

Are VQA Systems RAD? Measuring Robustness to Augmented Data with Focused Interventions

Daniel Rosenberg and Itai Gat and Amir Feder and Roi Reichart

Technion - Israel Institute of Technology

daniel.rnberg@gmail.com | {itaigat@ | feder@campus. | roiri@ }technion.ac.il

Abstract

Deep learning algorithms have shown promising results in visual question answering (VQA) tasks, but a more careful look reveals that they often do not understand the rich signal they are being fed with. To understand and better measure the generalization capabilities of VQA systems, we look at their robustness to counterfactually augmented data. Our proposed augmentations are designed to make a focused intervention on a specific property of the question such that the answer changes. Using these augmentations, we propose a new robustness measure, Robustness to Augmented Data (RAD), which measures the consistency of model predictions between original and augmented examples. Through extensive experimentation, we show that RAD, unlike classical accuracy measures, can quantify when state-of-the-art systems are not robust to counterfactuals. We find substantial failure cases which reveal that current VQA systems are still brittle. Finally, we connect between robustness and generalization, demonstrating the predictive power of RAD for performance on unseen augmentations.¹

1 Introduction

In the task of Visual Question Answering (VQA), given an image and a natural language question about the image, a system is required to answer the question accurately (Antol et al., 2015). While the accuracy of these systems appears to be constantly improving (Fukui et al., 2016; Yang et al., 2016; Lu et al., 2016), they are sensitive to small perturbations in their input and seem overfitted to their training data (Kafle et al., 2019).

To address the problem of overfitting, the VQA-CP dataset was proposed (Agrawal et al., 2018). It is a reshuffling of the original VQA dataset, such

¹Our code and data are available at: <https://danrosenberg.github.io/rad-measure/>




Figure 1: Predictions and attention maps of a state-of-the-art VQA-CP model over a VQA example (left) and its augmentation (right). A robust model should use the information it utilizes in the original example to correctly answer the augmented one.

that the distribution of answers per question type (e.g., “what color”, “how many”) differs between the train and test sets. Using VQA-CP, Kafle et al. (2019) demonstrated the poor out-of-distribution generalization of many VQA systems. While many models were subsequently designed to deal with the VQA-CP dataset (Cadene et al., 2019; Clark et al., 2019; Chen et al., 2020; Gat et al., 2020), aiming to solve the out-of-distribution generalization problem in VQA, they were later demonstrated to overfit the unique properties of this dataset (Teney et al., 2020). Moreover, no measures for robustness to distribution shifts have been proposed.

In this work we propose a consistency-based measure that can indicate on the robustness of VQA models to distribution shifts. Our robustness measure is based on counterfactual data augmentations (CADs), which were shown useful for both training (Kaushik et al., 2019) and evaluation (Garg et al., 2019; Agarwal et al., 2020). CADs are aimed at manipulating a specific property while preserving all other information, allowing us to evaluate the robustness of the model to changes to this property.

For example, consider transforming a “what color” question to a “yes/no” question, as depicted in Figure 1. The counterfactual reasoning for such a transformation is: “what would be the question if it had a yes/no answer?”. While VQA models

have seen many of both question types, their combination (yes/no questions about color) has been scarcely seen. If a model errs on such a combination, this suggests that to answer the original question correctly, the model uses a spurious signal such as the correlation between the appearance of the word “color” in the question and a particular color in the answer (e.g. here, color \Rightarrow white). Further, this example shows that some models cannot even identify that they are being asked a “yes/no” question, distracted by the word “color” in the augmented question and answering “green”.

Our robustness measure is named RAD: Robustness to (counterfactually) Augmented Data (Section 2.1). RAD receives (image, question, answer) triplets, each augmented with a triplet where the question and answer were manipulated. It measures the consistency of model predictions when changing a triplet to its augmentation, i.e., the robustness of the model to (counterfactual) augmentations. We show that using RAD with focused interventions may uncover substantial weaknesses to specific phenomenon (Section 3.2), namely, users are encouraged to precisely define their interventions such that they create *counterfactual* augmentations. As a result, pairing RAD values with accuracy gives a better description of model behavior.

In general, to effectively choose a model in complex tasks, complementary measures are required (D’Amour et al., 2020). Thus, it is important to have interpretable measures that are widely applicable. Note that in this work we manipulate only textual inputs - questions and answers, but RAD can be applied to any dataset for which augmentations are available. In particular, exploring visual augmentations would be beneficial for the analysis of VQA systems. Further, representation-level counterfactual augmentations are also valid, which is useful when generating meaningful counterfactual text is difficult (Feder et al., 2020).

Our augmentations (CADs) are generated semi-automatically (Section 2.2), allowing us to directly intervene on a property of choice through simple templates. As in the above example, our augmentations are based on compositions of two frequent properties in the data (e.g., “what color” and “yes/no” questions), while their combination is scarce. Intuitively, we would expect a model with good generalization capacities to properly handle such augmentations. While this approach can promise coverage of only a subset of the examples

in the VQA and VQA-CP datasets, it allows us to control the sources of the model’s prediction errors.

We conduct extensive experiments and report three key findings. First, for three datasets, VQA, VQA-CP, and VisDial (Das et al., 2017), models with seemingly similar accuracy are very different in terms of robustness, when considering RAD with our CADs (Section 3). Second, we show that RAD with alternative augmentation methods, which prioritize coverage over focused intervention, cannot reveal the robustness differences. Finally, we show that measuring robustness using RAD with our CADs predicts the accuracy of VQA models on unseen augmentations, establishing the connection between robustness to our controlled augmentations and generalization (Section 4).

2 Robustness to Counterfactuals

In this section, we first present RAD (Section 2.1), which measures model consistency on question-answer pairs and their augmented modifications. Then, we describe our template-based CAD generation approach (Section 2.2), designed to provide control over the augmentation process.

2.1 Model Robustness

We denote a VQA dataset with $\mathcal{U} = \{(x_v, x_q, y) \in \mathcal{V} \times \mathcal{Q} \times \mathcal{Y}\}$, where x_v is an image, x_q is a question and y is an answer. We consider a subset $\mathcal{D} \subseteq \mathcal{U}$ for which we can generate augmentations. For an example $(x_v, x_q, y) \in \mathcal{D}$, we denote an augmented example as $(x_v, x'_q, y') \in \mathcal{D}'$. In this paper we generate a single augmentation for each example in \mathcal{D} , resulting in a one-to-one correspondence between \mathcal{D} and the dataset of modified examples \mathcal{D}' . We further define $J(\mathcal{D}; f)$ as the set of example indices for which a model f correctly predicts y given x_v and x_q .

RAD assesses the proportion of correctly answered modified questions, among correctly answered original questions, and is defined as,

$$\text{RAD}(\mathcal{D}, \mathcal{D}'; f) = \frac{|J(\mathcal{D}; f) \cap J(\mathcal{D}'; f)|}{|J(\mathcal{D}; f)|}. \quad (1)$$

Note that RAD is in $[0, 1]$ and the higher the RAD of f is, the more robust f is.

As original examples and their augmentations may differ in terms of their difficulty to the model, it is important to maintain symmetry between \mathcal{D} and \mathcal{D}' . We hence also consider the backward view

of RAD, defined as $\text{RAD}(\mathcal{D}', \mathcal{D}; f)$. For example, “yes/no” VQA questions are easier to answer compared to “what color” questions, as the former have two possible answers while the latter have as many as eight. Indeed, state-of-the-art VQA models are much more accurate on yes/no questions compared to other question types (Yu et al., 2019). Hence, if “what color” questions are augmented with “yes/no” counterfactuals, we would not expect $\text{RAD}(\mathcal{D}', \mathcal{D}; f) = 1$ as generalizing from “yes/no” questions (\mathcal{D}') to “what color” questions (\mathcal{D}) requires additional reasoning capabilities.

RAD is not dependant on the accuracy of the model on the test set. A model may perform poorly overall but be very consistent on questions that it has answered correctly. Conversely, a model that demonstrates seemingly high performance may be achieving this by exploiting many dataset biases and be very inconsistent on similar questions.

2.2 Counterfactual Augmentations

In the VQA dataset there are three answer types: “yes/no”, “number” (e.g., ‘2’, ‘0’) and “other” (e.g., ‘red’, ‘tennis’), and 65 question types (e.g., “what color”, “how many”, “what sport”). In our augmentations, we generate “yes/no” questions from “number” and “other” questions.

For example, consider the question-answer pair “What color is the vehicle? Red”, this question-answer pair can be easily transformed into “Is the color of the vehicle red? Yes”. In general, “what color” questions can be represented by the template: “What color is the $\langle \text{Subj} \rangle$? $\langle \text{Color} \rangle$ ”. To generate a new question, we first identify the subject ($\langle \text{Subj} \rangle$) for every “what color” question, and then integrate it into the template “Is the color of the $\langle \text{Subj} \rangle$ $\langle \text{Color} \rangle$? Yes”. As the model was exposed to both “what color” and “yes/no” questions, we expect it to correctly answer the augmented question given that it correctly answers the original. Yet, this augmentation requires some generalization capacity because the VQA dataset contains very few yes/no questions about color.

Our templates are presented in Table 1 (see Table 6 in the appendix for some realizations). The augmentations are counterfactual since we intervene on the question type, a priori that many VQA systems exploit (Goyal et al., 2017), keeping everything else equal. The generation process is semi-automatic, as we had to first manually specify templates that would yield augmented questions that we can expect the model to answer correctly given

	Original	Augmented
Y/N \leftarrow C	What color is the $\langle S \rangle$? $\langle C1 \rangle$	Is the color of the $\langle S \rangle$ $\langle C2 \rangle$? Yes/No
Y/N \leftarrow HM	How many $\langle S \rangle$? $\langle N1 \rangle$	Are there $\langle N2 \rangle$ $\langle S \rangle$? Yes/No
Y/N \leftarrow WK	What kind of $\langle S \rangle$ is this? $\langle O1 \rangle$	Is this $\langle S \rangle$ $\langle O2 \rangle$? Yes/No

Table 1: Our proposed template-based augmentations.

that it succeeds on the original question.

To achieve this goal, we apply two criteria: **(a)** The template should generate a grammatical English question; and **(b)** The generated question type should be included in the dataset, but not in questions that address the same semantic property as the original question. Indeed, yes/no questions are frequent in the VQA datasets, but few of them address color (first template), number of objects (second template), and object types (third template). When both criteria are fulfilled, it is reasonable to expect the model to generalize from its training set to the new question type.

Criterion (a) led us to focus on yes/no questions since other transformations required manual verification for output grammaticality. While we could have employed augmentation templates from additional question types into yes/no questions, we believe that our three templates are sufficient for evaluating model robustness. Overall, our templates cover 11% of the VQA examples (Section 3.1).

3 Robustness with RAD and CADs

In the following, we perform experiments to test the robustness of VQA models to augmentations. We describe the experimental setup, and evaluate VQAv2, VQA-CPv2, VisDial models, each on our augmentations and on other alternatives.²

3.1 Experimental Setup

Baseline Augmentations We compare our augmentations to three alternatives: VQA-Rephrasings (Reph, Shah et al., 2019), ConVQA (Ray et al., 2019), and back-translation (BT, Sennrich et al., 2016). VQA-Rephrasings is a manual generation method, where annotators augment each validation question with three re-phrasings. ConVQA is divided into the L-ConVQA and CS-ConVQA subsets. In both subsets, original validation examples are augmented to create new question-answer pairs. L-ConVQA is automatically generated based

²The URLs of the software and datasets, and the implementation details are all provided in Appendices C and D.

Dataset	Model $\backslash\mathcal{D}'$	RAD($\mathcal{D}, \mathcal{D}'$) (%)							Acc.
		Y/N \leftarrow C	Y/N \leftarrow HM	Y/N \leftarrow WK	BT	Reph	L-ConVQA	CS-ConVQA	
VQA-CP	RUBi	64.92	57.15	62.59	85.57	77.73	78.02	65.93	46.66
	LMH	1.01	22.82	50.10	83.68	75.04	64.54	50.65	53.72
	CSS	0.94	11.73	39.95	77.54	68.89	10.67	38.64	58.47
VQA	BUTD	67.15	58.68	78.59	87.43	79.28	75.78	70.19	63.09
	BAN	74.40	62.45	82.51	88.17	81.14	79.37	70.18	65.92
	Pythia	65.00	60.61	81.60	88.42	82.86	77.02	69.45	64.56
	VisualBERT	79.99	68.29	85.98	88.52	84.09	82.09	71.75	65.62
VisDial	FGA	31.36	57.69	-	91.42	-	-	-	53.07
	VisDialBERT	62.08	56.06	-	94.04	-	-	-	55.78

Table 2: RAD over our proposed augmentations (Y/N \leftarrow C, Y/N \leftarrow HM, Y/N \leftarrow WK) and alternatives (BT, Reph, ConVQA). The rows correspond to state-of-the-art models on VQA-CP (top), VQA (middle) and Visual Dialog (bottom). Reph and ConVQA were not created for VisDial, and it does not have “what kind” questions. The last column corresponds to validation accuracy.

Dataset	Model $\backslash\mathcal{D}'$	Accuracy(\mathcal{D}) (%)						
		Y/N \leftarrow C	Y/N \leftarrow HM	Y/N \leftarrow WK	BT	Reph	L-ConVQA	CS-ConVQA
VQA-CP	RUBi	65.85	17.35	44.14	45.80	46.51	72.14	66.67
	LMH	68.87	44.24	50.58	52.35	53.78	65.07	61.76
	CSS	72.87	63.16	51.83	56.37	58.81	49.84	56.12
VQA	BUTD	79.44	54.43	63.49	60.37	62.23	75.05	62.42
	BAN	80.72	62.37	66.48	63.02	64.81	74.94	65.01
	Pythia	81.62	57.49	64.42	61.69	63.88	74.55	63.79
	VisualBERT	80.85	58.89	64.46	62.71	64.96	76.50	66.01
VisDial	FGA	55.62	40.00	-	61.53	-	-	-
	VisDialBERT	68.99	50.77	-	63.47	-	-	-

Table 3: Original accuracy over our proposed augmentations (Y/N \leftarrow C, Y/N \leftarrow HM, Y/N \leftarrow WK) and alternatives (BT, Reph, ConVQA). The rows correspond to state-of-the-art models on VQA-CP (top), VQA (middle) and Visual Dialog (bottom). Reph and ConVQA were not created for VisDial, and it does not have “what kind” questions.

on scene graphs attached to each image, and CS-ConVQA is manually generated by annotators. Finally, back-translation, translating to another language and back, is a high-coverage although low-quality approach to text augmentation. It was used during training and shown to improve NLP models (Sennrich et al., 2016), but was not considered in VQA. We use English-German translations.

Models The VQA-CP models we consider are RUBi (Cadene et al., 2019), LMH (Clark et al., 2019) and CSS (Chen et al., 2020). The VQA models we consider are BUTD (Anderson et al., 2018), BAN (Kim et al., 2018), Pythia (Jiang et al., 2018) and VisualBERT (Li et al., 2019). For VisDial we use FGA (Schwartz et al., 2019) and VisDialBERT (Murahari et al., 2020). We trained all the models using their official implementations.

3.2 Results

Table 2 presents our main results. RAD values for all of our augmentations are substantially lower than those of the alternatives, supporting the value

of our focused intervention approach for measuring robustness. The high RAD values for BT and Reph might indicate that VQA models are indeed robust to linguistic variation, as long as the answer does not change. Interestingly, our augmentations also reveal that VQA-CP models are less robust than VQA models. This suggests that despite the attempt to design more robust models, VQA-CP models still overfit their training data.

In VQA-CP, RUBi has the lowest accuracy performance in terms of its validation accuracy, even though it is more robust to augmentations compared with LMH and CSS. For VQA models, in contrast, BUTD has the lowest RAD scores on our augmentations and the lowest accuracy. VisualBERT, the only model that utilizes contextual word embeddings, demonstrates the highest robustness among the VQA models.

Finally, while both VisDial models have similar accuracy, they have significantly different RAD scores on our augmentations. Specifically, VisDialBERT performs better than FGA on Y/N \leftarrow C

augmentations. This is another indication of the value of our approach as it can help distinguish between two seemingly very similar models.

Complementary to the RAD values in Table 2 we also provide accuracies on original questions in Table 3. Note that across all the original questions, except ConVQA questions, RUBi has the lowest accuracy while CSS has the highest accuracy. This trend is reversed when looking at RAD scores - CSS has the lowest score while RUBi has the highest score. This emphasizes the importance of RAD as a complementary metric, since considering only accuracy in this case would be misleading. Namely, RAD provides additional critical information for model selection.

4 Measuring Generalization with RAD

To establish the connection between RAD and generalization, we design experiments to demonstrate RAD’s added value in predicting model accuracy on unseen modified examples. Concretely, we generate 45 BUTD (VQA) and LMH (VQA-CP) instances, differing by the distribution of question types observed during training (for each model instance we drop between 10% and 99% of each of the question types “what color”, “how many” and “what kind” from its training data; see Appendix E for exact implementation details). For each of the above models we calculate RAD values and accuracies in the following manner.

We split the validation set into two parts: \mathcal{D} (features) and \mathcal{T} (target). Consider a pool of four original question sets that are taken from their corresponding modifications: $Y/N \leftarrow C$, $Y/N \leftarrow HM$, $Y/N \leftarrow WK$, $Reph$. Then we have four possible configurations in which \mathcal{D} is three sets from the pool and \mathcal{T} is the remaining set. For each model and for each configuration, we compute model accuracy on \mathcal{D} ($Accuracy(\mathcal{D})$) and on the modifications of questions in \mathcal{T} (the predicted variable $y(\mathcal{T}) = Accuracy(\mathcal{T}')$) which are modified with the target augmentation of the experiment. We also compute the RAD values of the model on the modified questions in \mathcal{D} which are generated using the other three augmentations ($RAD(\mathcal{D}, \mathcal{D}')$, and $RAD(\mathcal{D}', \mathcal{D})$). Then, we train a linear regression model using $Accuracy(\mathcal{D})$, $RAD(\mathcal{D}, \mathcal{D}')$, and $RAD(\mathcal{D}', \mathcal{D})$, trying to predict $y(\mathcal{T})$. We perform this experiment four times, each using a different configuration (different augmentation type as our target), and average across the configurations.

Features\Model	R^2
	LMH
Accuracy(\mathcal{D}), RAD($\mathcal{D}, \mathcal{D}'$), RAD($\mathcal{D}', \mathcal{D}$)	0.917 ± 0.117
Accuracy(\mathcal{D})	0.829 ± 0.237
RAD($\mathcal{D}, \mathcal{D}'$)	0.899 ± 0.133
RAD($\mathcal{D}', \mathcal{D}$)	0.849 ± 0.213

Table 4: Linear regression experiments, predicting accuracy performance on unseen augmentation types.

Results Table 4 presents the average R^2 values and standard deviations over the four experiments. RAD improves the R^2 when used alongside the validation accuracy. Interestingly, a model’s accuracy on one set of augmentations does not always generalize to other, unseen augmentations. Only when adding RAD to the regression model are we able to identify a robust model. Notably, for LMH the usefulness of RAD is significant, as it improves the R^2 by 11%. It also predicts performance better than validation accuracy when used without it in the regression. The standard deviations further confirm that the above claims hold over all configurations. These observations hold when running the same experiment with respect to the BUTD model, however, the improvements are smaller since the regression task is much easier with respect to this model (R^2 of 0.995 with all features).

5 Conclusion

We proposed RAD, a new measure that penalizes models for inconsistent predictions over data augmentations. We used it to show that state-of-the-art VQA models fail on CADs that we would expect them to properly address. Moreover, we have demonstrated the value of our CADs by showing that alternative augmentation methods cannot identify robustness differences as effectively. Finally, we have shown that RAD is predictive of generalization to unseen augmentation types.

We believe that the RAD measure brings substantial value to model evaluation and consequently to model selection. It encourages the good practice of testing on augmented data, which was shown to uncover considerable model weaknesses in NLP (Ribeiro et al., 2020). Further, given visual augmentations, which we plan to explore in future work, or linguistic augmentations, RAD is applicable to any classification task, providing researchers with meaningful indications of robustness.

References

- Vedika Agarwal, Rakshith Shetty, and Mario Fritz. 2020. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In *CVPR*.
- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. Don't just assume; look and answer: Overcoming priors for visual question answering. In *CVPR*.
- 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 *CVPR*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *ICCV*.
- Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. 2019. Rubi: Reducing unimodal biases for visual question answering. In *NeurIPS*.
- Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. 2020. Counterfactual samples synthesizing for robust visual question answering. In *CVPR*.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. Don't take the easy way out: Ensemble based methods for avoiding known dataset biases. In *EMNLP*.
- Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yian Ma, Cory McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. 2020. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395*.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *CVPR*.
- Amir Feder, Nadav Oved, Uri Shalit, and Roi Reichart. 2020. Causalm: Causal model explanation through counterfactual language models. *arXiv preprint arXiv:2005.13407*.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*.
- Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *AAAI*.
- Itai Gat, Idan Schwartz, Alexander Schwing, and Tamir Hazan. 2020. Removing bias in multi-modal classifiers: Regularization by maximizing functional entropies. In *NeurIPS*.
- 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 *CVPR*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Yu Jiang, Vivek Natarajan, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2018. Pythia v0. 1: the winning entry to the vqa challenge 2018. *arXiv preprint arXiv:1807.09956*.
- Kushal Kafle, Robik Shrestha, and Christopher Kanan. 2019. Challenges and prospects in vision and language research. *Frontiers in Artificial Intelligence*.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2019. Learning the difference that makes a difference with counterfactually-augmented data. In *ICLR*.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear attention networks. In *NeurIPS*.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. In *NeurIPS*.
- Vishvak Murahari, Dhruv Batra, Devi Parikh, and Abhishek Das. 2020. Large-scale pretraining for visual dialog: A simple state-of-the-art baseline. In *ECCV*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *NeurIPS*.

- Arijit Ray, Karan Sikka, Ajay Divakaran, Stefan Lee, and Giedrius Burachas. 2019. Sunny and dark outside?! improving answer consistency in vqa through entailed question generation. In *EMNLP*.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of nlp models with checklist. In *ACL*.
- Idan Schwartz, Seunghak Yu, Tamir Hazan, and Alexander G Schwing. 2019. Factor graph attention. In *CVPR*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *ACL*.
- Meet Shah, Xinlei Chen, Marcus Rohrbach, and Devi Parikh. 2019. Cycle-consistency for robust visual question answering. In *CVPR*.
- Amanpreet Singh, Vedanuj Goswami, Vivek Nataraajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2020. Mmf: A multimodal framework for vision and language research. <https://github.com/facebookresearch/mmf>.
- Damien Teney, Kushal Kafle, Robik Shrestha, Ehsan Abbasnejad, Christopher Kanan, and Anton van den Hengel. 2020. On the value of out-of-distribution testing: An example of goodhart’s law. *NeurIPS*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *EMNLP*.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *CVPR*.
- Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep modular co-attention networks for visual question answering. In *CVPR*.

A Dataset Statistics

Please see Table 5 for the number of examples in each dataset that we use (VQA, VQA-CP and VisDial). We also report the number of augmentations we produce for each of our three augmentation types ($Y/N \leftarrow C$, $Y/N \leftarrow HM$ and $Y/N \leftarrow WK$), alongside previous augmentation approaches used in our experiments (BT, Reprh, L-ConVQA and CS-ConVQA).

B Our Augmentations

We describe the manual steps required to meet the desired standard for each augmentation type. For $Y/N \leftarrow C$, we filter out questions that start with “What color is the”. For $Y/N \leftarrow HM$, we use questions that starts with “How many”. For $Y/N \leftarrow WK$, we consider questions that match the pattern “What kind of $\langle S \rangle$ is this? $\langle OI \rangle$ ”. Table 6 presents several realizations of the templates we define (see Section 2.2 for a discussion of these templates).

In $Y/N \leftarrow HM$, we ensure that when the answer is ‘1’, we use “Is there ...” instead of “Are there ...”. We also ensure that the subsequent word to “How many” is a noun. We verify it is a noun using the part-of-speech tagger available through the spaCy library (Honnibal et al., 2020).

We allow the generation of both ‘yes’ and ‘no’ answers. Creating a modified question that is answered with a ‘yes’ requires a simple permutation of words in the original question-answer pair, e.g., for $Y/N \leftarrow C$, take “ $\langle C1 \rangle$ ” = “ $\langle C2 \rangle$ ” (see Table 1). Similarly, to generate a question that should be answered with a ‘no’, we repeat the above process and change “ $\langle C2 \rangle$ ”. In this case, we randomly pick an answer and replace it with the original answer with probability weighted with respect to the frequency in the data, among the pool of possible answers for the given augmentation type. When generating a new question, we first randomly decide whether to generate a ‘yes’ or ‘no’ question (with a probability of 0.5 for each). Then, for example, if we choose to generate a ‘no’, and “ $\langle C1 \rangle$ ” = “red”, we have a 63% chance of having “ $\langle C2 \rangle$ ” = “blue”.

C URLs of Data and Code

Data We consider three VQA datasets:

- The VQAv2 dataset (Goyal et al., 2017): <https://visualqa.org/>.
- The VQA-CPv2 dataset (Agrawal et al., 2018): <https://www.cc.gatech.edu/gr>

<ads/a/aagraval307/vqa-cp/>.

- The VisDial dataset (Das et al., 2017): <https://visualdialog.org/>

We also consider three previous augmentation methods:

- VQA-Rephrasings (Shah et al., 2019): <https://facebookresearch.github.io/VQA-Rephrasings/>.
- ConVQA (Ray et al., 2019): <https://arijitr1993.github.io/ConVQA/>.
- Back-translations (Sennrich et al., 2016). We have generated these utilizing the transformers library (Wolf et al., 2020), <https://github.com/huggingface/transformers>. Specifically, we used two pre-trained translation models, English to German, and German to English: <https://huggingface.co/Helsinki-NLP/opus-mt-en-de>, <https://huggingface.co/Helsinki-NLP/opus-mt-de-en>.

Models We consider nine models, where each model’s code was taken from the official implementation. All implementations are via PyTorch (Paszke et al., 2019).

The three VQA-CPv2 models:

- RUBi (Cadene et al., 2019): <https://github.com/cdancette/rubi.bootstrap.pytorch>.
- LMH (Clark et al., 2019): <https://github.com/chris36/bottom-up-attention-vqa>.
- CSS (Chen et al., 2020): <https://github.com/yanxinzju/CSS-VQA>.

The four VQAv2 models:

- BUTD (Anderson et al., 2018): <https://github.com/hengyuan-hu/bottom-up-attention-vqa>.
- BAN (Kim et al., 2018): <https://github.com/jnhwkim/ban-vqa>.
- Pythia (Jiang et al., 2018): Using the implementation in the MMF library (Singh et al., 2020), <https://github.com/facebookresearch/mmf>.
- VisualBERT (Li et al., 2019): Using the implementation in the MMF library.

And the two VisDial models:

Dataset	Augmentation Count							Validation Count
	Y/N ← C	Y/N ← HM	Y/N ← WK	BT	Reph	L-ConVQA	CS-ConVQA	
VQA-CP	12,910	13,437	1,346	149,329	39,936	127,924	423	219,928
VQA	12,835	10,233	1,654	138,043	121,512	127,924	1,365	214,354
VisDial	516	130	-	1,136	-	-	-	20,640

Table 5: Number of examples in each of the datasets we use.

Yes/No ← Colors	Yes/No ← How Many	Yes/No ← What Kind
What color is the cat? White Is the color of the cat white? Yes	How many athletes are on the field? 5 Are there five athletes on the field? Yes	What kind of food is this? Breakfast Is this food breakfast? Yes
What color is the court? Green Is the color of the court green? Yes	How many dogs are in the picture? 3 Are there two dogs in the picture? No	What kind of event is this? Skiing Is this a skiing event? Yes
What color is the vase? Blue Is the color of the vase red? No	How many giraffes are walking around? 2 Are there four giraffes walking around? No	What kind of animal is this? Cow Is this animal an elephant? No
What color is the man’s hat? Red Is the color of the man’s hat red? Yes	How many cakes are on the table? 0 Is there one cake on the table? No	What kind of building is this? Church Is this building a church? Yes
What color is the sky? Blue Is the color of the sky blue? Yes	How many dogs? 1 Are there zero dogs? No	What kind of floor is this? Wood Is this a wood floor? Yes

Table 6: Some realizations of our templates (defined in Table 1). The black text (top) is the original question-answer pair and the blue text (bottom) is the corresponding augmented question-answer pair.

- FGA (Schwartz et al., 2019): <https://github.com/idansc/fga>.
- VisDialBERT (Murahari et al., 2020): <https://github.com/vmurahari3/visdial-bert>.

D Model Settings

We have trained the VQAv2 and the VQA-CPv2 models that we use, as pre-trained weights were not available for our requirements. For our evaluations, we require a model that is trained solely on the VQAv2 train set, such that we match the VQA-CPv2 settings, where there are only two sets, train and validation. In contrast, pre-trained models that are built for VQAv2 are trained on the VQAv2 training set and on the VQAv2 validation set together, as the dataset contains a third development set that is commonly used for validation.

We have trained six VQA models using the default hyper-parameters from their official implementations (URLs in Appendix C): RUBi, LMH, CSS, BUTD, BAN and Pythia. We trained the above models on a single Nvidia GeForce RTX 2080 Ti GPU, when the training time for each of the models was less than 12 hours. In addition, inference in this setting took less than an hour for all models.

The VisualBERT model is more computationally intensive, and we had to reduce the default batch size from 480 to 54 to fit it on our resources. Using three Nvidia GeForce RTX 2080 Ti GPUs for

VisualBERT, training took 36 hours and inference took 4 hours.

For the VisDial models, FGA, and VisDialBERT, we have downloaded the pre-trained weights and used them solely for inference. On a single Nvidia GeForce RTX 2080 Ti GPU, inference took 15 minutes for FGA, and 8 hours for VisDialBERT.

All the models we consider have less than 200M parameters.

When accuracies are reported on VQAv2 and on VQA-CP (Tables 2 and 3) we use the VQA-accuracy metric (Antol et al., 2015). For VisDial we use the standard accuracy metric (denoted originally as R@1).

E Regression Experiments

We generate 45 BUTD (VQA) instances and 45 LMH (VQA-CP) instances. To generate different model instances, we create 45 new training sets by removing examples from the original train set. For each of the three question types, “what color”, “how many” and “what kind”, we remove the following 15 percentage values of examples from the original train set: [10%, 20%, 30%, 40%, 50%, 60%, 70%, 80%, 90%, 92%, 95%, 96%, 97%, 98%, 99%], resulting in 45 new training sets. Then, each model instance is created by training on one of the 45 training sets.

We split the validation set into two parts: \mathcal{D} and \mathcal{T} . \mathcal{D} is used to calculate the features in our linear

regression model. We denote with \mathcal{D}'_1 the questions in \mathcal{D} that can be modified using the Y/N \leftarrow C augmentation, after these questions were modified. Similarly, we define \mathcal{D}'_2 , \mathcal{D}'_3 , and \mathcal{D}'_4 for Y/N \leftarrow HM, Y/N \leftarrow WK, and Reph, respectively.

We average the R^2 of four linear regression experiments, when in each experiment we set a different i ($i \in \{1, 2, 3, 4\}$) for which $\mathcal{T} = \mathcal{D}'_i$ and use the remaining three templates to calculate our features. We denote the regression features with $x_1 = \text{Accuracy}(\mathcal{D})$, $x_2 = \text{RAD}(\mathcal{D}, \mathcal{D}')$, and $x_3 = \text{RAD}(\mathcal{D}', \mathcal{D})$, where $\text{RAD}(\mathcal{D}, \mathcal{D}')$ and $\text{RAD}(\mathcal{D}', \mathcal{D})$ are computed with respect to the three other templates ($j \in \{1, 2, 3, 4\}, j \neq i$). The predicted label is $y(\mathcal{T}) = \text{Accuracy}(\mathcal{T})$.

Thus the equation for our regression is:

$$y(\mathcal{T}) = b_1x_1 + b_2x_2 + b_3x_3 + \epsilon .$$

We also perform three regression experiment for each feature alone:

$$y(\mathcal{T}) = bx_k + \epsilon, \quad k = 1, 2, 3 ,$$

and compare the results of these experiments in [Table 4](#).