
Examining Cooperation in Visual Dialog Models

Mircea Mironenco *
University of Amsterdam
Amsterdam, Netherlands
mircea.mironenco@student.uva.nl

Dana Kianfar *
University of Amsterdam
Amsterdam, Netherlands
dana.kianfar@student.uva.nl

Ke Tran
ILPS
University of Amsterdam
Amsterdam, Netherlands
m.k.tran@uva.nl

Evangelos Kanoulas
ILPS
University of Amsterdam
Amsterdam, Netherlands
e.kanoulas@uva.nl

Efstratios Gavves
QUVA Lab
University of Amsterdam
Amsterdam, Netherlands
e.gavves@uva.nl

Abstract

In this work we propose a blackbox intervention method for visual dialog models, with the aim of assessing the contribution of individual linguistic or visual components. Concretely, we conduct structured or randomized interventions that aim to impair an individual component of the model, and observe changes in task performance. We reproduce a state-of-the-art visual dialog model and demonstrate that our methodology yields surprising insights, namely that both dialog and image information have minimal contributions to task performance. The intervention method presented here can be applied as a sanity check for the strength and robustness of each component in visual dialog systems.

1 Introduction

By combining vision and language, tasks that require high-level image understanding such as image captioning [2, 6, 11] and visual question answering [1, 8, 9], leverage the performance of deep neural networks in an attempt to simulate the way humans acquire and use information from different modalities in their environment. In more recent work, the availability of large-scale datasets has seen dialog models being proposed as a medium for communicating visual information [7, 10].

Two recent approaches have proposed models that aim to acquire natural language by multi-agent dialog on a downstream visual task [3, 5]. Both models are asymmetrically *primed* with initial textual and visual information, and leverage the information gap between two agents to simulate a human-like conversation.

Our goal is to dissect the contributions of linguistic and visual components and their interplay. We believe that progressing visually-grounded conversational artificial intelligence requires the understanding of communicative protocols exchanged by the agents and how they utilize language and visual information cooperatively. In this work, we present a simple black-box *intervention* method which aims to determine which source of linguistic and visual information agents exploit the most in order to complete the task. Our method is model-agnostic, and can be utilized as a sanity-check in the process of designing a new model. We demonstrate our method on the visual dialog model presented in [4]. Our method empirically tests the robustness of the model with respect to variations in inputs, and performs unit testing on specific components.

*Equal contribution by DK and MM. Order determined by coin toss.

2 Dialog Agents

We replicate the supervised hierarchical recurrent encoder-decoder model proposed by Das *et al.* [4] which engages in the cooperative image guessing game introduced in [3]. The game includes a question bot (Q-BOT) and an answer bot (A-BOT). Both bots are provided with a short description of an image and A-BOT additionally accesses the image itself. Without seeing the image, Q-BOT must communicate with A-BOT by asking questions in order to guess the image. Since the outcome of the game is decided based on the accuracy of Q-BOT’s prediction of the ground truth image, A-BOT must cooperate with Q-BOT to win the game. Choosing this model was primarily motivated by its structure which incorporates multiple components working towards the same goal, as well as the VisDial dataset which is a sound testbed for this type of investigation due to its diversity and size.

To quantify the performance of the model, Q-BOT is asked to rank the correct image among a set of candidates. Consistent with the original evaluation, we use mean-percentile rank (MPR) as our metric. A mean-percentile rank of 90% means that the prediction of the Q-BOT is closer to the ground truth image than 90% of the images in the set. In the guessing game, forcing two agents to communicate through natural language enables humans to inspect the behaviours of the agents painlessly by creating meaningful and interpretable interventions.

3 Methodology

We consider two types of interventions: (1) intervening on the initial condition (image and caption), and (2) intervening in the course of conversation by changing the responses generated by either agent. Understanding the initial condition is crucial in designing a conversational AI. For example, it sheds light on what data we should collect, or what is the natural interface between humans and AI. To understand the role of images and their captions, we perform interventions as follows.

- **Image:** We replace image feature vector by random noise $z \sim \text{Uniform}(0, 1)$. If the images are useful cues, then this intervention completely destroys a piece of essential information. We therefore expect a degradation in the evaluation performance.
- **Caption:** We replace a content word by a random word. Additionally we observe that many captions are poorly related to their corresponding images.

Intervening during the dialogs on other hand, allows us to glimpse into the model’s internal representations and its ability to exploit and exchange meaningful bits of information. We expect that intelligent systems should be sensitive to perturbations especially when they are not trained to cope against it. In this setting, we intervene on

- **Question:** with probability p , each token in the question is replaced by a random token before giving it to A-BOT.
- **Answer:** with probability p , each token in the answer of A-BOT is replaced by a random token before giving it to Q-BOT.

In addition to random noise, we propose using *negation* as a more principled approach to gaining insights into the cooperative behaviors of the two agents. Particularly we change the answer of A-BOT from *yes* to *no* and vice versa. If the Q-BOT behaves cooperatively its predictions should alternate dramatically. Here we choose to manipulate A-BOT and observe the outcomes of the VisDial game because it is easier to negate the answer than the question. Moreover *yes* and *no* answers make up 37% of the responses of A-BOT in training and 45% in validation data, therefore it is reasonable to expect both bots to learn the negation concept.

4 Experiments

All experiments² were performed on the model described in Das *et al.* [4]. We perform the four interventions described in Section 3 during inference on the validation set. We study the performance of each modified dialog system, on the same task of ranking the image inside a collection, and

²Our code will be available at <https://github.com/danakianfar/Examining-Cooperation-in-VDM>

compare it with the regular performance without interventions. Our aim is to understand whether the dialog or the image is being leveraged to provide further information, which the Q-BOT can use to make better predictions. A priori, we expect to observe a large decline performance for all intervention experiments, as they essentially involve replacing an information source with random noise.

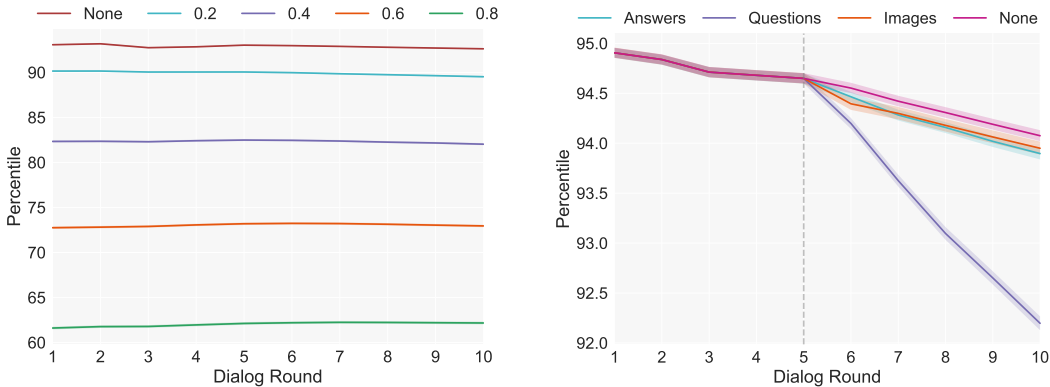
Specifically we intervene on the caption with different probabilities $p \in \{0.2, 0.4, 0.6, 0.8\}$ at the start of the dialog, and on the image, answers and questions at round 5 with $p = 0.8$. In the case of the image intervention, we replace the entire image with random noise.

5 Results

In this section we present the mean percentile rank (MPR) on the 40K validation set of the VisDial v0.9 dataset [3] for each intervention experiment (answers, captions, questions, and images) as well as the regular inference without interventions (called "None") as described in Section 3.

Caption Interventions As displayed in Figure 1a higher values of p correspond to poorer performance. Despite the fact that the caption is only seen once at the start of the dialog, it nevertheless plays a very important role in the predictive performance of the network across all rounds. Figure 2 in the appendix provides an example of a "positive" manual intervention where replacing the original image caption with a more informative one resulted in a much better ranking.

Image, Caption and Answer Interventions As seen in Figure 1b, the interventions have a negative effect on the performance after the intervention on round 5. However, the decrease in performance by each individual experiment is much smaller than the decrease with interventions on captions. It suggests that Q-BOT relies mainly on the caption as it contains most of the information needed to make predictions. We also note that randomly intervening on the questions affects the performance the most.



(a) **Caption:** MPR when a token in the caption is replaced by another with $p \in \{0.2, 0.4, 0.6, 0.8\}$. "None" represents no interventions.

(b) **Image, Question & Answer:** MPR when we intervene starting at round 5 with $p = 0.8$. "None" represents no interventions.

Figure 1: Comparison of rankings with and without interventions. Note that the y-axes in the two figures are not aligned.

We clearly observe a large discrepancy between rankings of caption interventions and the other experiments. Surprisingly, the decline in performance caused by interventions on answers and questions is less than expected, suggesting that the dialog itself is not used effectively for image identification. Interestingly replacing images with complete noise has minimal to no impact on the performance of the model. Q-BOT mainly relies on the caption provided at the beginning of the dialog to make prediction. This suggests that there is little cooperation between two bots. This effect is more clearly observed in the extreme case where we intervene with $p = 1$ on all rounds, shown in Table 1. We note that the ranking performance of the caption interventions is improved as the dialog continues, although it never recovers completely.

Table 1: MPR [%] of each intervention experiment with $p = 1.0$, where we intervened at each round during inference on the validation set. The difference between "None" and all other intervention rankings at the final round is shown in the final row.

Intervention	None	Images	Captions	Answers	Questions
Round 1	93.1	93.1	50.0	93.0	92.5
Round 2	93.2	93.1	50.3	93.1	92.6
Round 3	92.7	92.5	50.7	92.5	91.8
Round 4	92.8	92.5	51.3	92.5	91.4
Round 5	93.0	92.7	51.6	92.5	91.3
Round 6	93.0	92.6	51.9	92.4	90.9
Round 7	92.9	92.5	52.2	92.2	90.6
Round 8	92.8	92.4	52.3	92.0	90.2
Round 9	92.7	92.3	52.4	91.9	89.9
Round 10	92.6	92.2	52.5	91.7	89.6
Gap @10	0.0	0.4	40.1	0.8	3.0

Table 2: MPR [%] of Negation intervention results.

Round	1	2	3	4	5	6	7	8	9	10
8								94.3	94.1	93.9
6						94.6	94.4	94.2	94.1	93.9
4				94.7	94.6	94.5	94.3	94.2	94.0	93.8
2		94.8	94.7	94.6	94.6	94.4	94.3	94.1	94.0	93.8
0	94.8	94.8	94.7	94.6	94.5	94.4	94.2	94.1	93.9	93.7

Table 2 shows the results of negation intervention experiments. We start intervening from round 0, 2, 4, 6, and 8. Comparing the MPR across all rounds (columns in the table), we see that Q-BOT is sensitive to the non-cooperative behaviors of A-BOT, however only to a small degree ($< 0.3\%$). Our results suggest that in addition to downstream evaluation in a cooperative game setting, forcing one agent to play non-cooperatively could help researchers design a better experimental setup to understand the cooperative behaviours amongst agents in the system.

6 Discussion

We have presented a simple yet effective method for assessing the interaction of linguistic and visual components in visual dialog models. Using an example which combines multiple sources of information in a cooperative multi-agent setup, we have demonstrated that impairing individual components can reveal the extent to which each information source is exploited by the agents to accomplish their goals. A pitfall with designing multi-modal and blackbox systems is that the role of individual components can not be deduced from the overall performance of the model.

In a series of surprising results, we discovered that the role of images is minimal in the image retrieval performance. Furthermore, our evaluation suggests that the dialog itself was not exploited significantly by the bots in the cooperative setting. We argue that designing multi-modal systems requires careful evaluation, or *unit testing*, of each component. Grounding natural language is a difficult problem as models which combine modalities must account for the individual impact of each information source.

We encourage future researchers to account for the effect of each modality on the total performance of the model. An interesting research direction is how to learn interaction models of independent modalities which can be shown to generalize well in conjunction, and avoid the pitfall of overfitting to spurious correlations that optimize the surrogate learning objectives of each independent modality. Reinforcement learning (RL) has traditionally been used to train dialog models and has been motivated as a natural training paradigm for dialogue models, and has been applied to visual dialog models in [4, 5]. We did not include RL methods in our analysis and leave such assessment to future work.

7 Acknowledgements

This research was supported by the Google Faculty Research Award program and the Dutch national program COMMIT.

References

- [1] A. Agrawal, J. Lu, S. Antol, M. Mitchell, C. L. Zitnick, D. Parikh, and D. Batra. Vqa: Visual question answering. *International Journal of Computer Vision*, 123(1):4–31, May 2017.
- [2] X. Chen and C. Lawrence Zitnick. Mind’s eye: A recurrent visual representation for image caption generation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [3] A. Das, S. Kottur, K. Gupta, A. Singh, D. Yadav, J. M. Moura, D. Parikh, and D. Batra. Visual Dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [4] A. Das, S. Kottur, J. M. Moura, S. Lee, and D. Batra. Learning cooperative visual dialog agents with deep reinforcement learning. In *International Conference on Computer Vision (ICCV)*, 2017.
- [5] H. de Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. C. Courville. Guesswhat?! visual object discovery through multi-modal dialogue. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [6] J. Johnson, A. Karpathy, and L. Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [7] O. Lemon and O. Pietquin. *Data-Driven Methods for Adaptive Spoken Dialogue Systems: Computational Learning for Conversational Interfaces*. Springer Publishing Company, Incorporated, 2012.
- [8] M. Malinowski, M. Rohrbach, and M. Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [9] M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2953–2961. Curran Associates, Inc., 2015.
- [10] I. V. Serban, R. Lowe, P. Henderson, L. Charlin, and J. Pineau. A survey of available corpora for building data-driven dialogue systems. *CoRR*, abs/1512.05742, 2015.
- [11] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

8 Appendix

8.1 Manual interventions

We present examples of manual interventions below. All intervened values are displayed in blue. The original inference is displayed on the left column of each example, and the intervened dialog is presented on the right column. The normal inference rankings are displayed in dark blue and the intervened ranking in light blue.

Manual intervention on Caption

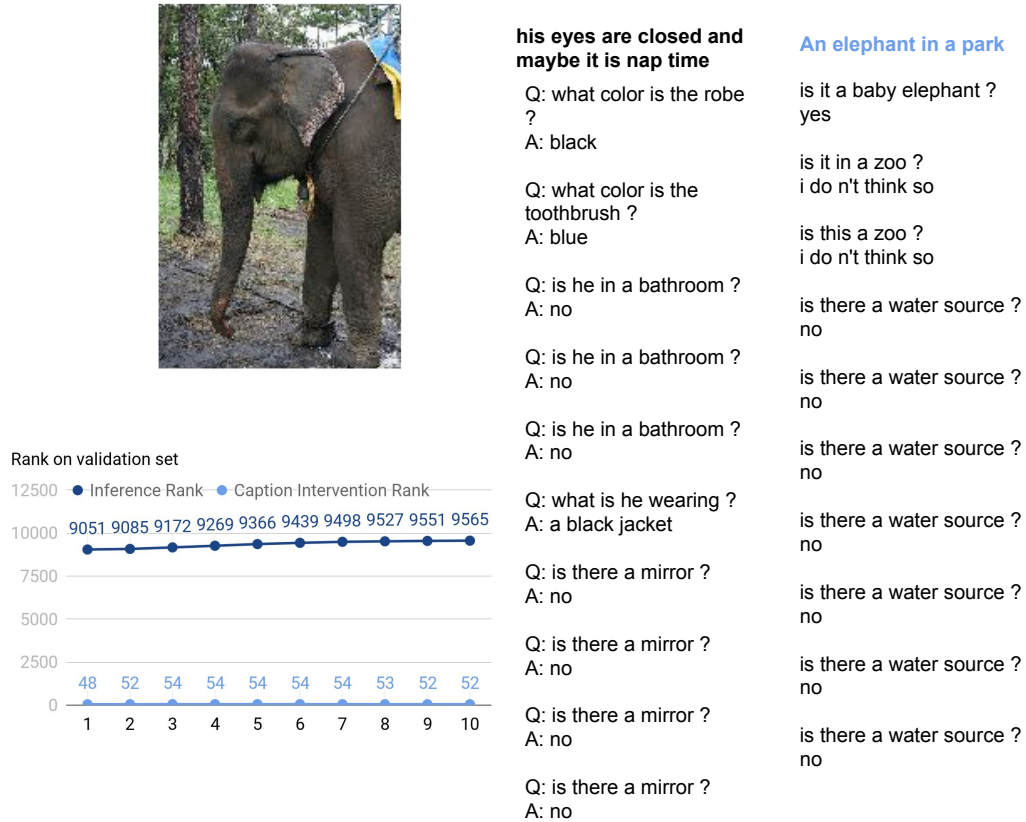
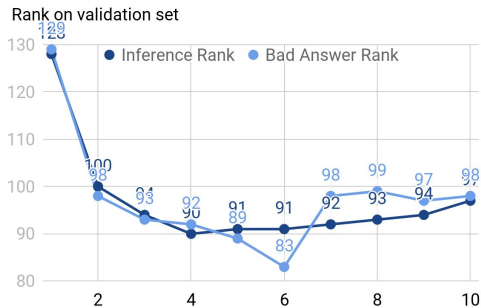


Figure 2: Manual 'positive' intervention on the caption of the image. The original caption (top of left column in bold) is uninformative and results in a poor rankings of approximately 9000 out of approximately 40000. Changing the caption to make it more descriptive (right column in blue) improves the rankings dramatically to around 50.

Adding wrong answers



a red stop sign that is on a pole

Q: what does the sign say ?

A: UNK

Q: are there cars ?

A: no

Q: is it daytime ?

A: yes

Q: any cars ?

A: no

Q: any cars ?

A: no

Q: any cars ?

A: no

Q: any people ?

A: no

Q: any people ?

A: no

Q: is it sunny ?

A: yes

Q: is it sunny ?

A: yes

a red stop sign that is on a pole

what does the sign say ?

Turn right

are there cars ?

yes

is it daytime ?

no

any cars ?

yes

any cars ?

yes

any cars ?

yes

any people ?

yes

any people ?

yes

is it sunny ?

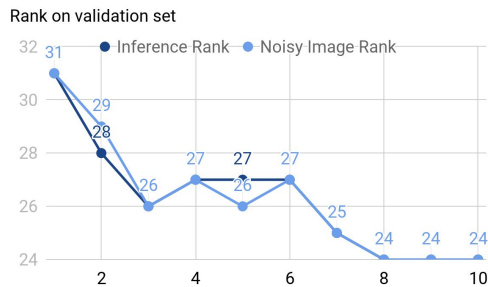
no

is it sunny ?

no

Figure 3: Manual intervention on the answers. We consistently provide answers that are either false or which negate the inferred dialog (displayed in blue). The original dialog achieves a final ranking of 97 out of approximately 40000. Surprisingly, the interventions do not cause a very large perturbation despite being misleading and uninformative.

Noisy image features



a man standing next to a child on a snow covered slope

Q: is the man wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

Q: is he wearing a hat ?
A: yes

a man standing next to a child on a snow covered slope

is the man wearing a hat ?
no

is he wearing a hat ?
no

is he wearing a hat ?
no

is he wearing a hat ?
no

is he wearing a hat ?
no

is he wearing a hat ?
no

is he wearing a hat ?
no


is he wearing a hat ?
no

is he wearing a hat ?
no

is he wearing a hat ?
no

Figure 4: Manual intervention where we replaced the image feature vector with random noise. Surprisingly, there is no noticeable change in image rankings. The decoded sequences by the answer bot are however different.

Manual intervention on questions



the sub from subway is ready to be cut and eaten

Q: is the crane on the hood ?
A: no

Q: is it a commercial truck ?
A: yes

Q: is the picture in color ?
A: yes

Q: is there a person in the picture ?
A: no

Q: is the water calm ?
A: yes

Q: is it a sunny day ?
A: yes

Q: is it a sunny day ?
A: yes

Q: is it sunny ?
A: yes

Q: is it sunny ?
A: yes

Q: is it sunny ?
A: yes

the sub from subway is ready to be cut and eaten

Are there people ?
no

is it food ?
no

Is it on a table ?
yes

What color is the table?
White

Is there a knife ?
no

Is it big ?
yes

Where is it located in the image ?
On a table

is it dark ?
no

Does it fit in one hand ?
Yes

Are there vegetables ?
yes

Figure 5: Manual intervention where we provided more meaningful questions (displayed in blue). Surprisingly, there is no change in image rankings. The two image rankings perfectly overlap.