

# Decoupled Box Proposal and Featurization with Ultrafine-Grained Semantic Labels Improve Image Captioning and Visual Question Answering

Soravit Changpinyo   Bo Pang   Piyush Sharma   Radu Soricut

Google AI

Venice, CA 90291

{schangpi, bopang, piyushsharma, rsoricut}@google.com

## Abstract

Object detection plays an important role in current solutions to vision and language tasks like image captioning and visual question answering. However, popular models like Faster R-CNN rely on a costly process of annotating ground-truths for *both* the bounding boxes *and* their corresponding semantic labels, making it less amenable as a primitive task for transfer learning. In this paper, we examine the effect of decoupling box proposal and featurization for down-stream tasks. The key insight is that this allows us to leverage a large amount of labeled annotations that were previously unavailable for standard object detection benchmarks. Empirically, we demonstrate that this leads to effective transfer learning and improved image captioning and visual question answering models, as measured on publicly-available benchmarks.

## 1 Introduction

Object detection has been employed extensively as a primitive task for vision and language tasks such as image captioning and visual question answering (VQA); see (Anderson et al., 2018) and the work that follows it. One motivation is that the ability to recognize salient regions and objects may be too difficult to learn from weakly-supervised *top-down* signals, in the form of captions and question-answer pairs. Indeed, *bottom-up* signals provided by object detection often correspond to semantic units of language such as words or phrases, making them suitable for text generation and image-text alignment.

However, object detection itself can be broken down into multiple subtasks (Liu et al., 2018). A family of “two-stage” object detectors first proposes category-agnostic bounding box candidates and then featurizes and classifies the cropped regions into one of the available semantic labels.



Q: How much money is this?

ERCNN (VG) A: 1 dollar

Ultra A: 20

Q: What is this?

A: beer

A: bbq sauce

Figure 1: Ultrafine-grained semantic labels (at “instance level”) provide transfer learning power to down-stream tasks like visual question answering.

Even “one-stage” object detection approaches, where these boxes become category-specific, can be formulated in a *bottom-up* manner as detecting and grouping extreme and center points (Zhou et al., 2019). Can we take advantage of this observation to learn to transfer more effectively? In this work, we take a step in this direction by examining the effect of decoupling box proposal and featurization *on downstream vision and language tasks*. In particular, we consider a two-stage object detector and set a goal of pushing the “featurization” aspect of the task further than before.

Our choice to break free from “featurization by object detection models” has at least two advantages. First, there is a larger amount of labeled data that can be leveraged to train a better featurization module, even if such data do not support learning box proposals. To put it another way, the quality of features directly provided by object detectors is limited by the fact that annotating ground-truths for *both* the bounding boxes *and* their corresponding semantic labels is costly and scales poorly. By separating them, we reintroduce the freedom to annotate for object-agnostic box segmentation, without the burden of baking in an-

notation decisions related to the granularity level of the semantic labels (i.e., do we use as semantic labels “money”, “euro”, or “20 euro”?). As illustrated in Figure 1, the granularity level of the semantic labels plays a crucial role for downstream tasks such as VQA.

Second, this approach is better suited to downstream tasks whose domains are different from the one the object detector is trained on. In other words, it allows us to benefit from transfer learning, which is a great advantage given the relatively modest amount of available supervised data for these downstream vision and language tasks.

We empirically demonstrate the above-mentioned advantages through a focused study of the effect of improved featurization on image captioning and VQA in transfer learning settings. In particular, we (i) leverage ultra-fine-grained semantic labels (e.g., “golden gate bridge” vs. “bridge”) for featurization (Juan et al., 2019); and, (ii) focus on scenarios in which object detection modules trained on Visual Genome (VG) (Krishna et al., 2017) are applied to out-of-domain images: image captioning on the Conceptual Captions dataset (Sharma et al., 2018), and VQA on the VizWiz dataset (Gurari et al., 2018). Our results indicate that there are ways to incorporate low-level pre-training tasks that benefit vision and language models via higher-quality bottom-up signals.

## 2 Related Work

Attention-based deep models are popular in image captioning and VQA. Early work used fixed partitions of images as candidate regions (Xu et al., 2015). However, variable sized regions that are better correlated with object boundaries have gained momentum (Fu et al., 2017; Pedersoli et al., 2017; Anderson et al., 2018). Indeed, Anderson et al. (2018) established new state-of-the-art performance over both image captioning and VQA tasks on the MSCOCO and VQA2 benchmarks using a Faster R-CNN detector trained on Visual Genome. As both Visual Genome and VQA2 were built on images from MSCOCO, the object detector was applied largely to in-domain images. In contrast, our work focuses on more realistic settings in which domains of different tasks may not be perfectly aligned (Chen et al., 2018).

We leverage image representations extracted from a network pre-trained over large amounts of

labeled data. Prior work demonstrated the power of pre-training with image classification at scale (Sun et al., 2017; Mahajan et al., 2018; Wu et al., 2019). However, we consider downstream vision *and language* tasks (image captioning and visual question answering), in contrast to less complex vision-only tasks explored in such work: object detection and in some cases semantic segmentation and human pose estimation. Furthermore, our transfer learning technique is based on decoupled region proposal and ultra-finegrained featurization, not fine-tuning the pre-trained network.

Another set of closely related work utilized additional data for scaling up either vision tasks (Hoffman et al., 2016; Tang et al., 2017; Redmon and Farhadi, 2017) or vision and language tasks (Venugopalan et al., 2017; Lu et al., 2018; Noh et al., 2019). For instance, YOLO9000 (Redmon and Farhadi, 2017) built a “WordTree” hierarchy based on the WordNet synsets (Miller et al., 1990), mapped categories in both COCO object detection and ImageNet classification datasets into the hierarchy, and proposed a joint detection and classification training framework. Our approach to transfer learning with ultrafine-grained featurization can similarly address the long-tail nature of target vocabulary (see Figure 2) while being simpler (e.g., not require carefully merging different sets of vocabulary as in YOLO9000). The number of classes we consider is also several orders of magnitude larger.

Incorporating object detection signals in downstream tasks appropriately is non-trivial and an active subject for research (Santoro et al., 2017; Zhang et al., 2018). In this work, we ask the orthogonal question of whether it is necessary to accept the object detector’s output as-is.

## 3 Features and Experimental Setup

Our starting point is a two-stage object detector, which consists of two core modules. One is responsible for category-agnostic box proposal, and the other for featurizing each cropped region for semantic label prediction. In this paper, we select Faster R-CNN (Ren et al., 2015b), a widely-used object detector in image captioning and VQA.

**Faster R-CNN Model** We reimplement the Faster R-CNN model, training it to predict both 1,600 object and 400 attribute labels in Visual Genome (Krishna et al., 2017), following the standard setting from Anderson et al. (2018). ResNet-

101 (He et al., 2016) pre-trained on ImageNet (Russakovsky et al., 2015) is used as the core featurization network<sup>1</sup>. We achieve a mAP@50 of 10.96 for object detection and 1.5 for attribute detection. Given an image, Faster R-CNN proposes  $K$  bounding box regions, each of which comes with a  $D$ -dimensional feature vector as well as object/attribute class predictions (along with their scores).  $K$  is set to 100 and  $D$  to 2048 in our experiments. Using output features on the task of VQA and our model described in Section 5, we obtain an accuracy of 66.9% on the validation set of the VQA2 dataset (Goyal et al., 2017). For comparison, this number already surpasses all validation accuracy numbers in Table 2 for a strong model by Peng et al. (2019), suggesting that our Faster R-CNN features are of high-quality.

**Decoupled Box Proposal and Featurization with Ultra-finegrained Semantic Labels** In standard use of object detectors following Anderson et al. (2018), downstream tasks receive “knowledge” merely about a few thousand classes and four hundred attributes. Here, we exploit the fact that box proposal and featurization can be decoupled, and work on improving the object representation (featurization).

More concretely, we conduct a study toward understanding the utility of improved featurization on downstream tasks. To this end, we exploit a graph-based, semi-supervised representation learning approach called Graph-Regularized Image Semantic Embedding (Graph-RISE) (Juan et al., 2019). Specifically, Graph-RISE is based on ResNet-101 where the  $10 \times 10 \times 2K$  feature map is first average pooled to  $4 \times 4 \times 2K$ , and then flattened and projected to a 64-dimensional embedding before the softmax layer. Learned from  $\mathcal{O}(260M)$  web images and  $\mathcal{O}(40M)$  (noisy) semantic labels, these compact 64-dimensional feature vectors are trained to capture a whole spectrum of semantic similarity, ranging from coarse-grained / category-level (e.g., “bridge”), fine-grained level (e.g., “steel red bridge”), to ultrafine-grained / instance-level (e.g., “golden gate bridge”).

**Our Objective** The main goal is to compare two approaches in using bottom-up signals: 1) FRCNN: use the default visual features from the Faster R-CNN detector; 2) Ultra: use bounding boxes from the Faster R-CNN detector, then featurize them using the much more compact representation from Graph-RISE that potentially reflects differentiation of ultrafine-grained semantic labels. Next, we evaluate this setup on downstream tasks for image captioning and visual question answering.

turize them using the much more compact representation from Graph-RISE that potentially reflects differentiation of ultrafine-grained semantic labels. Next, we evaluate this setup on downstream tasks for image captioning and visual question answering.

## 4 Image Captioning

**Dataset** We use the Conceptual Captions (CC) dataset (Sharma et al., 2018), consisting of 3.3 million training and 15,000 validation images/caption pairs. Another 12,000 image/caption pairs comprise the hidden test set. Official scores on the test set are obtained by submitting models to the CC Challenge server<sup>2</sup>. Unlike other image captioning datasets, images from CC are pulled from across the web and thus exhibit a wide variety of both images and image-caption styles. Most notably, the domain of images can be very different from Visual Genome, unlike in popular benchmarks such as MSCOCO (Lin et al., 2014).

**Model** We adopt the encoder-decoder model from (Sharma et al., 2018), whose basic building block is a Transformer Network (Vaswani et al., 2017). To convert multi-modal inputs to a sequence of encoder feature vectors, we use up to three types of image features:

- G** : Global features by Graph-RISE, a dense 64D vector extracted from the whole image;
- B** : Box-region features by Faster R-CNN (FRCNN, sparse 2048D), or Graph-RISE (Ultra, dense 64D), extracted from each cropped image region resized to  $224 \times 224$  (cf. Sec. 3);
- L** : Label embeddings, obtained by embedding predicted object semantic labels from Google Cloud Vision APIs<sup>3</sup> into a 512D feature vector. These semantic labels are then mapped to embeddings using an embedding layer pre-trained to predict label co-occurrences in web documents using a word2vec model (Mikolov et al., 2013).

For both **B** and **L**, we select the inputs with highest scores and order the sequence inputs based on such scores from high to low. Additionally for **B**, we remove box regions whose scores are lower than 0.001. We use beam search with width 5 for the decoder in all of our experiments<sup>4</sup>.

<sup>2</sup>ai.google.com/research/ConceptualCaptions

<sup>3</sup>cloud.google.com/vision

<sup>4</sup>See further details in the supplementary material.

<sup>1</sup>See further details in the supplementary material.

	dev	test		
	CIDEr	CIDEr	ROUGE-L	SPICE
Transf-Baseline	-	0.772	0.244	0.172
TTI-BIC (single)	-	0.980	<b>0.266</b>	0.186
<b>G</b> (Base)	0.868	-	-	-
<b>B</b> -FRCNN	0.667	-	-	-
<b>B</b> -Ultra	0.873	-	-	-
<b>L</b>	0.606	-	-	-
<b>G</b> + <b>B</b> -FRCNN	0.871	-	-	-
<b>G</b> + <b>B</b> -Ultra	0.912	-	-	-
<b>G</b> + <b>L</b>	0.888	-	-	-
<b>G</b> + <b>B</b> -FRCNN + <b>L</b>	0.892	0.944	0.261	0.190
<b>G</b> + <b>B</b> -Ultra + <b>L</b>	<b>0.937</b>	<b>0.984</b>	<b>0.265</b>	<b>0.195</b>

Table 1: Automatic metric scores for the image captioning task on Conceptual Captions. Ablation results are reported for our model using different sets of visual features. The top two baselines are from the Conceptual Captions Leaderboard as of August 30, 2019.

**Metrics** We adopt the standard automatic metrics for image captioning: CIDEr (Vedantam et al., 2015), ROUGE-L (Lin and Och, 2004), and SPICE (Anderson et al., 2016), as implemented in the COCO-caption evaluation toolkit<sup>5</sup>.

**Results** We report results on both the dev and test sets for Conceptual Captions in Table 1. “Base” uses the **G** feature only. We first compare the Base **G** against each of the feature types (**B**-FRCNN, **B**-Ultra, and **L**). We then perform ablations under the +**B** condition (FRCNN/Ultra) to the Base **G** or stronger **G** + **L** models.

According to dev CIDEr scores, global or box Graph-RISE features **G** and **B**-Ultra are (individually) clearly stronger than box features by Faster R-CNN **B**-FRCNN or label embeddings **L** features. Nevertheless, these features are considerably complementary. Specifically, box features **B**-Ultra complements the Base **G**, pushing the score from 0.868 to 0.912. It is also worth noting that, albeit their low individual scores, **B**-FRCNN or **L** improves upon each model they are added to.

Our models with Ultra features clearly outperform the ones with FRCNN. This is demonstrated in three conditions: when they are on their own, when they are added to the simple **G** model, and when they are added to the stronger **G** + **L** model. Manual inspection of the models’ predictions further supports this; a qualitative comparison of **B**-Ultra vs. **B**-FRCNN in Figure 2 suggests that ultra-finegrained featurization leads to an improved correspondence between visual inputs and caption tokens of *unfamiliar* objects (such as “monks” and “staircase”).

<sup>5</sup><https://github.com/tylin/coco-caption>.

To get test scores, we submit our best model using FRCNN and our best model using Ultra (based on dev CIDEr) to the CC Challenge server. Test scores for other models were not obtained due to the limited number of submissions per time period. As of August 30, 2019, the **G** + **B**-Ultra + **L** model outperforms all other single baselines<sup>6</sup>, for both CIDEr and SPICE (and tie on ROUGE-L).

## 5 Visual Question Answering

**Dataset** We use the recently-proposed VizWiz dataset (Gurari et al., 2018), in which both images and questions originate from visually-impaired or blind people. It consists of 20,000/3,173 (image, question, answers) triplets in the train/val splits, and additional 8,000 triplets for the test split. Each question is independently annotated with 10 answers. We choose the VizWiz benchmark specifically because it is a more suitable benchmark for measuring transfer learning effects. Other VQA datasets, including VQA1.0 (Antol et al., 2015), VQA2.0 (Goyal et al., 2017), Visual7W (Zhu et al., 2016), COCOQA (Ren et al., 2015a), and GQA (Hudson and Manning, 2019) are completely or partly based on MSCOCO or Visual Genome. As such, they may not provide unbiased grounds for measuring the impact of object-detection features based on Visual Genome versus alternative featurization techniques.

**Model** We follow the setting described in Pythia v0.1 (Jiang et al., 2018), the winning entry to the VQA challenge 2018. In particular, the architecture is a *simplified* “up-down” model from (Anderson et al., 2018)<sup>7</sup>. The featurization of the bounding boxes follows the description from Section 4. For the base condition, we use the box features based on Faster R-CNN (**B**-FRCNN), following the majority of previous work. For the test condition, we replace them with the Ultra-based features (**B**-Ultra).

**Metrics** As commonly done in previous work (Antol et al., 2015), we use as our accuracy metric the average score over 9 subsets of the ground-truth 10 answers, where each score is computed by the formula:  $\min(\# \text{ humans that provided that answer} / 3, 1)$ . Accuracy on the test-dev and test-standard splits is obtained by submitting the mod-

<sup>6</sup>[ai.google.com/research/ConceptualCaptions/leaderboard](https://ai.google.com/research/ConceptualCaptions/leaderboard)

<sup>7</sup>See further details in the supplementary material.






Ground-truth  
Box FRCNN (VG)  
Box Ultra

“monks clean a garden at a temple .”  
“a woman walks through the streets .”  
“monks walking in front of a temple”



“black sesame seeds on a white background”  
“a pile of dried flowers”  
“black chia seeds on a white background”



Ground-truth  
Box FRCNN (VG)  
Box Ultra

“a photo of a staircase inside a historic house”  
“the interior of the church”  
“the staircase of the house”



“car & tree ornament this heart of mine”  
“digital art selected for the #”  
“christmas tree in a toy car”

Figure 2: Qualitative results from our image captioning models using **B-FRCNN** vs. **B-Ultra** (see text for details), along with ground-truth captions. Ultra is more capable than FRCNN of dealing with images with unfamiliar objects, those that do not perfectly fall into the domain where the Faster R-CNN object detector is trained on.

	all	y/n	num	unans	other
VizWiz	46.9	59.6	21.0	80.5	27.3
BAN	51.6	<b>68.1</b>	17.9	<b>85.3</b>	31.5
Ours (FRCNN)	51.9	66.7	24.3	85.0	32.1
Ours (Ultra)	<b>53.7</b>	<b>68.1</b>	<b>28.8</b>	84.0	<b>35.4</b>

Table 2: Accuracy (%) on the test-standard split for the VQA task on the VizWiz dataset. Additionally, we provide accuracy per answer type: yes/no (y/n), number (num), unanswerable (unans), and the rest (other). The baselines include VizWiz (Gurari et al., 2018) and BAN (Kim et al., 2018).

els to the VizWiz Challenge server<sup>8</sup>.

**Results** We report results on the VizWiz benchmark in Table 2. Our model with FRCNN provides a strong baseline, slightly outperforming the previous-best model, BAN (Kim et al., 2018), a different architecture that also uses the FRCNN-based features for object bounding boxes. The model using Ultra features further improves upon this; at 53.7%, it outperforms the one using FRCNN by a significant margin (1.8% accuracy on “all” question types). Moreover, this 1.8% improvement is a weighted average across answer types; the per-answer-type numbers indicate that

our approach achieves even better improvements on two of the more difficult answer types, “number” (+4.5%) and “rest” (+3.3%). These improvements are illustrated by the examples provided in Figure 1.

This illustrates the effectiveness of decoupling bounding box proposal and featurization, and quantifies the impact of using transfer learning via large amounts of training data and ultrafine-grained semantic labels used for object representations.

## 6 Conclusion

In this work, we propose to (re)decouple box proposal and featurization. We show that this allows us to leverage additional signals and annotations, leading to more effective transfer learning for downstream vision and language tasks: image captioning and visual question answering. This result suggests that large-scale datasets with fine-grained image-level semantic labels, even when they do not dissect complex visual scenes, can benefit current state-of-the-art models – especially when applied to benchmarks where images are from diverse domains.

<sup>8</sup>evalai.cloudcv.org/web/challenges/challenge-page/102/overview

## References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: semantic propositional image caption evaluation. In *ECCV*.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of CVPR*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual question answering. In *Proceedings of ICCV*.
- Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. 2018. Domain adaptive Faster R-CNN for object detection in the wild. In *Proceedings of CVPR*.
- Kun Fu, Junqi Jin, Runpeng Cui, Fei Sha, and Changshui Zhang. 2017. Aligning where to see and what to tell: Image captioning with region-based attention and scene-specific contexts. *TPAMI*, 39(12):2321–2334.
- 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 CVPR*.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. VizWiz Grand Challenge: Answering visual questions from blind people. In *Proceedings of CVPR*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of CVPR*.
- Judy Hoffman, Deepak Pathak, Eric Tzeng, Jonathan Long, Sergio Guadarrama, Trevor Darrell, and Kate Saenko. 2016. Large scale visual recognition through adaptation using joint representation and multiple instance learning. *JMLR*, 17(1):4954–4984.
- Drew A. Hudson and Christopher D. Manning. 2019. GQA: a new dataset for compositional question answering over real-world images. In *Proceedings of CVPR*.
- Yu Jiang, Vivek Natarajan, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2018. Pythia v0.1: the winning entry to the VQA Challenge 2018. *arXiv preprint arXiv:1807.09956*.
- Da-Cheng Juan, Chun-Ta Lu, Zhen Li, Futang Peng, Aleksei Timofeev, Yi-Ting Chen, Yaxi Gao, Tom Duerig, Andrew Tomkins, and Sujith Ravi. 2019. Graph-RISE: Graph-regularized image semantic embedding. *arXiv preprint arXiv:1902.10814*.
- Jin-Hwa Kim, Yongseok Choi, Sungeun Hong, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear attention networks for VizWiz challenge. In *Proceedings of the ECCV Workshop on VizWiz Grand challenge*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael Bernstein, and Li Fei-Fei. 2017. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123(1):32–73.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of ACL*.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: common objects in context. In *Proceedings of ECCV*.
- Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. 2018. Deep learning for generic object detection: A survey. *arXiv preprint arXiv:1809.02165*.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2018. Neural baby talk. In *Proceedings of CVPR*.
- Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. 2018. Exploring the limits of weakly supervised pretraining. In *Proceedings of ECCV*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of NeurIPS*.
- George A. Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine Miller. 1990. Introduction to WordNet: An on-line lexical database. *International Journal of Lexicography*, 3(4):235–244.
- Hyeonwoo Noh, Taehoon Kim, Jonghwan Mun, and Bohyung Han. 2019. Transfer learning via unsupervised task discovery for visual question answering. In *Proceedings of CVPR*.
- Marco Pedersoli, Thomas Lucas, Cordelia Schmid, and Jakob Verbeek. 2017. Areas of attention for image captioning. In *Proceedings of ICCV*.
- Gao Peng, Hongsheng Li, Haoxuan You, Zhengkai Jiang, Pan Lu, Steven Hoi, and Xiaogang Wang. 2019. Dynamic fusion with intra-and inter-modality attention flow for visual question answering. In *Proceedings of CVPR*.

- Joseph Redmon and Ali Farhadi. 2017. YOLO9000: Better, faster, stronger. In *Proceedings of CVPR*.
- Mengye Ren, Ryan Kiros, and Richard Zemel. 2015a. Exploring models and data for image question answering. In *Proceedings of NeurIPS*.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015b. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Proceedings of NeurIPS*.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet large scale visual recognition challenge. *IJCV*, 115(3):211–252.
- Adam Santoro, David Raposo, David G. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. In *Proceedings of NeurIPS*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual Captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of ACL*.
- Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of ICCV*.
- Yuxing Tang, Josiah Wang, Xiaofang Wang, Boyang Gao, Emmanuel Dellandréa, Robert Gaizauskas, and Liming Chen. 2017. Visual and semantic knowledge transfer for large scale semi-supervised object detection. *TPAMI*, 40(12):3045–3058.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of NeurIPS*.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based image description evaluation. In *Proceedings of CVPR*.
- Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2017. Captioning images with diverse objects. In *Proceedings of CVPR*.
- Baoyuan Wu, Weidong Chen, Yanbo Fan, Yong Zhang, Jinlong Hou, Junzhou Huang, Wei Liu, and Tong Zhang. 2019. Tencent ML-Images: A large-scale multi-label image database for visual representation learning. *arXiv preprint arXiv:1901.01703*.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of ICML*.
- Yan Zhang, Jonathon Hare, and Adam Prügél-Bennett. 2018. Learning to count objects in natural images for visual question answering. In *Proceedings of ICLR*.
- Xingyi Zhou, Jiacheng Zhuo, and Philipp Krähenbühl. 2019. Bottom-up object detection by grouping extreme and center points. In *Proceedings of CVPR*.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7W: Grounded question answering in images. In *Proceedings of CVPR*.