20.13</b>	<b>23.63</b>
Llama2-13B-base				Llama2-13B-chat		
Greedy	68.34	25.04	32.71	66.32	19.82	30.03
+DoLa-low	68.67	28.64	28.78	65.54	17.82	29.14
+DoLa-high	62.08	26.12	25.68	61.86	16.32	27.42
+AD	67.67	17.91	30.80	64.50	<b>22.91</b>	<b>34.52</b>
+SLED	<u>71.47</u>	<u>30.53</u>	32.52	<u>66.40</u>	19.84	29.89
+NDAD	<b>71.66</b>	<b>30.64</b>	<b>32.88</b>	<b>66.48</b>	<u>19.85</u>	30.11

340  
341  
342 4.2 EVALUATION ON DIFFERENT BENCHMARKS  
343

344 **Multiple-Choices Tasks.** These tasks are designed to evaluate whether the decoding strategy can  
345 more effectively assign higher probabilities to correct answers or reasonable completions, while  
346 suppressing its preference for incorrect options. It should be noted that these tasks are essentially  
347 distribution-fitting problems, and overfitting to specific tasks often undermines the generalization  
348 capability of a decoding method. Since our goal is to enhance factuality and robustness while  
349 preserving broad applicability, even when the performance deviation is small or only marginal  
350 improvements are achieved, the results remain understandable and acceptable. We validated the  
351 effectiveness of the NDAD method through short-answer factuality tests on the TruthfulQA dataset  
352 and long-paragraph factuality tests on the FACTOR dataset. The corresponding experimental results  
353 are summarized in Table 1, and more detailed analyses are provided in Appendix B.1. Our NDAD  
354 method demonstrates strong generalization across different models and datasets, and largely achieves  
355 improvements over the baseline SLED. This suggests that the proposed decoding strategy is generally  
356 more effective at calibrating probability assignment between correct and incorrect answers.

357 **Chain-of-Thought Reasoning Tasks.** This task primarily focuses on evaluating how different  
358 decoding methods can be adapted to the CoT strategy to effectively handle complex reasoning  
359 problems. The detailed results can be found in Table 1. Our NDAD method consistently outperforms  
360 all baselines in decoding performance. At the same time, the limitations of AD become particularly  
361 evident on CoT datasets. AD constrains next-token probabilities by incorporating contextual entropy  
362 to enhance factuality. However, it falls short on reasoning tasks because tokens in CoT datasets exhibit  
363 strong logical dependencies, and relying solely on token-level activation entropy from the context  
364 may deviate from the original semantics. Moreover, some intermediate tokens lack contextual support  
365 and are prone to being misclassified as hallucinations, thereby impairing reasoning performance.

366 **Open-Ended Generation Tasks.** For open-ended tasks, we adopt TriviaQA, PopQA, and NQ-Open  
367 datasets. Our NDAD method consistently achieves further improvements over the baselines. Results  
368 are shown in Table 2. Since PopQA and NQ-Open are highly knowledge-intensive, models tend to  
369 rely more on contextual information during generation. The AD method, which is inherently designed  
370 to adjust decoding based on contextual attention, therefore shows exceptionally strong reasoning  
371 performance on the Llama-13B-chat model. However, when compared with the results on CoT tasks  
372 in Table 1, it becomes evident that AD exhibits substantial variability. Therefore, our NDAD method  
373 demonstrates the strongest robustness.

374  
375  
376  
377 4.3 EVALUATION ON DIFFERENT LLMs  
378

379 We further conduct experiments on a broader range of model architectures, including models from  
380 different families as well as different variants within the same family. As reported in Table 3, NDAD  
381 consistently delivers state-of-the-art results across all tested configurations, surpassing other baselines.  
382 This demonstrates that the proposed method is not only effective for a specific model class but

378  
379  
380 Table 3: Evaluation results on varying LLMs.  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395

Model	TruthfluQA(MC)				Factor Wiki	CoT GSM8K
	MC1	MC2	MC3	Avg.		
Qwen2.5-7B-instruct	41.00	64.59	38.17	47.92	54.54	84.46
	+DoLa-low	36.60	66.03	34.21	45.61	83.02
	+DoLa-high	34.64	2.37	34.51	23.84	76.95
	+SLED	45.04	70.37	39.88	51.76	84.91
	+NDAD	45.17	70.37	39.89	51.81	85.14
Mistral-7B-instruct	40.27	68.32	37.06	48.55	60.49	53.45
	+DoLa-low	39.53	68.44	36.16	48.04	53.22
	+DoLa-high	39.53	68.43	36.09	48.02	53.30
	+SLED	45.41	71.17	40.27	52.28	53.90
	+NDAD	45.53	71.31	40.46	52.43	54.36
Llama3-8B-instruct	38.92	68.16	36.56	47.88	59.22	75.97
	+DoLa-low	35.74	65.27	33.60	44.87	75.82
	+DoLa-high	35.99	65.04	33.72	44.92	75.51
	+SLED	41.37	68.46	37.61	49.15	75.82
	+NDAD	41.37	69.21	37.89	49.49	77.18

396  
397 Table 4: Evaluation results on Llama2-70B.  
398

Method	Factor	GSM8K
Llama2-70B	61.92	56.10
+DoLa-Low	74.05	57.01
+DoLa-High	62.53	38.21
+SLED	77.32	57.01
+NDAD	77.52	57.54

399  
400 Table 5: Runtime and memory overhead on  
401 Llama2-7B-base.  
402  
403  
404

Method	Runtime (s)	Memory (MB)
Greedy	1.11	13503.47
DoLa	1.17	15261.98
SLED	1.17	15452.88
NDAD	1.34	17779.01

405  
406 also generalizes well across diverse architectures. Moreover, the performance gains are particularly  
407 pronounced on CoT datasets such as GSM8K, where NDAD exhibits substantial improvements over  
408 the baselines. This finding highlights the robustness of NDAD in handling complex reasoning tasks.  
409 Consequently, these results confirm that NDAD achieves both cross-model generality and strong  
410 robustness, making it a versatile and effective decoding strategy.  
411

412  
413 4.4 EVALUATION ON LARGER-SCALE LLM

414 To assess the viability of the method on substantially larger models, we conducted additional  
415 experiments using Llama2-70B on the Factor dataset for multiple-choice tasks and GSM8K for  
416 chain-of-thought reasoning. The results, presented in Table 4, show that the method continues to  
417 deliver strong performance on generative tasks such as GSM8K. The second-best baseline improves  
418 by 0.91%, whereas our method achieves an improvement of 1.44%, corresponding to a relative gain  
419 of 58%. For the Factor dataset, as discussed in Section 4.2, this task essentially evaluates distribution  
420 fitting, where maintaining a smooth upward trend is sufficient. These results demonstrate that the  
421 method remains effective when scaled to much larger models and exhibits strong robustness across  
422 different model sizes.  
423

## 424 4.5 ABLATION STUDY

425 **Incorporation of Hallucination Signal.** We first demonstrate that our method indeed introduces  
426 hallucination signal into the model. To this end, we directly decode the logits obtained after masking  
427 the importance attention heads and evaluate their performance. The experimental results are shown in  
428 Figure 3. As can be observed, compared with the original decoding, performance consistently drops  
429 across different models and datasets, with the most significant decline occurring on the GSM8K  
430 dataset. This indicates that complex reasoning tasks heavily rely on the aggregation and inference  
431 of internal attention heads, and masking these heads introduces stronger hallucination signal. This  
observation is consistent with the analysis in Section 4.3, where our NDAD method achieves better




Figure 3: Results from Decoding Hallucination Signals.




Figure 4: Different head and layer parameters on the Llama-7B-base.

results on GSM8K, suggesting that stronger hallucination signal can provide more effective leverage for enhancing the NDAD decoding strategy. Moreover, the ablation experiments in Table 6 based on random head and layer selection further support that hallucination induction guided by head importance and layer-level entropy contributes to the performance gains of NDAD.

**Importance of Head and Layer Parameters.** To effectively introduce hallucination signal, it is necessary to mask more important attention heads. Using the Llama-7B-base model as an example, we present results on the FACTOR and GSM8K datasets under different parameter settings. Figure 4a illustrates the impact on accuracy when varying the number of masked heads while keeping the number of masked layers fixed. Conversely, Figure 4b shows the effect of varying the number of masked layers while fixing the number of masked heads. Overall, the trend generally follows a rising-then-falling pattern. Notably, throughout the experiments, the range of masked heads and layers remained between [6, 13], within which the model consistently achieved relatively strong performance across both datasets. More detailed results are provided in Appendix B.3.

**Global and Local Weights.** We further analyze the effectiveness of the global and local weighting components in our method. The ablation results based on Llama2-7B-base and Llama2-13B-base are reported in Table 6, and the more comprehensive results and analyses can be found in Appendix B.2. Specifically, w/o global weight indicates removing the measurement of directional consistency between hallucination signal and the original signals, while w/o local weight corresponds to excluding the measurement of consistency between the tail-token evolution and the transition from the premature to the mature state. From the results, it is clear that both weighting mechanisms play a crucial role in enhancing the decoding performance. For example, in the case of Llama2-7B-base, removing either global or local weights leads to a drop in performance. A similar trend is observed for Llama2-13B-base, where the absence of these weights consistently reduces accuracy across all benchmarks. Importantly, the GSM8K dataset again shows the largest degradation, underscoring that complex reasoning tasks are particularly sensitive to the loss of these weighting mechanisms. These results confirm that both global and local weights contribute complementary benefits, and together they enable NDAD to achieve robust and state-of-the-art performance.

#### 4.6 COMPUTATIONAL OVERHEAD ANALYSIS

To evaluate the computational overhead of our method, we measured runtime and memory usage on the Llama2-7B-base model using a single GSM8K sample, and the results are presented in Table 5. As shown, the additional cost introduced by NDAD is relatively lightweight, with the primary overhead arising from the incorporation of the negative-direction signal. Consistent with existing decoding-based approaches, NDAD only modifies the logits of the final layer, requires no additional training, and does not depend on high-quality external data, giving it strong plug-and-play capability. In many real-world applications, safety and factual reliability are often more critical than achieving

486  
487  
488 Table 6: Ablation study on the effectiveness of each component in the NDAD method.  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500

Method	TruthfluQA(MC)				Factor Wiki	CoT	
	MC1	MC2	MC3	Avg.		StrQA	GSM8K
Llama2-7B-base	26.58	41.88	18.96	29.14	58.42	60.74	13.95
random head	34.15	62.55	31.91	42.87	67.17	61.13	13.95
random layer	34.15	62.61	31.84	42.87	67.10	61.40	14.71
w/o global weight	34.27	62.57	31.93	42.92	67.20	61.09	14.63
w/o local weight	33.90	61.13	31.43	42.15	67.17	61.44	14.10
NDAD	<b>34.39</b>	<b>62.62</b>	<b>31.98</b>	<b>43.00</b>	<b>67.30</b>	<b>61.57</b>	<b>14.86</b>
Llama2-13B-base	27.59	43.14	19.53	30.09	63.79	65.98	28.81
random head	34.88	63.58	31.94	43.47	71.04	66.72	28.13
random layer	34.76	63.56	31.91	43.41	71.01	66.72	28.66
w/o global weight	34.88	63.59	31.93	43.47	70.98	65.41	28.73
w/o local weight	34.76	63.57	31.89	43.41	70.91	66.07	27.98
NDAD	<b>34.88</b>	<b>63.60</b>	<b>31.97</b>	<b>43.48</b>	<b>71.18</b>	<b>66.81</b>	<b>29.26</b>

501  
502  
503 the absolute lowest decoding latency; thus, a moderate amount of runtime and memory overhead is  
504 generally acceptable.505  
506 

## 5 CONCLUSION

508  
509 We present an innovative decoding strategy NDAD, which explicitly elicits hallucination signal by  
510 masking critical attention heads and leverages them as negative directions for contrastive decoding.  
511 To controllably leverage these signals, we design a dynamic weighting mechanism: the global  
512 weight measures the directional consistency between the hallucination signal and the original early-  
513 layer logits, thereby quantifying the referential value of the current hallucination signal; the local  
514 weight characterizes the tendency of low-probability tokens to evolve toward the mature distribution.  
515 By suppressing the output probabilities of hallucination-prone tokens through gradient-descent  
516 adjustments during decoding, NDAD consistently improves factual reliability across diverse models  
517 and benchmarks, demonstrating particularly strong robustness in complex reasoning tasks. In  
518 conclusion, NDAD provides a lightweight yet effective solution for optimizing LLM decoding.519  
520 

## ETHICAL STATEMENT

521  
522 This paper presents a decoding strategy designed to improve the factual reliability of LLMs. Our  
523 research does not involve human subjects, sensitive personal data, or potentially harmful datasets. All  
524 benchmark datasets employed in our experiments are publicly available and widely used within the  
525 Natural Language Processing research community.526  
527 

## REPRODUCIBILITY STATEMENT

528  
529 To ensure the reproducibility of our experiments, we have provided the source codes in the supple-  
530 mentary materials for review. Upon acceptance of this paper, we will release the codes as open source  
531 to enable researchers to replicate and extend our experiments.532  
533 

## REFERENCES

534  
535 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
536 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.  
537 *arXiv preprint arXiv:2303.08774*, 2023.538 Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos,  
539 Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv*  
540 *preprint arXiv:2305.10403*, 2023.

540 Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. Benchmarking large language models in  
 541 retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*,  
 542 volume 38, pp. 17754–17762, 2024a.

543

544 Shiqi Chen, Miao Xiong, Junteng Liu, Zhengxuan Wu, Teng Xiao, Siyang Gao, and Junxian He.  
 545 In-context sharpness as alerts: An inner representation perspective for hallucination mitigation.  
 546 *arXiv preprint arXiv:2403.01548*, 2024b.

547

548 Xin Cheng, Di Luo, Xiuying Chen, Lemao Liu, Dongyan Zhao, and Rui Yan. Lift yourself up:  
 549 Retrieval-augmented text generation with self-memory. *Advances in Neural Information Processing  
 550 Systems*, 36:43780–43799, 2023.

551

552 Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. Dola:  
 553 Decoding by contrasting layers improves factuality in large language models. *arXiv preprint  
 554 arXiv:2309.03883*, 2023.

555

556 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
 557 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve  
 558 math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

559

560 Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and  
 561 Qing Li. A survey on rag meeting llms: Towards retrieval-augmented large language models. In  
 562 *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pp.  
 563 6491–6501, 2024.

564

565 Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan,  
 566 Vincent Y Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, et al. Rarr: Researching and revising what  
 567 language models say, using language models. *arXiv preprint arXiv:2210.08726*, 2022.

568

569 Ariel Gera, Roni Friedman, Ofir Ariv, Chulaka Gunasekara, Benjamin Sznajder, Noam Slonim, and  
 570 Eyal Shnarch. The benefits of bad advice: Autocontrastive decoding across model layers. *arXiv  
 571 preprint arXiv:2305.01628*, 2023.

572

573 Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle  
 574 use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of  
 575 the Association for Computational Linguistics*, 9:346–361, 2021.

576

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

580

581 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv  
 582 preprint arXiv:1503.02531*, 2015.

583

584 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang,  
 585 Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM  
 586 Computing Surveys*, 55(12):1–38, 2023.

587

588 Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,  
 589 Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al.  
 590 Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.

591

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

594

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

598

599 Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. Latent retrieval for weakly supervised open  
 600 domain question answering. *arXiv preprint arXiv:1906.00300*, 2019.

594 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
 595 Heinrich Kütller, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented genera-  
 596 tion for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:  
 597 9459–9474, 2020.

598 Junyi Li, Jie Chen, Ruiyang Ren, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen.  
 599 The dawn after the dark: An empirical study on factuality hallucination in large language models.  
 600 *arXiv preprint arXiv:2401.03205*, 2024.

602 Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time  
 603 intervention: Eliciting truthful answers from a language model. *Advances in Neural Information  
 604 Processing Systems*, 36:41451–41530, 2023a.

606 Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke  
 607 Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization.  
 608 *arXiv preprint arXiv:2210.15097*, 2022.

609 Xiaonan Li, Changtai Zhu, Linyang Li, Zhangyue Yin, Tianxiang Sun, and Xipeng Qiu. Llatrieval:  
 610 Llm-verified retrieval for verifiable generation. *arXiv preprint arXiv:2311.07838*, 2023b.

612 Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human  
 613 falsehoods. *arXiv preprint arXiv:2109.07958*, 2021.

615 Matéo Mahaut, Laura Aina, Paula Czarnowska, Momchil Hardalov, Thomas Müller, and Lluís  
 616 Márquez. Factual confidence of llms: on reliability and robustness of current estimators. *arXiv  
 617 preprint arXiv:2406.13415*, 2024.

618 Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi.  
 619 When not to trust language models: Investigating effectiveness of parametric and non-parametric  
 620 memories. *arXiv preprint arXiv:2212.10511*, 2022.

622 Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer,  
 623 Luke Zettlemoyer, and Hannaneh Hajishirzi. Factscore: Fine-grained atomic evaluation of factual  
 624 precision in long form text generation. *arXiv preprint arXiv:2305.14251*, 2023.

625 Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov,  
 626 and Hannaneh Hajishirzi. Fine-grained hallucination detection and editing for language models.  
 627 *arXiv preprint arXiv:2401.06855*, 2024.

629 Dor Muhlgay, Ori Ram, Inbal Magar, Yoav Levine, Nir Ratner, Yonatan Belinkov, Omri Abend,  
 630 Kevin Leyton-Brown, Amnon Shashua, and Yoav Shoham. Generating benchmarks for factuality  
 631 evaluation of language models. *arXiv preprint arXiv:2307.06908*, 2023.

632 Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances  
 633 in neural information processing systems*, 32, 2019.

635 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
 636 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow  
 637 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–  
 638 27744, 2022.

639 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea  
 640 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances  
 641 in neural information processing systems*, 36:53728–53741, 2023.

643 Vipula Rawte, Amit Sheth, and Amitava Das. A survey of hallucination in large foundation models.  
 644 *arXiv preprint arXiv:2309.05922*, 2023.

646 William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike.  
 647 Self-critiquing models for assisting human evaluators, 2022. URL <https://arxiv.org/abs/2206.05802>,  
 2023.

648 Chufan Shi, Haoran Yang, Deng Cai, Zhisong Zhang, Yifan Wang, Yujiu Yang, and Wai Lam. A  
 649 thorough examination of decoding methods in the era of llms. *arXiv preprint arXiv:2402.06925*,  
 650 2024.

651 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,  
 652 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly  
 653 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

654 Qwen Team. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2, 2024.

655 Katherine Tian, Eric Mitchell, Huaxiu Yao, Christopher Manning, and Chelsea Finn. Fine-tuning  
 656 language models for factuality. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction*  
 657 *Following*, 2023.

658 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée  
 659 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and  
 660 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.

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

664 Chenguang Wang, Xiao Liu, and Dawn Song. Language models are open knowledge graphs. *arXiv*  
 665 *preprint arXiv:2010.11967*, 2020.

666 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-  
 667 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.  
 668 *arXiv preprint arXiv:2203.11171*, 2022.

669 Sean Welleck, Amanda Bertsch, Matthew Finlayson, Hailey Schoelkopf, Alex Xie, Graham Neubig,  
 670 Ilia Kulikov, and Zaid Harchaoui. From decoding to meta-generation: Inference-time algorithms  
 671 for large language models. *arXiv preprint arXiv:2406.16838*, 2024.

672 Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. Retrieval head mechanistically  
 673 explains long-context factuality. *arXiv preprint arXiv:2404.15574*, 2024.

674 Chang-Bin Zhang, Peng-Tao Jiang, Qibin Hou, Yunchao Wei, Qi Han, Zhen Li, and Ming-Ming  
 675 Cheng. Delving deep into label smoothing. *IEEE Transactions on Image Processing*, 30:5984–  
 676 5996, 2021.

677 Jianyi Zhang, Da-Cheng Juan, Cyrus Rashtchian, Chun-Sung Ferng, Heinrich Jiang, and Yiran Chen.  
 678 Sled: Self logits evolution decoding for improving factuality in large language models. *Advances*  
 679 *in Neural Information Processing Systems*, 37:5188–5209, 2024.

680 Yue Zhang, Leyang Cui, Wei Bi, and Shuming Shi. Alleviating hallucinations of large language  
 681 models through induced hallucinations. *arXiv preprint arXiv:2312.15710*, 2023a.

682 Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao,  
 683 Yu Zhang, Yulong Chen, et al. Siren’s song in the ai ocean: a survey on hallucination in large  
 684 language models. *arXiv preprint arXiv:2309.01219*, 2023b.

690 **A EXPERIMENTAL SETTINGS**

692 **A.1 PARAMETER SETTINGS**

694 For the parameters  $\alpha$  in Equation 10 and the  $I$  correctness distributions in Equation 5, we set the  
 695 default values to  $\alpha = 2$  and  $I = 10$ . However, due to dataset uncertainty, additional hyperparameter  
 696 tuning may be required in special cases. Following the work of (Zhang et al., 2024), we test  $\alpha$  from  
 697  $\{0.01, 0.1, 1, 2, 5, 10\}$  and  $I$  from  $\{5, 10, 20, 50\}$ . During the aforementioned tests, we guarantee  
 698 that the chosen parameters achieve performance better than greedy decoding. On this basis, we then  
 699 incorporate our hallucination signal to conduct adaptive negative-direction aware decoding. For the  
 700 number of masked heads and layers used in introducing hallucination signal, we partly explained  
 701 this in Section 4.5. In experiments, we usually set the range to  $[6, 13]$ , which generally yields strong  
 702 performance.




Figure 5: Different head and layer parameters on the Llama-7B-chat.

## 713 B ADDITIONAL EXPERIMENTAL RESULTS

### 714 B.1 EXTENDED ANALYSIS OF MULTIPLE-CHOICES TASKS

717 As shown in Table 1, the performance improvements of NDAD on multiple-choice tasks are slightly  
 718 smaller compared to other task types. This is consistent with the discussion in Section 4.2, where  
 719 multiple-choice problems essentially reduce to a logits-fitting task; as long as the model achieves  
 720 stable performance without large fluctuations and delivers moderate gains, the results remain reasonable.  
 721 Moreover, since the multiple-choice format inherently constrains the output space with a fixed  
 722 set of candidate answers, the likelihood of hallucination is substantially reduced, leading to weaker  
 723 hallucination signals and thus smaller benefits from NDAD’s decoding adjustments. Nevertheless,  
 724 our primary focus is on more complex open-ended generation tasks, where hallucinations are far  
 725 more prevalent and where NDAD demonstrates clear advantages in suppressing hallucination-prone  
 726 directions and enhancing factual reliability.

### 727 B.2 EXTENDED ABLATION ANALYSIS

729 We further conducted ablation experiments on Llama2-7B-chat and Llama2-13B-chat to examine the  
 730 effect of different components in NDAD, with the experimental setup summarized in Table 7.

732 **Hallucination Signal Induction.** During the stage of hallucination signal induction, we observed  
 733 that the random selection of attention heads or layers occasionally outperformed our guided masking  
 734 strategy based on head importance and layer-level entropy. This can be attributed to the inherently  
 735 greedy nature of the masking strategy: although generally effective, it does not fully explore the  
 736 extensive search space. Consequently, certain random configurations may fortuitously yield superior  
 737 outcomes. Nonetheless, such instances are expected and do not diminish the overall effectiveness of  
 738 a principled importance-guided approach.

739 **Global Weighting in Multiple-Choice Tasks.** For the global weighting component, the performance  
 740 on Llama2-7B-chat with the TruthfulQA dataset was slightly better when the global weighting was  
 741 not applied compared to the full NDAD method. As discussed in Section B.1, these multiple-choice  
 742 tasks essentially reduce to a logits-fitting problem with a small set of candidate answers. Since all  
 743 options are inherently more reliable than open-ended generations, the model is less vulnerable to  
 744 noisy hallucinations in this setting. Consequently, assessing the reliability of hallucination signals  
 745 becomes less critical, and the global weighting may even introduce unnecessary adjustments that  
 746 interfere with straightforward logits alignment. By contrast, in open-ended generation tasks, where  
 747 hallucination is more prevalent, the global and local weighting strategies play a much more important  
 748 role in enhancing factual reliability.

### 749 B.3 EXTENDED PARAMETER ANALYSIS

751 We further conducted hyperparameter experiments on Llama2-7B-chat. As shown in Figure 5, for  
 752 both the number of masked attention heads and the number of masked layers, performance exhibits a  
 753 general rising-then-falling trend: as the number of masked heads or layers increases, performance  
 754 initially improves but declines once the masking becomes excessive. The results suggest that the  
 755 optimal settings typically lie within the range of 6 to 13, where a better balance is achieved between  
 inducing hallucination signals and preserving the original representations.

756 Table 7: Additional ablation study on the effectiveness of each component in the NDAD method.  
757

758 759 Method	760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 799 800 801 802 803 804 805 806 807 808 809				760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 799 800 801 802 803 804 805 806 807 808 809				
	760 MC1	761 MC2	762 MC3	763 Avg.	764 Factor	765 CoT	766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 799 800 801 802 803 804 805 806 807 808 809		
Llama2-7B-chat	35.62	57.47	32.10	41.73	56.68	63.58	21.23		
random head	36.84	63.38	32.65	44.29	64.93	<b>64.72</b>	21.15		
random layer	36.47	62.99	32.59	44.02	65.00	<b>64.72</b>	20.62		
w/o global weight	<b>36.84</b>	<b>63.71</b>	<b>32.80</b>	<b>44.45</b>	64.93	64.37	21.00		
w/o local weight	36.47	60.82	32.40	43.23	64.96	63.58	20.62		
NDAD	36.84	63.27	32.76	44.29	<b>65.06</b>	64.67	<b>21.99</b>		
Llama2-13B-chat	36.47	63.06	32.77	44.10	61.96	69.65	36.69		
random head	37.45	63.61	32.95	44.67	67.47	69.91	35.63		
random layer	<b>37.70</b>	<b>63.58</b>	<b>33.07</b>	<b>44.78</b>	67.57	69.52	35.78		
w/o global weight	35.62	63.91	32.49	44.01	67.60	69.43	35.71		
w/o local weight	37.21	64.02	32.90	44.71	67.67	69.65	37.00		
NDAD	37.58	63.63	33.02	44.74	<b>67.74</b>	<b>69.96</b>	<b>37.30</b>		

772  
773 B.4 EXTENDED LINGUISTIC QUALITY EVALUATION

774 To assess whether NDAD introduces any degradation in linguistic quality, we conduct an additional  
775 evaluation focusing on fluency, coherence, and comprehensibility. These dimensions reflect whether  
776 the generated responses remain natural, logically organized, and easy to understand—qualities that  
777 are essential for real-world deployment but are often overlooked in factuality-oriented methods. We  
778 generate model outputs using Llama2-70B on GSM8K and obtain linguistic quality scores from the  
779 external evaluator Gemini-2.5-Pro. The results are presented in Table 8. As shown, the scores across  
780 all methods are highly consistent, and NDAD performs on par with or slightly better than existing  
781 decoding strategies, indicating that NDAD does not introduce noticeable negative effects on linguistic  
782 quality. This evaluation further demonstrates that NDAD improves factuality while preserving the  
783 naturalness and readability of generated text. Table 9 is the full evaluation prompt used for scoring  
784 with the Gemini model.

785 Table 8: Linguistic quality evaluation of different decoding methods using Gemini-2.5-Pro.  
786

787 788 Method	789 790 Fluency	791 792 Coherence	793 794 Comprehensibility
Greedy	9.37	7.96	8.65
DoLa	9.29	7.91	8.58
SLED	9.32	8.02	8.69
NDAD	9.31	8.04	8.67

## 795 C ALGORITHM OF NDAD

796 The entire algorithmic workflow of the NDAD method is presented in Algorithm 1 and 2.

## 797 D CASE STUDY

801 Table 10 reports the results of the Llama-7B-Base model on the GSM8K dataset under different  
802 decoding strategies. The examples demonstrate that our NDAD method is more effective in eliciting  
803 factual outputs from the model.

810  
811  
812  
813  
814  
815  
816

Table 9: Prompt for Gemini-2.5-Pro.

817 You are an advanced artificial intelligence review system  
818 specialized in evaluating the quality of model responses. Your  
819 task is to rate the quality from three perspectives: fluency,  
820 coherence, and comprehensibility. Please strictly follow the  
821 evaluation dimensions below to score each item (range: 0–10,  
822 with higher scores indicating better quality).

823 [Evaluation Criteria]

825 Fluency: Whether the sentence structure of the answer is clear  
826 and natural, with no obvious grammatical errors, inappropriate  
827 word usage, or issues affecting the reading experience. Higher  
828 scores indicate smooth language that can be read without  
829 difficulty.

830 Coherence: Whether the logical connections between parts of  
831 the answer are tight and information flows smoothly. Check  
832 for jumps, breaks, contradictions, or repetition that affect  
833 logical coherence. Higher scores indicate clear thinking and  
834 reasonable structure.

836 Comprehensibility: Whether the answer is easy for the target  
837 reader to understand. Higher scores indicate clear information  
838 delivery, easy understanding, and no ambiguity or obscure  
839 expressions.

841 [Output Format]

843 Please output in the following JSON format:

844 {

846 "Scores for Each Dimension": {  
847 "Fluency": score,  
848 "Coherence": score,  
849 "Comprehensibility": score  
850 },

852 "Reason for Scoring": Explain the reasons for scoring each  
853 dimension, and briefly summarize the overall evaluation

855 }

856 Please validate the question and return the result in JSON  
857 format, with no other content except the JSON.

859  
860  
861  
862  
863

864

865

866

867

**Algorithm 1** Hallucination Signal Induction

---

- 1: LLM with  $L$  layers,  $sequence$ , following the work of (Wu et al., 2024), a original score list of  $n$  attention head  $\{s_{l,1}, s_{l,2}, \dots, s_{l,n}\}$  in layer  $l$ , number of masked attention heads  $x$ , number of masked layer  $K$ .
- 2: **for**  $l < L$  **do**
- 3:     Normalize scores into probability distribution:  $p_{l,i} = \frac{s_{l,i}}{\sum_{j=1}^n s_{l,j}}$ ,  $i = 1, \dots, n$ .
- 4:     Compute attention head scores distribution entropy:  $E_l = -\sum_{i=1}^n p_{l,i} \log p_{l,i}$ .
- 5: **end for**
- 6: Obtain the set of distribution entropy  $\{E_1, E_2, \dots, E_L\}$ .
- 7: Select the set  $\mathcal{L}$  consisting of the  $K$  layers  $l$  corresponding to the largest entropy values.
- 8: **for**  $l \in \mathcal{L}$  **do**
- 9:     Set the weights of the top- $x$  scoring attention heads to 0.
- 10: **end for**
- 11: The  $sequence$  into the LLM to obtain the hallucination signals  $logits_l^{\text{mask}}$ , where  $l \leq L$ .
- 12: **Return:**  $\{logits_1^{\text{mask}}, logits_2^{\text{mask}}, \dots, logits_L^{\text{mask}}\}$

---

885

886

887

888

889

890

891

892

893

**Algorithm 2** Negative-Direction Aware Decoding

---

- 1: **Initialization:** LLM with  $L$  layers,  $sequence$ ,  $\alpha$  in Equation 10, number of correctness directions  $I$ ,  $\epsilon \rightarrow 0$ ,  $\varphi(\cdot)$  maps values into  $[0, 1]$ , the one-hot vectors  $\mathcal{T} = \{\mathcal{P}_{e_1}, \mathcal{P}_{e_2}, \dots, \mathcal{P}_{e_I}\}$  of correctness directions.
- 2: The  $sequence$  into the LLM to obtain the original logits  $logits_l$  and hallucination signal  $logits_l^{\text{mask}}$  given by Algorithm 1, the probabilities at each layer  $l$  denoted as  $\mathcal{P}_{logits_l} = \text{softmax}(logits_l)$  and  $\mathcal{P}_{logits_l^{\text{mask}}} = \text{softmax}(logits_l^{\text{mask}})$ , where  $l \leq L$ .
- 3: Identify the tokens with the top- $I$  largest probabilities in  $\mathcal{P}_{logits_L}$  and assign the value 1 to their indices and 0 to the remaining positions.
- 4: Set the indices of top- $I$  largest probabilities tokens in  $\mathcal{P}_{logits_l^{\text{mask}}}$  to  $\epsilon$ :  $\mathcal{P}_{logits_l^{\text{mask}}} \rightarrow \mathcal{P}_{logits_l^{\text{mask}}}^{\text{tail}}$ .
- 5: **for**  $l < L$  **do**
- 6:     Compute  $\mathcal{W}_l^{\text{global}} = \varphi(\text{cos\_sim}(logits_l, logits_l^{\text{mask}}))$ .
- 7:     Compute  $\mathcal{W}_{l,i}^{\text{local}} = \max(\text{cos\_sim}(logits_l^{\text{mask}} - logits_L, \mathcal{P}_{logits_l^{\text{mask}}}^{\text{tail}} - \mathcal{P}_{e_i}), 0)$ ,  $\mathcal{P}_{e_i} \in \mathcal{T}$ .
- 8:     Calculate  $\tilde{\mathcal{W}}_{l,i} = (\mathcal{W}_l^{\text{global}} \mathcal{W}_{l,i}^{\text{local}})^2$ ,  $i \in [1, I]$ .
- 9: **end for**
- 10: Obtain the current latent distribution  $\mathcal{P}_{\text{NDAD}} = \frac{\sum_{l=1}^L \tilde{\mathcal{W}}_{l,i}}{\sum_{l=1}^L \sum_{j=1}^{|I|} \tilde{\mathcal{W}}_{l,j}}$  by computing each  $i \in [1, I]$  across different layers.
- 11: **Return:**  $logits_L^{\text{new}} = logits_L - \alpha \nabla KL(\mathcal{P}_{\text{NDAD}}, \mathcal{P}_{logits_L})$

---

914

915

916

917

918  
919  
920  
921  
922  
923  
924

Table 10: Case study of Llama-7B-base on the GSM8K Dataset.

925	<b>Input:</b>	Q: There are 15 trees in the grove. Grove workers will plant trees in the grove today. After they are done, there will be 21 trees. How many trees did the grove workers plant today? A: There are 15 trees originally. Then there were 21 trees after some more were planted. So there must have been $21 - 15 = 6$ . The answer is 6.
926		Q: If there are 3 cars in the parking lot and 2 more cars arrive, how many cars are in the parking lot? A: There are originally 3 cars. 2 more cars arrive. $3 + 2 = 5$ . The answer is 5.
927		Q: Leah had 32 chocolates and her sister had 42. If they ate 35, how many pieces do they have left in total? A: Originally, Leah had 32 chocolates. Her sister had 42. So in total they had $32 + 42 = 74$ . After eating 35, they had $74 - 35 = 39$ . The answer is 39.
928		Q: Jason had 20 lollipops. He gave Denny some lollipops. Now Jason has 12 lollipops. How many lollipops did Jason give to Denny? A: Jason started with 20 lollipops. Then he had 12 after giving some to Denny. So he gave Denny $20 - 12 = 8$ . The answer is 8.
929		Q: Shawn has five toys. For Christmas, he got two toys each from his mom and dad. How many toys does he have now? A: Shawn started with 5 toys. If he got 2 toys each from his mom and dad, then that is 4 more toys. $5 + 4 = 9$ . The answer is 9.
930		Q: There were nine computers in the server room. Five more computers were installed each day, from monday to thursday. How many computers are now in the server room? A: There were originally 9 computers. For each of 4 days, 5 more computers were added. So $5 * 4 = 20$ computers were added. $9 + 20$ is 29. The answer is 29.
931		Q: Michael had 58 golf balls. On tuesday, he lost 23 golf balls. On wednesday, he lost 2 more. How many golf balls did he have at the end of wednesday? A: Michael started with 58 golf balls. After losing 23 on tuesday, he had $58 - 23 = 35$ . After losing 2 more, he had $35 - 2 = 33$ golf balls. The answer is 33.
932		Q: Olivia has 23. <i>She bought five bagels for 3 each.</i> How much money does she have left? A: Olivia had 23 dollars. 5 bagels for 3 dollars each will be $5 * 3 = 15$ dollars. So she has $23 - 15$ dollars left. $23 - 15$ is 8. The answer is 8.
933		Q: Marcell and Beatrice are having a contest to see who can eat the most fruit roll-ups, so they unroll as many as they can find. Unfortunately, someone makes a mistake and Beatrice's was two roll-ups wide and 24 rolls up long while Marcell's was 3 roll-ups wide and 14 roll-ups long. If they both ate their entire amount, how many did they eat on average?
934		A:
935	<b>Dreedy</b>	Marcell's was 3 wide and 14 long. Beatrice's was 2 wide and 24 long. So Marcell's was $3 * 14 = 42$ rolls. Beatrice's was $2 * 24 = 48$ rolls. So Marcell ate 42 rolls and Beatrice ate 48 rolls. $42 + 48 = 90$ . The answer is 90.
936	<b>SLED</b>	Marcell's was 3 wide and 14 long. Beatrice's was 2 wide and 24 long. So Marcell's was $3 * 14 = 42$ rolls. Beatrice's was $2 * 24 = 48$ rolls. So Marcell ate 42 rolls and Beatrice ate 48 rolls. $42 + 48 = 90$ rolls. The answer is 90.
937	<b>NDAD</b>	Marcell's was 3 roll-ups wide and 14 roll-ups long. Beatrice's was 2 roll-ups wide and 24 roll-ups long. So Marcell's was $3 * 14 = 42$ roll-ups long. Beatrice's was $2 * 24 = 48$ roll-ups long. So Marcell's was 42 roll-ups long and Beatrice's was 48 roll-ups long. So they both ate $42 + 48 = 90$ roll-ups. $90 / 2 = 45$ . The answer is 45.

967  
968  
969  
970  
971