

NLX-GPT: A Model for Natural Language Explanations in Vision and Vision-Language Tasks

Fawaz Sammani¹, Tanmoy Mukherjee^{1,2}, Nikos Deligiannis^{1,2}

¹ETRO Department, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium

²imec, Kapeldreef 75, B-3001 Leuven, Belgium

fawaz.sammani@vub.be, tmukherj@etrovub.be, ndeligia@etrovub.be

Abstract

Natural language explanation (NLE) models aim at explaining the decision-making process of a black box system via generating natural language sentences which are human-friendly, high-level and fine-grained. Current NLE models¹ explain the decision-making process of a vision or vision-language model (a.k.a., task model), e.g., a VQA model, via a language model (a.k.a., explanation model), e.g., GPT. Other than the additional memory resources and inference time required by the task model, the task and explanation models are completely independent, which disassociates the explanation from the reasoning process made to predict the answer. We introduce NLX-GPT, a general, compact and faithful language model that can simultaneously predict an answer and explain it. We first conduct pre-training on large scale data of image-caption pairs for general understanding of images, and then formulate the answer as a text prediction task along with the explanation. Without region proposals nor a task model, our resulting overall framework attains better evaluation scores, contains much less parameters and is 15× faster than the current SoA model. We then address the problem of evaluating the explanations which can be in many times generic, data-biased and can come in several forms. We therefore design 2 new evaluation measures: (1) explain-predict and (2) retrieval-based attack, a self-evaluation framework that requires no labels. Code is at: <https://github.com/fawazsammani/nlxgpt>.

1. Introduction

Deep learning models have enabled extraordinary breakthroughs in a variety of vision tasks (such as image classification [15, 17, 26]) and vision-language tasks (such as visual question answering [1, 3, 57], visual entailment [56], image

¹Throughout this paper, we refer to NLE models as Natural Language Explanation models aimed for vision and vision-language tasks.

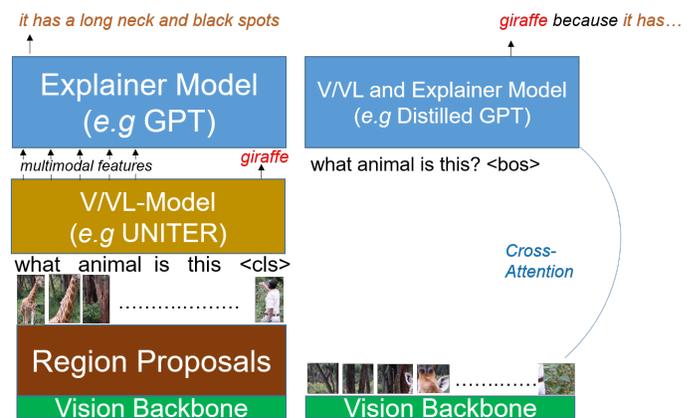


Figure 1. A comparison between previous models (left) and ours (right). Our model solely requires a visual encoder and a language model. We model the answer as a text prediction task along with the explanation. Best viewed in color.

captioning [11, 34, 43, 53], and more), achieving promising performance. However, they are black box systems. For these models to be deployed in everyday life, explaining their decision-making process becomes critical for several reasons such as trust, accountability, and model bias understanding and correctness. Different from visual or textual explanations which highlight regions or tokens in an image or sentence that lead to a specific prediction [5, 45, 48, 49], natural language explanation models [8, 33] explain the decision-making process of a model through natural language sentences. These sentences are easy to understand by humans and are much more detailed than highlighted regions or tokens. Recently, NLE for vision and vision-language (VL) tasks has been introduced [20, 32, 36, 54]. In this work, we focus on explaining models aimed for vision and vision-language tasks. Current NLE models [20, 32, 36, 54] first utilize a VL-model to get an answer for the task at hand (e.g., a visual question answering (VQA) model). The outputs of the VL-model (answer and mul-

timodal features) along with the question are then fed to language model (e.g., LSTM or Transformer) to get an explanation for the answer (see Figure 1). At training time, the language model is trained to produce explanations for the ground-truth answers with a NLE dataset. At test time, the output of a VL-model is utilized to predict an answer which is fed to the language model to get the explanation. This paradigm has two disadvantages. First, the addition of the task model requires higher storage and memory requirements (typically, approx. 80M and 300M parameters for small and large models, respectively). Second, the VL-model and language model are completely independent of each other, which disconnects the explanations from the reasoning process made to predict the answer. Finally, automatic NLE measures that evaluate the generated explanation do not always reflect the correctness, reasoning and semantic meaning of the explanations since explanations can come in different forms and may learn correlations and bias in the dataset.

In summary, we make the following contributions:

- We propose NLX-GPT, a model which can simultaneously predict an answer and explain it, by formulating the answer prediction as a text generation task *along with* the explanation. This eliminates the need for a VL-model to provide an answer and associates the explanation with the reasoning process made to predict the answer.
- Our method outperforms previous works on most metrics while being $15\times$ faster and requiring less memory resources. We further present an ablation analysis of the steps and components of our model, demonstrating that each aspect contributes non-trivially to the final performance of the model.
- We present two new evaluation frameworks for NLE which can reflect the correctness, reasoning, semantic meaning and the degree of biasness of the generated explanations

2. Related Work

The first work on NLE was proposed by [18] for vision tasks. It was then extended to video tasks [22], vision-language tasks [20, 27, 32, 36, 54] and NLP tasks [8, 19, 33, 40]. The authors of [18] propose a discriminative objective that can take into account class-discriminative image properties which justify a classification prediction. Later, [36] proposed two datasets: ACT-X and VQA-X. The former is used to explain decisions of activity recognition models, while the later is used to explain decisions of VQA models. The authors used the MCB VQA model [16] as the task model and an LSTM as the explanation model. The authors of [54] focus on generating more faithful multimodal

explanations and filters out the human textual explanations whose Grad-CAM visual explanation does not align with the Grad-CAM of the predicted answer. Their task model is the Up-Down VQA model [3] and their explanation model is the Up-Down LSTM model [3]. In [27], an automatic dataset of explanations for the VQA task is constructed, in which the explanations are essentially image captions describing the image. [32] uses different vision models to encode the image and then inputs them along with the ground-truth answer and question to a GPT-2 model [39] to generate an explanation. Finally, [20] corrects the eSNLI-VE dataset [14] for explanations of the visual entailment task. Their e-UG model is composed of UNITER [9] as the task model and GPT-2 [39] as the explanation model.

As observed, all these works rely on a task model, which brings the disadvantages discussed previously. Our proposed NLX-GPT tackles these problems by eliminating the task model and produces the answer as a text generation task along with the explanation.

3. NLX-GPT

Consider a task model M_T which performs a specific vision or vision-language task. For example, M_T can be an answering model for VQA, visual entailment or activity recognition. Also, consider an explainer model M_E which provides an explanation for the output of the vision-language model. Previous models [20, 32, 36, 54] stack both M_T and M_E on top of each other. That is, M_T first answers the task, and M_E provides an explanation for the answer. Drawing inspiration from [10, 40], in our work we eliminate M_T and allow M_E to address both objectives. In other words, M_E simultaneously predicts an answer and explains it. This is achieved by formulating the answer prediction as a text generation task *along with* the explanation. This has two advantages: First, it eliminates the high memory requirements of M_T , which can reach to over 300M parameters [9, 28], and reduces the inference time. Second, unlike previous models, where M_T and M_E are completely independent, the generated answer and explanation from our model become associated with each other, in the sense that the explanation is intrinsic, internally affiliated and connected to the reasoning process made to predict the answer. In fact, the answering task may perform even better than having a separate model M_T . For example, our VQA accuracy outperforms the SoA [20] which uses UNITER [9] as M_T by 3%. This shows the strong ability of modeling the answering task as a text-generation task *along with* its explanation and vice-versa. There are other reasons for why removing a VL-model is advantageous. We refer readers to Section 7 of the Appendix for more information. In our case, we choose M_T to be a distilled version [44] of the GPT-2 transformer language model [7]. Our model is an encoder-decoder architecture,

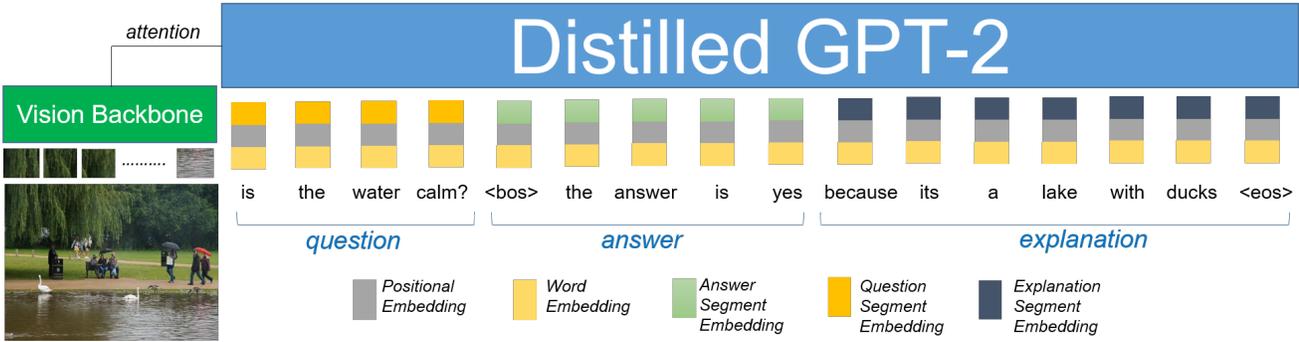


Figure 2. A schema of the proposed NLX-GPT model. At test time, we supply the question and `<bos>` token, and start generating the answer and explanation

as originally proposed in [51]. The encoder is a visual backbone that encodes the image and the decoder is the distilled GPT-2. We consider all sub-inputs (question/hypothesis, answer and explanation) as a single sequence to M_E and train M_E with the cross-entropy objective to generate a sequence $\mathbf{w} = w_1, w_2, \dots, w_T$ of T words (containing both the answer and explanation) by minimizing the negative log-likelihood:

$$\mathcal{L} = - \sum_{t=1}^T \log p_{\theta}(w_t | \mathbf{w}_{<t}), \quad (1)$$

where $\mathbf{w}_{<t}$ denotes the words before word t . A schema of our model is shown in Figure 2. Note that our model can still be used to explain the answer of any external VL-model by simply appending the answer output of that VL-model after the question at the input to our NLX-GPT, which conditions the explanation on that answer. In fact, this becomes a special case of our general model.

At the same time, we aim to deviate away from visual encoders biased towards a specific task (e.g., image classification) as well as from time-expensive bottom-up features [3]. We would instead like to fully rely on grid features. To this, we utilize the CLIP vision encoder [38] since the visual features are (1) non-biased and general and (2) close to the language embedding space of GPT-2 due to the contrastive learning objective of CLIP, which makes the fusing of visual and linguistic information easier. Without region proposals nor a VL-model, our model becomes significantly faster and more memory efficient than previous models, as shown in Table 1. However, we still demonstrate in later sections that even with a ResNet-101 visual encoder pretrained on ImageNet-1K, our model can still outperform previous models that also use a ResNet-101, even without bottom-up object level features [3]. With bottom-up object-level features, our model significantly outperforms the SoA even without using an additional BERT-based multimodal feature extractor (such as UNITER [9]). In the next sub-

sections, we will explain the two stages we conduct on our model: pretraining and finetuning. Please note that our visual encoder is fixed and not fine-tuned at any time during pretraining or finetuning.

3.1. Pretraining

Training an NLE model to explain the decision-making process of an answer given a particular image requires strong image understanding in the first place. Otherwise, the model may be susceptible to learning correlations and bias in the dataset, or overfitting due to the small-scale NLE dataset. In later sections, we make this evident through visualization. It is also shown in [12] how image understanding greatly helps image classification and object detection tasks. Following the trend of vision-language pretraining [9, 23, 28, 60], we pretrain our Distilled GPT-2 on a large-scale corpus of image-caption pairs. We choose image captioning as our pre-training task because (1) it is aligned with our downstream task of text generation and (2) image captions provide comprehensive image understanding by describing objects, their attributes and their relationships. Particularly, we use data from COCO captions [30], Flickr30k [37], visual genome (VG) [25] and image paragraph captioning [24]. In the case of VG region descriptions, we combine region descriptions per image to form a paragraph. More details can be found in the Appendix. We use the cross-entropy objective loss (as in (1)) to train our model to generate a caption c describing an image I in an autoregressive manner. Other pre-training objectives such as image feature regression, masked language modelling and image-text matching can be used; however, we find that the results are already satisfactory with cross-entropy training. It is also shown in [55] that these additional objectives only contribute a little to the overall performance. We pretrain the model with a batch size of 768; we refer to the Appendix for further implementation details.

Table 1. Comparison between other models in terms of the visual encoder, region proposals, VL-model, explanation generator, inference time and total number of parameters. Note that VL-models in previous works act as multi-modal feature extractors and answering models.

Model	Visual Encoder	Region Proposals	VL-model	Explanation Generator	Time (ms)	Parameters (M)
FME	ResNet-101	✓	Up-Down VQA	Up-Down LSTM	~910	142
RVT	ResNet-101	✓	BERT	GPT-2	~925	277
e-UG	ResNet-101	✓	UNITER	GPT-2	~929	277
NLX-GPT	ResNet-101	×	×	Distilled GPT-2	~93	138
NLX-GPT	ViT	×	×	Distilled GPT-2	~55	182

Table 2. Unfiltered Scores for VQA-X and ACT-X. B4, M, R, C, S are short for: BLEU-4, METEOR, ROUGE-L, CIDER and SPICE. Unfiltered scores on e-SNLI-VE are 11.9, 18.2, 32.5, 1.09 and 33.0, respectively.

Approach	VQA-X					ACT-X					Human
	B4	M	R	C	S	B4	M	R	C	S	
CAPS [36]	5.9	12.6	26.3	35.2	11.9	5.2	11.0	26.5	10.4	4.6	22.9
PJ-X [36]	19.5	18.2	43.4	71.3	15.1	15.3	15.6	40.0	22.0	7.2	38.2
FME [54]	24.4	19.5	47.4	88.8	17.9	-	-	-	-	-	-
NLX-GPT (w/o pretraining)	23.8	20.3	47.2	89.2	18.3	25.6	21.4	48.0	63.5	15.4	-
NLX-GPT (w/ pretraining)	25.6	21.5	48.7	97.2	20.2	28.1	22.6	49.7	74.9	17.6	89.0

3.2. Concept Detection

In the e-SNLI-VE dataset [20], we find that pre-training does not help our model generalize. We therefore propose to compensate for image understanding by predicting image concepts with a simple multilayer perceptron (MLP) layer attached on top of the visual encoder and trained on the Visual Genome dataset [25]. Particularly, we use binary cross-entropy as the loss function to train an N -way multi-label classifier, where N is the concept vocabulary:

$$\mathcal{L}_P = - \sum_{i \in \mathcal{B}} \sum_{j=1}^N p_{ij} \log \hat{p}_{ij} + (1 - p_{ij}) \log (1 - \hat{p}_{ij}), \quad (2)$$

where p_{ij} is the j -th target for the i -th datapoint in the batch \mathcal{B} and \hat{p}_{ij} is the sigmoid probability output. After detecting the concepts, we append them to the GPT-2 input and model a conditional language generation task [21]. Further implementation details can be found in the Appendix.

3.3. Finetuning

After the pretraining stage, we finetune our NLX-GPT model on NLE tasks using a batch size of 32. Particularly, we choose NLE for visual-question answering, visual entailment and visual commonsense reasoning (VQA-X [36], e-SNLI-VE [20] and VCR [58]) as vision-language tasks, and NLE for activity recognition (ACT-X) [36] as the vision task. See Section 4 for more details about these datasets.

3.4. Evaluation Measures

In order to evaluate our method, we use three types of evaluation measures. First, we consider the automatic natural language generation (NLG) metrics (BLEU [35], ME-

TEOR [6], ROUGE-L [29], CIDER [52], SPICE [2] and BERTScore [59]); all scores are computed with the publicly available code². Second, we use human evaluation as done in previous works. The process of human evaluation is identical to [20] for VQA-X and e-SNLI-VE, and identical to [36] for ACT-X. We still provide full details about the human evaluation process in the Appendix. One drawback of automatic NLG measures is that they do not always reflect the truthfulness and correctness of the explanations, since explanations can come in different forms as well as be very generic and data-biased. Taking VQA as an example, the data-biasness problem refers to models that can correctly answer a question without considering the image information, which is achieved by learning statistics and correlations in the training data. For example, their answer to the question “What is the color of the grass?” is always “green”. Or “What is the color of the banana?” is always “yellow”. When models are faced with a difficult learning problem, they prefer to learn from the statistical information that links the question with the most occurring answer, completely ignoring the image. By simply memorizing biases in the training data, they exhibit acceptable performance on the test set. As a result, when the model is faced with questions with different image information (e.g. gray grass rather than green grass, or a ripe black banana rather than a yellow banana), they show degraded performance, which also means they are not trustable. The best way to evaluate this phenomenon is conducting human evaluation. However, human evaluation is an expensive and tedious process. To this, automatic measures that better reflect the correctness, context, reasoning, semantic meaning

²<https://github.com/tylin/coco-caption>

Table 3. Filtered scores for VQA-X and e-SNLI-VE. BS stands for BERTScore. † Since the CIDER score is exceptionally high, we suspect a difference in the CIDER idf weights used by the e-VIL [20] authors. We therefore report our scores with a different evaluation API [46]. Our resulting scores on this API on NLX-GPT (w/ concepts) for B1, B4, R-L, M and C are: 35.3, 11.9, 33.1, 18.7, 112.5

VQA-X											
	B1	B2	B3	B4	R-L	M	C	S	BS	Human	Task Acc.
PJ-X [36]	57.4	42.4	30.9	22.7	46.0	19.7	82.7	17.1	84.6	65.4	76.4
FME [54]	59.1	43.4	31.7	23.1	47.1	20.4	87.0	18.4	85.2	63.2	75.5
RVT [32]	51.9	37.0	25.6	17.4	42.1	19.2	52.5	15.8	85.7	67.1	68.6
QA-only [20]	51.0	36.4	25.3	17.3	41.9	18.6	49.9	14.9	85.3	-	-
e-UG [20]	57.3	42.7	31.4	23.2	45.7	22.1	74.1	20.1	87.0	71.5	80.5
NLX-GPT	64.2	49.5	37.6	28.5	51.5	23.1	110.6	22.1	86.9	83.22	83.07
e-SNLI-VE											
	B1	B2	B3	B4	R-L	M	C	S	BS	Human	Task Acc.
PJ-X [36]	29.4	18.0	11.3	7.3	28.6	14.7	72.5	24.3	79.1	59.6	69.2
FME [54]	30.6	19.2	12.4	8.2	29.9	15.6	83.6	26.8	79.7	58.5	73.7
RVT [32]	29.9	19.8	13.6	9.6	27.3	18.8	81.7	32.5	81.1	59.4	72.0
QA-only [20]	29.8	19.7	13.5	9.5	27.0	18.7	80.4	32.1	81.1	-	-
e-UG [20]	30.1	19.9	13.7	9.6	27.8	19.6	85.9	34.5	81.7	68.9	79.5
NLX-GPT (w/o Concepts)	35.7	24.0	16.8	11.9	33.4	18.1	114.7	32.1	80.6	66.3	-
NLX-GPT (w/ Concepts)	37.0	25.3	17.9	12.9	34.2	18.8	117.4 †	33.6	80.8	67.4	73.91

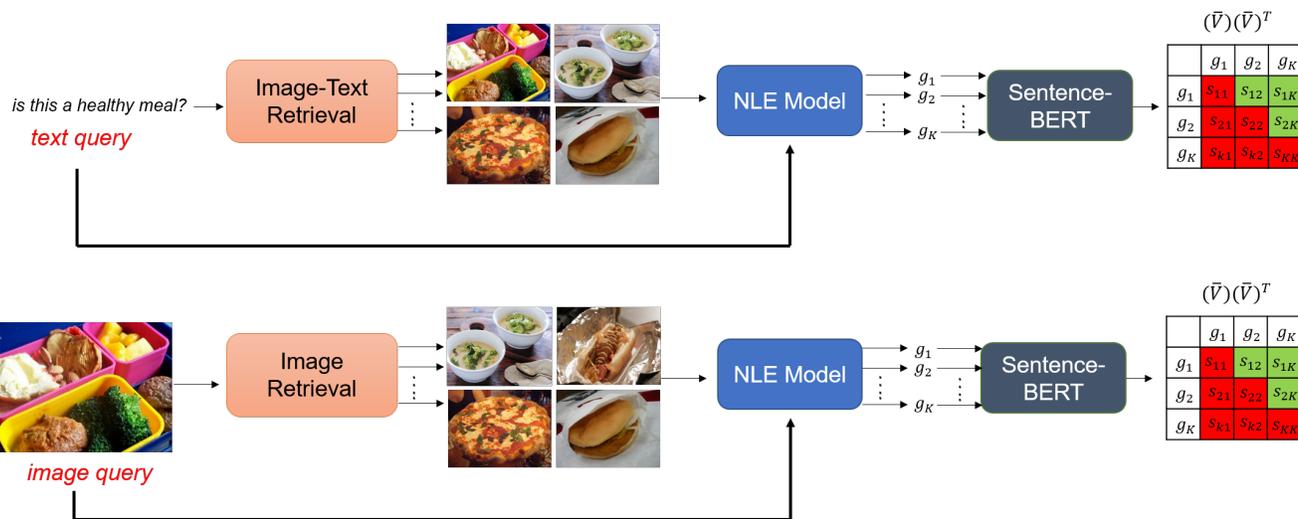


Figure 3. The retrieval-based attack evaluation framework. On top, we show how to measure the biasness to text in vision-language tasks. At the bottom, we show the process for vision tasks.

and the degree of biasness of the generated explanations are needed. We therefore propose two new automatic evaluation measures: (1) Explain-Predict, and (2) Retrieval-based attack. We elaborate on these measures below.

Explain-Predict: This paradigm is firstly introduced in [8]. While [8] uses it as the main model, we use it as an evaluation framework. In this evaluation framework, we measure how good the explanation justifies the answer. We input the question and the generated explanation (without the answer) to a language representation model, and aim to

predict the answer. This measure gives us a degree on the correlation between the explanation and the answer. If the explanation correctly justifies the answer, we assume the language representation model is able to predict it.

We choose DistilBERT [44] to be our language representation model. An example is given in Figure 4 where the input question is: "what animal is this?" and the generated explanation is "it has a long neck and black spots". Note that we append a CLS token at the beginning of the sequence (as in BERT [13]) and a SEP token in between the question and the generated explanation, in order to dis-

tinguish between them. We take the CLS token output as input to a classification layer over all the answers. In this example, the model predicts "giraffe" as an answer, which correctly justifies the generated explanation.

Retrieval-based Attack: In this evaluation framework, we measure the degree our model is susceptible to correlations and bias in the dataset by attacking it with similar inputs. For vision-language tasks, we measure the degree of biasness to both text and images. We will limit our explanation to the biasness to text, as the biasness to images is equivalent. To measure the biasness to text, we consider a text as a query to an image-text retrieval model, and retrieve K similar images close to the text query in the embedding space. For example, the text query could be "is this a healthy meal?" and the retrieved images would be images of different types of meals. Here, we would like to observe the following: given the same question, would the generated explanation always be the same for the different retrieved images? We assume that the same explanations are not due to model reasoning, but rather than correlations and bias in the dataset. In this case we would expect a high cosine intra-distance (distance between the generated explanations of all the retrieved elements). After retrieving the images, we supply our NLX-GPT model with the fixed text query (question) and vary the images (retrieved elements) to generate K different explanations. Given a query, let $\mathcal{G} = \{g_1 \dots g_K\}$ be the set of generated explanations for all the K retrieved elements. We feed each $g_k \in \mathcal{G}$ into a language representation model to get its encoded vector representation. By putting together all encoded representations, we obtain a matrix $V \in \mathbb{R}^{K \times d}$, where d is the dimension of the encoded representation. We first perform L2-normalization on each row of V ; that is, $\bar{v}_k = v_k / \|v_k\|_2$, $\forall k = 1, \dots, K$. We then compute the gram matrix $\bar{V}\bar{V}^T$ of the normalized matrix \bar{V} to find the average cosine distance between each sample with all the other samples as: $s_{\text{avg}} = \frac{1}{K} \sum_{i \in \mathcal{U}} [\bar{V}\bar{V}^T]_i$, and \mathcal{U} is the set of entries in the upper triangular part of $\bar{V}\bar{V}^T$ (see green part of the matrix in Figure 3). Note that the negative distances are clamped to 0. Thus, s_{avg} represents the average intra-distance between the generated explanations of the retrieved elements. The lower the distance is, the lower the bias will be. Therefore, a lower distance is better. We choose Sentence-BERT [41] as the language representation model, a BERT model fine-tuned to contrast between sentence pairs. We use CLIP [38] as the retrieval model. Note that this evaluation framework requires no ground-truth labels, which is advantageous. For vision tasks, since there is no question or hypothesis involved, we utilize an image-retrieval model (i.e., the image part of CLIP [38]) to retrieve similar images to a given image query, and the remaining process is the same. Figure 3 depicts the retrieval-based attack evaluation framework.

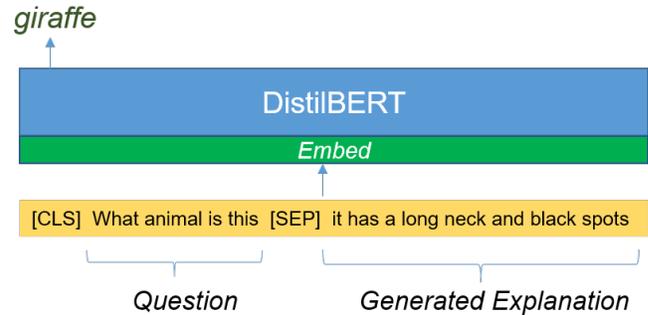


Figure 4. The explain-predict evaluation framework.

4. Experiments and Results

We mainly experiment with 3 different vision and vision-language NLE datasets: VQA-X [36], ACT-X [36] and e-SNLI-VE [20]. We also experiment with the VCR [58] dataset by re-formulating it as a text generation task. Since the setup of VCR is different from the previously mentioned NLE datasets, we exclude it from the main results and refer readers to the Appendix. VQA-X is a vision-language NLE dataset which extends the Visual Question Answering (VQA v2) dataset [1] and provides explanations for the answers. The images are obtained from the COCO dataset [30]. It contains 33K QA pairs (28K images). ACT-X is a vision NLE dataset which extends the activity recognition dataset [4] and provides explanations for the activity answers. It contains 18K images. Finally, e-SNLI-VE [20] is a recently introduced vision-language NLE dataset which provides explanations for the visual entailment prediction task (the task of predicting whether an image and hypothesis *entail*, *contradict* or are neutral to each other), and mainly corrects [14]. The images are obtained from Flickr30k [37]. It contains over 430K examples. There are currently two different ways previous works evaluate NLE with automatic NLG metrics. The first is evaluating all the explanations in the dataset, regardless of whether the predicted answer for the explanation is true or false. This has been used in [36, 54]. We refer to this variant as "unfiltered". The second way is to only consider the explanations for which the predicted answer is correct. This assumes that an explanation is wrong if it justifies a wrong answer, and is therefore not considered in the evaluation. This has been used in [20]. We refer to this variant as "filtered". We evaluate our method on both variants.

4.1. Quantitative Analysis

Table 2 shows our results on the "unfiltered" variant. As seen, our model outperforms all previous works on both VQA-X and ACT-X. Table 3 shows our results on the "filtered" variant. Our model also outperforms all previous

Table 4. Explain-Predict scores. GT indicates that the ground-truth explanations are fed. Scores are in %

Dataset	GT	NLX-GPT
VQA-X	86.35	77.82
e-SNLI-VE	93.10	73.43
ACT-X	65.09	48.14

Table 5. Retrieval-based attack scores. K indicates how many elements we retrieve for a given query. Lower is better.

	Image Biasness		Text Biasness	
	K@5	K@10	K@5	K@10
VQA-X	44.56	44.99	47.84	46.26
e-SNLI-VE	39.79	37.95	67.16	65.69
	Image Biasness			
	K@5	K@10	K@15	
ACT-X	63.78	55.08	48.19	

models on most of the scores on both VQA-X and e-SNLI-VE, while being much more memory efficient, a lot more faster, and requiring no regional and strong multi-modal features. We also show our task performance accuracy scores. For VQA-X, we consider an answer to be correct if it is included in the set of all possible ground-truth answers. In Table 4, we report our results on our proposed explain-predict evaluation framework. We also show the results when feeding the ground-truth explanations (column indicated by "GT"). This gives a measure on the top-performing score. In Table 5, we show our results on our proposed retrieval-based attack framework. These are computed by averaging over all the intra-distances scores for all the test data (images or text) in the respective dataset (lower is better). We report scores when we retrieve $K = 5, 10, 15$ elements for a given query.

4.2. Qualitative Analysis

In Figure 5 we visualize the attention maps for the predicted answer, which is modeled as a text prediction task along with the explanation. In Figure 6, we visualize the attention map for the predicted answer without and with model pretraining. This clearly shows how model pretraining helps the model to reason and visual-ground the correct answer on which the explanation will be conditioned on. In Figure 7, we show some qualitative results from our model on all three tasks. We refer to the Appendix for more qualitative examples, failure cases and retrieval-based attack examples.

4.3. Ablation Studies

We ablate 2 different aspects of our model: First, we wish to observe how much performance does pretraining (for VQA-X and ACT-X) or concept detection (for e-SNLI-



Figure 5. The attention maps for the predicted answer

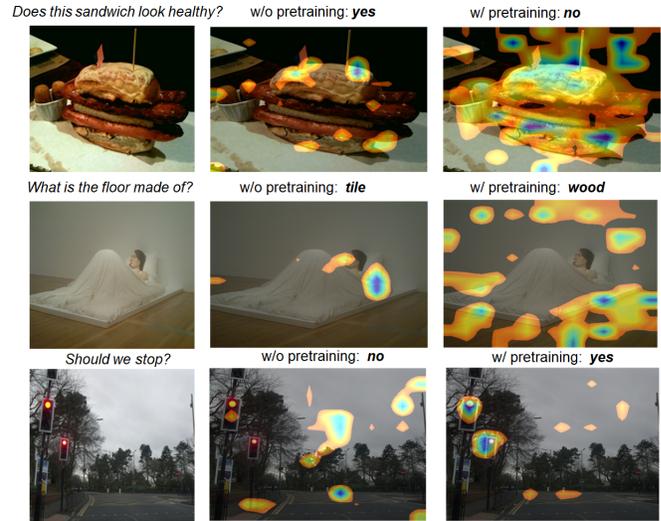


Figure 6. The attention maps for the predicted answer without pretraining (middle) and with pretraining (right)

VE) add to the overall performance. The last two rows in Tables 2 and 3 demonstrate this effect. Second, we ablate different image encoders. Particularly, we experiment with the vision transformer [15], ResNet-101 [17], DeiT [50] and bottom-up features [3]. Table 6 demonstrates our results on these image encoders. The best performance is obtained by the CLIP vision transformer [38]. Different from [47], we find the CLIP vision transformer to outperform the CLIP ResNet-101. We also show the superiority of our model even with a ResNet-101 pretrained on ImageNet-1K [26] or with bottom-up features [3]. Note that the scores in Table 6 are *without* image-caption pretraining.

5. Limitations

There is no free lunch – Although our model brings many advantages, it still has some limitations. It performs slightly worse than [20] on the e-SNLI-VE task for three automatic NLG metrics: METEOR, SPICE and BERTScore. These metrics have shown high correlation with human judgment in [20]. Our model thus favors N-gram (such as BLEU, ROUGE) and human consensus (such as CIDER) metrics

VQA-X



Is it raining?

yes because the ground is wet and people are using umbrellas



What kind of animals are on the man's lap?

cats because they are furry and have whiskers



What beverage is in the cup?

coffee because it is a dark liquid with a dark brown top

ACT-X



Skydiving because he is in the air with a parachute on his back and is falling off



cooking or food preparation because she is standing in a kitchen and cutting up food on a cutting board



Callisthenics because he is in the air doing a push up with his hands

e-SNLI-VE



Hypothesis: The animal is playing outside.
Detected Concepts: [frisbee, red, grass, brown, dog, black, white, leg, bare, orange, tail, tree, ear, dead, paw]
Prediction: entailment because a dog is an animal and field is outside



Hypothesis: The three children are playing hide and seek in the hay
Detected Concepts: [girl, child, shirt, blue, grass, boy, male child, hair, brown, trouser, hand, little, white, young]
Prediction: neutral because just because the children are playing in the hay does not mean they are playing hide and seek



Hypothesis: A girl is sitting at home watching TV.
Detected Concepts: [playground, resort_area, child, girl, shirt, little, slide, kid, blue, pink, white, tree, young]
Prediction: contradiction because the girl cannot be sitting at home and playing on a playground simultaneously

Figure 7. Some qualitative examples from our model on the three tasks

Table 6. Ablation studies on different image encoders. The scores are *filtered* as in Table 3. Also note that the scores are *without* image captioning pretraining. † indicates that we use the CLIP model [39] weights. * indicates the model is pretrained on ImageNet-1K. † means the vision backbone is fine-tuned. Otherwise, the vision backbone is fixed.

Image Backbone	BLEU-4	METEOR	CIDER	ROUGE
ViT†	28.1	22.6	108.5	50.9
ResNet-101†	26.7	22.1	102.3	50.6
DeiT*	26.4	21.8	97.6	49.7
ResNet-101*	22.4	19.5	81.6	45.7
ResNet-101*†	25.2	20.5	95.2	47.3
BU-feats	26.5	22.5	101.1	50.1

on the e-SNLI-VE task and gives more weight to those metrics rather than the others. We also observe that other NLE models on the e-SNLI-VE task give more weight to either the former or the later group of metrics, but not both. This is also observed in [20] where the BLEU-N, R-L and CIDER

scores are lower than ours.

6. Conclusion and Future Directions

We proposed NLX-GPT, a model which can simultaneously predict an answer and explain it by formulating the answer prediction as a text generation task along with the explanation. We also proposed two new evaluation frameworks to better evaluate the generated explanations. In the future, we would like to take advantage of powerful language understanding models through techniques such as distillation, or even by leveraging NLE datasets aimed for NLP tasks, which are much more diverse and of large-scale, and thus can greatly benefit NLE vision-language models. We also expect to see self-critical sequence training [42] or its better variants [31] incorporated.

Acknowledgement: This research has been supported by the Research Foundation - Flanders (FWO) under the Project G0A4720N.

References

- [1] Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Vqa: Visual question answering. *International Journal of Computer Vision*, 123:4–31, 2015. [1](#), [6](#)
- [2] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016. [4](#)
- [3] Peter Anderson, X. He, C. Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6077–6086, 2018. [1](#), [2](#), [3](#), [7](#)
- [4] Mykhaylo Andriluka, Leonid Pishchulin, Peter V. Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693, 2014. [6](#)
- [5] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, 10, 2015. [1](#)
- [6] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *IEEevaluation@ACL*, 2005. [4](#)
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020. [2](#)
- [8] Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. e-snli: Natural language inference with natural language explanations. In *NeurIPS*, 2018. [1](#), [2](#), [5](#)
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, A. E. Kholy, Faisal Ahmed, Zhe Gan, Y. Cheng, and Jing jing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. [2](#), [3](#)
- [10] Jaemin Cho, Jie Lei, Haochen Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *ICML*, 2021. [2](#)
- [11] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10575–10584, 2020. [1](#)
- [12] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11157–11168, 2021. [3](#)
- [13] J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. [5](#)
- [14] Virginie Do, Oana-Maria Camburu, Zeynep Akata, and Thomas Lukasiewicz. e-snli-ve-2.0: Corrected visual-textual entailment with natural language explanations. *ArXiv*, abs/2004.03744, 2020. [2](#), [6](#)
- [15] A. Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, M. Dehghani, Matthias Minderer, Georg Heigold, S. Gelly, Jakob Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2020. [1](#), [7](#)
- [16] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016. [2](#)
- [17] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [1](#), [7](#)
- [18] Lisa Anne Hendricks, Zeynep Akata, Marcus Rohrbach, Jeff Donahue, Bernt Schiele, and Trevor Darrell. Generating visual explanations. In *ECCV*, 2016. [2](#)
- [19] Myeongjun Jang and Thomas Lukasiewicz. Are training resources insufficient? predict first then explain! *ArXiv*, abs/2110.02056, 2021. [2](#)
- [20] Maxime Kayser, Oana-Maria Camburu, Leonard Salewski, Cornelius Emde, Virginie Do, Zeynep Akata, and Thomas Lukasiewicz. e-vil: A dataset and benchmark for natural language explanations in vision-language tasks. *ArXiv*, abs/2105.03761, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [21] Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. Ctrl: A conditional transformer language model for controllable generation. *ArXiv*, abs/1909.05858, 2019. [4](#)
- [22] Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John F. Canny, and Zeynep Akata. Textual explanations for self-driving vehicles. In *ECCV*, 2018. [2](#)
- [23] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. [3](#)
- [24] Jonathan Krause, Justin Johnson, Ranjay Krishna, and Li Fei-Fei. A hierarchical approach for generating descriptive image paragraphs. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3337–3345, 2017. [3](#)
- [25] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123:32–73, 2016. [3](#), [4](#)
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60:84 – 90, 2012. [1](#), [7](#)

- [27] Qing Li, Qingyi Tao, Shafiq R. Joty, Jianfei Cai, and Jiebo Luo. Vqa-e: Explaining, elaborating, and enhancing your answers for visual questions. *ArXiv*, abs/1803.07464, 2018. [2](#)
- [28] Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *ECCV*, 2020. [2](#), [3](#)
- [29] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *ACL 2004*, 2004. [4](#)
- [30] Tsung-Yi Lin, M. Maire, Serge J. Belongie, James Hays, P. Perona, D. Ramanan, Piotr Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [3](#), [6](#)
- [31] Ruotian Luo. A better variant of self-critical sequence training. *ArXiv*, abs/2003.09971, 2020. [8](#)
- [32] Ana Marasović, Chandra Bhagavatula, Jae Sung Park, Roman Le Bras, Noah A. Smith, and Yejin Choi. Natural language rationales with full-stack visual reasoning: From pixels to semantic frames to commonsense graphs. In *FINDINGS*, 2020. [1](#), [2](#), [5](#)
- [33] Sharan Narang, Colin Raffel, Katherine Lee, Adam Roberts, Noah Fiedel, and Karishma Malkan. Wt5?! training text-to-text models to explain their predictions. *ArXiv*, abs/2004.14546, 2020. [1](#), [2](#)
- [34] Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. X-linear attention networks for image captioning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10968–10977, 2020. [1](#)
- [35] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2001. [4](#)
- [36] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8779–8788, 2018. [1](#), [2](#), [4](#), [5](#), [6](#)
- [37] Bryan A. Plummer, Liwei Wang, C. Cervantes, Juan C. Caicedo, J. Hockenmaier, and S. Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *International Journal of Computer Vision*, 123:74–93, 2015. [3](#), [6](#)
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. [3](#), [6](#), [7](#)
- [39] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. [2](#), [8](#)
- [40] Nazneen Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. Explain yourself! leveraging language models for commonsense reasoning. In *ACL*, 2019. [2](#)
- [41] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. [6](#)
- [42] Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1179–1195, 2017. [8](#)
- [43] Fawaz Sammani and Luke Melas-Kyriazi. Show, edit and tell: A framework for editing image captions. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4807–4815, 2020. [1](#)
- [44] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108, 2019. [2](#), [5](#)
- [45] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128:336–359, 2019. [1](#)
- [46] Shikhar Sharma, Layla El Asri, Hannes Schulz, and Jeremie Zumer. Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation. *CoRR*, abs/1706.09799, 2017. [5](#)
- [47] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *ArXiv*, abs/2107.06383, 2021. [7](#)
- [48] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*, abs/1312.6034, 2014. [1](#)
- [49] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. *ArXiv*, abs/1703.01365, 2017. [1](#)
- [50] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. [7](#)
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, L. Kaiser, and Illia Polosukhin. Attention is all you need. *ArXiv*, abs/1706.03762, 2017. [3](#)
- [52] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4566–4575, 2014. [4](#)
- [53] Oriol Vinyals, Alexander Toshev, Samy Bengio, and D. Erhan. Show and tell: A neural image caption generator. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3156–3164, 2015. [1](#)
- [54] Jialin Wu and Raymond J. Mooney. Faithful multi-modal explanation for visual question answering. *ArXiv*, abs/1809.02805, 2019. [1](#), [2](#), [4](#), [5](#), [6](#)
- [55] Qiaolin Xia, Haoyang Huang, Nan Duan, Dongdong Zhang, Lei Ji, Zhifang Sui, Edward Cui, Taroon Bharti, Xin Liu, and

- Ming Zhou. Xgpt: Cross-modal generative pre-training for image captioning. In *NLPCC*, 2021. 3
- [56] Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. Visual entailment: A novel task for fine-grained image understanding. *ArXiv*, abs/1901.06706, 2019. 1
- [57] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6274–6283, 2019. 1
- [58] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6713–6724, 2019. 4, 6
- [59] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. *ArXiv*, abs/1904.09675, 2020. 4
- [60] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J. Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. *ArXiv*, abs/1909.11059, 2020. 3