

COIN: Counterfactual Image Generation for VQA Interpretation

Zeyd Boukhers Timo Hartmann Jan Jürjens
{boukhers,tihartmann,juerjens}@uni-koblenz.de
University of Koblenz-Landau
Koblenz, Germany

Abstract

Due to the significant advancement of Natural Language Processing and Computer Vision-based models, Visual Question Answering (VQA) systems are becoming more intelligent and advanced. However, they are still error-prone when dealing with relatively complex questions. Therefore, it is important to understand the behaviour of the VQA models before adopting their results. In this paper, we introduce an interpretability approach for VQA models by generating counterfactual images. Specifically, the generated image is supposed to have the minimal possible change to the original image and leads the VQA model to give a different answer. In addition, our approach ensures that the generated image is realistic. Since quantitative metrics cannot be employed to evaluate the interpretability of the model, we carried out a user study to assess different aspects of our approach. In addition to interpreting the result of VQA models on single images, the obtained results and the discussion provides an extensive explanation of VQA models' behaviour.

Keywords: ML Interpretability, VQA, GAN, UXE

1 Introduction

Over the past years, the task of Visual Question Answering (VQA) has been widely investigated taking advantage of the development strides of Natural Language Processing and Computer Vision. A VQA model aims to answer a natural language question about the content of an image or one of the appearing objects. Due to the complexity of the task, VQA models are still in the early stage of research and up to our knowledge, is not integrated into any running system. One of the inherent problems of VQA systems is the reliance on the correlation between the question and answer more than the content of the image [6, 31]. Furthermore, the available datasets are usually unbalanced regarding certain types of questions. In the VQAv1 dataset [2], for instance, simply answering tennis to any sports-related question without considering the image yields an accuracy of approximately 40%. This is because the dataset creators tend to generate questions about objects detectable in the image which make the dataset suffer from the so-called “visual priming bias”. For example, blindly answering “yes” to all questions starting

with “Do you see a...?” without considering anything else yields approximately 90% accuracy in the VQAv1 dataset [15, 31]. In practice, this bias is not distinctly perceivable because the users tend to ask similar questions related to the image of the appearing objects and they most likely know the correct answer. For more complex questions or when the user lacks the knowledge expertise of the questions or the image content (e.g. medical domain), it won't be possible to capture the behaviour of the VQA system and whether it is biased. Therefore, it is important to interpret the result of these systems and find what caused the model to output a given answer based on the image-question pair.

Although there is no uniform definition of “Interpretability”, researchers agree that the interpretability of ML models increases the users' trust in ML systems. For VQA models, there exist a limited number of papers that investigate the interpretation of their models. Existing attempts include (i) the identification of visual attributes that are relevant to the question [22, 51] and (ii) the generation of counterfactuals [31, 33, 43].

In the direction of (i), the method proposed by Zhang et al. [51] generates a heatmap over the input image to highlight the image regions that are relevant to the answer of the VQA model. However, as pointed out by Fernández-Loría et al. [10], this approach explains the model's prediction but does not provide a sufficient explanation of its decision. They suggest instead that counterfactual explanations offer a more sophisticated way to increase the interpretability of an ML model because they reveal the causal relationship between features in the input and the model's decision. For example, considering an ML model that classifies MRI images into *Malignant* or *Benign*, the generated heatmap can highlight the key ROI of the model's prediction. Although this heatmap answers the question *What did lead to this decision?*, it does not answer an important question *Why did it lead to this decision?* Consequently, this approach does not provide insight into how the model would behave under alternative conditions. To provide more in-depth interpretability, generating a counterfactual image that is minimally different from the original one but leads to a different model's output would indirectly answer the question *Why did the model take such a decision?* To the best of our knowledge, only three existing methods [6, 33, 43] aim at making VQA models interpretable by providing counterfactual images.

Chen et al. [6] introduce a method that generates counterfactual samples by applying masks to critical objects in images or words in questions. Similarly, Teney et al. [43] suggest a method that masks features in images whose bounding boxes overlap with human-annotated attention maps. Finally, in their ongoing research, Pan et al. [33] propose a framework to generate counterfactual images by editing the original image such that the VQA model returns an alternative answer for a given question. Due to the complexity of the problem, the approach is restricted to colour-based questions. Given a tuple (Image, Question and the VQA’s answer), the approach first finds the question-critical object and then changes its colour so that the VQA system gives a different answer. However, this change is not limited to the question-critical object but all regions with similar colours are changed. Given that the main goal of interpreting VQA systems is to help an expert user understand the behaviour of the VQA model, the approach presented in [33] requires that the user understands the relationship between the image, the question and the answer.

As the user needs interpretation mainly when he lacks the necessary knowledge to understand the relationship between the input and output, we propose to extend the work of Pan et al. [33] by restricting the change to the question-critical object. To this end, this paper introduces an attention mechanism that identifies question-critical objects in the image and guides the counterfactual generator to apply the changes on specific regions. Moreover, a weighted reconstruction loss is introduced in order to allow the counterfactual generator to make more significant changes to question-critical spatial regions than the rest of the image. For further improvement and future work, we made the entire implementation of the guided generator publicly available. In summary, this paper aims to best interpret the output of VQA models by answering the following research questions:

- **RQ1:** How to change the answer of a VQA model with the minimum possible edit on the input image?
- **RQ2:** How to alter exclusively the region in the image on which the VQA model focuses to derive an answer to a certain question?
- **RQ3:** How to generate realistic counterfactual images?

Following this section, Section 2 discusses the related works. Section 3 presents the proposed approach and Section 4 presents the conducted experiments and the obtained results that validate the effectiveness of the proposed approach. Finally, Section 5 concludes this paper and gives insight into future directions.

2 Related Work

This paper addresses the problem of interpreting the outcome of VQA systems. Therefore, we will review in this section the related works divided into three categories: (1)

Interpretable Machine Learning, (2) Visual Question Answering (VQA) and (3) Interpretable VQA.

2.1 Interpretable Machine Learning

Throughout the past decades, the notion of interpretability increasingly gained attention in the ML domain [5, 9, 12, 20, 39, 44, 49, 50]. According to Kim et al. [20], interpretability is particularly important with respect to systems whose decisions have a significant impact such as healthcare, criminal justice systems and finance. Interpretability serves several purposes, including protecting certain groups from being discriminated against, understanding the effect of parameter and input variation on the model’s robustness and increasing the user’s trust in automated intelligent systems [9]. Therefore, a model is considered interpretable if it allows a human to consistently and correctly classify its outputs [20] and understand the reason behind the model’s output [5]. ML models such as decision trees are inherently interpretable, as they provide explanations during training or while the output is generated. However, most of the sophisticated ML models used nowadays such as deep neural networks are not interpretable by nature and require post-hoc explanations [28]

One way to create such explanations is to use global surrogate models, where the aim is to approximate the prediction function f of a black-box and complex model (e.g. neural networks) to the best possible with a prediction function g such that g is the prediction function of an inherently interpretable model (e.g. decision trees or linear regression) [28, 32]. Another way is to use local surrogate models, which individually explain the predictions of a trained ML model to have an overview of its behaviour [32]. Ribeiro et al. [35] propose a Local Interpretable Model-agnostic Explanations (LIME) which approximates the output of black-box models by examining how variations in the training data affect its predictions. Particularly, LIME permutes a trained black-box model’s training samples to generate a new dataset. Based on the black-box model’s predictions on the permuted dataset, LIME trains an interpretable model, which is weighted by the proximity of the sampled instances to the instance of interest.

Feature visualization is another direction to increase the interpretability of black-box models. The goal is to visualize the features that maximize the activation of a NN’s unit [32]. This direction takes advantage of the structure of NNs, where the relevance is backpropagated from the output layer to the input layer [50]. Mainly, most of the approaches under this direction are dedicated to image classification tasks by providing a saliency map that highlights the pixels relevant to the model’s output [27, 28, 40]. A common interpretability direction suggests generating example-based explanations for complex data distributions. The aim is to find prototypes from the training dataset that summarize the prediction of the model [28]. Although this approach can satisfy the user

need for interpretation in simple tasks, it is not practical for most of the real-world data which are heavily complex and seldom contain representative prototypes [20]. Therefore, Kim et al. [20] propose to identify some criticism samples that deliver insights about those prototypes which is not covered by the model.

The problem of post-hoc interpretability methods is their incapability to answer how the model would behave under alternative conditions (e.g. different training data). Therefore, causal interpretability approaches aim at finding why did the model make a certain decision instead of another one or what would be the output of the model for a slightly different input [28]. Here, the goal is to extract causal relationships from the data by analysing whether changing one variable cause an effect in another one [29]. One way to achieve this is by finding a counterfactual input that affects the model’s prediction [8, 13, 16, 18, 41, 45]. For instance, Goyal et al. [16] propose a method that identifies how a given image “ I ” could change so that the image classifier outputs a different class by replacing the key discriminative regions in “ I ” with pixels from an identified “distractor” image “ I_0 ” that has a different class label. Pearl [34] suggests that generating counterfactuals allows for the highest degree of interpretability among all methods to explain black-box models.

2.2 Visual Question Answering (VQA)

Taking advantage of the remarkable advancement of computer vision and natural language processing, due to Deep Neural Networks, several research works have addressed the task of VQA in the past years [1, 2, 6, 11, 15, 17, 24, 42, 48]. According to Antol et al. [2], VQA methods aim at answering natural language questions about an input image. This combination of image and textual data makes VQA a challenging multi-modal task that involves understanding both the question and the image [42]. The answer’s format can be of several types: a word, a phrase, a binary answer, a multiple-choice answer, or a “fill in the blank” answer [17, 42].

In contrast to earlier contributions, recent VQA approaches aim at generating answers to free-form open-ended questions [46]. Agrawal et al. [2], for example, propose a system that classifies an answer to a given question about an image by combining a CNN-based model to extract features from the image and Long Short-Term Memory (LSTM)-based model to process the question. This model, referred to as *Vanilla VQA*, can be considered as a benchmark for DL-based VQA methods [42]. Yang et al. [47] introduce *Stacked Attention Networks (SANs)* that uses CNNs and LSTMs to compute an images’ regions related to the answer based on the semantic representation of a natural language question. Similarly, Anderson et al. [1] propose to narrow down the features in images by using top-down signals based on a natural language question to determine what to look for. These signals are combined with bottom-up signals stemming from a purely visual feed-forward attention mechanism.

Despite the continuous advancements in VQA, several papers suggest that VQA models tend to suffer from the language prior problem, where the models tend to achieve good superficial performances but do not truly understand the visual context [6, 15, 48, 53]. Specifically, Goyal et al. [15] found that in the VQAv1 dataset [2] blindly answering “yes” to any question starting with “Do you see a...?” without taking into account the rest of the question or the image results in an accuracy of 87%. To overcome this, they proposed a balanced dataset to counter language biases, such that for a given triplet (image I ; question Q ; answer A) from the VQAv1 dataset [2], humans were asked to identify a similar image I' for which the answer to question Q is different from A . Similarly, Zhang et al. [48] propose a balanced VQA dataset for binary questions, where for each question, pairs of images showing abstract scenes were collected so that the answer to the question is “yes” for one image and “no” for the other.

Moreover, Chen et al. [6] assume that existing VQA models capture superficial linguistic correlations between questions and answers in the training set and, hence, yields low generalizability. Therefore, they propose a model-agnostic Counterfactual Samples Synthesizing (CSS) training scheme that aims at improving VQA models’ visual-explainable and question-sensitive abilities. The CSS algorithm masks (i) objects relevant to answering a question in the original image to generate a counterfactual image and (ii) critical words to synthesize a counterfactual question. In the same context, Zhu et al. [53] propose a self-supervised learning framework that balances the training data but first, identifies whether a given question-image pair is relevant (i.e., the image contains critical information for answering the question) or irrelevant. This information is then fed to the VQA model to overcome language priors.

These above-discussed problems indicate that the empirical results of VQA models do not reflect their efficacy, especially when promoting VQA models to serve their intended purposes. Specifically, answering questions that the user cannot answer, such as in the healthcare domain. Therefore, it is important to make the output of the VQA model interpretable and not only rely on the evaluation results.

2.3 Interpretable VQA

In most real-world scenarios, human users want to get an explanation along with a VQA system’s output, especially if it fails to answer a question correctly or when the user does not know whether the answer is correct [22]. However, there exist only a few papers addressing the task of interpreting and explaining the outcome of VQA models [7, 22, 31, 33, 51]. Also, most of the existing approaches rarely provide human-understandable explanations regarding the mechanism that led to a given answer. Li et al. [22] introduce a method that simulates the human question-answering behaviour. First,

they apply pre-trained attribute detectors and image captioning to extract attributes and generate descriptions for the given image. Second, the generated explanations are used instead of the image data to infer an answer to a question. Consequently, providing critical attributes and captions to the end-user allows them to understand better what the system extracts from the image. Zhang et al. [51] introduce a heat map-based system to display the image’s regions relevant to the question to the user. To this end, they employed in their model region descriptions and object annotations provided in the Visual Genome dataset [21].

Pan et al. [33] introduce a method that provides counterfactual images along with a VQA model’s output. Precisely, for a given question-image pair, the system generates a counterfactual image that is minimally different from the original image and visually realistic but leads the VQA model to output a different answer for the given question. In its current form, their method is restricted to the context of colour questions. Furthermore, since their model makes edits on a pixel-by-pixel level, the counterfactual images contain changes also in areas irrelevant to a given question. Therefore, the counterfactual image does not provide any meaningful interpretability if the question-critical object has the same relative colour as other objects.

3 Method

In this paper, we propose a method named *COIN* that provides human-interpretable discriminatory explanations to increase a VQA system’s interpretability. The aim is to interpret the outcome of the VQA model by answering the question: “How would the image look like so that the VQA model gives a different outcome?”. Concretely, given an image-question pair $(I; Q)$ and a VQA model $f : (I, Q) \rightarrow A$, where A is its predicted answer, the goal is to train a model \mathcal{G} to generate a new image I' , such that:

$$\mathcal{G} : (I, Q, A) \rightarrow I' \quad | \quad f : (I', Q) \rightarrow A' \quad ; \quad \forall A' \neq A, \quad (1)$$

where I' is the counterfactual image of I . Since there are infinite possible images that can satisfy the above constraint, along the lines of Pan et al. [33], *COIN* aims to tackle the research question **RQ1** by constraining \mathcal{G} to (i) be different as minimally as possible from I , (ii) be visually realistic, (iii) contain semantically meaningful edits and (iv) be applied only on the question-relevant regions. Intuitively, with these constraints, we aim to change the output of the VQA model by applying as minimum as possible changes only on the semantically relevant object so that the user can perceive what can change the output of the VQA model. To this end, we propose to extend the counterfactual GAN introduced by Pan et al. [33] by allowing for semantic edits beyond colour modifications (e.g., editing shapes) and ensuring that only the question-critical regions in an image are altered while

Variable	Description
\mathcal{G}	The counterfactual generator proposed in this paper.
f	The VQA model (i.e. MUTAN [3] in this paper)
I	Original image
Q	Question about I
A	The answer of f to Q given I
h	I ’s height
w	I ’s width
I'	The counterfactual image of I , generated by \mathcal{G}
A'	The answer of f to Q given I'
\hat{I}	An image generated by \mathcal{G}
M	The attention map of I
M'	The attention map of I'

Table 1. Summary of variables used in this paper

retaining the background. The variables used for the system definition are summarized in Table 1.

Figure. 1 illustrates an overview of the proposed architecture. In the depicted example, given the image I , the answer of the VQA model to the question “ Q : What is the colour of the bird?” is “ A : Blue and White?”. To explain this output, \mathcal{G} has to abide by the above-mentioned constraints. To this end, \mathcal{G} consists of several components:

3.1 ROI Guide

To tackle the research question **RQ2**, \mathcal{G} has to be guided to primarily edit regions in I that are relevant to Q . To this end, *COIN* aims to identify the question-critical object in I , but, complex images may contain various objects, of which, usually, only one or a few are relevant when answering a given question. Therefore, an object can be considered to be question-critical if it is key to finding an answer to a given question. For example, given the question “What colour are the man’s shorts?”, the question-critical object in Figure. 2 is the man’s shorts. Therefore, *COIN* aims to guide the generator with a continuous attention map $M \in \mathbb{R}^{1 \times h \times w}$ in the range $[0, 1]$, where h and w correspond to the input image’s height and width, respectively. This map is supposed to highlight the discriminative ROIs of the image I that led $f(I)$ to output the answer A . Thus, instead of generating a counterfactual image I' based on the original image I , the latter is concatenated with the attention map M , such that the concatenation $[I; M]$ serves as an input to the generator \mathcal{G} .

To obtain M , an attention mechanism is used to determine each pixel’s importance regarding the VQA model’s decision. The intuition is to identify the spatial regions in an image that are most relevant to answer a given question. For this reason, *COIN* applies the Gradient-weighted Class




Figure 1. Overview of the proposed architecture inspired by Pan et al. [33]

Activation Mapping (Grad-CAM) algorithm [38] to the VQA model’s final convolution layer because convolution layers retain spatial information that is not kept by fully connected layers. Hence, CNN’s last convolution layer is supposed to have the finest balance between high-level semantics and fine-grained spatial information. Grad-CAM exploits this property by finding the gradient of the most dominant logit (i.e., in the case of a VQA model, this corresponds to the answer with the highest probability) that flows into the model’s final activation map. Furthermore, since Grad-CAM is suitable for various CNN-based models, it can be applied to most VQA models.

Intuitively, the algorithm computes the importance of each neuron activated in the CNN’s final convolutional layer with respect to its prediction. Computing the gradient y^a of the logit corresponding to the VQA model’s predicted answer a with respect to the k th feature map’s activations ϕ^k of a convolutional layer, i.e. $\frac{\delta y^a}{\delta \phi^k}$, reveals the localization map $L_{Grad-CAM}^c \in \mathbb{R}^{u \times v}$ of width u and height v . Next, channel-wise pooling with respect to the width and height dimensions is applied to the gradients. The pooled gradients are then used to weigh the activation channels. Finally, the weighted activations α_k^a reveal each channel’s importance with respect to the VQA model’s prediction [38]:

$$\alpha_k^a = \frac{1}{Z} \sum_i^u \sum_j^v \frac{\delta y^a}{\delta \phi_{i,j}^k} \quad (2)$$

Performing a weighted combination of forward activation maps followed by a ReLU finally yields a coarse saliency map of the same size as the convolutional feature maps [38]:



Figure 2. Example question-image pair from the VQAv1 dataset [2]. The red bounding box indicates the question-critical object.

$$L_{Grad-CAM}^c = \text{ReLU} \left(\sum_k \alpha_k^a \phi^k \right) \quad (3)$$

Finally, to obtain M , the feature maps $L_{Grad-CAM}^c$ are interpolated to match the size of the input image I . Furthermore, a gaussian filter with a mean $\mu = 0$ and a population standard deviation $\sigma = 2$ is applied for improved preservation of the selected image regions’ edges [4].

3.2 Language-Conditioned Counterfactual Image Generation

To drive \mathcal{G} generates a counterfactual image I' such that the corresponding answer $A' \neq A$, *COIN* follows Pan et al. [33] by adopting an architecture based on LingUNet [25], which is an encoder-decoder NN similar to the popular pixel-to-pixel UNet model [36] and maps conditioning language to key intermediate filter weights based on an embedding of natural language text.

Similarly to [33], *COIN* feed \mathcal{G} with language embedding which is the concatenation of the VQA model’s question encoding and answer encoding. The question encoding is represented by the question embedding \bar{q} , which stems from the VQA model’s language encoding for the question Q . The answer encoding is represented by the answer embedding \bar{a} , which is the VQA’s final logits weight vector w.r.t its prediction A for the image-question pair (I, Q) . The goal here is to train \mathcal{G} with the VQA model’s negated cross-entropy for A being the target. Consequently, the generated image I' should contain semantically meaningful differences compared to I , such that the VQA model outputs two different answers for I' and I given the same question Q .

Precisely, \mathcal{G} applies a series of operations to condition the image generation process on language. First, the question embedding \bar{q} and the answer embedding \bar{a} are concatenated to create a language representation \bar{x} . Second, \mathcal{G} applies a 2D 1×1 convolutional filter with weights K_k to each feature map F_j . Each K_k is computed by splitting \bar{x} into m equally sized vectors $\{\bar{x}\}_{j=1}^m$ and applying a 1×1 linear transformation to each of them. Applying the filter weights to each F_j yields the language-conditioned feature maps G_j [25].

Next, LingUNet performs a series of convolution and deconvolution operations to generate a new image \hat{I} . The final counterfactual I' is retrieved as follows:

$$I' = M \odot \hat{I} + (\mathbf{1} - M) \odot I, \quad (4)$$

where \odot denotes the element-wise multiplication and $\mathbf{1}$ is an all-ones matrix with the same dimension. Intuitively, I' is created by incorporating to the original image’s background, the foreground of \hat{I} , which is denoted by pixels with large attention values representing a higher intensity compared to those with low attention values.

3.3 Minimum change

Although I and I' should have distinct semantics with respect to a given question Q , the differences between the two images should be as minimal as possible. To this end, *COIN* incorporates a reconstruction loss, which penalizes the generator for creating outputs different from the input. To ensure that question-critical objects can change their semantic meaning, the generator should be allowed to make significantly more changes in the corresponding image regions (i.e. the foreground) than in the rest of the image (i.e., the background). A

modified ℓ_2 -loss adapted to this purpose, which incorporates the attention map M as a relative weighting term, acts as the reconstruction loss:

$$\ell_2 = \left[\left\| (\mathbf{1} - M) \odot I - (\mathbf{1} - M) \odot I' \right\|_2^2 \right]. \quad (5)$$

Applying a weighted reconstruction loss aims at contributing to the desired traits that (i) the model predominantly edits critical objects and (ii) a relatively ℓ_2 -loss constraint is applied to question-critical regions, allowing for more significant semantic edits. Contrarily, the stricter ℓ_2 -constraint for question-irrelevant regions ensures that the generator retains them nearly unchanged.

3.4 Realism

The counterfactual images generated by \mathcal{G} should be visually realistic. To this end and to tackle the research question (**RQ3**), *COIN* employs a PatchGAN discriminator as proposed by Isola et al. [19]. This discriminator learns to distinguish between real and fake images and penalizes unrealistic generated counterfactual images. The generator and the discriminator are trained in an adversarial manner as in GAN training [14].

Spectral Normalization for stabilize training: Training GANs can suffer from instability and be vulnerable to the problems of exploding and vanishing gradients [23]. In their approach, Pan et al. [33] applied gradient clipping to counter this problem, which requires extensive empirical fine-tuning of the training regime. To bypass this extensive procedure, *COIN* uses spectral normalization [23, 26] to counteract training instability as Miyato et al. [26] suggest that using spectral normalization in GANs can lead to the generated images having a higher quality relative to other training stabilization techniques, such as gradient clipping. Given a real function $\gamma : \mathbb{R} \rightarrow \mathbb{R}$, the Lipschitz constraint is followed if $|\gamma(x_1) - \gamma(x_2)| / |x_1 - x_2| \leq k$, where k is the Lipschitz constant (e.g., $k = 1$). Given a CNN CNN_θ with L layers and weights $\theta = \{w_1, w_2, \dots, w_L\}$, its output for an input x can be computed as [23]:

$$CNN_\theta = a_L \odot l_{w_L} \odot a_{L-1} \odot l_{w_{L-1}} \odot \dots \odot a_1 \odot l_{w_1}(x), \quad (6)$$

where $a_{i=1}^L$ denotes the activation function in the i th layer. Spectral normalization regularizes the convolutional kernels $w_i \in \mathbb{R}^{c_{out} \times c_{in} \times k_w \times k_h}$ with kernel width k_w and height k_h of the fully connected layers l_{w_i} and c_{in} and c_{out} be the input and output channels, respectively. To this end, w_i is first reshaped into a matrix $\hat{w}_i \in \mathbb{R}^{c_{out} \times (c_{in} \times k_w \times k_h)}$, which is then normalized such that the spectral norm $\|\hat{w}_i\|_{sp} = 1 \forall i = 1, \dots, L$. Thereby, the spectral norm is computed as follows [23]:

$$\|\hat{w}_i\|_{sp} = \frac{\hat{w}_i}{u_i^T \times \hat{w}_i \times v_i}, \quad (7)$$

where u_i and v_i denote the left and right singular vectors of \hat{w}_i with respect to its largest singular value.

4 Experiments

In this section, we evaluate the effectiveness of the proposed approach from different aspects, namely, the capability of \mathcal{G} to (i) generate counterfactual image I' such that $f(I'; Q) \neq f(I; Q)$ (**RQ1**), (ii) focus the changes on the object-critical object (**RQ2**), (iii) generate realistic images (**RQ3**). For result reproducibility and further improvements, we made our code and results publicly available under this link ¹

4.1 Dataset

For all our experiments, we used a subset of the VQAv1 dataset’s *Real Images* portion introduced by Agrawal et al. [5] The dataset covers (i) images of everyday scenes with a wide variety of questions about the images and the corresponding ground truth answers. Despite the dataset includes samples with several types of questions, for feasibility reasons, we focus in this experiment on color- and shape-based questions only. This yields a set of 23,469 tuples (Image, Question, Answer). The subset is publicly available under this link ¹

4.2 VQA model

As a VQA model, we employed MUTAN [3], which is trained on VQAv1 dataset. MUTAN achieved overall accuracy of approximately 67% and it performed particularly well on questions with binary *Yes/No* answers (accuracy $\approx 85.14\%$). For quantitative questions (e.g., “*How many ...?*”), it achieved an accuracy of 39.81%. For all other question types, including color and shape-based questions, it achieved an accuracy of 58.52%.

4.3 Evaluation and Results

Automatically and objectively assessing the quality of synthetically generated images is a challenging task [30, 37, 52]. Salimans et al. [37] suggest that there exists no objective function to assess a GAN’s performance. Furthermore, the goal of COIN is to provide an understandable interpretability to the VQA output. The quality of this interpretability can only be assessed by the satisfaction of the user. Therefore, we conducted a user study by presenting the output of COIN (i.e. I' and A') together with I , A and Q . For every sample, the user answers five questions divided in two phases:

- **Phase I:** We present the participant with I' , Q and A' . The participant is requested then to answer the following questions:
 1. *Is the answer correct?* with three possible answers: *Yes*, *No*, and *I am not sure*.
 2. *Does the picture look photoshopped: any noticeable edit or distortion (automatic or manual)?* with five possible answers (i.e. from *Very real* to *Clearly photoshopped*).

- **Phase II:** We present the participant with Q and both images I and I' together with the answers of the VQA model A and A' . The participant is requested then to answer the following questions:

1. *Which of the images is the original?* with three possible answers: *Image 1*², *Image 2*² and *I am not sure*
2. *Is the difference between both images related to the question-critical object?* with three possible answers: *Yes*, *No* and *I am not sure*
3. *Which pair (Image, Answer) is correct?* with four possible answers: *Image 1*², *Image 2*², *Both* and *None*

To make this experience easy for the participants, we built a web application ¹, where 94 participants have participated in the survey answering the above questions for 1320 unique samples. Note that some samples have been treated by more than one participant, but maximum by maximum three, which make the total number of samples 2001. In the following, we present the qualitative and quantitative (obtained from the user study) results of COIN w.r.t the above-mentioned evaluation aspects:

Semantic change (RQ1): Since the final goal of COIN is to generate an counterfactual image that drives the VQA model to give a different answer, we evaluate here the capability of \mathcal{G} to generate these images. Among 12096 counterfactual images generated by \mathcal{G} , 37.82% of them lead f to output an answer $A' \neq A$. In particular, f outputs a different answer for the counterfactual than for the original image for 38.05% of the color-based questions and 25.45% of the shape-based questions. Among the samples treated in our survey, the participants found that the VQA outputs a correct answer A given I and Q for only 45.8% of the presented images, while it outputs a wrong answer for 41.4% of them. In the remaining 12.8% of the images, the participants couldn’t decide because (i) the question was not understood, (ii) the correct answer is not unique or (iii) the answer is only partially correct. After presenting the participants with both images I and I' , the question Q and the answers A and A' , the participants changed their opinions about 484 samples, where they found that A is correct for 70.5% of the presented images. This indicates that the participants could understand the question and answer better after interpreting the result of f . For the generated images, the participants found that f outputs a correct answer A' for only 40.5%.

Figure 3 illustrates qualitative results of \mathcal{G} for color-based questions, where each row represents, from left to right, the original image I , the corresponding map M and the generated counterfactual image I' and its corresponding map M' . The rows 3a to 3c show examples of counterfactual images with different answers than their originals with realistic and understandable changes. As can be noticed in some examples such as Figures 3a and 3b, the VQA model f slightly

¹<https://coin.ai-research.net/>

²Note that we do not give any information about the image.

shifts its attention after the change is applied. This means, that theoretically, f can give a different answer because of focusing on another region after the edit and not because of the edit itself.

For the rest of the samples, f fails to alter their semantic meaning. Figure 3d depicts such an example, where f 's answer does not change when being provided with the counterfactual image. The cause of this failure is due to the difference between the original and counterfactual images is too small for f to change the prediction. This because the clashing constraints that \mathcal{G} has to obey. For example, \mathcal{G} is required to change the output of f but is at the same time restricted to apply as minimally as possible of changes. Another reason might be a wrong gaudiness of the attention map. Suppose the question-critical object accounts for a large portion of the image, or the question is about the image's background. In that case, \mathcal{G} often only edits those parts on which the attention mechanism focuses. As a result, the relevant image region is not modified in its entirety and thus f cannot perceive a semantically meaningful change or its attention is shifted towards other regions of the question-critical object.

While \mathcal{G} can generate semantically meaningful counterfactual images for a lot of color-based questions, it mostly fails in the context of shapes. Figure 4 shows two examples. While the VQA model predicts a different answer for the counterfactual than for the original image in both cases, the changes are not semantically meaningful from a human observer's perspective. In example 4a, the object's shape remains roughly unchanged, while the counterfactual generator slightly edits the sign's color. Furthermore, f predicts an incorrect answer regarding both the original and the counterfactual image. As M indicates, the changes to the original image appear to be significant enough for f to shift its focus slightly to the lower right portion of the sign when making an inference on the counterfactual image.

This shift seems to cause the model to change its prediction. Contrarily, the changes in example 4b are more dominant. \mathcal{G} produces an artifact covering the kite and a segment of the sky surrounding it. As the M show, f focuses on the same area in both the original and the counterfactual image, but the artifact seems to cause it to change its answer. These two instances are exemplary for most of the counterfactual images for shape-based questions: the generator (i) applies only a few edits that are barely noticeable but they are more likely to change the answer of f which is the task of \mathcal{G} or (ii) produces artifacts that are not semantically meaningful for a human observer.

Both the examples in Figure 3 and Figure 4 show that \mathcal{G} 's edits vary depending on the questions and answers. Since the attention maps pose a strong constraint for \mathcal{G} , its edits heavily depend on them. If M focuses on the question-critical object, such as in Figure 3c, \mathcal{G} successfully modifies it. Contrarily, if M focuses only on a small portion of the object in question (such as in Figure 3e), language conditioning does

not have the desired effect. In these cases, \mathcal{G} fails to modify the areas relevant to the question-answer pair sufficiently. Despite the failure to generate the counterfactual image, \mathcal{G} provides an understandable interpretability to the behaviour and result of the VQA model for a particular pair (Image, Question).

Question-critical object (RQ2): One of the main aims of *COIN* is to edit only the question-critical object. This is controlled by the attention map M , where the generator \mathcal{G} is restricted to apply changes predominantly in the image regions on which M focuses. Figure 5 depicts, for three example samples, the original image I (Left), the question Q , the answer A given by MUTAN (f) and the interpolation of the map M with I (Center). The right image is the *background* obtained by computing $(1 - M) \odot I$, where 1 denotes all-ones matrix and \odot is the element-wise multiplication.

In Figure 5a, the question is about the ball and as shown in the interpolated map, f focuses specifically on the ball's region, which makes \mathcal{G} restricted to make changes only on that region. When the region of interest is large and/or sparse such as in Figure 5b, f might not focus on the entire question-critical region but only a portion of it, which is sufficient to answer the question. Figure 5c indicates that f wrongly answered the question as *red and white* but the correct answer is clearly *black*. The obtained map M explains that f was focusing on a different region. Consequently, \mathcal{G} applies the changes on the wrong region. The user can understand the behaviour of the model based on this generated image, which is an edit to the original one w.r.t a wrong region.

According to our user study, \mathcal{G} applied the changes on the critical object in 50.1% of the samples. For 32.2% of the samples, the changes were applied on (1) completely different region, (2) the region of interest and other irrelevant regions or (3) a small part of the region of interest. For the remaining 17.7%, the users couldn't determine whether the changes were applied on the question-critical object or not. From these results, we can derive three main patterns of the attention mechanism:

1. If the object relevant to answering a question is relatively small compared to the rest of the image, the attention mechanism focuses on it completely in most cases. In other words, the computed intensities are higher for pixels belonging to the object than for the rest of the pixels. Under these circumstances, the generator can make larger changes to the entire object than to the rest of the image.
2. Contrarily, if the object is very large or MUTAN pays attention to the background, the projection usually focuses only on a part of it. Consequently, the information that the generator receives allows it to apply more significant changes to a segment of the object or the background than to the rest of it.

(a) **Question:** What color are the peppers in the bottom left corner?



f 's answer: yellow

f 's answer: red

(b) **Question:** What color is the large central flower?



f 's answer: yellow

f 's answer: orange

(c) **Question:** What color is the bird?



f 's answer: blue and white

f 's answer: brown and white


(d) **Question:** What color is the dog?



f 's answer: brown

f 's answer: brown

(e) **Question:** What is the color scheme of the picture?



f 's answer: black and white

f 's answer: orange and white

Figure 3. Example outputs of \mathcal{G} for color-based questions from the VQA_{v1} [2] validation set. Left: original image I and the corresponding attention map M . Right: Generated counterfactual image I' and the corresponding attention map M' .

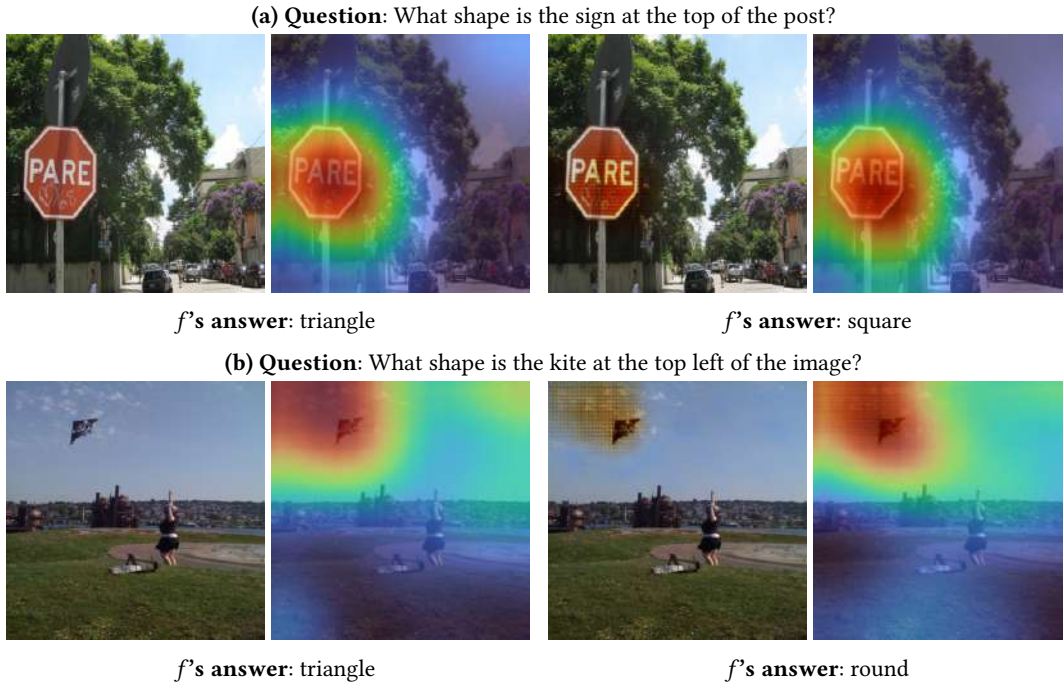


Figure 4. Example outputs of \mathcal{G} for shape-based questions from the VQAv1 [5] validation set. Left: original image I and the corresponding attention map M . Right: Generated counterfactual image I' and the corresponding attention map M' .

3. If MUTAN makes an incorrect prediction, this is often reflected by the projection not focusing on the question-critical object, but another element of the image, such as in Figure 5c.

In all these patterns, the applied changes of I w.r.t M is supposed to change the answer of f for the same question Q . This is because M is indicating where f is focusing to answer Q given I . However, when the VQA focuses on an irrelevant region in I , the applied changes might change the visual semantic of this irrelevant region such that when feeding f with I' and Q , f focuses on a different region than it did in I . This different region can also be the correct region.

Realism (RQ3): Generating realistic counterfactual images is very important to interpret the result of VQA models to users. As shown in Figure 3 and Figure 4, the degree of realism varies depending on the necessary edit that changes f ' answer and on the size difference of the question-critical region to the focused object. This is reflected also in the result of our user study that is demonstrated in Figure 6, where the users were presented only with the generated images and asked “Does the picture look photoshopped (any noticeable edit or distortion)?”. As Figure 6 indicates, the answers of the users vary depending on the generated images. When presenting the corresponding original image together with the generated one and asking “Which of the images is the original?”, the participants could correctly distinguish between the original and the generated one in $\sim 67.2\%$ of

the presented samples. In $\sim 12.9\%$ of the images, the participants selected the generated image as the original one and in the rest of image ($\sim 19.9\%$), the participants couldn't decide which of the images is the original and which one is the generated. This result indicates that \mathcal{G} could trick the human participants by generating counterfactual which look extremely realistic.

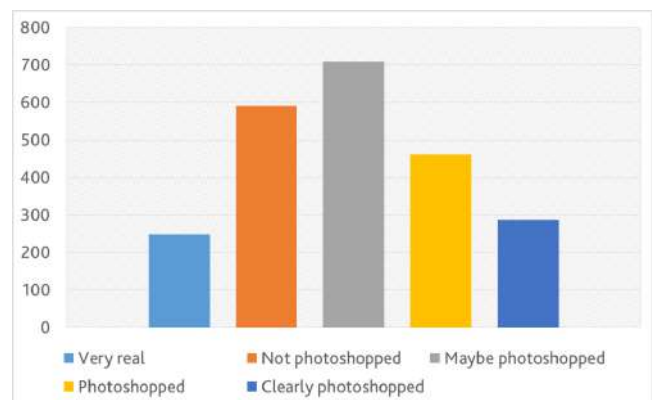


Figure 6. Frequency histogram of participants' answers to the question: “Does the picture look photoshopped?”

Minimality of Image edits: To evaluate \mathcal{G} 's performance on generating counterfactual images with minimum edits,

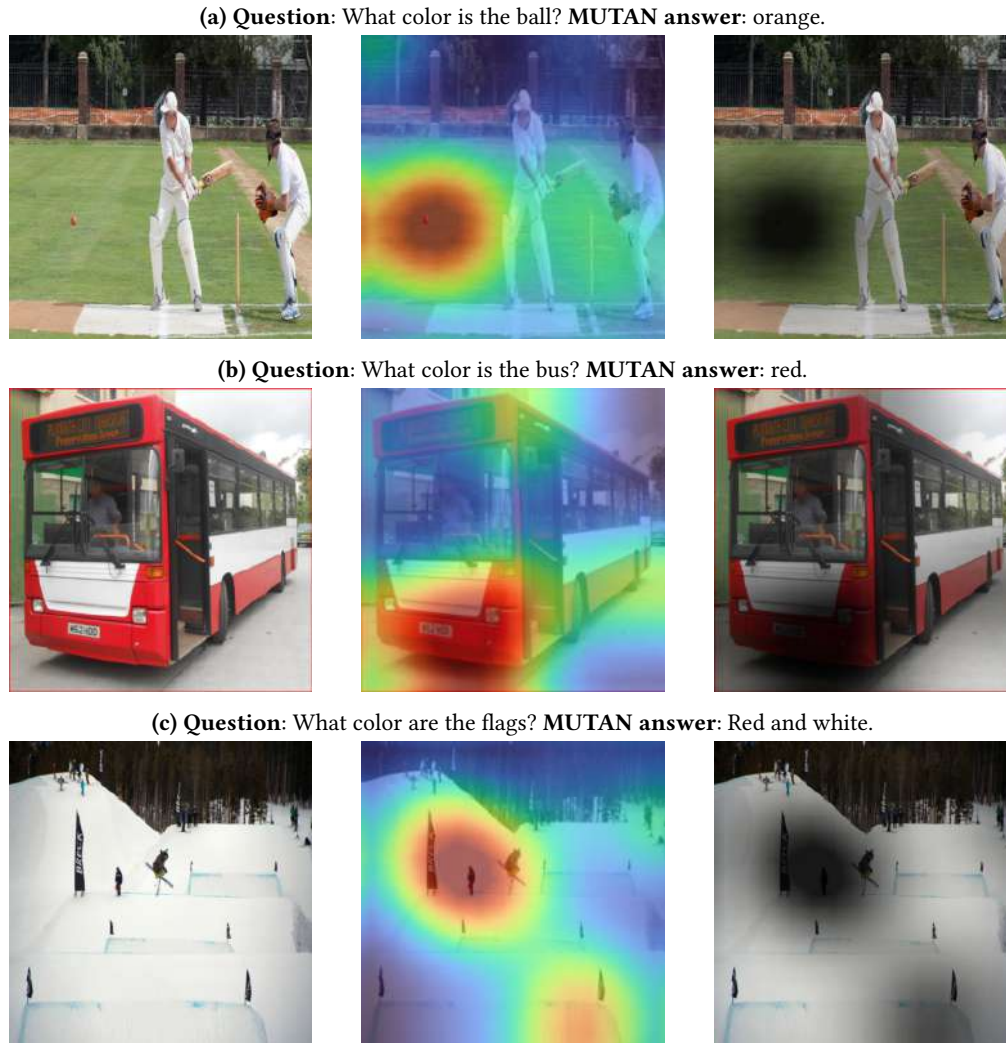


Figure 5. Example outputs of the Grad-CAM algorithm applied to MUTAN for color-based questions. Left: original image. Center: Interpolated attention map projected on the original image. Right: The background image.

we computed ℓ_1 -norm across both the training and the validation set. This measures the magnitude of changes in the counterfactual image relative to the original one, where lower values indicate fewer changes.

Table 2 summarizes the results of this evaluation, where the mean (denoted μ) and standard deviation (denoted σ) values are calculated for different splits of both the training and the validation sets. The first three columns represent the values computed across the entire dataset and for color-based and shape-based questions. The remaining six columns contain the same computations for the portion of pairs of original and counterfactual images for which \mathcal{G} predicts distinct or equal answers, respectively.

The results of Table 2 indicate that \mathcal{G} applied fewer changes when it comes to color-based questions compared to shape-related questions. This observation applies to both the training and the validation set and across all splits. Moreover,

overall, \mathcal{G} applied fewer changes in those cases where f 's predictions regarding the original and counterfactual image were distinct than if they were equal. This indicates that \mathcal{G} changed the original image to the maximum possible level but without successfully changing f 's answer.

5 Conclusion

In this paper, we introduced *COIN*, a GAN-based approach to interpret the output of VQA models by generating counterfactual images to drive the VQA model outputting different a different answer (*RQ1*). *COIN* is a modified implementation of LingUNet with incorporating a Grad-CAM-based attention mechanism that determines each pixel's importance regarding the VQA model's decision making process. With this, the counterfactual generator learns to apply modifications in an image predominantly to question-critical objects,

		All	Color	Shape	Same VQA Answers			Different VQA Answers		
					All	Color	Shape	All	Color	Shape
Training set	μ	0.0175	0.0174	0.0207	0.0177	0.0176	0.0212	0.0173	0.0172	0.0195
	σ	0.0039	0.0039	0.0048	0.0040	0.0039	0.0049	0.0039	0.0038	0.0046
Validation set	μ	0.0175	0.0174	0.0208	0.0177	0.0176	0.0212	0.0173	0.0173	0.0198
	σ	0.0041	0.004	0.0047	0.0041	0.0041	0.0048	0.0038	0.0038	0.0042

Table 2. Mean (μ) and standard deviation (σ) of the ℓ_1 -norm computed across the training and validation set and split across categories

while retaining the rest of the image (**RQ2**). The obtained results indicate that using an attention mechanism is an appropriate means to guide the modification process. Furthermore, the quality of the counterfactual images depended to a large extent on the attention maps. Extensive experiments on the challenging VQAv1 dataset have demonstrated that *COIN* achieves encouraging results for color questions by generating realistic counterfactual images (**RQ3**).

For future work, we will train *COIN* on a larger, more diverse dataset such as VQAv2 dataset, which contains multiple images per question rather than only a single one as in the VQAv1 dataset. Moreover, using an attention mechanism that focuses on the question-critical objects more accurately could also significantly improve the interpretability capabilities of *COIN*. To this end, we plan to employ super-pixel segmentation to extract concepts (e.g. color, texture, or a group of similar segments) and uses the Shapley Value algorithm to determine each concept’s contribution to a DNN’s decision. The generator will then be trained to alter the most important concept(s) in an image rather than providing it an attention map. Moreover, replacing an entire instance of a concept rather than editing an image on a pixel-by-pixel level could pave the way for semantic changes even larger than altering shapes.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6077–6086.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 2425–2433.
- [3] Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. 2017. Mutan: Multimodal tucker fusion for visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 2612–2620.
- [4] Fredrik Bergholm. 1987. Edge focusing. *IEEE transactions on pattern analysis and machine intelligence* 6 (1987), 726–741.
- [5] Supriyo Chakraborty, Richard Tomsett, Ramya Raghavendra, Daniel Harborne, Moustafa Alzantot, Federico Cerutti, Mani Srivastava, Alun Preece, Simon Julier, Raghuvver M Rao, et al. 2017. Interpretability of deep learning models: A survey of results. In *2017 IEEE smartworld, ubiquitous intelligence & computing, advanced & trusted computed, scalable computing & communications, cloud & big data computing, Internet of people and smart city innovation (smartworld/SCALCOM/UIC/ATC/CBDcom/IOP/SCI)*. IEEE, 1–6.
- [6] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. 2020. Counterfactual samples synthesizing for robust visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10800–10809.
- [7] Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. 2017. Human attention in visual question answering: Do humans and deep networks look at the same regions? *Computer Vision and Image Understanding* 163 (2017), 90–100.
- [8] Emily Denton, Ben Hutchinson, Margaret Mitchell, Timnit Gebru, and Andrew Zaldivar. 2019. Image counterfactual sensitivity analysis for detecting unintended bias. *arXiv preprint arXiv:1906.06439* (2019).
- [9] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).
- [10] Carlos Fernández-Loría, Foster Provost, and Xintian Han. 2020. Explaining data-driven decisions made by AI systems: the counterfactual approach. *arXiv preprint arXiv:2001.07417* (2020).
- [11] Donald Geman, Stuart Geman, Neil Hallonquist, and Laurent Younes. 2015. Visual Turing test for computer vision systems. *Proceedings of the National Academy of Sciences* 112, 12 (2015), 3618–3623.
- [12] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.
- [13] Oscar Gomez, Steffen Holter, Jun Yuan, and Enrico Bertini. 2020. ViCE: visual counterfactual explanations for machine learning models. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 531–535.
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [15] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6904–6913.
- [16] Yash Goyal, Ziyang Wu, Jan Ernst, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Counterfactual visual explanations. In *International Conference on Machine Learning*. PMLR, 2376–2384.
- [17] Akshay Kumar Gupta. 2017. Survey of visual question answering: Datasets and techniques. *arXiv preprint arXiv:1705.03865* (2017).
- [18] Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. 2018. Generating counterfactual explanations with natural language. *arXiv preprint arXiv:1806.09809* (2018).
- [19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [20] Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. 2016. Examples are not enough, learn to criticize! criticism for interpretability.

- Advances in neural information processing systems* 29 (2016).
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2016. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332* (2016).
- [22] Qing Li, Jianlong Fu, Dongfei Yu, Tao Mei, and Jiebo Luo. 2018. Tell-and-Answer: Towards Explainable Visual Question Answering using Attributes and Captions. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 1338–1346.
- [23] Zinan Lin, Vyas Sekar, and Giulia Fanti. 2020. Why Spectral Normalization Stabilizes GANs: Analysis and Improvements. *arXiv preprint arXiv:2009.02773* (2020).
- [24] Mateusz Malinowski and Mario Fritz. 2014. A multi-world approach to question answering about real-world scenes based on uncertain input. *Advances in neural information processing systems* 27 (2014), 1682–1690.
- [25] Dipendra Misra, Andrew Bennett, Valts Blukis, Eyvind Niklasson, Max Shatkhin, and Yoav Artzi. 2018. Mapping Instructions to Actions in 3D Environments with Visual Goal Prediction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2667–2678.
- [26] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. Spectral Normalization for Generative Adversarial Networks. In *International Conference on Learning Representations*.
- [27] Grégoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller. 2017. Explaining nonlinear classification decisions with deep taylor decomposition. *Pattern Recognition* 65 (2017), 211–222.
- [28] Raha Moraffah, Mansooreh Karami, Ruocheng Guo, Adrienne Raglin, and Huan Liu. 2020. Causal interpretability for machine learning-problems, methods and evaluation. *ACM SIGKDD Explorations Newsletter* 22, 1 (2020), 18–33.
- [29] W James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. 2019. Interpretable machine learning: definitions, methods, and applications. *arXiv preprint arXiv:1901.04592* (2019).
- [30] Seonghyeon Nam, Yunji Kim, and Seon Joo Kim. 2018. Text-adaptive generative adversarial networks: manipulating images with natural language. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. 42–51.
- [31] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12700–12710.
- [32] Caio Nóbrega and Leandro Marinho. 2019. Towards explaining recommendations through local surrogate models. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*. 1671–1678.
- [33] Jingjing Pan, Yash Goyal, and Stefan Lee. 2019. Question-conditioned counterfactual image generation for vqa. *arXiv preprint arXiv:1911.06352* (2019).
- [34] Judea Pearl. 2018. Theoretical Impediments to Machine Learning With Seven Sparks from the Causal Revolution. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 3–3.
- [35] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
- [37] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training gans. *Advances in neural information processing systems* 29 (2016), 2234–2242.
- [38] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*. 618–626.
- [39] Ravid Shwartz-Ziv and Naftali Tishby. 2017. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810* (2017).
- [40] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013).
- [41] Kacper Sokol and Peter A Flach. 2018. Glass-Box: Explaining AI Decisions With Counterfactual Statements Through Conversation With a Voice-enabled Virtual Assistant.. In *IJCAI*. 5868–5870.
- [42] Yash Srivastava, Vaishnav Murali, Shiv Ram Dubey, and Snehasis Mukherjee. 2019. Visual question answering using deep learning: A survey and performance analysis. *arXiv preprint arXiv:1909.01860* (2019).
- [43] Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. 2020. Learning what makes a difference from counterfactual examples and gradient supervision. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X* 16. Springer, 580–599.
- [44] Michael Tsang, Hanpeng Liu, Sanjay Purushotham, Pavankumar Murali, and Yan Liu. 2018. Neural interaction transparency (nit): Disentangling learned interactions for improved interpretability. *Advances in Neural Information Processing Systems* 31 (2018), 5804–5813.
- [45] Sahil Verma, John Dickerson, and Keegan Hines. 2020. Counterfactual explanations for machine learning: A review. *arXiv preprint arXiv:2010.10596* (2020).
- [46] Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2017. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding* 163 (2017), 21–40.
- [47] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 21–29.
- [48] Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2016. Yin and yang: Balancing and answering binary visual questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5014–5022.
- [49] Quanshi Zhang, Xin Wang, Ying Nian Wu, Huilin Zhou, and Song-Chun Zhu. 2020. Interpretable CNNs for object classification. *IEEE transactions on pattern analysis and machine intelligence* 43, 10 (2020), 3416–3431.
- [50] Quan-shi Zhang and Song-chun Zhu. 2018. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (2018), 27–39.
- [51] Yundong Zhang, Juan Carlos Niebles, and Alvaro Soto. 2019. Interpretable visual question answering by visual grounding from attention supervision mining. In *2019 IEEE winter conference on applications of computer vision (wacv)*. IEEE, 349–357.
- [52] Dawei Zhu, Aditya Mogadala, and Dietrich Klakow. 2021. Image manipulation with natural language using Two-sided Attentive Conditional Generative Adversarial Network. *Neural Networks* 136 (2021), 207–217.
- [53] Xi Zhu, Zhendong Mao, Chunxiao Liu, Peng Zhang, Bin Wang, and Yongdong Zhang. 2020. Overcoming language priors with self-supervised learning for visual question answering. *arXiv preprint arXiv:2012.11528* (2020).