

Going Beneath the Surface: Evaluating Image Captioning for Grammaticality, Truthfulness and Diversity

Huiyuan Xie¹, Tom Sherborne², Alexander Kuhnle¹, and Ann Copestake¹

¹Department of Computer Science and Technology, University of Cambridge

²School of Informatics, University of Edinburgh

{hx255,aok25,aac10}@cam.ac.uk, tom.sherborne@ed.ac.uk

Abstract


Image captioning as a multimodal task has drawn much interest in recent years. However, evaluation for this task remains a challenging problem. Existing evaluation metrics focus on surface similarity between a candidate caption and a set of reference captions, and do not check the actual relation between a caption and the underlying visual content. We introduce a new diagnostic evaluation framework for the task of image captioning, with the goal of directly assessing models for grammaticality, truthfulness and diversity (GTD) of generated captions. We demonstrate the potential of our evaluation framework by evaluating existing image captioning models on a wide ranging set of synthetic datasets that we construct for diagnostic evaluation. We empirically show how the GTD evaluation framework, in combination with diagnostic datasets, can provide insights into model capabilities and limitations to supplement standard evaluations.

Introduction

Automatically generating text to describe the content of images, also known as image captioning, is a multimodal task of considerable interest in both the computer vision and the NLP communities. Image captioning can be framed as a translation task from an image to a descriptive natural language statement. Many existing captioning models (Vinyals et al. 2015; Donahue et al. 2015; Yao et al. 2017; Aneja, Deshpande, and Schwing 2018) follow the typical encoder-decoder framework where a convolutional network is used to condense images into visual feature representations, combined with a recurrent network for language generation. While these models demonstrate promising results, quantifying image captioning performance remains a challenging problem, in a similar way to other generative tasks (Radev et al. 2003; Gatt and Kraemer 2018).

Evaluating candidate captions for human preference is slow and laborious. To alleviate this problem, many automatic evaluation metrics have been proposed, such as BLEU (Papineni et al. 2002), METEOR (Banerjee and Lavie 2005), ROUGE (Lin 2004) and CIDEr (Vedantam, Lawrence Zitnick, and Parikh 2015). These n-gram-based metrics eval-

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Caption 1: A circle is above a green rectangle.

Caption 2: A blue triangle is to the left of a semicircle.

Caption 3: A semicircle is below a gray triangle.

Caption 4: A semicircle is to the left of a triangle.

Figure 1: ShapeWorld example: spatial statements in the context of multiple shapes. The first three statements are truthful and diverse descriptions of the image. The fourth statement is wrong, but nonetheless exhibits a high degree of n-gram overlap with the true reference captions.

uate captioning performance based on surface similarity between a candidate caption and reference statements. A more recent evaluation metric for image captioning is SPICE (Anderson et al. 2016), which takes into account semantic propositional content of generated captions by scoring a caption based upon a graph-based semantic representation transformed from reference captions.

The rationale behind these evaluation metrics is that human reference captions serve as an approximate target and comparing model outputs to this target is a proxy for how well a system performs. Thus, a candidate caption is not directly evaluated with respect to image content, but compared to a set of human statements about that image.

However, in image captioning, visual scenes with multiple objects and relations correspond to a diversity of valid descriptions. Consider the example image and captions from the ShapeWorld framework (Kuhnle and Copestake 2017) shown in Figure 1. The first three captions are true statements about the image and express relevant ideas, but describe different objects, attributes and spatial relationships, while the fourth caption is wrong despite referring to the same objects as in the third caption. This casts doubt on the sufficiency of using a set of reference captions to approximate the content of an image. We argue that, while existing metrics have undeniably been useful for real-world captioning evaluation, their focus on approximate surface compar-

son limits deeper insights into the learning process and eventual behavior of captioning models.

To address this problem, we propose a set of principled evaluation criteria which evaluate image captioning models for grammaticality, truthfulness and diversity (GTD). These criteria correspond to necessary requirements for image captioning systems: (a) that the output is grammatical, (b) that the output statement is true with respect to the image, and (c) that outputs are diverse and mirror the variability of training captions.

Practical evaluation of GTD is currently only possible on synthetic data. We construct a range of datasets designed for image captioning evaluation. We call this diagnostic evaluation benchmark ShapeWorldICE (*ShapeWorld for Image Captioning Evaluation*). We illustrate the evaluation of specific image captioning models on ShapeWorldICE. We empirically demonstrate that the existing metrics BLEU and SPICE do not capture true caption-image agreement in all scenarios, while the GTD framework allows a fine-grained investigation of how well existing models cope with varied visual situations and linguistic constructions.

We believe that as a supplementary evaluation method to real-world metrics, the GTD framework provides evaluation insights that are sufficiently interesting to motivate future work.

Related work

Existing evaluation of image captioning

As a natural language generation task, image captioning frequently uses evaluation metrics such as BLEU (Papineni et al. 2002), METEOR (Banerjee and Lavie 2005), ROUGE (Lin 2004) and CIDEr (Vedantam, Lawrence Zitnick, and Parikh 2015). These metrics use n-gram similarity between the candidate caption and reference captions to approximate the correlation between a candidate caption and the associated ground truth. SPICE (Anderson et al. 2016) is a more recent metric specifically designed for image captioning. For SPICE, both the candidate caption and reference captions are parsed to scene graphs, and the agreement between tuples extracted from these scene graphs is examined. SPICE more closely relates to our truthfulness evaluation than the other metrics, but it still uses overlap comparison to reference captions as a proxy to ground truth. In contrast, our truthfulness metric directly evaluates a candidate caption against a model of the actual visual content.

Many researchers have pointed out problems with existing reference-based metrics including low correlations with human judgment (Elliott and Keller 2014; Anderson et al. 2016; Kilickaya et al. 2017) and strong baselines using nearest-neighbor methods (Devlin et al. 2015b) or relying solely on object detection (Wang, Madhyastha, and Specia 2018). Fundamental concerns have been raised with respect to BLEU, including variability in parameterization and precise score calculation leading to significantly different results (Post 2018). Its validity as a metric for tasks other than machine translation has been questioned (Reiter 2018), par-

ticularly for tasks for which the output content is not narrowly constrained, like dialogue (Liu et al. 2016).

Some recent work focuses on increasing the diversity of generated captions, for which various measures are proposed. Devlin et al. (Devlin et al. 2015a) explored the concept of caption diversity by evaluating performance on compositionally novel images. van Miltenburg et al (van Miltenburg, Elliott, and Vossen 2018) framed image captioning as a word recall task and proposed several metrics, predominantly focusing on diversity at the word level. However, this direction is still relatively new and lacks standardized benchmarks and metrics.

Synthetic datasets

Recently, many synthetic datasets have been proposed as diagnostic tools for deep learning models, such as CLEVR (Johnson et al. 2017) for visual question answering (VQA), the bAbI tasks (Weston et al. 2015) for text understanding and reasoning, and ShapeWorld (Kuhnle and Copestake 2017) for visually grounded language understanding. The primary motivation is to reduce complexity which is considered irrelevant to the evaluation focus, to enable better control over the data, and to provide more detailed insights into strengths and limitations of existing models.

In this work, we develop the evaluation datasets within the ShapeWorld framework. ShapeWorld is a controlled data generation framework consisting of abstract colored shapes (see Figure 1 for an example). We use ShapeWorld to generate training and evaluation data for two major reasons. ShapeWorld supports customized data generation according to user specification, which enables a variety of model inspections in terms of language construction, visual complexity and reasoning ability. Another benefit is that each training and test instance generated in ShapeWorld is returned as a triplet of $\langle \text{image, caption, world model} \rangle$. The world model stores information about the underlying microworld used to generate an image and a descriptive caption, internally represented as a list of entities with their attributes, such as *shape, color, position*. During data generation, ShapeWorld randomly samples a world model from a set of available entities and attributes. The generated world model is then used to realize a corresponding instance consisting of image and caption. The world model gives the actual semantic information contained in an image, which allows evaluation of caption truthfulness.

GTD Evaluation Framework

In the following we introduce GTD in more detail, consider it as an evaluation protocol covering necessary aspects of the multifaceted captioning task, rather than a specific metric.

Grammaticality

An essential criterion for an image captioning model is that the captions generated are grammatically well-formed. Fully accurate assessment of grammaticality in a general context is itself a difficult task, but becomes more feasible in a very constrained context like our diagnostic language data.







Type	Variant	Caption	Image
Existential	OneShape	There is a green cross. A rectangle is green. There is a cyan shape.	
	MultiShapes	A shape is a gray triangle. There is a square. There is a yellow shape.	
Spatial	TwoShapes	A square is above a red pentagon. A yellow square is above a yellow pentagon. A square is to the left of a pentagon.	
	MultiShapes	A blue triangle is to the left of a semicircle. A circle is above a green rectangle. A semicircle is to the left of a circle.	
Quantification	Count	Exactly two rectangles are green. Exactly one shape is a yellow circle. Exactly zero shapes are ellipses.	
	Ratio	A quarter of the shapes are rectangles. A third of the rectangles are magenta. Half the shapes are green.	

Table 1: Sample captions and images from ShapeWorldICE datasets (truthful captions in blue, false in red). Images from Existential-OneShape only contain one object, while images from Spatial-TwoShapes contain two objects. Images from the other four datasets follow the same distribution with multiple abstract objects present in a visual scene.

We take parseability with the English Resource Grammar (Flickinger 2000, ERG) as a surrogate for grammaticality, meaning that a sentence is considered grammatically well-formed if we obtain a parse using the ERG.

The ERG is a broad-coverage grammar based on the head-driven phrase structure grammar (HPSG) framework. It is linguistically precise: sentences only parse if they are valid according to its hand-built rules. It is designed to be general-purpose: verified coverage is around 80% for Wikipedia, and over 90% for corpora with shorter sentences and more limited vocabulary (for details see Flickinger (2011)). Since the ShapeWorld training data – the only language source for models to learn from – is generated using the same grammar, the ERG has $\sim 100\%$ coverage of grammaticality in the model output space.

Truthfulness

The second aspect we investigate is truthfulness, that is, whether a candidate caption is compatible with the content of the image it is supposed to describe. We evaluate caption truthfulness on the basis of a linguistically-motivated approach using formal semantics. We convert the output of the ERG parse for a grammatical caption to a Dependency Minimal Recursion Semantics (DMRS) graph using the `pydmrs` tool (Copestake et al. 2016). Each converted DMRS is a logical semantic graph representation corresponding to the caption. We construct a logical proposition from the DMRS graph, and evaluate it against the actual world model of the corresponding image. A caption can be said to agree with an image only if the proposition evaluates as true on the basis of the world model. By examining the logical agreement between a caption representation and a world model, we can check whether the semantics of this caption agrees with the

visual content which the world model represents. Thus we do not rely on a set of captions as a surrogate for the content of an image, but instead leverage the fact that we have the ground truth, thus enabling the evaluation of true image-caption agreement.

Diversity

While grammaticality and truthfulness are essential requirements for image captions, these criteria alone can easily be “gamed” by specializing on a small set of generic statements which are true most of the time. In the context of abstract shapes, such captions include examples like “*There is a shape*” or “*At least zero shapes are blue*” (which is technically true even if there is no blue shape). This motivates the third fundamental requirement of captioning output to be diverse.

As ShapeWorldICE exploits a limited size of open-class words, we emphasize the diversity in ShapeWorldICE at the sentence level rather than the word level. Since the ground-truth reference captions in ShapeWorld are randomly sampled, we take the sampled captions accompanying the test images as a proxy for optimal caption diversity, and compare it with the empirical output diversity of the evaluated model on these test images. Practically, we look at language constructions used and compute the corresponding diversity score as the ratio of observed number versus optimal number:

$$\text{diversity} = \frac{\#\{\text{model-generated}\}}{\#\{\text{ShapeWorld-generated}\}}$$

Language constructions here correspond to reduced caption representations which only record whether an object is described by shape (e.g., “*square*”), color (e.g., “*red shape*”) or color-shape combination (e.g., “*red square*”). So the

statement “A square is red” and “A circle is blue” are considered the same, while “A shape is red” is different.




Figure 2: Performance comparison of the Show&Tell model and the LRCN_{1u} model on Existential-MultiShapes. *SnT* represents the Show&Tell model while *LRCN* represents the LRCN_{1u} model. *Grammaticality*, *Truthfulness* and *Diversity* refer to the grammaticality ratio, the truthfulness ratio and the diversity ratio of generated captions, respectively.

Experimental Setup

Datasets

We develop a variety of ShapeWorldICE datasets, with a similar idea to the “skill tasks” in the bAbI framework (Weston et al. 2015). Table 1 gives an overview for different ShapeWorldICE datasets we use in this paper. We consider three different types of captioning tasks, each of which focuses on a distinct aspect of reasoning abilities. Existential descriptions examine whether a certain object is present in an image. Spatial descriptions identify spatial relationships among visual objects. Quantification descriptions involve count-based and ratio-based statements, with an explicit focus on inspecting models for their counting ability. We develop two variants for each type of dataset to enable different levels of visual complexity or specific aspects of the same reasoning type. All the training and test captions sampled in this work are in English.

Each dataset variant consists of around 200k training instances and 4,096 validation instances, plus 4,096 test instances. Each training instance consists of an image and a reference caption. At test time, only the test images are available to the evaluated models. Underlying world models are kept from the models and are used for later GTD evaluation. For each test instance, we sample ten reference captions of the same distribution as the training captions to enable the comparison of our proposed metrics to BLEU and SPICE. We fine-tune our model hyperparameters based on the performance on the validation set. All reported results are measured on the test split with the parameters yielding the best validation performance.

Models

We experiment with two image captioning models: the Show&Tell model (Vinyals et al. 2015) and the LRCN_{1u} model (Donahue et al. 2015). Both models follow the basic encoder-decoder architecture design that uses a CNN encoder to condense the visual information into an image embedding, which in turn conditions an LSTM decoder to generate a natural language caption. The main difference between the two models is the way they condition the decoder. The Show&Tell model feeds the image embedding as the “*predecessor word embedding*” to the first produced word, while the LRCN_{1u} model concatenates the image features with the embedded previous word as the input to the sequence model at each time step.

We follow the common practice in image captioning to use a CNN component pretrained on object detection and fine-tune its parameters on the image captioning task. The encoder and decoder components are jointly optimized with respect to the standard cross-entropy sequence loss on the respective ShapeWorldICE dataset. For all our experiments, we train models end-to-end for a fixed number of 100k iterations with a batch size of 64. We use Adam optimization (Kingma and Ba 2014) with a learning rate of 0.001. Word embeddings are randomly initialized and jointly trained during the training.

Results

We train and evaluate the Show&Tell and LRCN_{1u} models on the ShapeWorldICE datasets. Here we discuss in detail the diagnostic results of these experiments. During training, we periodically record model output on the test images, to be able to analyze the development of our evaluation metrics throughout the process. We also compute BLEU-4 scores and SPICE scores of generated captions for comparison, using 10 reference captions per test image.

LRCN_{1u} exhibits clearly superior performance in terms of truthfulness. We start off by comparing performance of the Show&Tell model and the LRCN_{1u} model, see Figure 2. While both models learn to produce grammatical sentences early on, it can be seen that LRCN_{1u} is clearly superior in terms of truthfulness, achieving 100% halfway through training, whereas Show&Tell only slowly reaches around 90% by the end of 100k iterations. This indicates that incorporating visual features at every generation step is beneficial for producing true captions. The diversity ratios of captions generated by two models both increase substantially as the training progresses, with LRCN_{1u} exhibiting a slightly greater caption diversity at the end of training.

We observed similar results on other ShapeWorldICE datasets that we experimented with, validating the superiority of LRCN_{1u} over Show&Tell on ShapeWorldICE. Consequently, we decided to focus on the LRCN_{1u} architecture in subsequent evaluations, where we report detailed results with respect to the GTD framework on a variety of datasets.

Correlation between the BLEU/SPICE scores and the ground truth. From the learning curves shown in Figure




Figure 3: Learning curves for $LRCN_{1u}$ on Existential-OneShape, Existential-MultiShapes and Spatial-MultiShapes. *Truthfulness* refers to the ratio of generated captions that are grammatical and agree with ground-truth world models. *BLEU* and *SPICE* denote average BLEU-4 scores and average SPICE scores across the test split, respectively.

3, we find low or no correlation between the BLEU/SPICE scores and caption truthfulness.

On Existential-OneShape, the BLEU curve follows the trend of the truthfulness curve in general, indicating that BLEU is able to capture caption truthfulness well in this simple scenario. However, while BLEU reports equivalent model performance on Existential-MultiShapes and Spatial-MultiShapes, the truthfulness metric demonstrates very different results. The BLEU score for generated Existential-MultiShapes captions increases rapidly at the beginning of training and then plateaus despite the continuous increase in truthfulness ratio. Captions generated on Spatial-MultiShapes attain a relatively high BLEU score from an early stage of training, but exhibit low agreement (<0.6 truthfulness ratio) with ground-truth visual scenes. In the case of Spatial-MultiShapes, spatial descriptors for two objects are chosen from a fixed set (“above”, “below”, “to the left of” and “to the right of”). It is very likely for a generated spatial descriptor to match one of the descriptors mentioned in reference captions. In this particular case, the model is apt to infer a caption which has high n-gram overlaps with reference captions, resulting in a relatively high BLEU score. Thus an increased BLEU score does not necessarily indicate improved performance.

While the truthfulness and BLEU scores in Figure 3a both increase rapidly early on and then stay stable at a high rate after training for 20k iterations, the SPICE curve instead shows a downward trend in the later stage of training. We examined the output SPICE score for each test instance. SPICE reports a precision score of 1.0 for most test instances after 20k iterations, which is consistent with the truthfulness and BLEU scores. However, SPICE forms the reference scene graph as the union of the scene graphs extracted from individual reference captions, thus introducing redundancies. SPICE uses the F1 score of scene graph matching between the candidate and reference and hence is lowered by imperfect recall.

Comparing SPICE curves for three datasets shown in Fig-

ure 3a-3c, they suggest an increase in task complexity, but they do not reflect the successively closing gap of caption truthfulness scores between two Existential datasets, or the substantial difference in caption truthfulness between captions on Existential-MultiShapes and Spatial-MultiShapes.

In the remainder of the paper we discuss in detail the diagnostic results of the $LRCN_{1u}$ model demonstrated by the GTD evaluation framework.




Figure 4: Ratio of grammatical sentences produced by $LRCN_{1u}$ for different ShapeWorldICE datasets in the first 20k training iterations (stays at 100% afterwards).

Perfect grammaticality for all caption types. As shown in Figure 4, generated captions for all types of ShapeWorldICE datasets attain quasi-perfect grammaticality scores in fewer than 5,000 iterations, suggesting that the model quickly learns to generate grammatically well-formed sentences.

Failure to learn complex spatial relationships. While CNNs can produce rich visual representations that can be used for a variety of vision tasks (Sermanet et al. 2013), it remains an open question whether these condensed visual representations are rich enough for mul-




Figure 5: Truthfulness ratio of sentences produced by $LRCN_{1u}$ for different ShapeWorldICE datasets.

timodal tasks that require higher-level abilities of scene understanding and visual reasoning. From Figure 5, we can see that while the model performs rather well on Existential datasets, it exhibits a worse performance on Spatial data. The caption agreement ratio in the simple Spatial-TwoShapes scenario is relatively high, but drops significantly on Spatial-MultiShapes, demonstrating the deficiencies of the model in learning spatial relationships from complex visual scenes.

The counting task is non-trivial. Counting has long been considered to be a challenging task in multimodal reasoning (Antol et al. 2015; Jabri, Joulin, and Van Der Maaten 2016). To explore how well the $LRCN_{1u}$ model copes with counting tasks, we generated two Quantification datasets. The Quant-Count captions describe the number of objects with certain attributes that appear in an image (e.g. “Exactly four shapes are crosses”), while the Quant-Ratio captions describe the ratio of certain objects (e.g. “A third of the shapes are blue squares”).

From Figure 5, we notice that the $LRCN_{1u}$ model performs poorly on these datasets in terms of truthfulness, reflected in the 0.50 and 0.46 scores achieved by the model on the Quant-Count and Quant-Ratio tasks respectively. The learning curve for Quant-Ratio exhibits a more gradual rise as the training progresses, suggesting a greater complexity for the ratio-based task.

Caption diversity benefits from varied language constructions in the training data. The diversity ratios of generated captions for different ShapeWorldICE datasets are illustrated in Figure 6. We can see that the diversity of inferred captions is largely sensitive to the caption variability in the dataset itself. For simple datasets (such as Existential-OneShape) where language constructions in the training set are less diverse, the output captions tend to have uniform sentence structures. The high diversity ratios of generated Spatial and Quantification captions suggest that caption diversity benefits from heterogeneous language constructions in complex datasets.




Figure 6: Diversity ratio of sentences produced by $LRCN_{1u}$ on different ShapeWorldICE datasets.

Discussions and Conclusions

Evaluation metrics are required as a proxy for performance in real applications. As such, they should, as far as possible, allow measurement of fundamental aspects of the performance of models on tasks. In this work, we propose the GTD evaluation framework as a supplement to standard image captioning evaluation which explicitly focuses on grammaticality, truthfulness and diversity. We developed the ShapeWorldICE evaluation suite to allow in-depth and fine-grained inspection of model behaviors. We have empirically verified that GTD captures different aspects of performance to existing metrics by evaluating image captioning models on the ShapeWorldICE suite. We hope that this framework will shed light on important aspects of model behaviour and that this will help guide future research efforts.

While performing the evaluation experiments on the $LRCN_{1u}$ model, we noticed that caption agreement does not always improve as the training loss decreases. Ideally, the training objective should be in accordance with how a model is eventually evaluated. In future work, we plan to investigate the feasibility of deliberately encoding the GTD signal in the training process, for instance, by implementing a GTD-aware loss. We also plan to extend the existing ShapeWorldICE benchmark to include more linguistic constructions (such as relative clauses, compound sentences and coreference). By doing so, we hope to reveal how well existing image captioning models cope with complex generation tasks.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. HX is grateful for being supported by the CSC Cambridge Scholarship. TS is supported in part by the EPSRC Centre for Doctoral Training in Data Science, funded by the EPSRC (grant EP/L016427/1) and the University of Edinburgh. AK is grateful for being supported by a Qualcomm Research Studentship and an EPSRC Doctoral Training Studentship.

References

- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. In *ECCV*, 382–398.
- Aneja, J.; Deshpande, A.; and Schwing, A. G. 2018. Convolutional image captioning. In *CVPR*, 5561–5570.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Lawrence Zitnick, C.; and Parikh, D. 2015. Vqa: Visual question answering. In *ICCV*, 2425–2433.
- Banerjee, S., and Lavie, A. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL*, 65–72.
- Copestake, A.; Emerson, G.; Goodman, M. W.; Horvat, M.; Kuhnle, A.; and Muszynska, E. 2016. Resources for building applications with dependency minimal recursion semantics. In *LREC*.
- Devlin, J.; Cheng, H.; Fang, H.; Gupta, S.; Deng, L.; He, X.; Zweig, G.; and Mitchell, M. 2015a. Language models for image captioning: The quirks and what works. *arXiv preprint arXiv:1505.01809*.
- Devlin, J.; Gupta, S.; Girshick, R.; Mitchell, M.; and Zitnick, C. L. 2015b. Exploring nearest neighbor approaches for image captioning. *arXiv e-prints 1505.04467*.
- Donahue, J.; Anne Hendricks, L.; Guadarrama, S.; Rohrbach, M.; Venugopalan, S.; Saenko, K.; and Darrell, T. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2625–2634.
- Elliott, D., and Keller, F. 2014. Comparing automatic evaluation measures for image description. In *ACL*, 452–457.
- Flickinger, D. 2000. On building a more efficient grammar by exploiting types. *Natural Language Engineering* 6(1):15–28.
- Flickinger, D. 2011. Accuracy vs. robustness in grammar engineering. *Language from a Cognitive Perspective: Grammar, Usage, and Processing* 31–50.
- Gatt, A., and Kraemer, E. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research* 61:65–170.
- Jabri, A.; Joulin, A.; and Van Der Maaten, L. 2016. Revisiting visual question answering baselines. In *ECCV*, 727–739.
- Johnson, J.; Hariharan, B.; van der Maaten, L.; Fei-Fei, L.; Lawrence Zitnick, C.; and Girshick, R. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2901–2910.
- Kilickaya, M.; Erdem, A.; Ikişler-Cinbis, N.; and Erdem, E. 2017. Re-evaluating automatic metrics for image captioning. In *EACL*, 199–209.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kuhnle, A., and Copestake, A. 2017. Shapeworld: A new test methodology for multimodal language understanding. *arXiv preprint arXiv:1704.04517*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Liu, C.-W.; Lowe, R.; Serban, I.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*, 2122–2132.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: A method for automatic evaluation of machine translation. In *ACL*, 311–318.
- Post, M. 2018. A call for clarity in reporting bleu scores. In *Proceedings of the Third Conference on Machine Translation*, 186–191.
- Radev, D. R.; Teufel, S.; Saggion, H.; Lam, W.; Blitzer, J.; Qi, H.; Celebi, A.; Liu, D.; and Drabek, E. 2003. Evaluation challenges in large-scale document summarization. In *ACL*, 375–382.
- Reiter, E. 2018. A structured review of the validity of BLEU. *Computational Linguistics* 44(3).
- Sermanet, P.; Eigen, D.; Zhang, X.; Mathieu, M.; Fergus, R.; and LeCun, Y. 2013. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*.
- van Miltenburg, E.; Elliott, D.; and Vossen, P. 2018. Measuring the diversity of automatic image descriptions. In *COLING*, 1730–1741.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *CVPR*, 4566–4575.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator. In *CVPR*, 3156–3164.
- Wang, J.; Madhyastha, P. S.; and Specia, L. 2018. Object counts! Bringing explicit detections back into image captioning. In *NAACL*, 2180–2193.
- Weston, J.; Bordes, A.; Chopra, S.; Rush, A. M.; van Merriënboer, B.; Joulin, A.; and Mikolov, T. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Yao, T.; Pan, Y.; Li, Y.; Qiu, Z.; and Mei, T. 2017. Boosting image captioning with attributes. In *ICCV*, 4894–4902.