

Deep learning evaluation using deep linguistic processing

Alexander Kuhnle

Department of Computer Science
and Technology
University of Cambridge
aok25@cam.ac.uk

Ann Copestake

Department of Computer Science
and Technology
University of Cambridge
aac10@cam.ac.uk

Abstract

We discuss problems with the standard approaches to evaluation for tasks like visual question answering, and argue that artificial data can be used to address these as a complement to current practice. We demonstrate that with the help of existing ‘deep’ linguistic processing technology we are able to create challenging abstract datasets, which enable us to investigate the language understanding abilities of multimodal deep learning models in detail, as compared to a single performance value on a static and monolithic dataset.

1 Introduction & related work

In recent years, deep neural networks (DNNs) have established a new level of performance for many tasks in natural language processing (NLP), speech, computer vision and artificial intelligence. Simultaneously, we observe a move towards simulated environments and artificial data, particularly in reinforcement learning (Bellemare et al., 2013; Brockman et al., 2016). As outlined by Kiela et al. (2016), simulated data is appealing for various reasons. Most importantly, it acts as a prototypical problem presentation, abstracted from its noisy and intertwined real-world appearance.

However, with a few notable exceptions (Scheffler and Young, 2001; Byron et al., 2007), artificial data is relatively little used in NLP. Only recently people started arguing for the use of simulated data, like the long-term research proposal of Mikolov et al. (2015) on learning to understand language from scratch in a virtual environment, and introduced benchmark datasets, like the bAbI tasks (Weston et al., 2015) or the VQA datasets discussed below. Here we focus on the problem of *visually grounded language understanding* in the context of visual question answering (VQA). In principle, this task is particularly interesting from

a semantic perspective, since it combines general language understanding, reference resolution and grounded language reasoning in a simple and clear task. However, recent work (Goyal et al., 2017; Agrawal et al., 2016) has suggested that the popular VQA Dataset (Antol et al., 2015) is inadequate, due to various issues which allow a system to achieve competitive performance without truly learning these abilities.

To address this, modifications to the existing VQA Dataset and several artificial VQA datasets have been released. The former include C-VQA (Agrawal et al., 2017), a new composition-focused split, and VQA 2.0 (Goyal et al., 2017), an extension based on minimal image pairs. Similar approaches have been proposed in the context of image captioning (Shekhar et al., 2017; Hodosh and Hockenmaier, 2016), which relate to our proposal in that they modify language in a principled way. However, despite ‘mild artificiality’, some issues with real-world data like the VQA Dataset remain.

On the other hand, examples of new artificial datasets include the SHAPES dataset (Andreas et al., 2016), the CLEVR dataset (Johnson et al., 2017a), the NLVR dataset (Suhr et al., 2017), and the ShapeWorld framework (Kuhnle and Copestake, 2017), which is our implementation of the proposal presented here. They all consist of images showing abstract scenes with colored objects and, except for NLVR, use artificially produced language. Language generation for SHAPES and CLEVR is template-based and dataset-specific, while ShapeWorld leverages an existing broad-coverage semantic grammar formalism.

These datasets are introduced with the motivation to provide a clear and challenging evaluation for VQA systems. Johnson et al. (2017a) and Kuhnle and Copestake (2017) investigated popular VQA systems on their datasets, and demonstrate how artificial data provides us with detailed

insights previously not possible. Despite its simplicity, they uncover fundamental shortcomings of current VQA models. Since then, CLEVR has been of great importance for the development of new VQA models based on dynamically assembled modules (Hu et al., 2017; Johnson et al., 2017b), a dedicated relational module (Santoro et al., 2017), or a general modulation technique (Perez et al., 2018), all of which achieve close-to-perfect accuracy on CLEVR.

The advantage of artificial data in this paper is not seen in its capacity to improve existing models by augmenting training data, although this would be conceivable. Instead we are interested in its capacity to provide data for targeted investigations of specific model capabilities. We argue that it constitutes a *necessary*, though not in itself *sufficient* benchmark for genuine language understanding abilities. The aforementioned models exhibit clearly superior understanding of the types of questions CLEVR contains. This paper proposes a principled way of continuing the incremental progress in multimodal language understanding initiated by CLEVR and its template-based generation approach, based on deep linguistic processing tools. Our initial experiments show that we can provide data that is challenging for state-of-the-art models, like the quantification examples presented in section 3.3. Note that while success on such narrower datasets may not directly translate to improved performance on broader datasets like the VQA Dataset, the underlying mechanisms are important for progress in the longer run.

Our aims in this paper are threefold. First, we provide a brief but systematic review of the problems surrounding current standard evaluation practices in deep learning. Secondly, we use this to motivate the potential of artificial data from simulated microworlds to investigate DNNs for visually grounded language understanding. Thirdly, we present an evaluation methodology based on linguistic processing resources, and show why compositional semantic representations from a symbolic grammar are particularly suitable for the production of artificial datasets.

2 Problems of real-world datasets

In the following, we review a variety of issues related to the practice of evaluating DNNs on popular real-world datasets for tasks like VQA. We emphasize that our arguments are mainly based on

large-scale and broad-coverage datasets obtained in a relatively unconstrained way. Some of the points do not (fully) apply to more specific and carefully obtained data, like the NLVR dataset.

2.1 Issues with crowdsourced real-world data

The fact that DNNs require immense amounts of data for successful training led to the practice of adopting online data, such as the Flickr photo sharing platform, and leveraging crowdsourcing, usually via Amazon Mechanical Turk (AMT). For instance, the image captioning dataset MS COCO (Lin et al., 2014) contains more than 300,000 images annotated with more than 2 million human-written captions, while the popular VQA Dataset (Antol et al., 2015) is based on MS COCO.

Data obtained this way tends to be comparatively simple in terms of syntax and compositional semantics, despite exhibiting a high degree of lexical complexity due to its real-world breadth. Moreover, ‘re-purposed’ photos do not – and were never intended to – reflect the visual complexity of every-day scenarios (Pinto et al., 2008). Humans given the task of captioning such images will mostly produce descriptions which are syntactically simple. The way that workers on crowdsourcing platforms are paid gives them an incentive to come up with captions quickly, and hence further increases the tendency to simplicity. Note also that, while this is a form of real-world data, it has very little relationship to the way that a human language learner perceives the world, from the fact that image/question pairs are presented in no meaningful order to the impossibility of any kind of interaction with a particular scene.

Natural language follows Zipf’s law for many aspects (sentence length, syntactic complexity, word usage, etc), and consequently has an inbuilt simplicity bias when considered in terms of probability mass. The contents of image datasets based on photos also have a Zipfian distribution, but with biases which relate to what people choose to photograph rather than to what they see. Animal images in the VQA Dataset are predominantly cats and dogs, sport images mainly baseball and tennis (see Antol et al. (2015) for more statistics). Considering all these biases both in language and vision, the common evaluation measure – simple accuracy of questions answered correctly – is not a good reflection of a system’s general ability to understand visually grounded language.

2.2 The Clever Hans effect

Crowdsourced visual questions have other unexpected properties. Goyal et al. (2017) and Mahendru et al. (2017) note how questions rarely talk about objects that are not present in the image, hence an existential question like “*Do you see a...?*” is mostly true. Agrawal et al. (2016) also give the example of questions like “*What covers the ground?*”, which can confidently be answered with “*snow*” because of biases in common real-world scenes, or, more precisely, biases in the photographs of real-world scenes. Such biases help to explain why some *text-only* systems turn out to perform well on *visual* question answering when evaluated on the VQA Dataset.

Sturm (2014) compared such unexpected cues when evaluation machine learning systems to the story of ‘Clever Hans’, a horse exhibited in the early 20th century which was claimed to understand German and have extensive arithmetical and reasoning abilities. Hans was eventually found to be picking up on very subtle cues which were given completely unconsciously by his owner and which were not noticed by ordinary observers. Some of the recent findings for DNNs, particularly in NLP, suggest similarly problematic conclusions, like the surprisingly strong performance of a bag-of-words model for sequential information (Adi et al., 2017) or of text-only systems for multimodal information (Jabri et al., 2016).

A more fundamental form of this effect is illustrated by recent investigations in image recognition. Szegedy et al. (2014) and Nguyen et al. (2015) have shown surprisingly odd system behavior when confronted with either only minimally modified images or almost random noise. This behavior seems due to the specific interplay of a few parameters which dominate the model’s decision, and have led to an entire research subfield on adversarial instances in vision. Such investigations are not yet as prominent in the NLP community, although see, e.g., Jia and Liang (2017), Sproat and Jaitly (2016) and Arthur et al. (2016).

The ability to work with raw input data and to pick up correlations/biases, which humans cannot always manifest in explicit symbolic rules, is precisely the strength of DNNs as feature extractors. But given the often millions of parameters and large number of unstructured input values, it is difficult to avoid unexpected hidden cues. Real-world data with its enormous ‘sample space’,

which is necessarily only sparsely reflected, is hence particularly prone to this effect.

The immediate problem is that a system trained this way may not generalize appropriately to other situations. The longer-term problem is that, while we do not expect that DNNs will simulate human capabilities in a fine-grained way, there has to be some degree of comparability if they are ever to be capable of justifying or explaining their behavior. The ‘Clever Hans effect’ thus refers to situations where we wrongly and prematurely attribute such human-like reasoning mechanisms to trained models, when more careful and systematic investigations would have revealed our misjudgement.

2.3 Guiding principles for DNN evaluation

Compositionality is a fundamental aspect of language, and consequently a necessary prerequisite for any claim about ‘*understanding*’ natural language. Besides being required for proper generalization to novel utterances, it constitutes a far more efficient way of learning in the form of a structural prior, it leads to more interpretable inference results by forcing more systematic processing, and it results in more robust behavior, promising to reduce vulnerability to adversarial examples. However, Lake and Baroni (2017) recently gave reason to doubt the compositional capabilities of recurrent DNNs, which are at the heart of virtually all state-of-the-art NLP models. We conclude from this that a different kind of test data than existing benchmarks is required for more conclusive evaluations, and propose three simple principled ways to reduce the risk of encountering such problems:

- Avoid solely evaluating a system on a single and supposedly representative set, but design test instances with the aim of specifically investigating and confirming the system’s intended improvement over other models.
- Instead of keeping training and test data distributions similar, focus on the true compositional generalization abilities required by dissimilar distributions. A more asymmetric dataset represents a harder, but hence potentially more interesting task.
- Do at least some experiments with clean data, which reduces the likelihood of hidden biases or correlations compared to more ‘realistic’ and complex data. For instance, the relationship between image and text should be explicitly controlled in multimodal data.

A pentagon is above a green ellipse, and no blue shape is an ellipse.

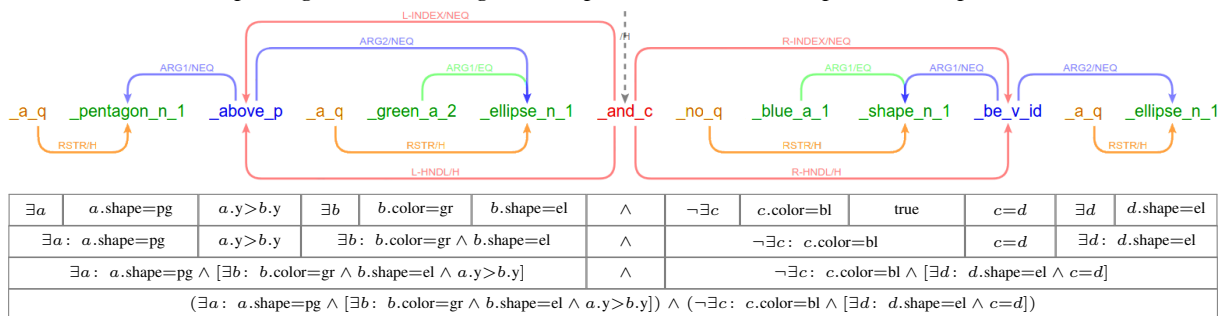


Figure 1: A caption with DMRS graph and semantic interpretation, illustrating how compositionality enables us to generate combinatorially large amounts of non-trivial captions and infer their semantics from atomic elements.

3 Automatic generation of language data

In the following, we describe our approach for automatic generation of artificial VQA data using existing deep linguistic processing technology, based on our implementation in the ShapeWorld framework (Kuhnle and Copestake, 2017)¹. We argue that a compositional semantic approach using a bidirectional grammar gives us precisely the sort of data as outlined by the above principles. We propose this approach as a complementary evaluation step, since it is not intended to replace real-world evaluation, but instead aims to cover aspects which existing datasets cannot provide.

3.1 Abstract microworlds

The generation process we use is based on randomly sampled abstract world models, i.e. values which specify a microworld, entities and all their attributes. In the case of our framework these include the number of entities, their shape and color, position, rotation, shade, etc. Such a world model can be visualized straightforwardly.

In this context, datasets are generators which can create an unlimited amount of data instances, hence making multiple iterations over a fixed set of training instances obsolete. Importantly, different datasets constrain the general sampling process in different ways by, for instance, restricting the number of objects, the attribute values available, the positioning of entities, and more. This addresses the point of specifying different data distributions for training and testing. Moreover, it makes it possible to partition evaluation data as desired, which facilitates the detailed investigation of system behavior for specific instances and hence the discovery of systematic shortcomings.

¹<https://github.com/AlexKuhnle/ShapeWorld>

3.2 Syntactically rich language generation


Of the recent abstract datasets mentioned in the introduction, Suhr et al. (2017) use human-written captions, the SHAPES dataset (Andreas et al., 2016) a minimalist grammar, and the CLEVR dataset (Johnson et al., 2017a) a more complex one based on functional building blocks, both template-based and specifically designed for their data. For our approach we leverage technology made available by the DELPH-IN (Deep Linguistic Processing with HPSG) consortium. More specifically, we make use of the broad-coverage, bidirectional², high-precision English Resource Grammar (Flickinger, 2000), which builds on the compositional semantic framework of Minimal Recursion Semantics (Copestake et al., 2005). For our system we use one of its variant, Dependency MRS (DMRS, Copestake (2009), Copestake et al. (2016)), and generate natural language sentences from abstract DMRS graphs using Packard’s parser-generator ACE³.

We have found that DMRS graphs can easily be enriched with appropriate semantics to be evaluated on a given world model. This means that the internals of the language system are essentially using a form of model-theoretic semantics. However, the external presentation of our task is still ‘natural’, i.e. only consists of image and language. Compositional representations like DMRS further enable us to produce an infinite number of captions of arbitrary syntactic complexity.

Figure 1 shows an example of a non-trivial caption with corresponding DMRS graph and logical representation over a world model. Both the

²Bidirectional grammars can be used for generation as well as parsing, of which the latter might be useful here, for instance, in investigating ambiguity effects.

³<http://sweaglesw.org/linguistics/ace/>



- Less than one triangle is cyan.
- At least half the triangles are red.
- More than a third of the shapes are cyan squares.
- Exactly all the five squares are red.
- More than one of the seven cyan shapes is a cyan square.
- Twice as many red shapes as yellow shapes are circles.

Figure 2: An image with several example captions focusing on quantification. The task is image caption agreement, that is, to decide whether the caption agrees with the image (green) or not (red), similar to yes/no questions in VQA.

abstractness and compositionality of the semantic representation are essential to allow us to scale beyond toy examples. The abstract scenario puts an emphasis on experiments with closed-class vocabulary and syntax, as compared to open-class dominated real-world datasets. However, the same approach can be extended to more complex domains, like the clip-arts of Zitnick et al. (2016).

In the future, we plan to implement two interesting extensions for our framework: First, paraphrase rules can be expressed on grammar-level and integrated into the generation process as post-processing step for increased linguistic variety. Second, (D)MRS-based grammars for other languages, such as the JACY grammar for Japanese (Siegel et al., 2016), can be used simply by translating the internal mapping of atomic DMRS components to corresponding semantic elements.

3.3 Quantification example

Figure 2 presents an example image accompanied by a variety of correct and incorrect captions focusing on quantification. We produce both count-based (“three”) and fraction-based (generalized) quantifiers (“half”) in various modifications (“less than three”), optionally with additional number restriction (“at least three of the five”) or comparative (“half as many as”).

We decide to focus on quantifiers here because, on the one hand, they can exhibit a high degree of structural complexity, which can only be resolved by using visual information. On the other hand, categories like ‘number’ in VQA 2.0 or ‘count’ and ‘compare integer’ in CLEVR imply that count-based quantification is specifically covered by these datasets. As the various captions in figure 2 illustrate, this is not fully the case. Note that we so far do not consider scope ambiguity of nested quantifiers, although our approach can be extended accordingly, since the (D)MRS formalism supports scope underspecification, which is one of the reasons for choosing DMRS.

4 Conclusion: Why use artificial data?

Challenging test data. The interplay of abstract world model and semantic language representation enables us to generate captions requiring non-trivial multimodal reasoning. In fact, the resulting captions can be more complex than the sort of captions we could plausibly obtain from humans, and do not suffer from a Zipfian tendency to simplicity on average (unless configured accordingly).

Avoid Clever Hans effect. The simple, abstract domain and the controlled generation process based on randomly sampling microworlds makes such data comparatively unbiased and greatly reduces the possibility of hidden complex correlations. We can be confident that we cover the data space both relatively uniformly and more exhaustively than this is the case in real-world datasets.

Flexibility & reusability. Real-world and/or human-created data essentially has to be obtained again for every change/update, like for VQA v2.0 (Goyal et al., 2017). In contrast to that, modularity and detailed configurability make our approach easily reusable for a wide range of potentially unforeseen changes in evaluation focus.

Rich evaluation. Ultimately, our goal in providing datasets is to enable detailed evaluations (of DNNs). By creating atomic test datasets specifically evaluating instance types individually (e.g., counting, spatial relations, or even more fine-grained), we can unit-test a DNN for specific sub-tasks. We believe that such a modular approach is a better way to establish trust in the understanding abilities of DNNs than a monolithic dataset and a single accuracy number to assess performance.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. AK is grateful for being supported by a Qualcomm Research Studentship and an EPSRC Doctoral Training Studentship.

References

- Yossi Adi, Einat Kermary, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *Proceedings of the International Conference on Learning Representations*, ICLR 2017.
- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. Analyzing the behavior of visual question answering models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP 2016, pages 1955–1960.
- Aishwarya Agrawal, Aniruddha Kembhavi, Dhruv Batra, and Devi Parikh. 2017. C-VQA: A compositional split of the visual question answering (VQA) v1.0 dataset. *CoRR*, abs/1704.08243.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, CVPR 2016.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, ICCV 2015.
- Philip Arthur, Graham Neubig, and Satoshi Nakamura. 2016. Incorporating discrete translation lexicons into neural machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP 2016.
- Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47(1):253–279.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym.
- Donna Byron, Alexander Koller, Jon Oberlander, Laura Stoia, and Kristina Striegnitz. 2007. Generating instructions in virtual environments (GIVE): A challenge and an evaluation testbed for NLG. In *Proceedings of the Workshop on Shared Tasks and Comparative Evaluation in Natural Language Generation*.
- Ann Copestake. 2009. Slacker semantics: Why superficiality, dependency and avoidance of commitment can be the right way to go. In *Proceedings of the 12th Meeting of the European Chapter of the Association for Computational Linguistics*, EACL 2009, pages 1–9.
- Ann Copestake, Guy Emerson, Michael W. Goodman, Matic Horvat, Alexander Kuhnle, and Ewa Muszyńska. 2016. Resources for building applications with Dependency Minimal Recursion Semantics. In *Proceedings of the 10th International Conference on Language Resources and Evaluation*, LREC 2016, pages 1240–1247.
- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A. Sag. 2005. Minimal Recursion Semantics: An introduction. *Research on Language and Computation*, 3(4):281–332.
- Dan Flickinger. 2000. On building a more efficient grammar by exploiting types. *Natural Language Engineering*, 6(1):15–28.
- 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, CVPR 2017, pages 6325–6334.
- Micah Hodosh and Julia Hockenmaier. 2016. Focused evaluation for image description with binary forced-choice tasks. In *Proceedings of the 5th Workshop on Vision and Language*, pages 19–28.
- Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. 2017. Learning to reason: End-to-end module networks for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, ICCV 2017.
- Allan Jabri, Armand Joulin, and Laurens van der Maaten. 2016. Revisiting visual question answering baselines. In *Proceedings of the 14th European Conference on Computer Vision*, ECCV 2016, pages 727–739.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP 2017, pages 2021–2031.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. 2017a. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, CVPR 2017.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017b. Inferring and executing programs for visual reasoning. In *Proceedings of the IEEE International Conference on Computer Vision*, ICCV 2017.
- Douwe Kiela, Luana Bulat, Anita L. Vero, and Stephen Clark. 2016. Virtual embodiment: A scalable long-term strategy for artificial intelligence research. *CoRR*, abs/1610.07432.
- Alexander Kuhnle and Ann Copestake. 2017. Shapeworld - A new test methodology for multimodal language understanding. *CoRR*, abs/1704.04517.

- Brenden M. Lake and Marco Baroni. 2017. Still not systematic after all these years: On the compositional skills of sequence-to-sequence recurrent networks. *CoRR*, abs/1711.00350.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *Proceedings of the European Conference on Computer Vision, ECCV 2014*.
- Aroma Mahendru, Viraj Prabhu, Akrit Mohapatra, Dhruv Batra, and Stefan Lee. 2017. The promise of premise: Harnessing question premises in visual question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP 2017*.
- Tomas Mikolov, Armand Joulin, and Marco Baroni. 2015. A roadmap towards machine intelligence. *CoRR*, abs/1511.08130.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015*.
- Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. 2018. FiLM: Visual reasoning with a general conditioning layer. In *AAAI*.
- Nicolas Pinto, David D. Cox, and James J. DiCarlo. 2008. Why is real-world visual object recognition hard? *PLOS Computational Biology*, 4(1):1–6.
- Adam Santoro, David Raposo, David G. T. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy P. Lillicrap. 2017. A simple neural network module for relational reasoning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems, NIPS 2017*, pages 4974–4983.
- Konrad Scheffler and Steve Young. 2001. Corpus-based dialogue simulation for automatic strategy learning and evaluation. In *Proceedings of the NAACL Workshop on Adaptation in Dialogue Systems*, pages 64–70.
- Ravi Shekhar, Sandro Pezzelle, Yauhen Klimovich, Aurélie Herbelot, Moin Nabi, Enver Sangineto, and Raffaella Bernardi. 2017. FOIL it! Find one mismatch between image and language caption. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, pages 255–265.
- Melanie Siegel, Emily M. Bender, and Francis Bond. 2016. *Jacy: An Implemented Grammar of Japanese*. CSLI Studies in Computational Linguistics. CSLI Publications, Stanford, California, USA.
- Richard Sproat and Navdeep Jaitly. 2016. RNN approaches to text normalization: A challenge. *CoRR*, abs/1611.00068.
- Bob L. Sturm. 2014. A simple method to determine if a music information retrieval system is a “horse”. *IEEE Transactions on Multimedia*, 16(6):1636–1644.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. *CoRR*, abs/1312.6199.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. Towards AI-complete question answering: A set of prerequisite toy tasks. *CoRR*, abs/1502.05698.
- C. Lawrence Zitnick, Ramakrishna Vedantam, and Devi Parikh. 2016. Adopting abstract images for semantic scene understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(4):627–638.