

# Defeasible Visual Entailment: Benchmark, Evaluator, and Reward-Driven Optimization

Yue Zhang, Liqiang Jing, Vibhav Gogate

The University of Texas at Dallas  
{yue.zhang,vibhav.gogate}@utdallas.edu, jingliqiang6@gmail.com

## Abstract

We introduce a new task called Defeasible Visual Entailment (DVE), where the goal is to allow the modification of the entailment relationship between an image premise and a text hypothesis based on an additional update. While this concept is well-established in Natural Language Inference, it remains unexplored in visual entailment. At a high level, DVE enables models to refine their initial interpretations, leading to improved accuracy and reliability in various applications such as detecting misleading information in images, enhancing visual question answering, and refining decision-making processes in autonomous systems. Existing metrics do not adequately capture the change in the entailment relationship brought by updates. To address this, we propose a novel inference-aware evaluator designed to capture changes in entailment strength induced by updates, using pairwise contrastive learning and categorical information learning. Additionally, we introduce a reward-driven update optimization method to further enhance the quality of updates generated by multimodal models. Experimental results demonstrate the effectiveness of our proposed evaluator and optimization method.

## Introduction

Natural Language Inference (NLI) (Bos and Markert 2005; Dagan, Glickman, and Magnini 2005; MacCartney and Manning 2009; Bowman et al. 2015) is a fundamental task that involves determining the logical relationship between two sentences, specifically identifying whether one sentence entails, contradicts, or is neutral with respect to the other. To further investigate the logical relationship across modalities, researchers have introduced a new inference task called Visual Entailment (VE). In VE, the premise is provided by an image and the hypothesis by a sentence and the task is to determine whether the image supports, contradicts, or is unrelated to the statement in the sentence (Xie et al. 2019).

Existing approaches to the VE task typically leverage pre-trained vision-language models, such as OFA (Wang et al. 2022), UNITER (Chen et al. 2020c), FGAIF (Jing and Du 2024), and CoCa (Yu et al. 2022). These models are designed to understand and reason across visual and textual modalities and have greatly improved our ability to accurately link and interpret images and text.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: An example of defeasibility in visual entailment.

Despite progress in this area, existing works on VE have mostly focused on clear, definite relationships and have not fully considered the uncertainties that can affect how images and text relate to each other. These uncertainties stem from factors such as incomplete information, unseen elements, image complexity, ambiguity, varying interpretations, differing perspectives, context, and the inherent subjectivity in visual perception.

To address this gap, we introduce the concept of *Defeasible Visual Entailment* (DVE). The aim of DVE is to provide additional textual information that can either strengthen or weaken the relationship between an image premise and a text hypothesis. As illustrated in Figure 1, the premise shows a brown dog running through a grassy field. A strengthener could argue that “The dog is a hunting dog,” which strengthens the premise because a hunting dog is more likely to chase a rabbit. On the other hand, a weakener might state “The ball bounces once,” suggesting the dog is more likely chasing the ball than the rabbit.

A key challenge with the DVE task is that existing datasets used in visual entailment research are not suitable for evaluating and benchmarking methods designed to solve the DVE task. More specifically, previous benchmarks in visual entailment have primarily focused on definite relationships, often overlooking the role of defeasibility when uncertainties are present. Therefore, to fully harness the potential of defeasible inference in visual entailment, we introduce a new benchmark. To create this benchmark, we developed a new dataset for the DVE task by replacing the premises in the  $\delta$ -NLI dataset (Rudinger et al. 2020) with images from the Flickr30k dataset (Young et al. 2014). This approach minimizes costs while maximizing the use of existing resources. In our dataset, each premise-hypothesis pair

includes multiple strengtheners and weakeners.

Building upon this dataset, we propose two specific DVE tasks: (1) *Classification Task*: predicting whether a provided update (sentence) acts as a strengthener or a weakener for the premise-hypothesis pair. (2) *Generation Task*: given a premise-hypothesis pair as input, generate an update sentence that weakens or strengthens the hypothesis. While the classification task can be easily evaluated using accuracy, the generation task lacks a metric that effectively captures the change in entailment strength introduced by the generated update. An ideal metric should reflect how the update influences the increase or decrease in entailment strength.

To address this issue, we propose a new *evaluator* capable of measuring the change of entailment strength brought by the generated update. We also introduce a learning scheme that employs pairwise contrastive learning and categorical information learning to train the evaluator in an unsupervised manner. Our evaluator outputs a value representing the entailment strength for a given triplet (update, premise, hypothesis). We conducted a human evaluation to compare the performance of our evaluator with existing metrics, such as ROUGE-L (Lin 2004), and CLIPScore (Hessel et al. 2021). Our experimental results demonstrate that our metric achieves the best correlation with human evaluation results and existing metrics are unable to accurately capture the change of reasoning relationship brought by the update.

In our experiments, we found that directly adapting existing VE methods for the DVE task (baseline approaches) results in low-quality updates, which often fail to alter the entailment relationship between the premise and hypothesis. To address this, we developed a reward-driven update optimization technique that leverages the evaluation results from our evaluator to further refine the generated updates. Our experimental results demonstrate that this new method produces higher-quality updates compared to baseline approaches. In summary, our contributions are:

1. We propose a defeasible visual entailment task and build the first benchmark<sup>1</sup> for it. This benchmark enables a thorough investigation of the fine-grained multimodal understanding capabilities of state-of-the-art models.
2. We devise a novel inference-aware evaluator that leverages advanced pairwise contrastive learning and categorical information learning for capturing the change of entailment strength brought by the update.
3. We propose a new reward-driven update optimization method and demonstrate experimentally that our method significantly enhances the quality of generated updates, outperforming state-of-the-art models.

## Task Definition

In this paper, we follow the definition of the defeasible task (Rudinger et al. 2020)

Given an image premise  $I$ , a hypothesis  $H$  is defeasible if there exists an update  $U$  (consistent with  $I$ ) such that a human would find  $H$  less likely to be true

<sup>1</sup>Our code and data are available at [https://github.com/skywalkerzhang/Defeasible\\_Visual\\_Entailment](https://github.com/skywalkerzhang/Defeasible_Visual_Entailment).

after learning  $U$ . Specifically, an update  $U$  is called a weakener if given a premise  $I$  and hypothesis  $H$ , a human would most likely find  $H$  less likely to be true after learning  $U$ ; if they find  $H$  more likely to be true, then we call  $U$  a strengthener.

## Classification Task

**Formulation** The goal of the classification task is to find a classification model  $\mathcal{M}_c$  which predicts the update type based on the premise  $I$ , hypothesis  $H$ , and update  $U$  as follows,

$$\hat{L} = \mathcal{M}_c(I, H, U), \quad (1)$$

where  $\hat{L} \in \{w, s\}$  denotes the predicted update type.  $\hat{L} = s$  (strengthener) is assigned if  $U$  makes the hypothesis  $H$  more likely given the image  $I$  while  $\hat{L} = w$  (weakener) is assigned if  $U$  makes the hypothesis  $H$  less likely given the image  $I$ .

**Evaluation Metric** We use accuracy as the metric.

## Generation Task

**Formulation** In this task, the model aims to generate an update based on the input premise  $I$ , hypothesis  $H$ , and goal  $G \in \{w, s\}$  (i.e., weakener or strengthener) as follows,

$$\hat{U} = \mathcal{M}_g(I, H, G), \quad (2)$$

where  $\hat{U}$  is the generated (textual) update.

**Evaluation Metric** To comprehensively assess the quality of the generation model  $\mathcal{M}_g$ , we utilize a variety of evaluation metrics, including traditional evaluation metrics: ROUGE-L (Lin 2004), BLEU-4 (Papineni et al. 2002), deep learning-based metrics: BERTScore (Zhang et al. 2020) and CLIPScore (Hessel et al. 2021), and our custom-designed reference-free Inference-aware Evaluator, which is detailed in the later section.

## Defeasible Visual Entailment Dataset

### Dataset Construction

In this section, we describe the construction of our dataset for the DVE task, which leverages three existing datasets: Flickr30k (Young et al. 2014), SNLI (Bowman et al. 2015) and the  $\delta$ -NLI dataset (Rudinger et al. 2020).

Flickr30k is a well-known image captioning dataset comprising 31,783 images and 158,915 captions, depicting everyday activities, events, and scenes. Each image in the dataset is annotated with five captions generated through crowdsourcing, providing diverse descriptions of the visual content. This dataset is essential for developing models that can understand and generate natural language descriptions of images, as it offers a rich set of image-caption pairs that cover a broad range of scenarios and objects.

The SNLI dataset is a large annotated textual entailment dataset. It comprises approximately 570,000 premise-hypothesis ( $T, H$ ) pairs, as well as their corresponding label categorized into three classes: entailment, neutral, and contradiction. The premise was originally collected from the captions in Flickr30k. The hypothesis was written via Amazon Mechanical Turk for each class.

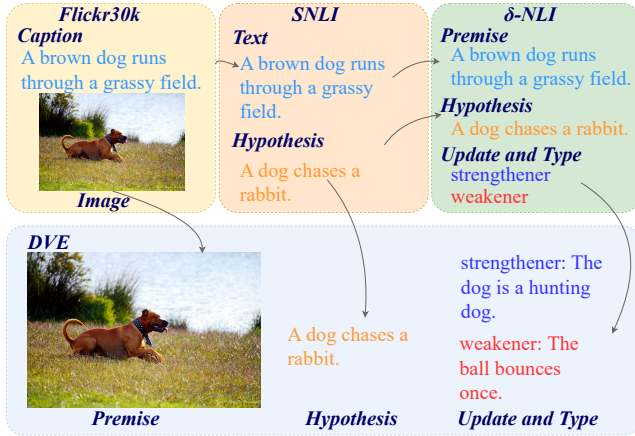


Figure 2: The workflow of generating the DVE dataset by integrating premises and hypotheses from SNLI with images from Flickr30k and updates from  $\delta$ -NLI.

The  $\delta$ -NLI dataset is designed to collect strengtheners and weakeners for the NLI task, which can be used to further investigate the semantic understanding ability in models. The new dataset was devised for the defeasible inference tasks in natural language. This dataset contains 10,000 neutral premise-hypothesis pairs derived from the SNLI dataset. In the context of SNLI, neutral premise-hypothesis pairs are those where the hypothesis is neither entailed nor contradicted by the premise, thereby making it easy to issue additional information to strengthen or weaken the statement under appropriate conditions. The premise is from the captions in the Flickr30k dataset. Crowdsourced workers were assigned the task of writing updates, including both strengtheners and weakeners.

Although existing datasets have been successful in assessing the semantic entailment capability of models, defeasibility in the visual domain has not been explored. Therefore, our work focuses on creating a novel dataset for DVE, which consists of image premises, text hypotheses, and updates (including weakeners and strengtheners) for premise-hypothesis pairs. To simplify and save cost, we constructed our new dataset based on the Flickr30k, SNLI, and  $\delta$ -NLI datasets. Specifically, for each premise-hypothesis pair  $(T, H)$  in the SNLI dataset, we replace the text premise with its corresponding image in Flickr30k, with the premise-hypothesis pair formulated as  $(I, H)$ . Thereafter, we incorporate the update from  $\delta$ -NLI into our DVE dataset. We only retain the premise-hypothesis pair that has an update in the  $\delta$ -NLI dataset. The overall workflow is shown in Figure 2.

### Statistics of DVE

In this section, we present the statistical overview of the DVE dataset, divided into training, development, and test sets. The statistics are summarized in Table 1. Overall, the DVE dataset’s balanced and diverse data support comprehensive training and evaluation of models on visual defeasible inference tasks. We further compare our DVE dataset with the related datasets in the supplementary material.

Statistics	Train set	Validation set	Test set
Total samples	93,082	1,888	1,972
Update type dist.			
Weakener	46,541	944	986
Strengtheners	46,541	944	986
Average premise length	12.83	13.82	13.21
Average hypothesis length	8.27	8.41	8.23
Unique premises	9,293	191	200
Unique hypotheses	9,438	195	203
Average updates per image	9.79	9.68	9.71
Unique images	9,507	195	203

Table 1: Statistics of the DVE dataset.

## Estimating the Impact of Updates on Multimodal Defeasible Reasoning

For the generation task, we utilize standard generation evaluation metrics such as BLEU, ROUGE, and BERTScore to measure the quality of the generated updates. These metrics assess the lexical or semantic similarity between the generated update and the reference updates, but it is not realistic to collect a comprehensive set of ground-truth references for such open-domain tasks, where answers can vary widely. We also employ the reference-free metric CLIPScore, which primarily evaluates the similarity between the answer and the image. While these metrics provide some insight into the quality of the updates, they are not well-suited to accurately capture the changes in entailment strength brought about by strengtheners and weakeners. To address this, we propose a new reference-free evaluation approach utilizing contrastive learning to train an unsupervised model capable of representing the entailment strength of the changes caused by updates.

### Inference-aware Evaluator

As mentioned before, for the generation task, we designed a novel reference-free evaluation method that leverages contrastive learning to capture the impact of updates on inference strength. Our model consists of the following components. The overall architecture of the model is illustrated in Figure 3, which consists of three modules: Multimodal Embedding, Feature Fusion, and Multitask Learning.

**Multimodal Embedding** The input data for the model consists of both images and text. To feed the multimodal data into our model, we first get the embeddings of the text and image as follows.

**Visual Embedding** Since ResNet (He et al. 2016) has shown great success on vision tasks, such as image classification (Russakovsky et al. 2015; Krizhevsky, Hinton et al. 2009), object detection (Everingham et al. 2010; Lin et al. 2014), semantic segmentation (Zhou et al. 2017), we also use it to extract the visual embedding. Specifically, we use the pretrained ResNet-50 model to extract the visual embedding as follows,

$$\mathbf{i} = \text{ResNet}(I), \quad (3)$$

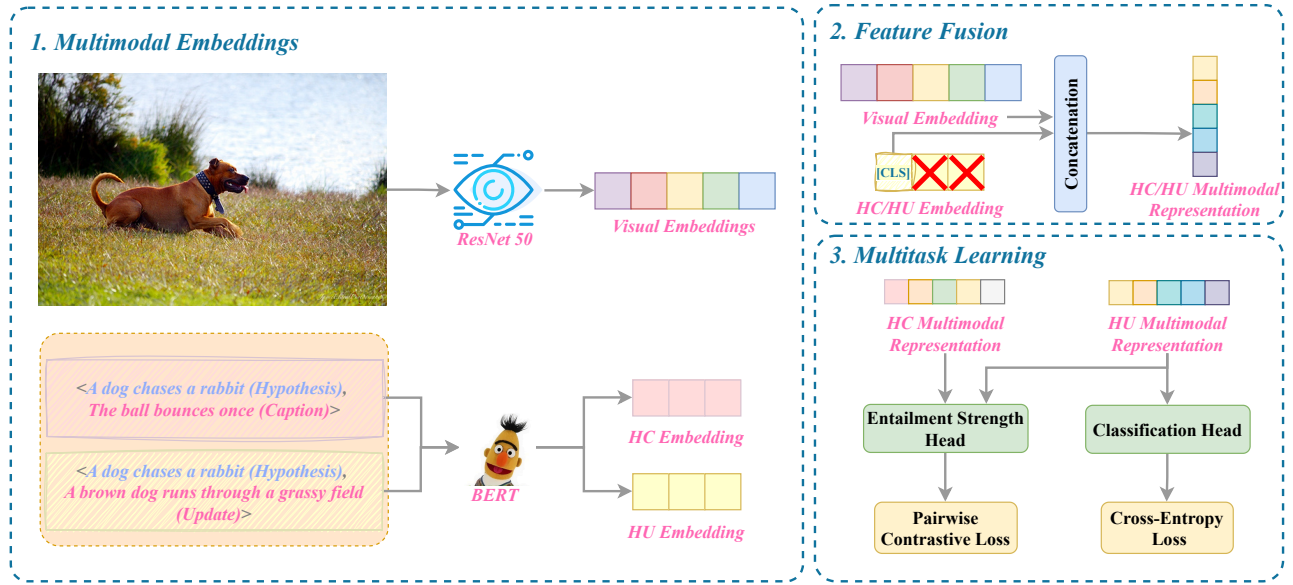


Figure 3: The architecture of our Inference-aware Evaluator, including three modules: Multimodal Embedding, Feature Fusion, and Multitask Learning. HC/HU Embedding means the embedding of the hypothesis-caption/hypothesis-update pair. Similarly, HC/HU Multimodal Representation stands for the multimodal representation of the hypothesis-caption/hypothesis-update pair.

where  $\mathbf{i} \in \mathbb{R}^{d_1}$  denotes the image embedding of the image premise  $I$ .  $d_1$  is the embedding size.  $\text{ResNet}(\cdot)$  refers to the ResNet-50 model.

**Textual Embedding** It is known that BERT (Devlin et al. 2019) achieves superior performance on various natural language models, such as Language Understanding (Wang et al. 2019), Question Answering (Rajpurkar et al. 2016) and Commonsense Inference (Zellers et al. 2018). Therefore, we use BERT to extract the textual features. In particular, we encode a pair of text inputs: the hypothesis and update with BERT as follows,

$$\mathbf{e} = \text{BERT}([H, U]), \quad (4)$$

where  $\mathbf{e} \in \mathbb{R}^{d_2}$  represent the embedding of the [CLS] token output from BERT, which is used to represent the overall semantics of the text pairs.  $\text{BERT}(\cdot)$  refers to the BERT model.

**Feature Fusion** We propose to concatenate the extracted visual and textual features to form a combined multimodal feature representation, denoted as  $\mathbf{m}$ :

$$\mathbf{m} = [\mathbf{i}, \mathbf{e}], \quad (5)$$

where  $[\cdot, \cdot]$  denotes the concatenation operation. In this context,  $\mathbf{m} \in \mathbb{R}^{d_1+d_2}$  represents the combined features that integrate the multimodal information, enabling the model to leverage both visual and textual contexts effectively.

**Multitask Learning** Our evaluator employs a multitask learning framework to jointly perform classification and inference strength tasks, utilizing shared representations to improve overall performance. The inference strength score is ultimately used to represent the strength of visual entailment brought by updates.

**Pairwise Contrastive Learning** Since the existing entailment datasets only label update classes without indicating

entailment strength, they cannot be used to train a model that predicts this strength for a given (premise, hypothesis, and update) triplet. While human scoring could be an option, it is impractical due to its difficulty, cost, and lack of scalability. Instead, motivated by the contrastive learning framework (Chen et al. 2020b), we develop an unsupervised method to train our evaluator by comparing the entailment strength between pairs, requiring only knowledge of which pair has stronger entailment.

Specifically, we first devise an entailment strength head to output a numerical score  $s$  representing the impact of the update on the hypothesis as follows,

$$s = \mathbf{W}_s \mathbf{m} + \mathbf{b}_s, \quad (6)$$

where  $\mathbf{W}_s \in \mathbb{R}^{d_1+d_2}$  and  $\mathbf{b}_s \in \mathbb{R}^1$  are the trainable weights and bias of the entailment strength layer respectively. The entailment strength score  $s$  is used as the final measure of the strength of entailment inference, indicating how the update affects the hypothesis. A higher score indicates that the update makes the hypothesis more likely in relation to the premise. In contrast, a lower score indicates that the update makes the hypothesis less likely in relation to the premise.

To train the evaluator, we design a custom pairwise contrastive loss function that can capture the change in entailment strength by comparing triplets (update, premise, and hypothesis). It is evident that the the entailment strength of the triplet (strengtheners, premise, and hypothesis) is bigger than the triplet (caption, premise, and hypothesis) and the the entailment strength of the triplet (weakeners, premise, and hypothesis) is smaller than the triplet (caption, premise, and hypothesis). Therefore, we devise the pairwise contrastive

loss function as follows,

$$\mathcal{L}_p = -\frac{1}{N} \sum_{i=1}^N \log(\sigma((s_u^i - s_c^i) \cdot l^i)), \quad (7)$$

where  $s_u^i$  is the score computed by the Eqn.(6) for the triplet (update, premise, and hypothesis).  $s_c^i$  is the score computed by the Eqn.(6) for the triplet (caption, premise, and hypothesis).  $l^i \in \{-1, 1\}$ , where  $-1$  represents the update is a weakener and  $1$  represents the update is a strengthener.  $\sigma(\cdot)$  is the sigmoid function, and  $N$  is the number of samples.

**Categorical Information Learning** To further learn the category information of the update, we devise a categorical information loss function. Specifically, we first design a classification head that aims to classify the update as either a strengthener or a weakener as follows,

$$\hat{y} = \sigma(W_c \mathbf{m}_u + b_c), \quad (8)$$

where  $\sigma$  is the sigmoid activation function.  $W_c \in \mathbb{R}^{d_1+d_2}$  and  $b_c \in \mathbb{R}^2$  are the trainable weight and bias of the classification layer.  $\mathbf{m}_u$  is the combined multimodal feature representation by Eqn.(5) for the triplet (update, premise, hypothesis).  $\hat{y} \in \mathbb{R}^2$  is the corresponding predicted label (*i.e.*, strengthener and weakener) for the above triplet. Thereafter, we utilize the cross-entry loss function to learn the categorical information as follows,

$$\mathcal{L}_c = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C y_{ij} \log \hat{y}_{ij}, \quad (9)$$

where  $N$  is the number of samples,  $C$  is the number of classes,  $y_{ij}$  is the ground truth label for the  $i$ -th sample and the  $j$ -th class (1 if the sample belongs to the class, otherwise 0), and  $\hat{y}_{ij}$  is the predicted probability for the  $i$ -th sample and the  $j$ -th class.

## Training

The overall loss function for multitask learning is a weighted sum of the classification loss and the pairwise contrastive loss as follows,

$$\mathcal{L} = (1 - \alpha)\mathcal{L}_p + \alpha\mathcal{L}_c, \quad (10)$$

where  $\mathcal{L}_c$  is the binary cross-entropy loss for the classification task,  $\mathcal{L}_p$  is the pairwise contrastive loss, and  $\alpha$  is a hyper-parameter to balance their contributions.

## Meta-evaluate Evaluator for Automatic Evaluation

To verify the effectiveness of our automatic evaluator, we conduct human evaluations on the whole test dataset. For the evaluation model, we select answers from LLaVA-1.5 (Liu et al. 2023) and GPT-4o. More implementation details can be found in the supplementary material.

## Human Evaluations

As we mentioned before, we select LLaVA-1.5 and GPT-4o to generate strengthener and weakener for the 203 images in the testing set, a total  $203 \times 2 \times 2 = 812$  samples.

Metric	$r(\%)$	$\rho(\%)$	$\tau(\%)$
GPT-4o			
ROUGE-L	-0.1265	-0.1631	-0.1180
BLEU	-0.0081	-0.0295	-0.0265
BERTScore	-0.0566	-0.0821	-0.0558
CLIPScore	0.1068	0.1179	0.0853
<b>Ours</b>	<b>0.8262</b>	<b>0.8037</b>	<b>0.6552</b>
LLaVA-1.5			
ROUGE-L	-0.1964	-0.2047	-0.1502
BLEU	-0.0777	-0.0601	-0.0528
BERTScore	-0.0950	-0.1069	-0.0795
CLIPScore	0.2760	0.2690	0.2035
<b>Ours</b>	<b>0.7733</b>	<b>0.7368</b>	<b>0.6024</b>

Table 2: Correlation between each evaluation metric and human judgment on VDI, measured by Pearson’s  $r$ , Spearman’s  $\rho$ , and Kendall’s  $\tau$ . The best metrics for each correlation coefficient are highlighted in bold.

We employed 3 workers for annotation, with each person annotating 812 testing samples. For each test example, we meticulously designed an annotation process to evaluate the scores of the models’ generated answers. The score was conducted on a 5-point scale, ranging from “weakens a lot” to “strengthens a lot,” with a middle category of “neutral” for updates that have no effect. Each worker was paid 15-20 USD per hour. After the annotation process, we calculated the inter-annotator agreement rate using Fleiss’  $\kappa$ , achieving a result of 80.4%, which involved all annotators. This level of concordance among the evaluators suggests that the human evaluation results are reliable.

## Correlations with Human Evaluations

We evaluate our proposed metric against traditional metrics commonly used for generation tasks, such as ROUGE-L, BLEU, BERTScore, and CLIPScore. These metrics are widely recognized for their effectiveness in evaluating text generation (Narayan, Cohen, and Lapata 2018; Lin et al. 2020) and vision-language tasks (Lin et al. 2014; Sidorov et al. 2020). To quantify the alignment between human annotations and model-generated evaluations, we employed three different correlation coefficients: Pearson’s  $r$ , Spearman’s  $\rho$ , and Kendall’s  $\tau$ .

We show the correlation between automatic evaluation and human evaluation in Table 2. Except for our metric and CLIPScore, other evaluation metrics (*e.g.*, ROUGE-L and BLEU) show negative correlations across both GPT-4o and LLaVA-1.5’s results. This indicates that these metrics do not align well with human judgments. One potential reason is that these metrics pay more attention to the text overlap but this is not suitable for open-domain generation, especially when the answer is not a fixed one. Notably, our proposed metric consistently outperforms other metrics across all correlation measures, which indicates the effectiveness of our metric. In addition, a key virtue of our method is that compared with some traditional metrics (*e.g.*, BLEU






Figure 4: An overview of Reward-driven Update Optimization, which includes three steps: Initial Response Generation, Critique, and Refinement.

and ROUGE-L), our metric is reference-free. We provide the case study for the evaluator in the supplementary material.

### Reward-driven Update Optimization

In the generation task, our objective is to generate updates based on given premises to either strengthen or weaken hypotheses. In our experiments, we observed that the initial updates may suffer from quality issues, such as simply captioning the images rather than effectively achieving the intended goal.

To address this issue, we propose a new reward-driven update optimization method, which leverages the entailment strength of the generated update. Figure 4 presents an overview of the proposed method. Our method consists of the following steps:

- 1. Initial Response Generation:** We submit the user request to the Large Vision-Language Model (LVLM), to generate the initial response. This response serves as the baseline for subsequent comparisons with the refined responses produced by our method.
- 2. Critique:** Our inference-aware evaluator serves as the critique, assessing the entailment strength of the generated updates. If the critique assigns a low score, we proceed to the next step (*i.e.*, Refinement) to improve the response. We establish a threshold  $\eta$  to evaluate the quality of the generated update. Specifically, if the score of a generated strengthener is less than  $\eta$  or the score of a generated weakener exceeds  $-\eta$ , we classify the update as low-quality. Conversely, if the score indicates the update is of high quality, we output the current response as the final result.
- 3. Refinement:** In this step, we feed the score along with the current generation result into the LVLM to refine the response. After generating a new update, we return to the critique step to obtain a new score. This process is repeated until the model produces a high-quality update (as defined in the critique step) or until the loop reaches a maximum iteration count of  $M$ .

### Evaluate Models on DVE Tasks

#### Experimental Setup

For the *Classification Task*, we selected seven models, categorized into two types: finetuning-based methods and models evaluated in the zero-shot setting. The finetuning-based

Model	ROUGE-L	BLEU	BERTScore	CLIPScore	Ours
Strengthener					
InstructBLIP	0.0601	0.0141	0.1891	0.2111	0.7211
Multimodal-GPT	0.0541	0.0033	0.7774	0.2426	1.0690
MiniGPT-4	0.1376	0.0180	0.7696	0.2705	1.4998
mPLUG-Owl	<b>0.3308</b>	<b>0.0781</b>	0.8815	0.2733	2.2887
LLaVA-1.5	0.3163	0.0612	0.8847	0.2788	2.7868
GPT-4o	0.2702	0.0423	<b>0.8954</b>	0.2867	3.6413
GPT-4o (Optimized)	0.2653	0.0410	0.8787	<b>0.2872</b>	<b>4.0679</b>
Weaker					
InstructBLIP	0.0817	0.0280	0.2614	0.2161	0.4231
Multimodal-GPT	0.0481	0.0032	0.7748	0.2406	0.7194
MiniGPT-4	0.1193	0.0128	0.7183	0.2639	0.9776
mPLUG-Owl	<b>0.3386</b>	<b>0.0858</b>	0.8842	0.2732	1.0274
LLaVA-1.5	0.3438	0.0773	0.8865	0.2702	0.4834
GPT-4o	0.2800	0.0451	<b>0.8957</b>	0.2782	-2.5212
GPT-4o (Optimized)	0.2768	0.0440	0.8798	<b>0.2762</b>	<b>-2.9240</b>

Table 3: Evaluation metrics for strengtheners and weakeners across different models. For our metric, a higher value represents a higher entailment strength brought by updates, while a lower value indicates a lower entailment strength. Therefore, for strengtheners, a higher value reflects a stronger entailment strength update. Conversely, for weakeners, a lower value indicates a more effective weakening update. The best results in each category are highlighted in bold.

models include VILT (Kim, Son, and Kim 2021), FLAVA (Singh et al. 2022), and CLIP (Radford et al. 2021). We fine-tuned these models on our training set with standard cross-entropy classification loss function. The models under the zero-shot setting include InstructBLIP (Dai et al. 2023), LLaVA-1.5, mPLUG-Owl (Ye et al. 2023), and GPT-4o. We directly prompt these pretrained LVLMs to generate a prediction for classification results. For the *Generation Task*, we selected six widely used LVLMs in a zero-shot setting as baselines: 1) InstructBLIP; 2) Multimodal-GPT (Gong et al. 2023); 3) MiniGPT-4 (Zhu et al. 2023); 4) mPLUG-Owl; 5) LLaVA-1.5; 6) GPT-4o. We select GPT-4o as the LVLM in reward-driven update optimization. More details of the experiments can be found in the supplementary material.

### Results and Analysis

**Classification Task** Table 4 presents the accuracy of the various models on the classification task. From this table, we observe that: (1) Among the finetuning-based models, CLIP achieves the highest accuracy at 71.10%, followed by FLAVA at 70.03%, and VILT at 68.10%. The likely reason is that the pretraining dataset for CLIP is larger than those used for the other models. (2) The closed-source GPT-4o signif-

Model	Accuracy (%)
VILT	68.10
FLAVA	70.03
CLIP	71.10
InstructBLIP	31.32
mPLUG-Owl	31.16
LLaVA-1.5	52.07
GPT-4o	<b>81.76</b>

Table 4: Performance comparison among different methods in the classification task.

icantly outperforms all other open-source models with an accuracy of 81.76%, demonstrating its robust capability. (3) The fine-tuning-based models outperform most LVLMS in the zero-shot setting, except for GPT-4o. This suggests that despite being trained on large-scale datasets, current LVLMS still lack sufficient knowledge for our classification task.

**Generation Task** Table 3 presents the performance of supporter and defeater generation across various assessment metrics. Notably, even GPT-4o does not achieve high scores according to existing generation metrics, highlighting the limitations of current metrics in accurately evaluating the quality of generated updates. This underscores the necessity of our proposed metric. MiniGPT-4 and Multimodal-GPT outperform InstructBLIP in BERTScore, likely due to their more fluent and coherent outputs. This advantage can be attributed to the more advanced language models used by MiniGPT-4 and Multimodal-GPT, which are better equipped to generate contextually appropriate sentences.

Among all baselines, GPT-4o achieved the best performance, demonstrating its robustness. Our proposed framework, GPT-4o (Optimized), performs even better than GPT-4o alone. This improvement is due to our framework’s ability to provide feedback to GPT-4o, enabling it to refine low-quality responses. Additionally, it is evident that all models perform worse in generating weakeners, with only GPT-4o-based models being able to produce effective weakeners. This is likely because most models tend to default to simple image captioning rather than generating nuanced defeaters.

We also assessed human performance based on our evaluator. The average score for the strengthener is 5.0998, and the weakener score is -4.5412. This demonstrates that there is a significant gap between the model’s performance and human performance. Finally, we provide a case study of our proposed optimization method in the supplement.

## Related Work

**Natural Language Inference** Textual entailment (Bowman et al. 2015; Williams, Nangia, and Bowman 2018; Nie et al. 2020), defined as determining whether a human would typically consider a hypothesis to be likely true given a premise, has become a cornerstone task in natural language processing. However, the task of textual entailment has faced criticism, studies have shown significant variability in human agreement on entailment judgments (Pavlick

and Kwiatkowski 2019), leading to the proposal of alternative approaches that use ordinal or numeric values to represent plausibility (Zhang et al. 2017; Sakaguchi and Durme 2018; Chen et al. 2020a; Talman et al. 2023). This shift aims to capture the nuanced nature of entailment more accurately. In recent years, the focus has shifted towards the defeasibility of textual entailments, which involves revising or overturning conclusions based on new evidence. The  $\delta$ -NLI dataset extends existing NLI datasets by including scenarios where new information can alter inferences, providing a more realistic evaluation of models’ reasoning abilities (Rudinger et al. 2020). Similarly, the BoardgameQA dataset measures the reasoning capacity of language models when faced with contradictory information, guided by source preferences and implicit background knowledge, better reflecting real-world reasoning challenges (Kazemi et al. 2023). However, the defeasible entailment inference in the multimodal setting is still unexplored.

**Visual Understanding Tasks** Visual Question Answering (VQA), image captioning, and visual reasoning are common visual understanding tasks. VQA aims to answer natural language questions based on provided visual information. The VQA-v1.0 dataset (Antol et al. 2015) was one of the first to address this task, focusing on the basic interaction between visual content and natural language questions. However, it faced issues related to biases and limited reasoning capabilities (Xie et al. 2019). To address these limitations, several datasets (Johnson et al. 2017; Goyal et al. 2017; Han et al. 2023; Mathew et al. 2022; Lu et al. 2021) have been developed to reduce biases and enhance reasoning capabilities. While VQA focuses on understanding and answering questions about visual content, image captioning involves generating natural language descriptions of an image’s content (Lin et al. 2014; Young et al. 2014; Sidorov et al. 2020; Jing et al. 2024; Chang et al. 2024; Jing, Zuo, and Zhang 2024). In addition, visual reasoning involves understanding relationships and interactions between visual elements, enhancing comprehension of visual content (Thrush et al. 2022; Wu et al. 2023; Qiao et al. 2023). However, these tasks can not capture fine-grained semantics reasoning relation change brought by the new information.

## Conclusion

In this paper, we present a novel defeasible visual entailment task and a new benchmark for studying defeasibility in visual entailment. We also propose a novel inference-aware evaluator for capturing the change of entailment strength brought by the update and a new reward-driven update optimization method to further improve the quality of the update generated by the multimodal model. Our experimental results clearly show the effectiveness of our proposed inference-aware evaluator and reward-driven update optimization method.

## Acknowledgments

We want to thank our anonymous reviewers for their feedback. This work was supported in part by the DARPA

Perceptually-Enabled Task Guidance (PTG) Program under contract number HR00112220005, the DARPA Assured Neuro Symbolic Learning and Reasoning (ANSR) Program under contract number HR001122S0039, the National Science Foundation grant IIS-1652835, the AFOSR award FA9550-23-1-0239, and OpenAI Researcher Access Program 0000006384.

## References

- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. VQA: Visual Question Answering. In *ICCV*, 2425–2433. IEEE Computer Society.
- Bos, J.; and Markert, K. 2005. Recognising Textual Entailment with Logical Inference. In *HLT/EMNLP*, 628–635. The Association for Computational Linguistics.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*, 632–642. The Association for Computational Linguistics.
- Chang, Y.; Jing, L.; Zhang, X.; and Zhang, Y. 2024. A Unified Hallucination Mitigation Framework for Large Vision-Language Models. *CoRR*, abs/2409.16494.
- Chen, T.; Jiang, Z.; Poliak, A.; Sakaguchi, K.; and Durme, B. V. 2020a. Uncertain Natural Language Inference. In *ACL*, 8772–8779. Association for Computational Linguistics.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. E. 2020b. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML*, 1597–1607. PMLR.
- Chen, Y.; Li, L.; Yu, L.; Kholy, A. E.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020c. UNITER: UNiversal Image-TExt Representation Learning. In *ECCV*, 104–120. Springer.
- Dagan, I.; Glickman, O.; and Magnini, B. 2005. The PASCAL Recognising Textual Entailment Challenge. In *MLCW*, volume 3944, 177–190. Springer.
- Dai, W.; Li, J.; Li, D.; Tiong, A. M. H.; Zhao, J.; Wang, W.; Li, B.; Fung, P.; and Hoi, S. C. H. 2023. InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning. In *NeurIPS*.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*, 4171–4186. Association for Computational Linguistics.
- Everingham, M.; Gool, L. V.; Williams, C. K. I.; Winn, J. M.; and Zisserman, A. 2010. The Pascal Visual Object Classes (VOC) Challenge. *IJCV*, 88: 303–338.
- Gong, T.; Lyu, C.; Zhang, S.; Wang, Y.; Zheng, M.; Zhao, Q.; Liu, K.; Zhang, W.; Luo, P.; and Chen, K. 2023. MultiModal-GPT: A Vision and Language Model for Dialogue with Humans. *CoRR*, abs/2305.04790.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *CVPR*, 6325–6334. IEEE Computer Society.
- Han, X.; You, Q.; Liu, Y.; Chen, W.; Zheng, H.; Mrini, K.; Lin, X.; Wang, Y.; Zhai, B.; Yuan, J.; Wang, H.; and Yang, H. 2023. InfiMM-Eval: Complex Open-Ended Reasoning Evaluation For Multi-Modal Large Language Models.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. IEEE Computer Society.
- Hessel, J.; Holtzman, A.; Forbes, M.; Bras, R. L.; and Choi, Y. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *EMNLP*. Association for Computational Linguistics.
- Jing, L.; and Du, X. 2024. FGAIF: Aligning Large Vision-Language Models with Fine-grained AI Feedback. *CoRR*, abs/2404.05046.
- Jing, L.; Li, R.; Chen, Y.; and Du, X. 2024. FaithScore: Fine-grained Evaluations of Hallucinations in Large Vision-Language Models. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12–16, 2024*, 5042–5063. Association for Computational Linguistics.
- Jing, L.; Zuo, J.; and Zhang, Y. 2024. Fine-grained and Explainable Factuality Evaluation for Multimodal Summarization. *CoRR*, abs/2402.11414.
- Johnson, J.; Hariharan, B.; van der Maaten, L.; Fei-Fei, L.; Zitnick, C. L.; and Girshick, R. B. 2017. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. In *CVPR*, 1988–1997. IEEE Computer Society.
- Kazemi, M.; Yuan, Q.; Bhatia, D.; Kim, N.; Xu, X.; Imbrasaitė, V.; and Ramachandran, D. 2023. BoardgameQA: A Dataset for Natural Language Reasoning with Contradictory Information. In *NeurIPS 2023*.
- Kim, W.; Son, B.; and Kim, I. 2021. ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 5583–5594. PMLR.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. Technical report, Toronto, ON, Canada.
- Lin, B. Y.; Shen, M.; Zhou, W.; Zhou, P.; Bhagavatula, C.; Choi, Y.; and Ren, X. 2020. CommonGen: A Constrained Text Generation Challenge for Generative Commonsense Reasoning. In *AKBC*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *ECCV 2014*, 740–755. Springer.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. In *NeurIPS*.
- Lu, P.; Qiu, L.; Chen, J.; Xia, T.; Zhao, Y.; Zhang, W.; Yu, Z.; Liang, X.; and Zhu, S. 2021. IconQA: A New Benchmark for Abstract Diagram Understanding and Visual Language Reasoning. In *NeurIPS*.



- MacCartney, B.; and Manning, C. D. 2009. An extended model of natural logic. In *IWCS*, 140–156. Association for Computational Linguistics.
- Mathew, M.; Bagal, V.; Tito, R.; Karatzas, D.; Valveny, E.; and Jawahar, C. V. 2022. InfographicVQA. In *WACV*, 2582–2591. IEEE.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018. Don’t Give Me the Details, Just the Summary! Topic-Aware Convolutional Neural Networks for Extreme Summarization. In *EMNLP*, 1797–1807. Association for Computational Linguistics.
- Nie, Y.; Williams, A.; Dinan, E.; Bansal, M.; Weston, J.; and Kiela, D. 2020. Adversarial NLI: A New Benchmark for Natural Language Understanding. In *ACL*, 4885–4901. Association for Computational Linguistics.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *ACL*, 311–318. ACL.
- Pavlick, E.; and Kwiatkowski, T. 2019. Inherent Disagreements in Human Textual Inferences. *TACL*, 7: 677–694.
- Qiao, Y.; Jing, L.; Song, X.; Chen, X.; Zhu, L.; and Nie, L. 2023. Mutual-Enhanced Incongruity Learning Network for Multi-Modal Sarcasm Detection. In Williams, B.; Chen, Y.; and Neville, J., eds., *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 9507–9515. AAAI Press.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, volume 139, 8748–8763. PMLR.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100, 000+ Questions for Machine Comprehension of Text. In *EMNLP*, 2383–2392. The Association for Computational Linguistics.
- Rudinger, R.; Shwartz, V.; Hwang, J. D.; Bhagavatula, C.; Forbes, M.; Bras, R. L.; Smith, N. A.; and Choi, Y. 2020. Thinking Like a Skeptic: Defeasible Inference in Natural Language. In *EMNLP*, 4661–4675. Association for Computational Linguistics.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115: 211–252.
- Sakaguchi, K.; and Durme, B. V. 2018. Efficient Online Scalar Annotation with Bounded Support. In *ACL*, 208–218. Association for Computational Linguistics.
- Sidorov, O.; Hu, R.; Rohrbach, M.; and Singh, A. 2020. TextCaps: A Dataset for Image Captioning with Reading Comprehension. In *ECCV*, 742–758. Springer.
- Singh, A.; Hu, R.; Goswami, V.; Couairon, G.; Galuba, W.; Rohrbach, M.; and Kiela, D. 2022. FLAVA: A Foundational Language And Vision Alignment Model. In *CVPR*, 15617–15629. IEEE.
- Talman, A.; Çelikkanat, H.; Virpioja, S.; Heinonen, M.; and Tiedemann, J. 2023. Uncertainty-Aware Natural Language Inference with Stochastic Weight Averaging. In Alumäe, T.; and Fishel, M., eds., *NoDaLiDa*, 358–365. University of Tartu Library.
- Thrush, T.; Jiang, R.; Bartolo, M.; Singh, A.; Williams, A.; Kiela, D.; and Ross, C. 2022. Winoground: Probing Vision and Language Models for Visio-Linguistic Compositionality. In *CVPR*, 5228–5238. IEEE.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *ICLR*. OpenReview.net.
- Wang, P.; Yang, A.; Men, R.; Lin, J.; Bai, S.; Li, Z.; Ma, J.; Zhou, C.; Zhou, J.; and Yang, H. 2022. OFA: Unifying Architectures, Tasks, and Modalities Through a Simple Sequence-to-Sequence Learning Framework. In *ICML*, 23318–23340. PMLR.
- Williams, A.; Nangia, N.; and Bowman, S. R. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *NAACL-HLT*, 1112–1122. Association for Computational Linguistics.
- Wu, R.; Ma, X.; Li, Q.; Wang, W.; Zhang, Z.; Zhu, S.; and Wang, Y. 2023. Bongard-OpenWorld: Few-Shot Reasoning for Free-form Visual Concepts in the Real World. *CoRR*, abs/2310.10207.
- Xie, N.; Lai, F.; Doran, D.; and Kadav, A. 2019. Visual Entailment: A Novel Task for Fine-Grained Image Understanding. *CoRR*, abs/1901.06706.
- Ye, Q.; Xu, H.; Xu, G.; Ye, J.; Yan, M.; Zhou, Y.; Wang, J.; Hu, A.; Shi, P.; Shi, Y.; Li, C.; Xu, Y.; Chen, H.; Tian, J.; Qi, Q.; Zhang, J.; and Huang, F. 2023. mPLUG-Owl: Modularization Empowers Large Language Models with Multimodality. *CoRR*, abs/2304.14178.
- Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *TACL*, 2: 67–78.
- Yu, J.; Wang, Z.; Vasudevan, V.; Yeung, L.; Seyedhosseini, M.; and Wu, Y. 2022. CoCa: Contrastive Captioners are Image-Text Foundation Models. *TMLR*.
- Zellers, R.; Bisk, Y.; Schwartz, R.; and Choi, Y. 2018. SWAG: A Large-Scale Adversarial Dataset for Grounded Commonsense Inference. In *EMNLP*, 93–104. Association for Computational Linguistics.
- Zhang, S.; Rudinger, R.; Duh, K.; and Durme, B. V. 2017. Ordinal Common-sense Inference. *TACL*, 5: 379–395.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *ICLR*. OpenReview.net.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene Parsing through ADE20K Dataset. In *CVPR*, 5122–5130. IEEE Computer Society.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. *CoRR*, abs/2304.10592.