
Visual Entailment Task for Visually-Grounded Language Learning

Ning Xie*
Wright State University
Dayton, OH 45435
xie.25@wright.edu

Farley Lai
NEC Laboratories America
Princeton, NJ 08540
farleylai@nec-labs.com

Derek Doran
Wright State University
Dayton, OH 45435
derek.doran@wright.edu

Asim Kadav
NEC Laboratories America
Princeton, NJ 08540
asim@nec-labs.com

Abstract

We introduce a new inference task - **Visual Entailment (VE)** - which differs from traditional Textual Entailment (TE) tasks whereby a premise is defined by an image, rather than a natural language sentence as in TE tasks. A novel dataset *SNLI-VE* is proposed for VE tasks based on the Stanford Natural Language Inference corpus and Flickr30K. We introduce a differentiable architecture called the Explainable Visual Entailment model (EVE) to tackle the VE problem. EVE and several other state-of-the-art visual question answering (VQA) based models are evaluated on the SNLI-VE dataset, facilitating grounded language understanding and providing insights on how modern VQA based models perform.

1 Introduction

Multimodal inference, reasoning, and fact entailment across image data and text have the potential to solve problems where the veracity of a text statement is drawn from visual facts. Representative applications involve the fake news detection and court cross-examination. The former aims to detect contradictions between the text news and visual facts such as an image or video clip in order to reduce the influence of misleading news. The latter intends to validate the testimony in case of any contradictions to visual evidence for a fair judgment.

Recent progress in visual reasoning using datasets such as the Visual Question Answering (VQA) dataset [Antol et al., 2015] and CLEVR [Johnson et al., 2017a] has been encouraging. However, the high accuracy in these datasets is often because of the bias in these datasets. For the VQA dataset, there is a question-conditioned bias [Goyal et al., 2017] where questions may hint at the answers such that the correct answer may be inferred without even considering the visual information. The following version of the VQA dataset [Goyal et al., 2017] reduces the bias by pairing questions with similar images that lead to different answers. Even so, the sentence structures in the VQA dataset remain simple and the yes/no questions are insufficient for training entailment tasks that include the neutral case. CLEVR on the other hand is designed for fine-grained reasoning but its synthetic nature introduces the uniformity in image and text structures, resulting in very high accuracy models [Hudson and Manning, 2018] that may not generalize well to real world settings. Hence, we need a more challenging inference task that requires learning grounded representations from cross-modal (image, text) pairs, where the same image is used for multiple natural language sentences,

*Work performed as a NEC Labs intern




Figure 2: EVE architecture. EVE determines if a hypothesis (text input) is entailed by an image premise (image input). The bottom half shows two methods on image feature extraction, either from the CNN feature maps or object detection ROIs.

represents an image region. In contrast, EVE-ROI considers regions of interest (ROI) proposals from MASK-RCNN [He et al., 2017] to locate prominent objects in images. The image regions either from EVE-ROI or EVE-Image are also self-attended and further weighted by the text-image attention. Both the text and image features are finally fused for later prediction.

We apply **self-attention** to capture the hidden relations between elements in the text and the image feature spaces respectively. The intuition of using self-attention is, under a long and complex hypothesis, it is increasingly necessary for the model to be able to attend to only the most relevant words. The effect of self-attention on the image is similar: image regions that jointly benefit the current prediction receive more attention. On the other hand, the **text-image attention** allows the model to select relevant image regions conditioned on the given text hypothesis.

Lastly, the **split-transform-merge** technique follows from the VQA 2017 winner [Teney et al., 2017]. The fused features are split and transformed through the Text MLP and the Image MLP respectively in expectation of more representational power to make better conclusions. On the contrary, there are no similar pre-trained MLP weights corresponding to the answers as in VQA tasks because the VE class labels are not associated with particular images or text. Therefore, those MLP weights are simply initialized randomly.

3 Evaluation on SNLI-VE

Model Name	Val Acc (%)	Val Acc Per Class (%)			Test Acc (%)	Test Acc Per Class (%)		
		C	N	E		C	N	E
Hypothesis Only	67.04	65.45	63.36	72.31	67.01	65.85	63.78	71.40
Image Captioning	68.14	67.3	63.12	73.99	67.47	66.75	63.56	72.07
Relational Network	67.81	68.01	63.94	71.49	68.39	69.13	65.58	70.45
Attention Top-Down	70.59	72.94	66.88	71.96	70.3	72.94	66.63	71.34
Attention Bottom-Up	69.79	71.56	64.25	73.57	69.34	70.56	64.49	72.96
EVE-Image*	71.40	70.48	66.88	76.83	71.36	70.61	67.17	76.31
EVE-ROI*	71.11	66.41	68.2	78.69	70.21	65.63	68.83	76.16

Table 1: Model Performance on SNLI-VE dataset

We evaluate the performance of EVE against several other baselines over SNLI-VE including the existing state-of-the-art VQA based models. Details about the dataset and our experiments are

discussed in the supplemental materials. The performance results, as listed in Table 1, involve comparisons between the following models:

Hypothesis Only: This model uses hypotheses only without image premises. Based on no premises, the model was expected to make random guesses but the resulting accuracy is up to 67%, as reproduced by others [Gururangan et al., 2018, Vu et al., 2018]. This indicates the performance of our model must exceed the 67% lower bound to make sense.

Image Captioning: Before VE, there are many captioning models [Karpathy and Fei-Fei, 2015, Vinyals et al., 2017, Chen et al., 2017] which can serve as a useful baseline by generating an image caption as the premise and then apply existing TE models for classification. For this baseline, we use a PyTorch implementation [Choi] which extracts the image features with a pre-trained ResNet152 backbone and generates the captions using an LSTM. The generated text premise is encoded with the input text hypothesis. Both text features are concatenated for classification. The model performance achieves a marginally higher accuracy of 68.14% and 67.47% on the validation and test sets respectively, implying that the generated image caption premise does not help much. After reviewing the generated captions, it is possible that the quality of the generated captions are too poor or missing the necessary information for the TE classifier. To address this problem, the captioning may be improved by using sophisticated models such as the dense captioning [Johnson et al., 2016] but there is no guarantee that every detail in the image potentially described by the hypothesis would be covered. Nevertheless, the TE classifier could still perform poorly due to the increase in the length of the caption premises.

Relational Network: The Relational Network (RN), proposed to tackle the CLEVR dataset considers pairwise feature fusions between different image regions in the CNN feature maps and the question embedding [Santoro et al., 2017]. Although RN provides high accuracy on CLEVR, only a marginal improvement is achieved at the accuracy of 67.81% and 68.39% on the validation and test splits of SNLI-VE.

Attention Top-Down: We also adopt the model from the winner [Anderson et al., 2018] of VQA challenge 2017, which applies text-image attention to the image regions in the CNN feature maps based on the question embedding. The weighted image features are then projected and fused with the question embedding using dot-product for classification. This attention based VQA model achieves the best accuracy so far, with 70.59% and 70.3% on the validation and test splits, respectively, implying attention can effectively use image premise features.

Attention Bottom-Up: The model design for Attention Bottom-Up is quite similar to Attention Top-Down, except the image features used are the ROIs extracted by a Mask-RCNN [He et al., 2017] implementation [Matterport]. The best performance achieved is 69.79% and 69.34% accuracy on the validation and testing splits respectively. Though we also evaluate the model with more than 10 ROIs, we observe no significant improvement.

EVE-Image and EVE-ROI: We finally evaluate our model, EVE, as described in Section 2. EVE-Image achieves the best performance of 71.4% and 71.36% accuracy on the validation and test partitions. EVE-ROI achieves a slightly lower accuracy of 71.11% and 70.21% but still better than the counterpart Attention Bottom-Up. The improvement, even just marginal, is likely attributed to the introduction of self-attention that captures the hidden relations in the same feature space. We do not find evidence that the split-transform-merge construct contributes much to VE tasks.

4 Conclusion

This work introduces *visual entailment*, a novel multimodal task to determine if a text hypothesis is entailed based on the visual information in the image premise. We build the SNLI-VE dataset providing real-world images from Flickr30K as premises, and the corresponding text hypotheses from SNLI. To address VE, we develop EVE and demonstrate its performance over several baselines, including the existing state-of-the-art VQA based models. The inherent language-bias induced by SNLI [Gururangan et al., 2018] serves as a strong baseline. The SNLI-VE dataset is scheduled to be publicly available.

Acknowledgments

Ning Xie and Derek Doran were supported by the Ohio Federal Research Network project *Human-Centered Big Data*. Any opinions, findings, and conclusions or recommendations expressed in this article are those of the author(s) and do not necessarily reflect the views of the Ohio Federal Research Network.

References

- P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, volume 3, page 6, 2018.
- S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Zitnick, and D. Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*, 2015.
- L. Chen, H. Zhang, J. Xiao, L. Nie, J. Shao, W. Liu, and T.-S. Chua. Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6298–6306. IEEE, 2017.
- Y. Choi. Image captioning pytorch implementation. https://github.com/yunjey/pytorch-tutorial/tree/master/tutorials/03-advanced/image_captioning. Accessed: 2018-10-30.
- A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*, 2016.
- Y. Gong, L. Wang, M. Hodosh, J. Hockenmaier, and S. Lazebnik. Improving image-sentence embeddings using large weakly annotated photo collections. In *European Conference on Computer Vision*, pages 529–545. Springer, 2014.
- Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, volume 1, page 3, 2017.
- S. Gururangan, S. Swayamdipta, O. Levy, R. Schwartz, S. R. Bowman, and N. A. Smith. Annotation artifacts in natural language inference data. *arXiv preprint arXiv:1803.02324*, 2018.
- K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2980–2988. IEEE, 2017.
- D. A. Hudson and C. D. Manning. Compositional attention networks for machine reasoning. *arXiv preprint arXiv:1803.03067*, 2018.
- J. Johnson, A. Karpathy, and L. Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4565–4574, 2016.
- J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 1988–1997. IEEE, 2017a.
- J. Johnson, B. Hariharan, L. van der Maaten, J. Hoffman, L. Fei-Fei, C. L. Zitnick, and R. B. Girshick. Inferring and executing programs for visual reasoning. In *ICCV*, pages 3008–3017, 2017b.
- A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015.
- J.-H. Kim, J. Jun, and B.-T. Zhang. Bilinear attention networks. *arXiv preprint arXiv:1805.07932*, 2018.
- I. Matterport. Mask rcnn pytorch implementation. <https://github.com/multimodallearning/pytorch-mask-rcnn>. Accessed: 2018-10-30.
- J. Pennington, R. Socher, and C. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In *Advances in neural information processing systems*, pages 2953–2961, 2015.

- A. Santoro, D. Raposo, D. G. Barrett, M. Malinowski, R. Pascanu, P. Battaglia, and T. Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976, 2017.
- D. Teney, P. Anderson, X. He, and A. van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. *arXiv preprint arXiv:1708.02711*, 2017.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):652–663, 2017.
- H. T. Vu, C. Greco, A. Erofeeva, S. Jafaritazehjan, G. Linders, M. Tanti, A. Testoni, R. Bernardi, and A. Gatt. Grounded textual entailment. *arXiv preprint arXiv:1806.05645*, 2018.
- P. Young, A. Lai, M. Hodosh, and J. Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014.
- Y. Zhu, O. Groth, M. Bernstein, and L. Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4995–5004, 2016.

Supplementary Materials

Dataset statistics. The original SNLI dataset split does not consider the arrangement of the original caption images. Therefore, the same image may appear in both training and test sets if directly adapted to VE. To address the issue, we disjointedly partition SNLI-VE by images following the splits in [Gong et al., 2014] and make sure that each class instances are balanced across the training, validation, and test sets as shown in Table 2.

	Training	Validation	Testing
#Images	29,783	1,000	1,000
#Entailment	176,932	5,959	5,973
#Neutral	176,045	5,960	5,964
#Contradiction	176,550	5,939	5,964
Vocabulary Size	29,550	6,576	6,592

Table 2: SNLI-VE statistics: number of images, per class examples and vocabulary size by split.

Implementation details. The proposed EVE model is implemented in PyTorch. We use the pre-trained GloVe.6B.300D [Pennington et al., 2014] for word embedding, where 6B is the corpus size and 300D is the embedding dimension. The image features used for EVE-Image are generated from a pre-trained ResNet101. The ROI features used for EVE-ROI are extracted using the Mask-RCNN implementation [Matterport]. The Adam optimizer is used for training with a batch size of 64. Adaptive learning rate is applied with both initial value and weight decay set to be 0.0001.