

IMPROVING USERS' MENTAL MODEL WITH ATTENTION-DIRECTED COUNTERFACTUAL EDITS

Short title: Attention-based counterfactuals for VQA

AUTHORS

KAMRAN ALIPOUR¹, ARIJIT RAY², XIAO LIN², MICHAEL COGSWELL², JURGEN P. SCHULZE¹, YI YAO², GIEDRIUS T. BURACHAS²

¹ University of California San Diego, La Jolla, California

² SRI International, Princeton, New Jersey

ACKNOWLEDGMENTS

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA) under the Explainable AI (XAI) program. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

Abstract:

In the domain of Visual Question Answering (VQA), studies have shown improvement in users' mental model of the VQA system when they are exposed to examples of how these systems answer certain Image-Question (IQ) pairs¹. In this work, we show that showing controlled counterfactual image-question examples are more effective at improving the mental model of users as compared to simply showing random examples. We compare a generative approach and a retrieval-based approach to show counterfactual examples. We use recent advances in generative adversarial networks (GANs) to generate counterfactual images by deleting and inpainting certain regions of interest in the image. We then expose users to changes in the VQA system's answer on those altered images. To select the region of interest for inpainting, we experiment with using both human-annotated attention maps and a fully automatic method that uses the VQA system's attention values. Finally, we test the user's mental model by asking them to predict the model's performance on a test counterfactual image. We note an overall improvement in users' accuracy to predict answer change when shown counterfactual explanations. While realistic retrieved counterfactuals obviously are the most effective at improving the mental model, we show that a generative approach can also be equally effective.

Keywords: Explainable AI, Visual Question Answering, Counterfactual

1 | INTRODUCTION

With the growing application of AI in high-risk domains, it is important for human users to understand the extent and limits of AI system competencies to ensure efficient and safe deployment of such systems. While deep neural networks have made impressive strides, they are notorious for being unpredictable to a human user as to when they succeed or fail in producing correct outputs. Hence, we need effective approaches to improve the end users' mental model of the deep neural network-based AI systems.

There has been work in literature² that shows humans can improve their mental models by mere exposure to the system predictions for a variety of inputs. A mental model is a person's internal representation of the AI system she is interacting with and ideally builds a correct understanding of the way that system works³. In this paper, we explore the various ways we can present such explanatory additional input-output examples to a user to maximize their mental model improvement. We ask the question: are certain examples of how the machine behaves better than other examples to teach humans when to trust the model and when not to? We examine the effect of exposing the users to explanatory examples where the inputs are changed in a controlled manner in order to better observe how the machine output changes to controlled changes in the input. We call these

controlled changes in input, “counterfactuals”. We hypothesize that such controlled changes in the examples shown are better for mental model improvement than showing random examples.

Many approaches^{4,5,2} to improving mental models also focus on using explanations that aid the user in understanding how a deep network arrived at a certain conclusion. While many existing explanation approaches such as attention maps attempt to provide insights into the inner working of AI machines, they don’t necessarily convey the causal chain of inference that happens in the algorithm⁶. As a result, the research community actively seeks novel explanation modalities to probe the causality of AI as this form of explanation can resonate better with human logic. Humans tend to learn better from explanations that easily convey when a machine is about to be correct and when not⁶. Among different techniques, showing counterfactual examples are considered *human-friendly* explanations because they are contrastive and also selective when showing the feature changes⁷. Counterfactuals provide the opportunity for the user to explore the range of responses from AI as they manipulate certain features of the inputs and the conditions.

In this paper, we focus on improving users’ mental models by generating counterfactual explanations for the task of visual question answering (VQA) - answering natural language questions asked about images. Specifically, we compare various methods of generating counterfactual examples to maximize a user’s accuracy in predicting when a model is about to fail or succeed.

In this setting, given an image-question pair, a counterfactual explanation is showing the output of the model for the same question but on a different image where the answer should be different. For example, as shown in Figure 1, for the question “what sport is being played?” on the original image of playing tennis, the counterfactual examples could be showing the answer of the model on an image where someone is playing baseball (middle image), or where a tennis racket is absent (rightmost image).

Specifically, we compare a retrieval-based approach and a GAN-based approach to generate counterfactual images for a given question. For example, as shown in Figure 1, we can generate a counterfactual image (an image where the answer may be different from the original image) by either retrieving an image where the answer is different (middle image) or by removing the tennis racket using a GAN network (rightmost image). Our automated approach using a GAN provides the opportunity to produce counterfactuals at scale and evaluate their effectiveness on a large population of AMT workers.

One major challenge of automatically generating counterfactual images is that we are limited by the capability of current GAN models. We chose to use an in-filling network⁸ to remove parts of the image since we observed that current networks can achieve this with a reasonable performance. Limited by the capability of only being able to remove parts of images, we need to decide the most effective parts of images to remove to generate counterfactual examples that help users to learn the idiosyncrasies of the model in order to improve their mental model. To this regard, we experiment with using attention maps (heatmaps that point to where a machine looks at while answering the question) to decide relevant and irrelevant parts of the image to remove to generate counterfactual images.

In summary, our contributions include:

1) proposing effective ways to generate counterfactual examples: We outline several ways of generating counterfactual images. Specifically, we compare a retrieval-based method and an automated GAN-based method to generate counterfactual images.

2) we evaluate empirically the effectiveness of counterfactual examples relative to providing random examples. We show an improvement in the mental model of users when showing controlled counterfactual examples as compared to simply showing random examples or no examples at all.

In the following sections, we first take a look at the related work. We then discuss the methodology behind our approach. We then cover the details for our hypotheses and experimental designs. Finally, we provide the results from our studies and discuss our interpretations.

2 | RELATED WORK

VQA/Explanations Our approach is based on interactions with a visual question answering (VQA)⁹ machine. The use of attention-based layers and explanations in VQA has been a highly popular approach^{10,11,12,4}. Previous work in the attention-based VQA includes attempts to improve the attention mechanism through co-attention between image and question¹³, or a combined bottom-up and top-down to compute object-level attentions¹⁴. In recent work, Peng *et al.*¹⁵ propose a Multi-modal Relation Attention Network (MRA-Net) model with textual and visual relation attention for higher performance and interpretability. Patro

et al. utilize adversarial training of the attention regions as a two-player game between attention and explanation¹⁶. We adopted a VQA model similar to what was proposed by Alipour *et al.*⁴ where the attention is derived from a transformer model¹⁷.

Counterfactuals Counterfactual examples have also been used to explain image classifiers^{8,18}. They have also been used in an optimization process where¹⁹ proposed a loss function to find the minimum changes in the input that results in a change in the output of a classifier. Using counterfactual images as explanations can also be thought of as the visual equivalent to observing VQA behavior by rephrasing the question and checking if the model responds consistently^{20,21,22}. Hence, such counterfactual images hint at how consistent these models are to users and that aids in their mental model improvement.

Mental model evaluation Some of the previous studies introduce metrics to measure trust with users^{23,24}, or the role of explanations to achieve a goal^{25,26,5}. Dodge *et al.* investigated the fairness aspect of explanations through empirical studies²⁷. Lai and Tan²⁸ examined the role of explanations in user success within a spectrum from human agency to full machine agency. Lage *et al.* proposed a method to evaluate and optimize human-interpretability of explanations based on measures such as size and repeated terms in explanations²⁹. Other approaches measured the effectiveness of explanations in improving the predictability of a VQA model^{2,1}.

In this work, we develop a series of user studies with a subject population of lay users with minimum knowledge about AI. The experiments are designed to investigate effective methods to produce counterfactuals that can improve the user's mental model of a VQA system.

3 | METHOD

In this section, we first describe our VQA model and then explain how we generate counterfactual images using a GAN.

3.1 | VQA Model

Our VQA model is trained based on the VQA 2.0 dataset³⁰ and is capable of answering questions about images in textual format. The model is a transformer-based neural network that can parse a combination of visual and textual embeddings from an image and question. The model encodes the image into a 49×512 feature map with the help of ResNet152³¹. The objects in the image are also encoded separately into a 36×512 feature map using a Region Proposal Network³². The model accepts questions with a maximum length of 30 words and all questions below this limit are padded with 0's. The question array is also embedded to a 30×512 vector of features.

The model employs transformer-based attention layers that receive all the visual, object, and textual features in the concatenated shape of 115 ($30 + 36 + 49$) tokens. The transformer is comprised of four layers with 12 heads in each layer. Consequently, the model can provide an attention tensor between these tokens with a $4 \times 12 \times 115 \times 115$ dimension. The model provides its prediction as a softmax probability distribution over 3129 answer choices from the attention-weighted feature values.

For our experiments, we use a subset of the VQA 2.0 validation dataset. We first show the VQA model's answer to the original images and questions from this subset. For each example, we also show the answer to two counterfactual images for the question to the user. We finally test the user's mental model by asking the user to predict the correctness of the answer on a test image for the same question. We will now describe how we generate counterfactual images.

3.2 | Generating Counterfactual Images

We generate counterfactual images to serve as examples of VQA behavior under differing inputs to improve the mental model of users. For example, a user who sees a VQA not counting oranges properly when changing the number of oranges in a picture and asking "How many oranges?" will learn that the VQA model has a low accuracy for counting oranges. This sort of mental model improvement might not have been noticed if we presented the user with only one counting example and other random examples of images and questions. In our study, we focus on altering objects in the image for a certain question. Specifically, we use a GAN which has been trained to in-paint areas of the image such that it looks natural⁸. When asked to in-paint an area of a certain object in an image, such a GAN would usually omit the object and in-paint its content that matches the background/surrounding scenery. We use such an approach to remove objects from the scene. However, such approaches are currently noisy and we often note artifacts in the image that make it seem unnatural. Hence, we limit the size of all in-painting bounding boxes to 10% to 20% of the whole image area.

How to choose objects to in-paint

The ultimate goal for this algorithm is to generate counterfactual explanations that are helpful to the users in predicting AI's response. Given the diverse combinations and interactions between objects in a real scene, it is not obvious how to define an algorithm to select and alter the objects from images to maximize mental model improvement. In our approach, we use attention maps - heatmaps that convey the important regions of the image for answering the question - to decide objects to remove in the image. We use two different sources of attention to identify the in-paint candidates and then produce the counterfactual images based on them:

- Human annotated attentions for the image-question pair, which come from the human attention dataset.
- The attention layers from the AI system. As described in Section 3.1, our VQA system has multiple layers of attention that weigh the image and question features. We select the weights from the last layer (averaging over the transformer heads) to display the attention values over the image regions. The attention values over the image regions are also computed as the average weight over all question tokens.

Based on the above-computed attention maps, we generate two counterfactual images- 1) we remove a box that falls in a region of high attention, and 2) we remove a box that falls in a region of low attention. This ensures we remove a relevant and irrelevant object in the image to introspect how the VQA model's answer changes. Based on this observation, a user can hypothetically learn whether the VQA model is behaving rationally or not. To select the low and high attention boxes, we employ a threshold that first segments the high and low attention regions from the attention map. The bounding boxes surrounding these regions provide the in-paint area. In cases where the bounding boxes are outside the limits (10% - 20% of the image area), the proposed box is scaled to a size within the range.

4 | EXPERIMENTAL SETTINGS

We conduct experiments to quantify the improvement in the mental model for users after being exposed to counterfactual explanations. We measure the mental model of a user by asking them to predict the answer change or the correctness of the answer for a given image-question (IQ) pair, similar to concurrent studies on user mental model evaluation^{2,1}. We use the Amazon Mechanical Turk platform to recruit users for our study. We recruit workers who are located in the United States (due to IRB regulations) and who exceed 98% approval rating on over at least 50 such human-intelligence tasks (HITs).

In our study, each user goes through 1 HIT which consists of 20 episodes of IQ pairs. In each episode, the users first see the VQA model's response to the original Image-Question (IQ) pair. Based on their group configuration, then they may or may not see two counterfactual forms of the original image and AI's response to the original question for those counterfactual images. In the evaluation section of each episode, users attempt to predict AI's response to a test image for the same question. We quantify user's mental model states based on their accuracy in predicting AI's response.

We use two tasks to measure a user's mental model - a) answer change prediction to see if users can predict if the answer will change when a certain object is removed, and b) answer correctness prediction on a real test image based on the lessons learned from counterfactual examples.

4.1 | Answer change prediction

In this setting, we show an IQ pair to a user, the VQA model's answer on the original IQ. One group of users - Counterfactual Group (CF Group) - sees two examples of objects being removed from the image and the VQA model's answer on these two altered images. The Control Group of users see no such altered examples. Finally, both the groups of users are presented with another new object removed from the same image and are asked to predict if the VQA model's answer for that image will change from the original image or not.

In the experiment, the counterfactuals were generated based on human attention³³ annotations from the VQA-HAT dataset.

Note that the object removed in the test image is always different from the objects removed in the counterfactual examples shown.

We do this by choosing separate regions of minimum (min), maximum (max), or medium (mid) attention based on the human-attention values on the image. While the Min and Max regions are determined by extracting the areas from the two sides of the attention spectrum (see figure 2), Mid area is identified by avoiding the overlap with Min and Max and also maintaining the minimum attention possible. This procedure of in-painting assures the minimum overlap among the sample and the test in-paintings and therefore minimizes the chance of overlap between counterfactual samples and test images. While testing the

users, we randomly choose to show two of min, mid, and max as counterfactual examples and test on the unseen third. The group **CF-MinAtt** shows mid and max attention as samples and tests on the image with min attention region removed. Similarly, **CF-MidAtt** tests on image with mid attention region removed, and **CF-MaxAtt** tests on the image with max attention region removed.

4.2 | Answer correctness prediction

In the second set of experiments, we provide a more realistic setting to evaluate the user mental model. Instead of predicting an answer change for a counterfactual test image, the users now attempt to predict whether the model will answer the same question correctly for a different test image. Since the test images are also selected from the IQ pairs in the VQA dataset, that guarantees that the test image is relevant for the question asked.

We define four groups (shown in table 1) to check whether counterfactual examples improve users' mental models to be able to predict the model's correctness on an unseen test image:

- Control Group (**CG-NoExp**) sees no explanations and just the VQA model's answer on an IQ pair.
- the counterfactual group either based on counterfactual images generated using human-annotated attention (**CF-HAT**) or VQA model's attention (**CF-AIAtt**). These groups are to examine how the process of generating counterfactual images affects the mental model.
- a group that sees retrieved real counterfactual images (**CF-AltImg**). We retrieve images that are relevant to the question but have a different answer from the VQA dataset. We can think of these as ideal counterfactual examples. The performance of this group compared to the CF-HAT and CF-AIAtt groups would tell us if generated counterfactuals (CF-HAT and CF-AIAtt) can be used in place of real counterfactual data to reduce dataset collection costs.
- a group that sees random IQ pairs instead of counterfactuals (**CG-RandIQ**). This group is to understand how much we gain from simply presenting two samples of random IQ pairs instead of two samples of counterfactual IQ examples.

Note that in all cases, we make sure all images are relevant to the question asked since VQA models are not trained to answer irrelevant questions about images³⁴. Figure 4 visualizes the different interfaces used for the CG-NoExp group and the counterfactual groups.

5 | RESULTS AND DISCUSSION

In this section, we cover the results from the user studies conducted for the two major tasks described previously: answer change prediction and answer correctness prediction.

5.1 | Answer Change Prediction

Table 2 provides the detailed numbers on the user accuracy in all groups. For each group, the users collectively predict a certain number of episodes which is outlined as N in the table. We show the accuracy of users correctly predicting the system would be INCORRECT for the cases when the VQA model is INCORRECT and similarly for when the VQA model is CORRECT. In the last column, we finally present the normalized accuracy which is the average of the accuracy for the CORRECT and INCORRECT cases. Since the number of correct cases is more than the number of incorrect cases for a VQA model, a normalized accuracy score mitigates potential spurious increases of accuracy simply because a user always predicted a model would be correct. If a user always predicted a model would be correct, the recall for CORRECT cases would be 100% and 0% for the INCORRECT cases, resulting in a normalized accuracy of only 50%.

Counterfactual examples help over not providing examples for predicting answer change Users that are exposed to the counterfactual samples can predict the answer change better than the users in the baseline group. Moreover, we observe a consistent improvement in all subgroups regardless of the in-painting patterns. Even in CF-MinAtt and CF-MidAtt users tend to do better when exposed to the counterfactuals although predicting an answer change for those cases can be inherently harder. These findings suggest a positive impact by the counterfactual samples on the mental model independent of the testing scenario.

5.2 | Answer Correctness Prediction

Here, we evaluate the user's accuracy in predicting whether a VQA model will be correct or not on an unseen test image. We conduct most of our experiments on this task since this task is more realistic and challenging. Our ultimate goal is to see if users can learn from counterfactual explanations to predict the model's performance on unseen images.

We divide the results in correctness prediction results into two subgroups based on cases where the VQA model is CORRECT and INCORRECT as shown in Table 3. Users tend to have an initial optimistic bias towards AI accuracy and as a result, they are more inclined to predict that the AI machine would be correct. As described previously, to prevent a spurious accuracy increase simply due to a user predicting a model will be correct more often, we compute the normalized user accuracy as an average between AI correct and AI incorrect cases.

We now summarize our findings:

Counterfactual examples help over showing no examples All counterfactual groups - CF-AltImg, CF-HAT, CF-AIAtt - show improvement over the control group where no explanations are shown for users' mental model as shown in Table 3 row a vs row c,d,e. This is hardly a surprising result since counterfactual examples provide more information.

Counterfactual examples help over showing random examples To check how much we gain in the mental model from simply providing more information, we check the performance of users when we show two random examples to the users. We see that the counterfactual groups CF-AltImg and CF-HAT both improve the mental model over simply showing random examples as shown in Table 3 row b vs row c and e. This shows that counterfactuals are indeed an effective form of showing examples of how a model behaves to users.

Generated counterfactual images can be a close substitute to realistic counterfactual images We see that a generated counterfactual image using an in-painting network⁸ based on human-annotated attention (row c of Table 3) can be almost as effective as a real retrieved counterfactual image from the VQA dataset (row e). While human-attention annotation is still currently needed, it is a step towards automating the counterfactual generation process.

Fully automating the generation process for counterfactual images can be tricky and currently doesn't seem to help mental model improvement As seen from row d of Table 3, if we use the model's attention values to decide objects to remove, the counterfactual images generated do not improve the user's mental model significantly over no example cases or when random cases are shown. This suggests that further research is needed to effectively automate the counterfactual generation process.

6 | CONCLUSION

In this work, we demonstrated that showing counterfactual images is helpful for the mental model improvement of users in predicting a VQA model's performance. We showed that counterfactual examples are more effective than showing random examples or not showing any examples at all. We also showed that a generative approach to generate counterfactual images can also be effective at improving the mental model of users. We hope these results can serve as a foundation to improve generative models for producing effective counterfactual explanations to improve user mental models for the safe and effective deployment of AI systems in the wild.

References

1. Alipour K, Schulze JP, Yao Y, Ziskind A, Burachas G. A study on multimodal and interactive explanations for visual question answering. *arXiv preprint arXiv:2003.00431* 2020.
2. Chandrasekaran A, Prabhu V, Yadav D, Chattopadhyay P, Parikh D. Do explanations make VQA models more predictable to a human?. *arXiv preprint arXiv:1810.12366* 2018.
3. Rutjes H, Willemsen M, IJsselstein W. Considerations on explainable AI and users' mental models. In: Association for Computing Machinery, Inc. ; 2019.
4. Alipour K, Ray A, Lin X, Schulze JP, Yao Y, Burachas GT. The Impact of Explanations on AI Competency Prediction in VQA. In: IEEE. ; 2020: 25–32.

5. Ray A, Burachas G, Yao Y, Divakaran A. Lucid Explanations Help: Using a Human-AI Image-Guessing Game to Evaluate Machine Explanation Helpfulness. *arXiv preprint arXiv:1904.03285* 2019.
6. Ray A, Cogswell M, Lin X, et al. Knowing What VQA Does Not: Pointing to Error-Inducing Regions to Improve Explanation Helpfulness. *arXiv preprint arXiv:2103.14712* 2021.
7. Molnar C. *Interpretable machine learning*. Lulu.com . 2020.
8. Chang CH, Creager E, Goldenberg A, Duvenaud D. Explaining image classifiers by counterfactual generation. *arXiv preprint arXiv:1807.08024* 2018.
9. Antol S, Agrawal A, Lu J, et al. Vqa: Visual question answering. In: ; 2015: 2425–2433.
10. Teney D, Anderson P, He X, Hengel v. dA. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In: ; 2018: 4223–4232.
11. Xu H, Saenko K. Ask, Attend and Answer: Exploring Question-Guided Spatial Attention for Visual Question Answering. In: Leibe B, Matas J, Sebe N, Welling M., eds. *Computer Vision – ECCV 2016* Springer International Publishing; 2016; Cham: 451–466.
12. Fukui A, Park DH, Yang D, Rohrbach A, Darrell T, Rohrbach M. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847* 2016.
13. Lu J, Yang J, Batra D, Parikh D. Hierarchical Question-Image Co-Attention for Visual Question Answering. *CoRR* 2016; abs/1606.0.
14. Anderson P, He X, Buehler C, et al. Bottom-up and top-down attention for image captioning and visual question answering. In: ; 2018: 6077–6086.
15. Peng L, Yang Y, Wang Z, Huang Z, Shen HT. MRA-Net: Improving VQA via Multi-modal Relation Attention Network. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2020.
16. Patro B, Namboodiri V, others . Explanation vs attention: A two-player game to obtain attention for VQA. In: . 34. ; 2020: 11848–11855.
17. Devlin J, Chang MW, Lee K, Toutanova K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* 2018.
18. Goyal Y, Wu Z, Ernst J, Batra D, Parikh D, Lee S. Counterfactual Visual Explanations. In: Chaudhuri K, Salakhutdinov R., eds. *Proceedings of the 36th International Conference on Machine Learning*. 97 of *Proceedings of Machine Learning Research*. PMLR; 2019: 2376–2384.
19. Wachter S, Mittelstadt B, Russell C. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech*. 2017; 31: 841.
20. Ray A, Sikka K, Divakaran A, Lee S, Burachas G. Sunny and dark outside?! improving answer consistency in vqa through entailed question generation. *arXiv preprint arXiv:1909.04696* 2019.
21. Selvaraju RR, Tendulkar P, Parikh D, et al. SQuINTing at VQA Models: Introspecting VQA Models With Sub-Questions. In: ; 2020: 10003–10011.
22. Agarwal V, Shetty R, Fritz M. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In: ; 2020: 9690–9698.
23. Cosley D, Lam SK, Albert I, Konstan JA, Riedl J. Is seeing believing?: how recommender system interfaces affect users' opinions. In: ACM. ; 2003: 585–592.
24. Ribeiro MT, Singh S, Guestrin C. " Why should i trust you?" Explaining the predictions of any classifier. In: ; 2016: 1135–1144.

25. Kulesza T, Stumpf S, Burnett M, Kwan I. Tell me more? The effects of mental model soundness on personalizing an intelligent agent. In: ; 2012: 1–10.
26. Narayanan M, Chen E, He J, Kim B, Gershman S, Doshi-Velez F. How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682* 2018.
27. Dodge J, Liao QV, Zhang Y, Bellamy RKE, Dugan C. Explaining Models: An Empirical Study of How Explanations Impact Fairness Judgment. In: IUI '19. Association for Computing Machinery; 2019; New York, NY, USA: 275–285
28. Lai V, Tan C. On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In: ; 2019: 29–38.
29. Lage I, Chen E, He J, et al. An evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1902.00006* 2019.
30. Goyal Y, Khot T, Summers-Stay D, Batra D, Parikh D. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In: ; 2017: 6904–6913.
31. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: ; 2016: 770–778.
32. He K, Gkioxari G, Dollár P, Girshick R. Mask r-cnn. In: ; 2017: 2961–2969.
33. Das A, Agrawal H, Zitnick L, Parikh D, Batra D. Human attention in visual question answering: Do humans and deep networks look at the same regions?. *Computer Vision and Image Understanding* 2017; 163: 90–100.
34. Ray A, Christie G, Bansal M, Batra D, Parikh D. Question relevance in VQA: identifying non-visual and false-premise questions. *arXiv preprint arXiv:1606.06622* 2016.
35. Kment B. Counterfactuals and explanation. *Mind* 2006; 115(458): 261–310.
36. Ruben DHEE. London–New York.; 2004.
37. Woodward J. *Making things happen: A theory of causal explanation*. Oxford university press . 2005.
38. Liang Z, Jiang W, Hu H, Zhu J. Learning to Contrast the Counterfactual Samples for Robust Visual Question Answering. In: ; 2020: 3285–3292.

TABLES

TABLE 1 User study groups for answer correctness prediction.

Group	Samples		Attention source
	Images	Questions	
CG-NoExp	—	—	—
CG-RandIQ	Random images	Random questions	—
CF-HAT	Original image in-painted over least/most attended areas	Original question	Human annotations
CF-AIAtt	Original image in-painted over least/most attended areas	Original question	AI attention
CF-AltImg	Alternative real images	Original question	—

TABLE 2 Normalized user accuracies in answer change prediction task.

Group	Baseline	Counterfactual
CF-MinAtt	54.23%	62.52%*
CF-MidAtt	56.29%	66.38%**
CF-MaxAtt	62.91%	66.01%**
All	61.30%	66.98%***

TABLE 3 User accuracy in answer correctness prediction task.

Group	AI correct		AI incorrect		Norm. Acc.
	N	User Acc.	N	User Acc.	
a. CG-NoExp	2995	70.42%	2605	45.60%	58.01%
b. CG-RandIQ	2921	78.26%	2659	46.30%	62.28%
c. CF-HAT	3005	73.14%	2625	54.48%	63.81%
d. CF-AIAtt	2942	77.16%	2558	42.81%	59.98%
e. CF-AltImg	1643	85.88%	917	43.84%	64.86%

FIGURE LEGENDS

FIGURE 1 While alternative real images may present a convincing counterfactual case for a VQA model, they are expensive to harvest and also often incapable of selecting specific features. In this sample, while the real image counterfactual may suggest that the AI agent is correctly capturing the type of sport, the in-painted counterfactual suggests that the change in the answer is not necessarily correlated to the changes in the input.

FIGURE 2 Generating counterfactual images based on human annotation attentions. The algorithm first identifies the most attended and least attended bounding boxes in the image and then applies the GAN to in-paint those bounding boxes and produce the counterfactual images.

FIGURE 3 The interfaces for the experiments that evaluate the impact of in-painted counterfactuals for the task of answer change prediction. Users in both groups are evaluated based on the same in-painting patterns. While the users in the counterfactual groups can utilize the counterfactual samples in their prediction, the baseline group attempts to predict the answer change merely based on the original IQ response. For the input and sample images, users see AI's top answer along with its probability (blue bar beneath the answers).

FIGURE 4 The workflow for different groups of the study. While steps 1 & 3 are shared among groups, the explanation step differentiates between them. In in-painted counterfactuals, sample 1 and sample 2 are in-painted over the least attended and most attended areas respectively. The real counterfactuals images are sampled from the VQA dataset.

