

# InEx: Hallucination Mitigation via Introspection and Cross-Modal Multi-Agent Collaboration

Zhongyu Yang<sup>1,2\*</sup>, Yingfang Yuan<sup>2\*</sup>, Xuanming Jiang<sup>1†</sup>, Baoyi An<sup>1</sup>, Wei Pang<sup>2†</sup>

<sup>1</sup>Xi'an Jiyun Technology Co., Ltd., Xi'an, China

<sup>2</sup>Bio-inspired Computing and Machine Learning Lab, Heriot-Watt University, Edinburgh, the United Kingdom  
{zy4028, y.yuan, w.pang}@hw.ac.uk, jiangxm24@163.com, anby2024@163.com

## Abstract

Hallucination remains a critical challenge in large language models (LLMs), hindering the development of reliable multimodal LLMs (MLLMs). However, existing solutions often rely on human intervention or underutilize the agent's ability to autonomously mitigate hallucination. To address these limitations, we draw inspiration from the way humans make reliable decisions in the real world. In particular, they begin with introspective reasoning to reduce uncertainty and form an initial judgment, then rely on external verification from diverse perspectives to reach a final decision. Motivated by this cognitive paradigm, we propose **InEx**, a training-free, multi-agent framework designed to autonomously mitigate hallucination. InEx introduces internal introspective reasoning, guided by entropy-based uncertainty estimation, to improve the reliability of the decision agent's reasoning process. The agent first generates a response, which is then iteratively verified and refined through external cross-modal multi-agent collaboration with the editing agent and self-reflection agents, further enhancing reliability and mitigating hallucination. Extensive experiments show that InEx consistently outperforms existing methods, achieving 4%-27% gains on general and hallucination benchmarks, and demonstrating strong robustness.

## Introduction

Large language models (LLMs) have enabled substantial progress in natural language processing, leading to the development of multimodal LLMs (MLLMs) capable of understanding both visual and textual information (Zhang et al. 2025b, 2024; Yang et al. 2025c). Despite their remarkable capabilities, MLLMs face the critical and unresolved challenge of hallucination. This phenomenon refers to generating responses that, while linguistically plausible, are factually inaccurate relative to the input image. Hallucination significantly compromises the reliability of MLLMs and presents a substantial barrier to their safe and effective deployment in real-world applications (Yang et al. 2025b).

Hallucination mitigation has been approached through three primary paradigms, each constrained by specific limitations that collectively motivate our work. Pre-processing

methods, such as fine-tuning or reinforcement learning from human feedback (Chen et al. 2025; Yu et al. 2024a, 2025), require substantial human intervention, including supervision and task-specific annotations, thus limiting scalability and adaptability. In-processing methods guide generation using hallucination-related heuristics or metrics (Farquhar et al. 2024; Zhang et al. 2025c). However, because these methods are tightly coupled with the model's internal reasoning and lack external verification, they remain vulnerable to meta-hallucination (i.e., meta-cognitive confidence in claims derived from incorrect knowledge) due to the model's inherent stereotype biases. Post-processing methods often employ an external verifier to detect or correct errors after generation (Madaan et al. 2023; Wu et al. 2025). In some cases, these verifiers require retraining, incurring additional computational overhead. Furthermore, we argue that treating in-processing and post-processing methods in isolation ultimately underutilizes the autonomous capabilities of agents for hallucination mitigation. In real-world decision-making, iterative reasoning and verification are inherently interconnected processes that mutually reinforce one another.

To address these challenges, we draw inspiration from the cognitive paradigm that humans make reliable decisions. This process typically begins with introspective internal reasoning, forming a preliminary judgment by reducing uncertainty. This decision is then verified and refined through external feedback from diverse perspectives, ultimately leading to a well-grounded decision. Motivated by this observation, we propose **InEx**, a training-free multi-agent framework designed to autonomously mitigate hallucination.

In the proposed InEx, a decision agent leverages unsupervised uncertainty estimation to perform internal introspective reasoning (**In**), proactively retrieving and analyzing visual information to generate an initial response without extra human intervention. To further mitigate hallucination, InEx introduces external cross-modal multi-agent collaboration (**Ex**) for iterative verification and refinement. In particular, the initial response from the decision agent undergoes repeated evaluation by self-reflection agents and refinement by the decision agent, guided by cross-modal evidence. This collaborative process of alternating verification and introspective refinement continues until the response achieves cross-modal consensus or meets a predefined stopping criterion, ensuring autonomy and reliability in response generation.

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

<sup>1\*</sup> Equal contribution

<sup>2†</sup> Corresponding Author

InEx unifies In and Ex into a cohesive process. Their interaction enables proactive uncertainty reduction and iterative cross-modal verification and refinement, with each component reinforcing the effectiveness of the other. Together, these modules enable fully autonomous hallucination mitigation, requiring neither human supervision nor model retraining. As a result, InEx consistently outperforms existing baselines across both general-purpose and hallucination benchmarks, while delivering well-calibrated and trustworthy responses.

Our main contributions can be summarized as follows:

- We propose **InEx**, a training-free multi-agent framework that mitigates hallucination via internal introspective reasoning and external cross-modal collaboration.
- We demonstrate that cross-modal consensus can effectively mitigate hallucination with the support of multi-modal, multi-perspective self-reflection agents.
- Extensive experiments show that InEx outperforms prior baselines by 3.8%–16.2% on general-purpose benchmarks and 6.5%–26.7% on hallucination benchmarks.

## Related Works

**Hallucination Mitigation in MLLMs:** Previous research follows three paradigms, all possessing inherent limitations: **(i) Data-centric tuning** leverages contrastive learning or reinforcement learning on curated adversarial prompts or hard negatives (Jiang et al. 2024; Liu et al. 2024a; Xiao et al. 2025). While effective on specific tasks, such methods incur high labeling and retraining costs, limiting scalability and generalization. **(ii) Post-hoc revision** modifies generated outputs using external verifiers, regenerators, or retrieval-augmented modules (Sahu, Sikka, and Divakaran 2024; Yin et al. 2024; Zhang et al. 2025a; Zhou et al. 2024). However, these methods are reactive rather than preventive, leading to latency, fragile pipelines, and ultimately failing to address the root causes of hallucination. **(iii) Decoding-time control** regulates token-level generation through logit calibration or uncertainty-based penalties (Chen et al. 2024; Deng, Chen, and Hooi 2024; Huo et al. 2025; Yang et al. 2025d). Despite their efficiency, these techniques often generate overconfident yet unreliable outputs due to shallow, modality-unaware reasoning and poor uncertainty modeling. Our **InEx** bridges this gap by unifying proactive, modality-aware uncertainty modeling with structured cross-modal verification.

**Self-reflection:** Recent studies have explored self-reflection mechanisms in LLMs, in which models iteratively refine their outputs through internal feedback loops to enhance reasoning and interpretability. These approaches can be categorized into three types: (1) self-evaluation within the textual modality (Shinn et al. 2023; Renze and Guven 2024); (2) feedback from an external trained textual critic (Gou et al. 2024; Yang et al. 2025a); and (3) integration of external knowledge sources such as Wikipedia or web search (Madaan et al. 2023; Asai et al. 2024). Although effective for text-only reasoning, these approaches lack visual grounding, limiting their applicability to multimodal contexts. To overcome this universal constraint, the proposed **InEx** extends self-reflection to the multimodal domain via structured collaboration among com-




Figure 1: ■ Hallucination ■ No Hallucination  
Uncertainty calibration across different methods, with each column showing uncertainty score distributions (top) and corresponding hallucination rates (bottom) across bins.

plementary visual and textual agents, enabling peer-review-style verification that enhances reasoning reliability.

## Preliminaries

**MLLM Architecture:** An MLLM  $f_\theta$ , parameterized by  $\theta$ , generates a response sequence  $\mathbf{y} = [y_1, \dots, y_T]$ , where each  $y_t$  is a text token, conditioned on an input image  $v$  and textual query  $\mathbf{x}$ . The image  $v$  is first processed by a vision encoder, then transformed into a set of visual tokens  $\mathbf{z} = [z_1, \dots, z_M]$  via a vision-language alignment module. These visual tokens, along with the query tokens from  $\mathbf{x}$ , are jointly fed into an LLM to generate the response  $\mathbf{y}$ . The LLM adopts a standard decoder-style architecture with  $L$  transformer layers. To predict token  $y_t \in \mathbf{y}$ , the hidden state is updated sequentially layer by layer. Given the hidden representation  $\mathbf{H}_t^{\ell-1}$  from layer  $\ell-1$ , the forward update at layer  $\ell$  is defined as follows:

$$\bar{\mathbf{H}}_t^\ell = \text{MHA}_\ell(\mathbf{H}_t^{\ell-1}) + \mathbf{H}_t^{\ell-1}, \mathbf{H}_t^\ell = \text{FFN}_\ell(\bar{\mathbf{H}}_t^\ell) + \bar{\mathbf{H}}_t^\ell, \quad (1)$$

where  $\text{MHA}_\ell(\cdot)$  and  $\text{FFN}_\ell(\cdot)$  denote the multi-head self-attention and feed-forward network at layer  $\ell$ , respectively. Each  $\text{MHA}_\ell$  consists of  $H$  attention heads as follows:

$$\text{MHA}_\ell(\mathbf{H}_t^{\ell-1}) = \text{Concat}(\text{Head}_{\ell,1}, \dots, \text{Head}_{\ell,H})\mathbf{W}_\ell^o, \quad (2)$$

where  $\mathbf{W}_\ell^o$  is a learnable parameter matrix at layer  $\ell$ .

A projection head  $\phi(\cdot)$  predicts the logits of the token  $y_t$ :  $\phi(\mathbf{H}_t^L) = \mathbf{H}_t^L \mathbf{W}_{\text{out}} + \mathbf{b}$ . Finally, the model produces the probability distribution over the vocabulary as:

$$p_\theta(y_t | v, \mathbf{x}, \mathbf{y}_{<t}) = \text{softmax}(\phi(\mathbf{H}_t^L))_{y_t}. \quad (3)$$

**Text-to-Visual Entropy Ratio (TVER):** A number of methods (Farquhar et al. 2024; Xiong et al. 2024; Wan et al. 2025) detect hallucination by uncovering patterns in the model’s internal reasoning process. In this work, we systematically investigate these methods and empirically find that TVER is the most effective, achieving the highest AUROC and the lowest ECE (Figure 1), thereby indicating superior calibration and a strong correlation with hallucination.

For each attention head at layer  $\ell$ , TVER first computes the raw attention distribution as follows:

$$\mathbf{a}_\ell = \text{softmax}\left(Q_\ell K_\ell^\top / \sqrt{d_k}\right). \quad (4)$$

The  $\mathbf{a}_\ell$  are partitioned based on textual indices  $\mathcal{T}$  and visual indices  $\mathcal{V}$ , and the entropy for each subset is computed as follows:

$$\text{Entropy}(\mathcal{T}||\mathcal{V}) = - \sum_{|\mathcal{T}||\mathcal{V}|} \mathbf{p}_{\ell,i} \log \mathbf{p}_{\ell,i}, \quad (5)$$

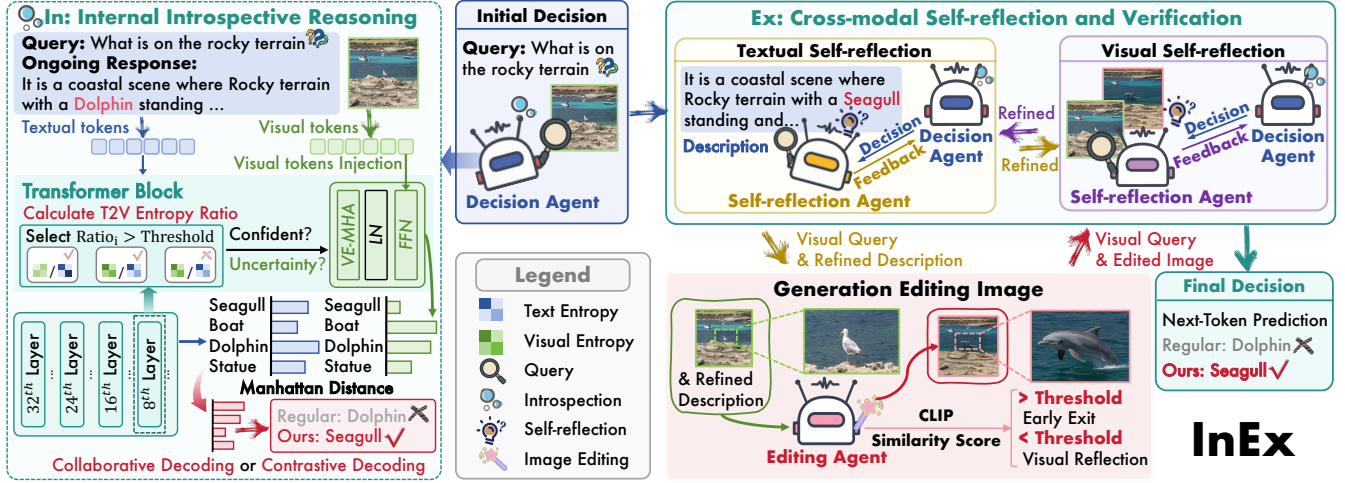


Figure 2: Overview of **InEx**. **In**: A decision agent initiates introspective reasoning guided by unsupervised uncertainty estimation, producing an initial response grounded in internal uncertainty signals. **Ex**: The response is then iteratively refined through alternating cross-modal verification and introspective updates, where self-reflection agents assess consistency with visual and textual evidence, and the decision agent continues to revise its output accordingly.

where  $i \in \mathcal{T}$  or  $i \in \mathcal{V}$ , and  $p_{\ell,i} = \text{softmax}(\mathbf{a}_{\ell,i})$ . The TVER at layer  $\ell$  for an attention head is defined as:

$$\text{TVER}_{\ell} = \frac{\text{Entropy}(\mathcal{T})}{\text{Entropy}(\mathcal{V})}. \quad (6)$$

As shown in Figure 1, a higher TVER score indicates a greater likelihood of hallucination, corresponding to increased entropy in  $\mathcal{T}$  and decreased entropy in  $\mathcal{V}$ . Since entropy reflects uncertainty, lower entropy in  $\mathcal{V}$  suggests confidence in the visual information. In such cases, elevated uncertainty in the textual modality may stem from misleading yet confident visual cues. To guide reasoning and mitigate hallucination, attention heads  $h \in H$  with TVER scores above the threshold  $\gamma_{\text{TVER}}$  are masked as follows:

$$\tilde{\mathbf{a}}_{\ell,h} = \begin{cases} 0, & \text{TVER}_{\ell} \geq \gamma_{\text{TVER}}, \\ \mathbf{a}_{\ell,h}, & \text{otherwise.} \end{cases} \quad (7)$$

Finally, we propose the Vision-Enhanced Multi-Head Attention (VE-MHA) as follows:

$$\text{VE-MHA}(\bar{\mathbf{H}}_t^{\ell}) = \text{Concat}(\tilde{\mathbf{a}}_{\ell,1}\bar{\mathbf{H}}_1^{\ell}, \dots, \tilde{\mathbf{a}}_{\ell,H}\bar{\mathbf{H}}_H^{\ell})\mathbf{W}_{\ell}^o, \quad (8)$$

where  $\bar{\mathbf{H}}_t^{\ell}$  denotes the output of MHA( $\cdot$ ) as defined in Eq. 1, and  $[\bar{\mathbf{H}}_1^{\ell}, \dots, \bar{\mathbf{H}}_H^{\ell}]$  represents the outputs from individual original attention heads. VE-MHA is applied after the original multi-head attention layer to refine its output through Self-Introspective Visual Augmentation, Enhanced Logits, and Decoding, which are discussed in the following section.

## Methodology

Our proposed **InEx** framework mitigates hallucination from two complementary perspectives: **In**, which enhances internal reasoning via TVER-guided introspection, and **Ex**, which conducts external cross-modal multi-agent collaboration.

## In: Internal Introspective Reasoning

As shown in Figure 2 and Algorithm 1, the MLLM decision agent  $\mathcal{A}_d$  generates a response  $\mathbf{y}$  conditioned on both visual and textual inputs. In real-world decision-making scenarios, humans typically engage in internal reasoning. This process is supported by introspection, which is guided by empirically validated cues to avoid errors and reduce uncertainty. Inspired by this cognitive mechanism, we introduce TVER as an introspective signal for  $\mathcal{A}_d$ , triggering an internal refinement process to improve decision reliability. When  $\mathcal{A}_d$  encounters ambiguous or underdetermined contexts (i.e., when  $\text{TVER}_{\ell} \geq \gamma_{\text{TVER}}$ ), it engages in self-introspective internal reasoning composed of three components: self-introspective augmentation, self-introspective enhanced logits, and self-introspective decoding. These mechanisms work together to reinforce relevant visual information, recalibrate confidence levels, and refine the response generation process.

For simplicity, in the following paragraphs, we omit the subscript  $t$  for hidden representations  $\mathbf{H}$ . By default, each  $\mathbf{H}$  is assumed to correspond to a token  $y_t$ .

**Self-Introspective Visual Augmentation:** As indicated by  $\text{TVER}_{\ell} \geq \gamma_{\text{TVER}}$ ,  $\mathcal{A}_d$  proactively retrieves relevant visual information to augment its reasoning trajectory through a memory-style modulation of the feed-forward network (FFN) at layer  $\ell$ . Given a hidden state  $\bar{\mathbf{H}}^{\ell}$  as shown in Eq. 1 and visual tokens  $\mathbf{z} = \{\mathbf{z}_1, \dots, \mathbf{z}_{|\mathbf{z}|}\}$ , we at first define a similarity-weighted retrieval:  $\Delta(\mathbf{z} | \bar{\mathbf{H}}^{\ell}) = \sum_{i=1}^{|\mathbf{z}|} \sigma(\langle \bar{\mathbf{H}}^{\ell}, \mathbf{z}_i \rangle) \cdot \mathbf{z}_i$ , where  $\sigma(\cdot)$  is an activation function (e.g., ReLU or SiLU). To incorporate retrieval into the reasoning process and augment it, the FFN( $\cdot$ ) in Eq. 1 is redefined as an enhanced FFN( $\cdot$ ):  $\text{FFN}_{\ell}^{\text{introspect}}(\bar{\mathbf{H}}^{\ell}) = \alpha \Delta(\mathbf{z} | \bar{\mathbf{H}}^{\ell}) + (1 - \alpha) \text{FFN}_{\ell}(\bar{\mathbf{H}}^{\ell})$ , where  $\alpha \in [0, 1]$  reflects the introspection strength. Through this process, the model introspectively “revisits” its visual grounding to recalibrate its

internal semantic space, thereby reinforcing alignment between modalities under uncertainty.

**Self-Introspective Enhanced Logits:** Furthermore, since the logits play a key role in determining the generated response, we incorporate VE-MHA at the final Transformer layer  $L$  to enhance the reliability of the information used in computing the logits. This is achieved by filtering out uncertain hidden representations while retaining the visual evidence introduced through  $\text{FFN}_\ell^{\text{introspect}}$ . The computation of the hidden representation used for generating the logits is defined by Eqs. 9-11. The masked heads are determined at layer  $\ell$ . Given the introspective representation  $\hat{\mathbf{H}}^L$ , the corresponding enhanced logits are defined as  $\phi(\hat{\mathbf{H}}^L)$ .

$$\bar{\mathbf{H}}^L = \text{MHA}_L(\mathbf{H}^{L-1}) + \mathbf{H}^{L-1}, \quad (9)$$

$$\mathbf{H}^L = \text{VE-MHA}_L(\bar{\mathbf{H}}^L) + \bar{\mathbf{H}}^L, \quad (10)$$

$$\hat{\mathbf{H}}^L = \text{FFN}_L(\mathbf{H}^L) + \mathbf{H}^L. \quad (11)$$

**Self-Introspective Decoding:** In the real world, when new evidence contradicts prior conclusions, humans typically compare the alternatives and adopt an appropriate strategy to reach a final decision. Inspired by this behavior, we propose self-introspective decoding to reconcile the original and enhanced predictions. Specifically, we compute the Manhattan distance between the two logits for each generated text token  $y_t$  using  $d_t = \|\phi(\mathbf{H}^L) - \phi(\hat{\mathbf{H}}^L)\|_T$ . To improve efficiency, the distance  $d_t$  is typically calculated over the top- $k$  logits (e.g.,  $k = 20$ ). The final fused logits are obtained as:

$$\text{Fused logits} = \begin{cases} \phi(\mathbf{H}^L) + \alpha_1 \phi(\hat{\mathbf{H}}^L), & \text{if } d_t < \gamma_d, \\ (1 + \alpha_2)\phi(\mathbf{H}^L) - \alpha_2 \phi(\hat{\mathbf{H}}^L), & \text{if } d_t \geq \gamma_d, \end{cases} \quad (12)$$

where  $\gamma_d$  is a confidence threshold, and  $\alpha_1, \alpha_2$  are weight factors. When  $d_t < \gamma_d$ , the original and enhanced logits are similar, indicating agreement. We apply a collaborative mode by combining both representations to reinforce the decision. When  $d_t \geq \gamma_d$ , the predictions diverge, signaling potential conflict. We then use a contrastive mode, emphasizing the original logits and using the enhanced ones only to calibrate them. Finally, the fused logits are used with softmax to generate  $y_t$  in an autoregressive manner.

## Ex: Cross-modal Multi-Agent Collaboration

Inspired by real-world decision-making, we introduce **Ex**, an external cross-modal multi-agent collaboration mechanism for alternating verification and refinement (see Algorithm 1). The process begins with the decision agent  $\mathcal{A}_d$  generating an initial response  $\mathbf{y}_0$  based on the input image  $v$  and text query  $\mathbf{x}$  (Line 2). The textual self-reflection agent  $\mathcal{A}_t$  verifies whether the response  $\mathbf{y}_0$  is supported by the dense caption  $\mathbf{c}$ , serving as a validation filter. If it is supported, the process returns  $\mathbf{y} \leftarrow \mathbf{y}_0$ . Otherwise,  $\mathcal{A}_t$  sends textual feedback  $\mathbf{f}_t$  to  $\mathcal{A}_d$ , prompting it to refine the response through introspection and initiate cross-modal verification. To mimic a rigorous human self-reflection process,  $\mathcal{A}_t$  performs verification and generates feedback through multi-perspective reflection and ensemble aggregation. The multi-perspective approach enables  $\mathcal{A}_t$  to focus on different aspects mentioned in the text, such

---

### Algorithm 1: InEx in Visual Question Answering

---

**Input:** image  $v$ , text query  $\mathbf{x}$ , decision agent  $\mathcal{A}_d$ , visual self-reflection agent  $\mathcal{A}_v$ , textual self-reflection agent  $\mathcal{A}_t$ , editing agent  $\mathcal{A}_i$

**Output:** Answer  $\mathbf{y}$

```

1 Initialisation:  $\mathbf{c}$  // dense caption for  $v$ 
2  $\mathbf{y}_0 = \mathcal{A}_d(v, \mathbf{x})$  // initial generation
3 if  $\mathcal{A}_t(\mathbf{c}, \mathbf{y}_0) == \text{True}$  then
4    $\mathbf{y} \leftarrow \mathbf{y}_0$ 
5   return  $\mathbf{y}$ 
6  $\mathbf{f}_t \leftarrow \mathcal{A}_t(\mathbf{c}, \mathbf{y}_i)$  // feedback from text
7 for  $i = 1$  to  $I$  do
8    $\mathbf{y}_i \leftarrow \mathcal{A}_d(v, \mathbf{x}, \mathbf{f}_t)$  // decision from text
9    $v' \leftarrow \mathcal{A}_i(\mathbf{y}_i, v)$  // image editing
10  if  $\mathcal{A}_v(v, v') == \text{True}$  then
11     $\mathbf{y} \leftarrow \mathbf{y}_i$ 
12    return  $\mathbf{y}$  // verification from vision
13   $\mathbf{f}_v \leftarrow \mathcal{A}_v(v, v')$  // feedback from vision
14   $\mathbf{y}_i \leftarrow \mathcal{A}_d(v, \mathbf{x}, \mathbf{f}_v)$  // decision from vision
15  if  $\mathcal{A}_t(\mathbf{c}, \mathbf{y}_i) == \text{True}$  then
16     $\mathbf{y} \leftarrow \mathbf{y}_i$ 
17    return  $\mathbf{y}$  // verification from text
18   $\mathbf{f}_t \leftarrow \mathcal{A}_t(\mathbf{c}, \mathbf{y}_i)$  // feedback from text
19   $t \leftarrow t + 1$ 

```

---

as actions, objects, and colors. The ensemble mechanism repeats the verification and feedback generation multiple times. The verification results and their supporting reasoning evidence are then aggregated through a self-aggregation process to form  $\tilde{\mathbf{y}}$ . If  $\tilde{\mathbf{y}} \neq \mathbf{y}'$ , the aggregated evidence is returned as  $\mathbf{f}_t$ . With the additional information  $\mathbf{f}_t$ ,  $\mathcal{A}_d$  generates a revised response  $\mathbf{y}_i$  (Line 8), which is then verified through the visual modality. In Line 9,  $\mathbf{y}_i$  and the original image  $v$  are input to the image editing agent  $\mathcal{A}_i$ , which performs image augmentation to produce a new image  $v'$ . In Line 10, the visual self-reflection agent  $\mathcal{A}_v$  compares  $v'$  with the original image  $v$ . If the similarity exceeds the threshold  $\gamma_{\text{CLIP}}$ ,  $\mathbf{y} \leftarrow \mathbf{y}_i$  is returned. Otherwise, in Line 13,  $\mathcal{A}_v$  generates visual feedback  $\mathbf{f}_v$  and sends it to  $\mathcal{A}_d$ .

The decision agent then revises its response based on visual feedback  $\mathbf{f}_v$ , which is subsequently verified by  $\mathcal{A}_t$ . This process repeats until a final response  $\mathbf{y}$  is returned or a preset iteration limit  $I$  is reached. Lines 12 and 17 together define the cross-modal consensus.

The design of Ex is guided by four core principles: (i) We assume that a response can be verified through a different modality (Lines 10 and 15), serving as a safeguard against hallucination. This mirrors how humans validate information by consulting multiple independent sources or by seeking external verification from individuals with different expertise and backgrounds. (ii) During text-based verification (Line 15), the decision must be consistent with the dense caption set  $\mathbf{c}$ , serving as a grounded textual representation of the visual content. This consistency ensures that the response is grounded in observable evidence, rather than being fabricated.

Evaluation	Methods	Random		Popular		Adversarial		Average	
		Accuracy $\uparrow$	F1-score $\uparrow$	Accuracy $\uparrow$	F1-score $\uparrow$	Accuracy $\uparrow$	F1-score $\uparrow$	Accuracy $\uparrow$	F1-score $\uparrow$
MSCOCO	—	83.49 $\uparrow$ 0.0	82.28 $\uparrow$ 0.0	79.98 $\uparrow$ 0.0	79.34 $\uparrow$ 0.0	76.03 $\uparrow$ 0.0	76.26 $\uparrow$ 0.0	79.83 $\uparrow$ 0.0	79.29 $\uparrow$ 0.0
	OPERA (Huang et al. 2024)	87.63 $\uparrow$ 4.2	86.45 $\uparrow$ 4.2	83.91 $\uparrow$ 4.0	83.50 $\uparrow$ 4.1	80.88 $\uparrow$ 4.9	80.69 $\uparrow$ 4.4	84.14 $\uparrow$ 4.3	83.55 $\uparrow$ 4.3
	ICD (Wang et al. 2024)	84.97 $\uparrow$ 1.5	83.23 $\uparrow$ 1.0	82.83 $\uparrow$ 2.9	81.39 $\uparrow$ 2.1	81.17 $\uparrow$ 5.1	80.06 $\uparrow$ 3.7	82.97 $\uparrow$ 3.1	81.56 $\uparrow$ 2.3
	VCD (Leng et al. 2024)	86.84 $\uparrow$ 3.4	86.83 $\uparrow$ 4.6	82.65 $\uparrow$ 2.7	83.37 $\uparrow$ 4.0	78.31 $\uparrow$ 2.3	79.28 $\uparrow$ 3.0	82.60 $\uparrow$ 2.7	83.16 $\uparrow$ 3.9
	<b>InEx (ours)</b>	<b>89.48</b> $\uparrow$ 6.0	<b>88.24</b> $\uparrow$ 6.0	<b>88.77</b> $\uparrow$ 8.7	<b>87.72</b> $\uparrow$ 8.4	<b>85.97</b> $\uparrow$ 9.9	<b>85.68</b> $\uparrow$ 9.4	<b>88.73</b> $\uparrow$ 8.9	<b>86.48</b> $\uparrow$ 7.2
A-OKVQA	—	83.45 $\uparrow$ 0.0	82.56 $\uparrow$ 0.0	79.90 $\uparrow$ 0.0	79.59 $\uparrow$ 0.0	74.04 $\uparrow$ 0.0	75.15 $\uparrow$ 0.0	79.13 $\uparrow$ 0.0	79.10 $\uparrow$ 0.0
	OPERA (Huang et al. 2024)	88.32 $\uparrow$ 4.9	87.64 $\uparrow$ 5.1	85.59 $\uparrow$ 5.7	84.64 $\uparrow$ 5.1	79.17 $\uparrow$ 5.1	79.97 $\uparrow$ 4.8	84.36 $\uparrow$ 5.2	84.08 $\uparrow$ 5.0
	ICD (Wang et al. 2024)	85.41 $\uparrow$ 2.0	84.99 $\uparrow$ 2.6	81.79 $\uparrow$ 1.9	81.98 $\uparrow$ 2.4	77.73 $\uparrow$ 3.7	78.79 $\uparrow$ 3.2	81.64 $\uparrow$ 2.5	82.00 $\uparrow$ 2.8
	VCD (Leng et al. 2024)	86.15 $\uparrow$ 2.7	86.34 $\uparrow$ 3.8	81.85 $\uparrow$ 2.0	82.82 $\uparrow$ 3.2	74.97 $\uparrow$ 0.9	77.73 $\uparrow$ 2.6	80.99 $\uparrow$ 1.9	82.30 $\uparrow$ 3.2
	<b>InEx (ours)</b>	<b>92.63</b> $\uparrow$ 9.2	<b>91.18</b> $\uparrow$ 8.6	<b>88.41</b> $\uparrow$ 8.5	<b>87.99</b> $\uparrow$ 8.4	<b>80.91</b> $\uparrow$ 6.9	<b>82.86</b> $\uparrow$ 7.7	<b>87.32</b> $\uparrow$ 8.2	<b>87.34</b> $\uparrow$ 8.3
GQA	—	83.73 $\uparrow$ 0.0	82.95 $\uparrow$ 0.0	78.17 $\uparrow$ 0.0	78.37 $\uparrow$ 0.0	75.08 $\uparrow$ 0.0	76.06 $\uparrow$ 0.0	78.99 $\uparrow$ 0.0	79.13 $\uparrow$ 0.0
	OPERA (Huang et al. 2024)	88.27 $\uparrow$ 4.5	87.52 $\uparrow$ 4.6	83.07 $\uparrow$ 4.9	82.93 $\uparrow$ 4.6	80.77 $\uparrow$ 5.7	81.05 $\uparrow$ 5.0	84.04 $\uparrow$ 5.1	83.83 $\uparrow$ 4.7
	ICD (Wang et al. 2024)	84.90 $\uparrow$ 1.2	84.22 $\uparrow$ 1.3	78.37 $\uparrow$ 0.2	78.81 $\uparrow$ 0.4	75.97 $\uparrow$ 0.9	76.93 $\uparrow$ 0.9	79.75 $\uparrow$ 0.8	79.99 $\uparrow$ 0.9
	VCD (Leng et al. 2024)	86.65 $\uparrow$ 2.9	86.99 $\uparrow$ 4.0	80.73 $\uparrow$ 2.6	82.24 $\uparrow$ 3.9	76.09 $\uparrow$ 1.0	78.78 $\uparrow$ 2.7	81.16 $\uparrow$ 2.2	82.67 $\uparrow$ 3.5
	<b>InEx (ours)</b>	<b>90.12</b> $\uparrow$ 6.4	<b>90.32</b> $\uparrow$ 7.4	<b>86.37</b> $\uparrow$ 8.2	<b>85.98</b> $\uparrow$ 7.6	<b>82.91</b> $\uparrow$ 7.8	<b>83.82</b> $\uparrow$ 7.8	<b>86.72</b> $\uparrow$ 7.7	<b>86.81</b> $\uparrow$ 7.7

Table 1: Performance on POPE benchmark using LLaVA-1.5-7B. The best results are in **blue**. We report accuracy and F1-score under three settings, e.g., Random, Popular, Adversarial, and Average, to show the robustness of the different methods directly.

(iii) For decisions inferred from text, we assess their validity by editing the image based on the generated output (Line 10). If the edited image remains coherent with the original, it confirms that the decision is anchored in visual reality rather than speculative reasoning. (iv) Cross-modal feedback is incorporated into  $\mathcal{A}_d$  to enhance its internal introspective reasoning, enabling it to generate a more reliable response, analogous to how humans iteratively refine their decisions by drawing on diverse sources of information.

In summary, **InEx** unifies internal introspective reasoning with external cross-modal collaboration to jointly verify and improve responses in an iterative loop, jointly enhancing response reliability and mitigating hallucination.

## Theoretical Analysis

To further illustrate the principles behind **InEx** and to comprehensively justify its effectiveness in mitigating hallucinations and ensuring robust performance, we present three theorems below, where  $\mathbf{H}$  denote the hidden states derived from the final original Transformer layer,  $\hat{\mathbf{H}}$  represent the updated states obtained after injecting visual evidence  $\mathbf{z}$  through Self-Introspective Visual Augmentation and Enhanced Logits.

**Theorem 1. (Mutual Information (MI) Increase)** *MI quantifies the amount of information one variable contains about another. InEx increases the MI between  $\hat{\mathbf{H}}$  and  $\mathbf{z}$ :*

$$I(\hat{\mathbf{H}}; \mathbf{z}) \geq I(\mathbf{H}; \mathbf{z}). \quad (13)$$

**Theorem 2. (Conditional Entropy Reduction)** *Let  $\mathbf{y}$  be the target output dependent on hidden states. An increase in MI leads to a corresponding reduction in the conditional entropy of the predicted output:*

$$E(\mathbf{y} | \hat{\mathbf{H}}) \leq E(\mathbf{y} | \mathbf{H}). \quad (14)$$

**Theorem 3. (IB Objective Improvement)** *Under the Information Bottleneck (IB) framework, InEx with Self-Introspective Visual Augmentation and Enhanced Logits optimizes the objective function by reducing the IB loss:*

$$\mathcal{L}(\hat{\mathbf{H}}) \leq \mathcal{L}(\mathbf{H}), \quad (15)$$

where  $\mathcal{L}(\mathbf{H}) = I(\mathbf{H}; \mathbf{x}) - \beta I(\mathbf{H}; \mathbf{y})$ ,  $\mathbf{x}$  denotes the text query embedding and  $\beta$  is a trade-off parameter.

We ground our theoretical framework in the Information Bottleneck theory, specifically the Data Processing Inequality and the contraction properties of stochastic mappings (Achille and Soatto 2018). Building on these validated principles, the proposed InEx framework is designed to enhance mutual information between internal representations and visual input while minimizing latent-state uncertainty.

Through different multi-agent collaboration, InEx simultaneously maximizes mutual information with visual input and minimizes latent-state uncertainty, thereby mitigating hallucination and yielding more reliable representations.

## Experiments

### Experiment Setup

**Datasets.** To rigorously assess the effectiveness of InEx, we conduct a comprehensive set of experiments across a diverse set of benchmarks, as detailed below:

- **Hallucination Benchmarks:** POPE benchmark (Li et al. 2023), CHAIR (Rohrbach et al. 2018), HallusionBench(H-Bench) (Guan et al. 2024).
- **General-purpose Benchmarks:** VizWiz-VQA (Gurari et al. 2018), MME (Fu et al. 2023), MMBench (Liu et al. 2024b), MM-Vet (Yu et al. 2024b), LLaVA-Bench (L-Bench) (Liu et al. 2023).

**Baselines and Models:** InEx is evaluated against three representative hallucination mitigation baseline methods: OPERA (Huang et al. 2024), ICD (Wang et al. 2024), and VCD (Leng et al. 2024). The evaluation is conducted using Qwen-VL (Bai et al. 2023), LLaVA-1.5 (Liu et al. 2023), and GLM-4V (GLM et al. 2024), covering 7B to 10B scales.

**Implementation Details:** We set  $\gamma_{\text{TVER}} = 0.55$ ,  $\gamma_d = 0.2$ ,  $\gamma_{\text{CLIP}} = 0.9$ ,  $\alpha_1 = 1$ , and  $\alpha_2 = 1$ . Self-reflection uses temperature-controlled decoding with a value of 0.7, combined with ensemble aggregation across three outputs and up to four specialized perspectives. The editing agent is implemented via IC-Edit (Zhang et al. 2025d). All experiments are conducted on NVIDIA H100 GPUs with FP16 precision.




Figure 3: InEx Parameter Analysis on MMBench using LLaVA-1.5-7B.

Methods	H-Bench		L-Bench		CHAIR					
	Accuracy ↑	Accuracy ↑	CHAIR <sub>S</sub> ↓	CHAIR <sub>T</sub> ↓	Recall ↑					
—	41.5	↑0.0	63.4	↑0.0	50.0	↑0.0	15.4	↑0.0	77.1	↑0.0
OPERA	41.2	↓0.3	59.8	↓3.6	47.9	↓2.1	14.5	↓0.9	76.8	↓0.3
ICD	38.2	↓3.3	49.8	↓13.	56.2	↑6.2	16.3	↑0.9	76.3	↓0.8
VCD	41.1	↓0.4	59.1	↓4.3	48.6	↓1.4	14.9	↓0.5	77.3	↑0.2
<b>InEx</b>	<b>44.2</b>	<b>↑2.7</b>	<b>66.5</b>	<b>↑3.1</b>	<b>45.1</b>	<b>↓4.9</b>	<b>13.3</b>	<b>↓4.1</b>	<b>83.2</b>	<b>↑6.1</b>

Table 2: Performance evaluation on H-Bench, L-Bench, and CHAIR using LLaVa-1.5-7B.

In	TEXT	VIS	POPE	CHAIR	VizWiz	MME	MM-Vet	L-Bench
✓	✗	✗	79.83	77.10	50.00	643.3	31.10	64.80
✓	✗	✗	86.43	79.40	51.30	648.5	32.60	65.40
✗	✓	✗	83.20	78.20	45.90	641.5	32.10	64.90
✗	✗	✓	85.20	80.20	49.70	645.5	31.70	65.30
✓	✓	✗	86.39	81.90	51.95	649.3	33.20	65.20
✓	✗	✓	87.77	82.50	52.25	645.8	35.10	66.10
✗	✓	✓	86.93	82.40	51.90	650.5	33.00	65.70
✓	✓	✓	<b>88.73</b>	<b>83.20</b>	<b>53.80</b>	<b>653.8</b>	<b>36.00</b>	<b>66.50</b>

Table 3: Ablation studies on LLaVA-1.5-7B.

### Results on Hallucination Benchmarks

Table 1 shows results on POPE, where InEx consistently outperforms all baselines across three sub-datasets and settings (random, popular, adversarial). On the MSCOCO subset, it boosts LLaVA-1.5-7B accuracy from 79.83 to 88.73(+ 8.9%). InEx also achieves the best results on HallusionBench and CHAIR across multiple metrics, as shown in Table 2.

### Results on General-purpose Benchmarks

We further evaluate InEx on general-purpose benchmarks. As shown in Table 2, InEx is the only method applied to LLaVA-1.5 on LLaVA-Bench, improving accuracy from 63.4% to 66.5%. In Table 4, InEx consistently outperforms existing methods for LLaVA-1.5, Qwen-VL, and GLM-4V across all datasets and metrics. For LLaVA-1.5, it achieves the highest score on MME-Hall (673.3), with strong results on object-level (Existence: 199.3, Count: 157.7) and attribute-level (Position: 137.0, Color: 179.3) metrics. It also leads on MM-Vet (36.00), VizWiz (53.80), and MMBench (67.17). InEx also demonstrates notable improvements for Qwen-VL and GLM-4V, particularly on MME-Hall, achieving scores of 677.3 (+59) and 732.3 (+35), respectively. These results highlight InEx’s robust capability in object and attribute understanding.




Figure 4: We conduct an ablation study to analyze the effect of 4 different image editing models using LLaVa-1.5-7B.

### Ablation Studies

We conduct ablation studies to analyze the individual contributions of In and Ex on seven datasets. For Ex, we evaluate both the textual (TEXT) and visual (VIS) self-reflection components. Table 3 shows that each component contributes to performance gains when used individually. However, the best performance is achieved when all three are combined, confirming the complementary effects of introspective reasoning and cross-modal verification in mitigating hallucination.

### Different Image Editing Model

Furthermore, the image editing model is a key component in the design of external verification. To evaluate its impact, we experiment with four image editing models with different architectures (diffusion-based and FLUX-based). As shown in Figure 4, InEx with In-Context Edit achieves the best overall performance. Nevertheless, InEx consistently improves upon the original model across all editing variants.

### Parameter Analysis

**Agents:** For the text self-reflection agent, we evaluate two key parameters: perspective, which denotes the number of different viewpoints used to verify a decision and generate feedback, and ensemble size, which indicates how many times the process is repeated. Both parameters lead to stable performance gains up to a value of four, effectively capturing diverse textual evidence through repetition. For the external collaboration process, we analyze the effect of the maximum number of iterations and observe that performance stabilizes when this value reaches four. This suggests that additional

Models	Methods	MME-Hall		Object-Level		Attribute-Level		MM-Vet	Vizwiz	MMBench
		Total $\uparrow$	Existence $\uparrow$	Count $\uparrow$	Position $\uparrow$	Color $\uparrow$	Accuracy $\uparrow$	Accuracy $\uparrow$	Accuracy $\uparrow$	
LLaVA-1.5-7B	—	643.3 $\uparrow$ 0.0	190.0 $\uparrow$ 0.0	155.0 $\uparrow$ 0.0	128.3 $\uparrow$ 0.0	170.0 $\uparrow$ 0.0	31.10 $\uparrow$ 0.0	50.00 $\uparrow$ 0.0	62.80 $\uparrow$ 0.0	
	OPERA (Huang et al. 2024)	610.0 $\downarrow$ 33.	195.0 $\uparrow$ 5.0	128.3 $\downarrow$ 26.	121.7 $\downarrow$ 6.6	165.0 $\downarrow$ 5.0	32.00 $\uparrow$ 0.9	50.76 $\uparrow$ 0.8	62.80 $\uparrow$ 0.0	
	ICD (Wang et al. 2024)	583.3 $\downarrow$ 60.	185.0 $\downarrow$ 5.0	130.0 $\downarrow$ 25.	121.7 $\downarrow$ 6.6	146.7 $\downarrow$ 23.	25.90 $\downarrow$ 5.2	37.62 $\downarrow$ 12.	39.78 $\downarrow$ 23.	
	VCD (Leng et al. 2024)	648.3 $\uparrow$ 5.0	190.0 $\uparrow$ 0.0	155.0 $\uparrow$ 0.0	133.3 $\uparrow$ 5.0	170.0 $\uparrow$ 0.0	30.20 $\downarrow$ 0.9	44.90 $\downarrow$ 5.1	54.21 $\downarrow$ 8.6	
	<b>InEx (Ours)</b>	<b>673.3</b> $\uparrow$ 30.	<b>199.3</b> $\uparrow$ 9.3	<b>157.7</b> $\uparrow$ 2.7	<b>137.0</b> $\uparrow$ 8.7	<b>179.3</b> $\uparrow$ 9.3	<b>36.00</b> $\uparrow$ 4.9	<b>53.80</b> $\uparrow$ 3.8	<b>67.17</b> $\uparrow$ 4.4	
Qwen-VL-10B	—	618.3 $\uparrow$ 0.0	175.0 $\uparrow$ 0.0	140.0 $\uparrow$ 0.0	123.3 $\uparrow$ 0.0	180.0 $\uparrow$ 0.0	49.00 $\uparrow$ 0.0	66.05 $\uparrow$ 0.0	56.53 $\uparrow$ 0.0	
	ICD (Wang et al. 2024)	616.7 $\downarrow$ 1.7	170.0 $\downarrow$ 5.0	138.3 $\downarrow$ 1.7	148.3 $\uparrow$ 25.	160.0 $\downarrow$ 20.	31.70 $\downarrow$ 17.	29.37 $\downarrow$ 36.	13.32 $\downarrow$ 43.	
	VCD (Leng et al. 2024)	648.3 $\uparrow$ 30.	175.0 $\uparrow$ 0.0	130.0 $\downarrow$ 10.	<b>153.3</b> $\uparrow$ 30.	190.0 $\uparrow$ 10.	34.60 $\downarrow$ 14.	34.54 $\downarrow$ 31.	39.18 $\downarrow$ 17.	
	<b>InEx (Ours)</b>	<b>677.3</b> $\uparrow$ 59.	<b>189.0</b> $\uparrow$ 13.	<b>151.0</b> $\uparrow$ 11.	143.3 $\uparrow$ 20.	<b>194.0</b> $\uparrow$ 14.	<b>52.30</b> $\uparrow$ 3.3	<b>68.56</b> $\uparrow$ 2.5	<b>59.74</b> $\uparrow$ 3.2	
GLM-4V-9B	—	697.3 $\uparrow$ 0.0	195.0 $\uparrow$ 0.0	165.0 $\uparrow$ 0.0	153.3 $\uparrow$ 0.0	184.0 $\uparrow$ 0.0	63.40 $\uparrow$ 0.0	57.39 $\uparrow$ 0.0	82.39 $\uparrow$ 0.0	
	ICD (Wang et al. 2024)	680.0 $\downarrow$ 17.	185.0 $\downarrow$ 10.	160.0 $\downarrow$ 5.0	170.0 $\uparrow$ 16.	165.0 $\downarrow$ 19.	59.40 $\downarrow$ 4.0	48.07 $\downarrow$ 9.4	78.01 $\downarrow$ 4.4	
	VCD (Leng et al. 2024)	714.0 $\uparrow$ 16.	195.0 $\uparrow$ 0.0	173.3 $\uparrow$ 8.3	160.7 $\uparrow$ 7.4	<b>185.0</b> $\uparrow$ 1.0	57.70 $\downarrow$ 5.7	50.01 $\downarrow$ 7.4	80.58 $\downarrow$ 1.8	
	<b>InEx (Ours)</b>	<b>732.3</b> $\uparrow$ 35.	<b>199.0</b> $\uparrow$ 4.0	<b>175.0</b> $\uparrow$ 15.	<b>173.3</b> $\uparrow$ 20.	<b>185.0</b> $\uparrow$ 1.0	<b>68.80</b> $\uparrow$ 5.4	<b>60.36</b> $\uparrow$ 3.0	<b>85.65</b> $\uparrow$ 3.3	

Table 4: Performance evaluation on MME Hallucination subset, MM-Vet, Vizwiz, and MMBench, where OPERA does not support Qwen-VL and GLM-4V. The best results are in blue.




Figure 5: InEx Parameter Analysis.

iterations yield diminishing returns, indicating that the verification mechanism has reached its effective capacity. All corresponding results are shown in Figure 3 (a). Figure 3 (b) shows the impact of image and text guidance scales, two key parameters of the image editing agent. Image guidance peaks at a scale of 7.5, suggesting that moderate constraints help preserve visual cues, while higher values introduce artifacts that reduce alignment. In contrast, text guidance improves performance steadily up to 10, where stronger semantic signals better align generated images with intended content. Furthermore, Figure 5 (left) illustrates that increasing the number of inference steps for the editing agent improves accuracy up to 100 steps by reducing noise and enhancing clarity. Beyond this point, however, performance gains taper off while computational cost continues to increase.

**Threshold:** InEx includes three thresholds:  $\gamma_{TVER}$ ,  $\gamma_d$ , and  $\gamma_{CLIP}$ . As shown in Figure 5 (right), increasing  $\gamma_{TVER}$  results in decreased accuracy, highlighting the importance of TVER in triggering introspective internal reasoning. A lower  $\gamma_{TVER}$  makes the system more sensitive to ambiguity and more likely to activate introspection. In contrast, increasing  $\gamma_{CLIP}$  enhances accuracy by strengthening the verification capacity of the visual self-reflection agent, further highlighting the benefits of cross-modal alignment. The final threshold,  $\gamma_d$ , determines the decoding mode. Accuracy fluctuates as  $\gamma_d$  changes, indicating that collaborative and contrastive decoding modes can complement each other effectively.

**Layer selection:** Self-introspective visual augmentation requires selecting a Transformer layer to inject visual features. Figure 3 (c) compares fixed-layer selection with dynamic selection informed by the TVER value. The dynamic approach

Method	Mean (%)	StdDev	$t$ -test $p$	Wilcoxon $p$
OPERA	84.14	0.62	$8.72 \times 10^{-25}$	$9.54 \times 10^{-7}$
ICD	82.97	0.41	$2.85 \times 10^{-37}$	$9.54 \times 10^{-7}$
VCD	82.60	0.54	$7.95 \times 10^{-30}$	$9.54 \times 10^{-7}$
<b>InEx</b>	<b>88.73</b>	<b>0.33</b>	-	-

Table 5: Statistical comparison on POPE benchmark. **Std-Dev** denotes standard deviation. Paired one-sided  $t$ -test and Wilcoxon signed-rank test  $p$ -values are reported ( $\alpha = 5\%$ ).

outperforms fixed-layer selection, demonstrating the benefit of adapting injection depth based on reasoning uncertainty.

Notably, InEx consistently enhances performance across all parameter configurations, as demonstrated in Figure 3.

**Statistical Significance Study:** We conduct 20 independent runs with different random seeds and compare our method against baseline models under identical settings. To assess statistical significance, we perform paired one-sided  $t$ -tests and Wilcoxon signed-rank tests. The null hypothesis ( $H_0$ ) states that InEx performs equally or worse than the baseline, while the alternative hypothesis ( $H_1$ ) states that our method performs better. As indicated in Table 5, our method achieves the highest mean score ( $88.73 \pm 0.33$ ), compared to OPERA ( $84.14 \pm 0.62$ ), ICD ( $82.97 \pm 0.41$ ), and VCD ( $82.60 \pm 0.54$ ). All obtained  $p$ -values are significantly below the threshold  $\alpha = 0.05$  (e.g.,  $t$ -test:  $8.72 \times 10^{-25}$  against OPERA), thereby allowing us to confidently reject  $H_0$  in favor of  $H_1$ . Thus, these statistical results confirm that our method consistently and significantly outperforms all baselines.

## Conclusion

We introduced **InEx**, a training-free multi-agent framework that reduces hallucination by combining internal introspection (**In**) with external cross-modal collaboration (**Ex**). Experiments validate that InEx consistently surpasses strong baselines across diverse hallucination and general-purpose benchmarks. Ablation studies confirm the contributions of both In and Ex. Beyond mitigating hallucination, InEx highlights how transitioning from single-agent reasoning to coordinated multi-agent interaction enables more reliable, autonomous behavior with minimal human intervention.

## References

- Achille, A.; and Soatto, S. 2018. Information dropout: Learning optimal representations through noisy computation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12): 2897–2905.
- Asai, A.; Wu, Z.; Wang, Y.; Sil, A.; and Hajishirzi, H. 2024. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. In *International Conference on Learning Representations*, 9112–9141.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.
- Chen, C.; Liu, M.; Jing, C.; Zhou, Y.; Rao, F.; Chen, H.; Zhang, B.; and Shen, C. 2025. PerturboLLaVA: Reducing multimodal hallucinations with perturbative visual training. In *International Conference on Learning Representations*, 61823–61858.
- Chen, Z.; Zhao, Z.; Luo, H.; Yao, H.; Li, B.; and Zhou, J. 2024. HALC: Object hallucination reduction via adaptive focal-contrast decoding. In *Proceedings of the 41st International Conference on Machine Learning*, 7824–7846.
- Deng, A.; Chen, Z.; and Hooi, B. 2024. Seeing is believing: Mitigating hallucination in large vision-language models via clip-guided decoding. *arXiv preprint arXiv:2402.15300*.
- Farquhar, S.; Kossen, J.; Kuhn, L.; and Gal, Y. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017): 625–630.
- Fu, C.; Chen, P.; Shen, Y.; Qin, Y.; Zhang, M.; Lin, X.; Yang, J.; Zheng, X.; Li, K.; Sun, X.; et al. 2023. MME: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- GLM, T.; Zeng, A.; Xu, B.; Wang, B.; Zhang, C.; Yin, D.; Rojas, D.; Feng, G.; Zhao, H.; Lai, H.; Yu, H.; Wang, H.; Sun, J.; Zhang, J.; Cheng, J.; Gui, J.; Tang, J.; Zhang, J.; Li, J.; Zhao, L.; Wu, L.; Zhong, L.; Liu, M.; Huang, M.; Zhang, P.; Zheng, Q.; Lu, R.; Duan, S.; Zhang, S.; Cao, S.; Yang, S.; Tam, W. L.; Zhao, W.; Liu, X.; Xia, X.; Zhang, X.; Gu, X.; Lv, X.; Liu, X.; Liu, X.; Yang, X.; Song, X.; Zhang, X.; An, Y.; Xu, Y.; Niu, Y.; Yang, Y.; Li, Y.; Bai, Y.; Dong, Y.; Qi, Z.; Wang, Z.; Yang, Z.; Du, Z.; Hou, Z.; and Wang, Z. 2024. ChatGLM: A family of large language models from GLM-130B to GLM-4 all tools. *arXiv preprint arXiv:2406.12793*.
- Gou, Z.; Shao, Z.; Gong, Y.; Shen, Y.; Yang, Y.; Duan, N.; and Chen, W. 2024. CRITIC: Large language models can self-correct with tool-interactive critiquing. In *International Conference on Learning Representations*, 57734–57811.
- Guan, T.; Liu, F.; Wu, X.; Xian, R.; Li, Z.; Liu, X.; Wang, X.; Chen, L.; Huang, F.; Yacoob, Y.; et al. 2024. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14375–14385.
- Gurari, D.; Li, Q.; Stangl, A. J.; Guo, A.; Lin, C.; Grauman, K.; Luo, J.; and Bigham, J. P. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3608–3617.
- Huang, Q.; Dong, X.; Zhang, P.; Wang, B.; He, C.; Wang, J.; Lin, D.; Zhang, W.; and Yu, N. 2024. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13418–13427.
- Huo, F.; Xu, W.; Zhang, Z.; Wang, H.; Chen, Z.; and Zhao, P. 2025. Self-introspective decoding: Alleviating hallucinations for large vision-language models. In *International Conference on Learning Representations*, 24272–24295.
- Jiang, C.; Xu, H.; Dong, M.; Chen, J.; Ye, W.; Yan, M.; Ye, Q.; Zhang, J.; Huang, F.; and Zhang, S. 2024. Hallucination augmented contrastive learning for multimodal large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27036–27046.
- Leng, S.; Zhang, H.; Chen, G.; Li, X.; Lu, S.; Miao, C.; and Bing, L. 2024. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13872–13882.
- Li, Y.; Du, Y.; Zhou, K.; Wang, J.; Zhao, W. X.; and Wen, J.-R. 2023. Evaluating object hallucination in large vision-language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 292–305.
- Liu, F.; Lin, K.; Li, L.; Wang, J.; Yacoob, Y.; and Wang, L. 2024a. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *International Conference on Learning Representations*, 57689–57733.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. In *Advances in Neural Information Processing Systems*, volume 36, 34892–34916.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; Chen, K.; and Lin, D. 2024b. Mmbench: Is your multi-modal model an all-around player? In *Computer Vision – ECCV 2024*, 216–233.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; Gupta, S.; Majumder, B. P.; Hermann, K.; Welleck, S.; Yazdanbakhsh, A.; and Clark, P. 2023. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, volume 36, 46534–46594.
- Renze, M.; and Guven, E. 2024. The benefits of a concise chain of thought on problem-solving in large language models. In *2024 2nd International Conference on Foundation and Large Language Models*, 476–483.
- Rohrbach, A.; Hendricks, L. A.; Burns, K.; Darrell, T.; and Saenko, K. 2018. Object hallucination in image captioning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4035–4045.
- Sahu, P.; Sikka, K.; and Divakaran, A. 2024. Pelican: Correcting hallucination in vision-LLMs via claim decomposition and program of thought verification. In *Proceedings of the*

2024 Conference on Empirical Methods in Natural Language Processing, 8228–8248.

Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, 8634–8652.

Wan, Z.; Zhang, C.; Yong, S.; Ma, M. Q.; Stepputtis, S.; Morency, L.-P.; Ramanan, D.; Sycara, K.; and Xie, Y. 2025. ONLY: One-layer intervention sufficiently mitigates hallucinations in large vision-language models. *arXiv preprint arXiv:2507.00898*.

Wang, X.; Pan, J.; Ding, L.; and Biemann, C. 2024. Mitigating hallucinations in large vision-language models with instruction contrastive decoding. In *Findings of the Association for Computational Linguistics ACL 2024*, 15840–15853.

Wu, T.-H.; Lee, H.; Ge, J.; Gonzalez, J. E.; Darrell, T.; and Chan, D. M. 2025. Generate, but verify: Reducing hallucination in vision-language models with retrospective resampling. *arXiv preprint arXiv:2504.13169*.

Xiao, W.; Huang, Z.; Gan, L.; He, W.; Li, H.; Yu, Z.; Shu, F.; Jiang, H.; and Zhu, L. 2025. Detecting and mitigating hallucination in large vision language models via fine-grained AI feedback. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(24): 25543–25551.

Xiong, M.; Hu, Z.; Lu, X.; LI, Y.; Fu, J.; He, J.; and Hooi, B. 2024. Can LLMs express their uncertainty? An empirical evaluation of confidence elicitation in LLMs. In *International Conference on Learning Representations*, 23650–23678.

Yang, Z.; Chen, J.; Xu, D.; Fei, J.; Shen, X.; Zhao, L.; Feng, C.-M.; and Elhoseiny, M. 2025a. WikiAutoGen: Towards multi-modal wikipedia-style article generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15532–15541.

Yang, Z.; Song, J.; Song, S.; Pang, W.; and Yuan, Y. 2025b. MERMAID: Multi-perspective self-reflective agents with generative augmentation for emotion recognition. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 24650–24666.

Yang, Z.; Wang, S.; Zhang, K.; Wu, K.; Leng, S.; Zhang, Y.; Qin, C.; Lu, S.; Li, X.; and Bing, L. 2025c. LongVT: Incentivizing “thinking with long videos” via native tool calling. *arXiv preprint arXiv:2511.20785*.

Yang, Z.; Xu, D.; Pang, W.; and Yuan, Y. 2025d. Script: Graph-structured and query-conditioned semantic token pruning for multimodal large language models. *Transactions on Machine Learning Research*.

Yin, S.; Fu, C.; Zhao, S.; Xu, T.; Wang, H.; Sui, D.; Shen, Y.; Li, K.; Sun, X.; and Chen, E. 2024. Woodpecker: Hallucination correction for multimodal large language models. *Science China Information Sciences*, 67(12): 220105.

Yu, T.; Yao, Y.; Zhang, H.; He, T.; Han, Y.; Cui, G.; Hu, J.; Liu, Z.; Zheng, H.-T.; Sun, M.; et al. 2024a. Rlhv-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13807–13816.

Yu, T.; Zhang, H.; Li, Q.; Xu, Q.; Yao, Y.; Chen, D.; Lu, X.; Cui, G.; Dang, Y.; He, T.; Feng, X.; Song, J.; Zheng, B.; Liu, Z.; Chua, T.-S.; and Sun, M. 2025. Rlaif-v: Aligning mllms through open-source ai feedback for super gpt-4v trustworthiness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19985–19995.

Yu, W.; Yang, Z.; Li, L.; Wang, J.; Lin, K.; Liu, Z.; Wang, X.; and Wang, L. 2024b. MM-vet: Evaluating large multimodal models for integrated capabilities. In *Proceedings of the 41st International Conference on Machine Learning*, 57730–57754.

Zhang, C.; Wan, Z.; Kan, Z.; Ma, M. Q.; Stepputtis, S.; Ramanan, D.; Salakhutdinov, R.; Morency, L.-P.; Sycara, K.; and Xie, Y. 2025a. Self-correcting decoding with generative feedback for mitigating hallucinations in large vision-language models. In *International Conference on Learning Representations*, 5494–5524.

Zhang, K.; Wu, K.; Yang, Z.; Hu, K.; Wang, B.; Liu, Z.; Li, X.; and Bing, L. 2025b. OpenMMReasoner: Pushing the frontiers for multimodal reasoning with an open and general recipe. *arXiv preprint arXiv:2511.16334*.

Zhang, P.; Zhang, K.; Li, B.; Zeng, G.; Yang, J.; Zhang, Y.; Wang, Z.; Tan, H.; Li, C.; and Liu, Z. 2024. Long context transfer from language to vision. *arXiv preprint arXiv:2406.16852*.

Zhang, R.; Zhang, H.; Fei, H.; and Zheng, Z. 2025c. Uncertainty-o: One model-agnostic framework for unveiling uncertainty in large multimodal models. *arXiv preprint arXiv:2506.07575*.

Zhang, Z.; Xie, J.; Lu, Y.; Yang, Z.; and Yang, Y. 2025d. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. *arXiv preprint arXiv:2504.20690*.

Zhou, Y.; Cui, C.; Yoon, J.; Zhang, L.; Deng, Z.; Finn, C.; Bansal, M.; and Yao, H. 2024. Analyzing and mitigating object hallucination in large vision-language models. In *International Conference on Learning Representations*, 56969–56998.