



HAL
open science

Egoshots, an ego-vision life-logging dataset and semantic fidelity metric to evaluate diversity in image captioning models

Pranav Agarwal, Alejandro Betancourt, Vana Panagiotou, Natalia Díaz-Rodríguez

► To cite this version:

Pranav Agarwal, Alejandro Betancourt, Vana Panagiotou, Natalia Díaz-Rodríguez. Egoshots, an ego-vision life-logging dataset and semantic fidelity metric to evaluate diversity in image captioning models. ICLR 2020 - 8th International Conference on Learning Representations, Apr 2020, Addis Ababa / Virtual, Ethiopia. hal-02864865

HAL Id: hal-02864865

<https://hal.science/hal-02864865v1>

Submitted on 11 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Egoshots, AN EGO-VISION LIFE-LOGGING DATASET AND SEMANTIC FIDELITY METRIC TO EVALUATE DIVERSITY IN IMAGE CAPTIONING MODELS

Pranav Agarwal
INRIA RITS, Paris, France
pranav2109@hotmail.com

Alejandro Betancourt
LANDING AI,
Colombia

Vana Panagiotou
Signal Ocean,
Greece

Natalia Díaz-Rodríguez*
U2IS, ENSTA, Institut Polytechnique Paris and INRIA Flowers, France
natalia.diaz@ensta-paris.fr

ABSTRACT

Image captioning models have been able to generate grammatically correct and human understandable sentences. However most of the captions convey limited information as the model used is trained on datasets that do not caption all possible objects existing in everyday life. Due to this lack of prior information most of the captions are biased to only a few objects present in the scene, hence limiting their usage in daily life. In this paper, we attempt to show the biased nature of the currently existing image captioning models and present a new image captioning dataset, *Egoshots*, consisting of 978 real life images with no captions. We further exploit the state of the art pre-trained image captioning and object recognition networks to annotate our images and show the limitations of existing works. Furthermore, in order to evaluate the quality of the generated captions, we propose a new image captioning metric, object based *Semantic Fidelity* (SF). Existing image captioning metrics can evaluate a caption only in the presence of their corresponding annotations; however, SF allows evaluating captions generated for images without annotations, making it highly useful for real life generated captions.

1 INTRODUCTION

Humans have a great ability to comprehend any new scene captured by their eyes. With the recent advancement of deep learning, the same ability has been shared with machines. This ability to describe any image in the form of a sentence is also well known as *image captioning* and has been at the forefront of research for both computer vision and natural language processing (Vinyals et al., 2014; Karpathy & Feifei, 2014; Venugopalan et al., 2016; Selvaraju et al., 2019; Vedantam et al., 2017). The combination of Convolutional Neural Networks (CNN) and Recurrent Neural Network (RNN) has played a major role in achieving close to human like performance, where the former maps the high dimensional image to efficient low dimensional features, and the latter use this low dimensional features to generate captions. However, most of the image captioning models have mostly been trained on MSCOCO (Lin et al., 2014) or Pascal-VOC (Everingham et al.), which consists of 80 and 20 object classes respectively. All the images are captioned taking into consideration only these classes. Thus, even though current models have been successful in generating grammatically correct sentences, they still give a poor interpretation of a scene because of the lack of knowledge about various other kinds of objects present in the world, along with those seen in the dataset. *Egoshots* dataset¹ has a wide variety of images ranging from both indoor to outdoor scenes and takes into consideration diverse situations encountered in real life which are hardly found in

*Work partially done during Díaz-Rodríguez and Panagiotou’s internship at Philips Research, and Betancourt’s stay at Eindhoven University of Technology, Netherlands.

¹*Egoshots* dataset available at <https://github.com/NataliaDiaz/Egoshots>




Figure 1: The overview of our approach for captioning the Egoshots dataset and further evaluating the caption using our proposed Semantic Fidelity (SF) metric. The image to be captioned is pre-processed and passed through a pre-trained image captioning model. The caption generated is processed to output only nouns. The same image is also input to an Object Detector (OD), which detects all the object classes in the given image. Both the noun and object classes detected are then used to compute the SF for a generated caption.

the MSCOCO or Pascal-VOC dataset. Hence our dataset can act as a benchmark to evaluate the performance and robustness of image captioning models on real-life images.

Image captioning has a wide range of applications from supporting a visually impaired person to the recommendation system. Guiding a visually challenged person (Gurari et al., 2018) is a highly sensitive task and a small error can lead to catastrophic accidents. A recent work (Weiss et al., 2019) has shown the ability to successfully complete this task in a simulated environment using reinforcement learning. Image captioning models can also play a major role in further improving these systems. However, in order to make these systems more reliable, the network trained should be able to predict considerably more descriptive captions taking into consideration all the objects present in the scene. In this work, along with releasing the captioned Egoshots dataset, we propose a metric of semantic fidelity to the actual scene being rendered. For images with complex settings, we show that the current image captioning networks lack robustness and are biased towards a few annotated object classes and the kind of controlled scenes present only in training datasets, hence they are not reliable and cannot be deployed for real-life applications without significant finetuning.

Current metrics are limited to evaluate the quality only in the presence of labels rendering them useless for captions generated for real-life scenes. We aim at tackling both these problems by proposing the Egoshots dataset and SF metric. Egoshots consists of 978 real-life ego-vision images captioned using state of the art image captioning models, and aims at evaluating the robustness, diversity, and sensitivity of these models, as well as providing an on-the-wild life-logging dataset that can aid the task of evaluating real life settings. Images were randomly taken by the *Autographer* camera worn by 2 female computer science interns (aged 25 and 29) for 1 month each in Eindhoven (Netherlands), doing regular activities (biking, office working, socializing...) during May-Jul 2015. We further propose a new metric SF so that captions generated by pre-trained models can be evaluated both on their relevance to the image being captioned, and the number of different objects the caption takes into consideration.

2 ANNOTATION PIPELINE AND SEMANTIC FIDELITY METRIC

Annotation Pipeline: Transfer learning has proven to be an effective approach to perform a known task in a new environment efficiently (Tan et al., 2018). The idea is to use a pre-trained model on a new dataset, in order to prevent learning features from scratch, which requires excessive computational resources and time. To caption Egoshots dataset we follow this methodology and use pre-trained weights of state of the art image captioning models without any finetuning. We restrict our work to three models, namely *Show Attend And Tell (SAT)*, *nocaps: novel object captioning at*

scale (NOC), and *Decoupled Novel Object Captioner* (DNOC), as they were able to achieve the best results for the real-life Egoshots dataset. Since pre-trained neural networks are strictly constrained to the size of input images on which they were trained, we use the exact pre-processing for images used in each of the models. Once processed, the images are mapped to their corresponding captions using the learnt models. Among the captions predicted, the one having maximum SF is used as the final caption for the dataset. Section A.1.3 further describes in detail each of the image captioning models.

Object based Semantic Fidelity Metric: To render caption-less datasets such as Egoshots useful, the aim is to map each image to the richest caption, i.e., the one covering as many objects as relevant in the scene, i.e., a caption as descriptive and detailed as possible. With current image captioning metrics, this task renders challenging, since all existing metrics evaluate the quality of generated captions using labeled captions from a dataset collected in well controlled and clean conditions (very different for real life first person vision images). In order to counterbalance this assumption not present in Egoshots, we propose a new image captioning metric called *Semantic Fidelity*. SF takes into consideration two elements: 1) the semantic closeness of the generated caption to the objects detected in the image, and 2) the object diversity with respect to the amount of object instances detected. Assuming a state-of-the-art quasi perfect object detector, by taking into account the semantic closeness among these two sets of (captioned and detected) entities (i.e. objects), we penalize the model when it predicts a caption containing objects completely different from the objects present in the scene.

For each caption generated by a captioning model, we eliminate all words except nouns. This is done as a simplification, assuming nouns convey the largest information on the number of different objects present (Wang et al., 2018). Let $C = \{c_1, \dots, c_m\}$ be a list of captions generated by the pretrained networks, and $W = \{w_1, \dots, w_k\}$ the generated words for a given caption c_i (further processed to keep only nouns for the given caption i $N_i = \{n_1, \dots, n_z\}$). In order to predict all objects present in the image we use a state of the art object detector (OD)². The objects predicted by the OD are represented as $O_{OD} = \{o_1, \dots, o_y\}$. Thus, for every image we obtain a list of noun words from its predicted caption, and a list of objects detected by the OD. A similarity metric among these word sets³ is calculated. The semantic similarity among these two word sets takes into account their semantic closeness using word embeddings. Recent works (Mikolov et al., 2013; Conneau et al., 2017) show the ability of word embeddings that is transforming a word into its vector form efficiently capture the semantic closeness of two given words. The SF metric uses this approach to calculate such semantic similarity between the noun words and objects in an image, for each caption, described as $S = \{s_1, \dots, s_m\}$. The cosine similarity as such does not take into account the diverse nature of the predicted captions in terms of the number of different objects present, since the final embedding calculated for each of the list averages out multiple objects of the same class as a single entity. Hence, to further penalise the quality of captions for its object count we compute the ratio of the number of noun words in a given caption to the number of objects predicted. The SF metric score thus measures the quality of a caption, as a proxy of the diversity of knowledge the network has regarding both the presence and diversity of different objects present in the image:

$$SF_i = s_i \cdot \frac{\#N}{\#O} \tag{1}$$

where, for image i , s_i is the semantic similarity among noun words in its predicted caption c_i and object nouns detected by the OD, $\#O$ is the cardinal of O_{OD} , and $\#N$ the number of nouns (representing objects in N_i) present in c_i . SF ranges in $[0, 1]$: captions having an SF closer to 1 convey more information and are semantically closer to the scene being captioned, in terms of the objects involved in the caption.

For a predicted caption, information about an average of 2-3 objects is generally conveyed (as shown in Fig. 2a), while for the OD we assume an ideal condition in which the OD acts as an oracle and predicts all different objects in the given image, so that $\#O \geq \#N$ (Assumption 1) for all images. Thus, the larger the number of noun entities in the caption, the more the ratio will approach 1,

²We tested several different ODs and settled with YOLO9000 (or Y9) in this particular showcase of SF metric because of its ability to detect the largest (9000) number of classes, while other detectors are restricted to either 80 or 20 classes only, as shown in Fig. 2b.

³Such as the cosine similarity (as we use here) of the mean of the embeddings of the words in each of these two sets.

and the closer SF will be to 1 for the best captions. This approach to compute SF will work only assuming robust object detectors satisfying enough scene annotation granularity. For the SF values to be reliable, the OD needs to detect correct object classes as present in the image. With completely different objects detected with respect to those in the image, the similarity metric for image i , s_i , will be inaccurate, and therefore, SF remains unreliable. In order for SF to be applicable, we will also assume Assumption 2: $\#O \neq 0$ (i.e., the OD can at least detect one object in the image). In scenarios when the object detector predicts less objects than nouns in the sentence (Assumption 1 broken), we skip the object diversity ratio and use $SF = s_i$ instead.

3 DISCUSSION AND CONCLUSION

Table 1 compares the SF performance on different image captioning methods using various object detectors (OD). To make SF reliable, we assume an OD to predict all the relevant different classes given in each image. Except Y9, all other ODs are trained only on either 80 or 20 classes, leading to larger SF values for weak detectors, as they find it difficult to penalize an inefficient caption given their lack of knowledge about various different object classes. Also ODs trained on 20 classes of VOC tend to show larger SF than those trained on 80 classes of COCO because of their inability to detect new classes not seen in the training dataset⁴. Better OD models will make SF more reliable, as is also reflected for the captions generated in Table 2. Table 2 compares the performance of each of the pre-trained image captioning models on the SF metric. OD Y9 for image 3 in Table 2 predicted a single incorrect class leading to inaccurate SF value due to using a similarity metric among wrong terms. However, for most of the other images, Y9 is able to predict the correct objects; therefore SF is able to penalize the captions which do not take into consideration all the objects present in a given scene. Hence, a well generalized and robust object detection model plays the most important role if the evaluation of captions is performed using SF.

Most images are captioned using 8-10 words only. Even if longer sentences may not always be preferred, critical applications, e.g., supporting the visually impaired or autonomous driving, where the agent does not have any prior information, the richer the interpretation of a scene, the more reliable the system will be.

We presented the *Egoshots* dataset, consisting of 978 real life first person vision images of everyday activities. Although the number of images in the dataset is far less than those in MSCOCO or Pascal VOC, our aim, along with presenting the dataset, is to analyze the performance of the pre-existing image captioning models and their reliability. To do so we propose the first image captioning metric (SF) that allows to evaluate captions from unlabelled images. Since previously existing metrics are limited to caption-labelled datasets, there is no way to analyze the quality of captions of real life images without costly annotated labels. We show on pre-trained models that, despite being able to successfully generate grammatically correct sentences, their captions are often misaligned with the objects present in the scene or hallucinate objects (Rohrbach et al., 2018). This is due to the presence of training bias towards objects imposed by the datasets these networks were trained with: they are unable to convey complete or rich information about the scene. The aim of *Egoshots* dataset and SF metric is thus to facilitate assessment, diversity and deployability of image captioning deep models.

Despite the major improvements in recent image captioning models, we show that there is room to achieve more semantically faithful and relevant caption generation systems; e.g. models could be affected by a significant amount of blurring, making object detection in the wild a lot harder. Thus, there are plenty of room for future work on OD and IC models on real life-logging situations, not only for assisted technology or telepresence, but for any autonomous system. Since the SF score’s aim is to measure the semantic quality of a caption, and it cannot evaluate the grammatical correctness of a sentence, future work should better assess both detector and captioning models quality to better map the desirable properties of such models at a finer grained domain-specific resolution. Such studies should include larger scale assessments with a positive control group (i.e., human annotations of real images on the wild).

⁴Or classes not always visually *observable*. An example of such words is, e.g. the class *entrepreneur* from YOLO9000.

REFERENCES

- Harsh Agrawal, Karan Desai, Yufei Wang, Rishabh Jain, Xinlei Chen, Dhruv Batra, Mark S Johnson, Devi Parikh, Peter Anderson, and Stefan Lee. nocaps: novel object captioning at scale. *International Conference on Computer Vision*, abs/1812.08658, 2019.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6077–6086, 2017.
- Satanjeev Banerjee and Alon Lavie. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *IJEEvaluation@ACL*, 2005.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word Translation Without Parallel Data. *ArXiv*, abs/1710.04087, 2017.
- Berkan Demirel, Ramazan Gokberk Cinbis, and Nazli Ikizler-Cinbis. Image Captioning with Unseen Objects. *ArXiv*, abs/1908.00047, 2019.
- Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Trevor Darrell, and Kate Saenko. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.
- Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. CenterNet: Keypoint Triplets for Object Detection. *ArXiv*, abs/1904.08189, 2019.
- M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. URL <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>.
- Chenyou Fan and David J. Crandall. DeepDiary: Automatic Caption Generation for Lifelogging Image Streams. *ArXiv*, abs/1608.03819, 2016.
- Yang Feng, Lin Ma, Wei Liu, and Jiebo Luo. Unsupervised Image Captioning. In *CVPR*, 2018.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. VizWiz Grand Challenge: Answering Visual Questions from Blind People. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3608–3617, 2018. URL <https://arxiv.org/abs/1802.08218>.
- Andrej Karpathy and Li Feifei. Deep Visual-Semantic Alignments for Generating Image Descriptions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3128–3137, 2014.
- Gal Sadeh Kenigsfeld and Ran El-Yaniv. Leveraging Auxiliary Text for Deep Recognition of Unseen Visual Relationships. *ArXiv*, abs/1910.12324, 2019.
- Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper R. R. Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and Vittorio Ferrari. The Open Images Dataset V4: Unified image classification, object detection, and visual relationship detection at scale. *ArXiv*, abs/1811.00982, 2018. URL <https://arxiv.org/abs/1811.00982>.
- Linghui Li, Sheng Tang, Lixi Deng, Yongdong Zhang, and Qi Tian. Image Caption with Global-Local Attention. In *AAAI*, 2017.
- Zhijun Liang, Yi-Sheng Guan, and Juan Rojas. Visual-Semantic Graph Attention Network for Human-Object Interaction Detection. 2020.
- Chin-Yew Lin. ROUGE: A Package For Automatic Evaluation Of Summaries. In *ACL 2004*, 2004.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common Objects in Context. In *ECCV*, 2014. URL <http://arxiv.org/abs/1405.0312>.

- Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: Single Shot MultiBox Detector. In *ECCV*, 2016.
- Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing When to Look: Adaptive Attention via a Visual Sentinel for Image Captioning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3242–3250, 2016.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural Baby Talk. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7219–7228, 2018.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*, 2013.
- Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. EGO-TOPO: Environment Affordances from Egocentric Video. *ArXiv*, abs/2001.04583, 2020.
- Silvio Olivastri, Gurkirt Singh, and Fabio Cuzzolin. End-to-End Video Captioning. 2019.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a Method for Automatic Evaluation of Machine Translation. In *ACL*, 2001.
- Joseph Paul Cohen, Genevieve Boucher, Craig A. Glastonbury, Henry Z. Lo, and Yoshua Bengio. Count-ception: Counting by Fully Convolutional Redundant Counting. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2017.
- Joseph Redmon and Ali Farhadi. YOLO9000: Better, Faster, Stronger. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6517–6525, 2016.
- Joseph Redmon and Ali Farhadi. YOLOv3: An Incremental Improvement. *ArXiv*, abs/1804.02767, 2018.
- Zhou Ren, Xiaoyu Wang, Ning Zhang, Xutao Lv, and Li-Jia Li. Deep Reinforcement Learning-Based Image Captioning with Embedding Reward. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1151–1159, 2017.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object Hallucination in Image Captioning. *CoRR*, abs/1809.02156, 2018. URL <http://arxiv.org/abs/1809.02156>.
- Ramprasaath R. Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Dhruv Batra, and Devi Parikh. Taking a HINT: Leveraging Explanations to Make Vision and Language Models More Grounded. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2591–2600, 2019.
- Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A Survey on Deep Transfer Learning. *ArXiv*, abs/1808.01974, 2018.
- Satoshi Tsutsui, Dian Zhi, Md Alimoor Reza, David J. Crandall, and Chen Yu. Active Object Manipulation Facilitates Visual Object Learning: An Egocentric Vision Study. *ArXiv*, abs/1906.01415, 2019.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based image description evaluation. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4566–4575, 2014.
- Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, and Gal Chechik. Context-Aware Captions from Context-Agnostic Supervision. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1070–1079, 2017.
- Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, Raymond J. Mooney, Trevor Darrell, and Kate Saenko. Captioning Images with Diverse Objects. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1170–1178, 2016.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3156–3164, 2014.

- Josiah Wang, Pranava Swaroop Madhyastha, and Lucia Specia. Object Counts! Bringing Explicit Detections Back into Image Captioning. In *NAACL-HLT*, 2018.
- Martin Weiss, Simón Chamorro, Roger Girgis, Margaux Luck, Samira Ebrahimi Kahou, Joseph Paul Cohen, Derek Nowrouzezahrai, Doina Precup, Florian Golemo, and Chris Pal. Navigation Agents for the Visually Impaired: A Sidewalk Simulator and Experiments. *ArXiv*, abs/1910.13249, 2019.
- Yuehua Wu, Linchao Zhu, Lu Jiang, and Yi Yang. Decoupled Novel Object Captioner. In *ACM Multimedia*, 2018.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *ICML*, 2015.
- Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image Captioning with Semantic Attention. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4651–4659, 2016.
- Peter Young, Alice Lai, Micah Hodosh, and Julia 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.
- Zehua Zhang, Ashish Tawari, Sujitha Martin, and David J. Crandall. Interaction Graphs for Object Importance Estimation in On-road Driving Videos. *ArXiv*, abs/2003.06045, 2020.

A APPENDIX

A.1 RELATED WORK

A.1.1 IMAGE CAPTIONING

The problem of scene understanding and image captioning has been vastly explored in the literature (Zhang et al., 2020; Nagarajan et al., 2020; Tsutsui et al., 2019; Liang et al., 2020; Kenigsfeld & El-Yaniv, 2019). Most of the work follows a sequence learning approach using an encoder and decoder (Vinyals et al., 2014; Olivastrri et al., 2019; Donahue et al., 2015; Fan & Crandall, 2016). The encoder consists of several stacked CNNs to achieve efficient latent representations of images, while the decoder has recurrent layers to map these latent vectors to a caption. In order to make the captions more diverse, the attention mechanism integrated into the encoder-decoder framework (Zhang et al., 2020; You et al., 2016; Li et al., 2017; Lu et al., 2018; Anderson et al., 2017; Selvaraju et al., 2019; Lu et al., 2016), which takes into consideration long-range dependencies. Attention gives importance to different spatial regions of the image by weighing them differently. The decoder then generates each word by taking into consideration the relative importance of each spatial region. Most approaches follow this methodology and train the models using datasets such as MS-COCO (Lin et al., 2014), PASCAL-VOC (Everingham et al.) and Flickr 30k (Young et al., 2014) to name a few. These datasets have millions of images with human labelled captions for a predefined number of object categories. COCO dataset has captions for 80 different object classes, while PASCAL-VOC has 20 different classes. Some of the recent works (Venugopalan et al., 2016; Wu et al., 2018; Demirel et al., 2019) have tried to break this limitation by integrating different object classes while training their image captioning model. They predict a placeholder in the generated caption for the object classes not present in the MSCOCO dataset and further replace this placeholder with the object classes. This approach was able to generate more diverse captions, to some extent, than previous models; however, it still lacks the ability to produce captions able to take into consideration all different objects present in a scene.

There have been very few works which have used unsupervised (Feng et al., 2018) or reinforcement learning (Ren et al., 2017) for captioning images. Unsupervised image captioning model use generative adversarial networks in order to generate captions close to those created by humans. They use an external corpus of sentences in order to train the model. The reinforcement learning model uses an actor-critic approach with a reward model for mapping images to their corresponding captions. The actor predicts the confidence of predicting the next word given the image, while the critic takes in the current state and the words predicted, and adjusts the goal to produce captions close to the ground truth.


A.1.2 IMAGE CAPTIONING METRICS

In order to evaluate the quality of the generated captions most of the models use automatic image captioning metrics such as BLEU (Papineni et al., 2001), Meteor (Banerjee & Lavie, 2005), Rouge (Lin, 2004) and CIDEr (Vedantam et al., 2014). These metrics evaluate the quality of the captions based on pre-existing labels. To the best of our knowledge there is no metric which can evaluate the quality of a captions without labels. Thus, testing the performance of an image captioning model on a real-life image not present in the dataset becomes a major limitation. On the other side, human evaluation is extensively time-consuming and not reliable, as it varies from person to person. In this work, we presented the new metric, SF, which allows to evaluate an unlabelled caption (i.e., captions for which its ground truth is not available).


A.1.3 IMAGE CAPTIONING AND OBJECT DETECTION MODELS ASSESSED

This section details the existing state of the art captioning models used to annotate Egoshots dataset and validate the SF metric.


Show Attend and Tell (SAT) (Xu et al., 2015) uses a CNN followed by an LSTM in order to caption the image. CNNs take the input image to extract efficient low dimensional features which are then further decoded by the LSTMs to generate captions in a sequence to sequence manner. In order to predict realistic captions, they integrate the attention module into LSTMs in order to focus on salient objects. Two kind of attention modules are used; namely, soft and hard attention.



(a)



(b)



(c)

Figure 2: a) Number of nouns per image for SAT, NOC, DNOC, and number of object categories for YOLO9000. b) Total number of different object classes present in the Egoshots dataset as predicted by each of the pre trained image captioning and object detection models. c) Caption length per image for SAT, NOC and YOLO9000.

Image Captioning Method	S-V	S-Co	Y3-V	Y3-Co	C-V	C-Co	Y9
<i>Show Attend And Tell</i>	0.35	0.34	0.34	0.33	0.30	0.36	0.28
<i>Novel Object Captioning at Scale</i>	0.40	0.39	0.39	0.37	0.34	0.40	0.33
<i>Decoupled Novel Object Captioner</i>	0.41	0.41	0.40	0.39	0.35	0.44	0.32

Table 1: Mean Semantic Fidelity of different image captioning models using various object detectors: S: SSD (Liu et al., 2016), Y3: YOLOv3 (Redmon & Farhadi, 2018), C: Center Net (Duan et al., 2019), Y9: YOLO9000 trained on ImageNet and COCO, V: trained on VOC, Co: trained on COCO.

Even though attention helps improving captions, at the same time, attention is responsible for less descriptive captions, as with attention the network focuses on only important objects in the image while filtering away a large number of objects.

Novel object captioning at scale (NOC) (Agrawal et al., 2019) tries to tackle the problem of having fewer object classes present in the captions of the COCO dataset by incorporating the Open Image dataset (Kuznetsova et al., 2018) which has 600 classes but still far less than YOLO9000. It tries to disentangle object detection and image captioning and claims to have a greater number of object classes in the generated captions in comparison to the MSCOCO but if we compare the performance on Egoshots dataset with respect to the total number of different object classes used for captioning it still lags behind YOLO9000 as shown in Fig. 2a.

Decoupled Novel Object Captioner (DNOC) (Wu et al., 2018) follows a two-step process for generating sentences. They start with predicting captions with placeholders for every novel object not seen previously in the dataset. In the second step, they use an object memory to replace the placeholder with the correct object word based on the visual features. Pre-trained object detection is used to predict novel objects. DNOC also uses an encoder-decoder architecture with a slight variation in the decoder part. The encoder is a network pre-trained on ImageNet to extract the low dimensional features. The decoder uses LSTMs and the output from the encoder to generate word by word a sentence. They predict novel objects in the captions not seen in the dataset but still, their captions are not descriptive with respect to the number of object classes per image far less than NOC and YOLO9000.

YOLO9000 (Y9 for short) (Redmon & Farhadi, 2016). We use a pre-trained YOLO9000 object detector model in the Egoshots dataset as state of the art OD in order to evaluate the diversity of generated captions. YOLO9000 achieved state of the art results for object detection, and can detect around 9000 different object classes. It integrates both object detection and image classification into a single module (trained on MSCOCO and ImageNet datasets). Since ImageNet has a greater number of object classes than MSCOCO, it detects more objects than those present in MSCOCO dataset. Despite being far from being 100% accurate, it is used as gold standard OD in our experiments to illustrate the use of SF. However, any state of the art object detector can be used/adapted to each domain specific problem, as state of the art OD.

In Fig. 2a we show that most of the captions generated convey information about max. 2-4 objects for a given image. At the same time, Y9 is able to detect 5-7 objects for most of the images, and is also able to predict a maximum of 12 classes for few images. In addition to this, Fig. 2b shows that for the Egoshots dataset, Y9 can detect 370 unique objects, while the captions predicted use approximately half the number of objects. The inability of captions to take into consideration all the objects present in the image is also shown in Fig. 2c.

Baselines implemented are available online⁵.

A.2 VALIDATING SEMANTIC FIDELITY METRIC WITH HUMAN SEMANTIC FIDELITY

SF metric can use different similarity metrics. In our case we use cosine similarity in order to compute each word embedding and then calculate the semantic similarity. We use spaCy NLP toolkit⁶ implementation.

⁵https://github.com/Pranav21091996/Semantic_Fidelity-and-Egoshots

⁶<https://spacy.io/>




Figure 3: Linear fitting test for SF and Human SF (HSF). Pearson correlation test for 100 MSCOCO dataset manually annotated images gives positive correlation with $\rho = 0.93$.

To further validate our SF metric, we perform a linear regression analysis by comparing SF scores with those SF scores provided by a human labeller (Human Semantic Fidelity, HSF). We use 100 MSCOCO dataset images. For SF we use MSCOCO ground truth (GT) captions, while for HSF we use human labelled image captions taken from MSCOCO caption, i.e.: we set HSF to be the number of objects in the caption divided by the number of real (GT) objects observed in image i :



$$HSF_i = \frac{\#N}{\#O_{GT}} \quad (2)$$

For both SF and HSF, a ground truth (GT) human annotated object detector is used to annotate the images, hence assumption 1 ($\#O \geq \#N$) is always valid. Pearson’s correlation ρ parameter (to compare the (co)relation between two independent variables) analysis showed a positive correlation among SF and HSF of 0.93, with a p-value of 1e-44. The coefficient of determination R^2 showed that approximately 76% of the observed variation can be explained by the linear model of the 100 manually labelled datapoints⁷.

A.3 COMPARING PRE-TRAINED IMAGE CAPTIONING MODELS USING SEMANTIC FIDELITY

The validity of the SF metric is subject to the generalization and robustness properties of the OD to new images (i.e., the higher the F-measure of the OD, the more reliable the SF will be).

The following examples illustrate the use of the Semantic Fidelity proposed metric on the Egoshots dataset, evaluating different image captioning models.

1)		YOLO9000	Model	Caption	SF
			SAT	A man is standing in front of a television.	0.31
			NOC	A man in a kitchen with a large mirror.	0.22
			DNOC	A man in a kitchen with a bottle.	0.19
2)		YOLO9000	Model	Caption	SF
			SAT	A man and a woman are riding a bike.	0.36
			NOC	A man is standing on a skateboard in the middle of a street.	0.46
			DNOC	A man and woman sitting on a bicycle.	0.38

⁷Available in the dataset folder in https://github.com/Pranav21091996/Semantic_Fidelity-and-Egoshots

3)		YOLO9000	Model	Caption	SF
		[Cellular Telephone]	SAT	A person riding a skateboard down a street.	0.24
			NOC	A woman is sitting on a bench with her bike.	0.18
			DNOC	A bicycle parked on a sidewalk near a street.	0.21
4)		YOLO9000	Model	Caption	SF
		[Evening bag, Albanian, Physical therapist]	SAT	A woman holding a nintendo wii game controller.	0.44
			NOC	A woman standing in a bathroom holding a wii remote.	0.53
			DNOC	A man in a black shirt and a cell phone.	0.54
5)		YOLO9000	Model	Caption	SF
		[Instigator, Panelist, Old lady, Vintner, Mixing faucet, Waiter]	SAT	A group of people playing a video game.	0.14
			NOC	A man in a kitchen preparing food in a kitchen.	0.39
			DNOC	A group of people standing around a table.	0.13
6)		YOLO9000	Model	Caption	SF
		[Entrepreneur, Wineglass, Vintner, Dinner table]	SAT	A group of people sitting at a table with wine glasses.	0.36
			NOC	A group of people sitting at a table with food.	0.27
			DNOC	A man and woman sitting at a table with food.	0.38
7)		YOLO9000	Model	Caption	SF
		[Rabbi, Fishmonger]	SAT	A woman sitting at a table with a glass of wine.	0.18
			NOC	A man is standing in the middle of a table with a bowl of food..	0.19
			DNOC	A man in a white shirt is holding a bowl.	0.18
8)		YOLO9000	Model	Caption	SF
		[Cereal Bowl, Salesperson, Tourist, Hairdresser]	SAT	a couple of people standing around a table.	0.23
			NOC	A man standing next to a woman in a city.	0.33
			DNOC	a group of people standing around a table.	0.16
9)		YOLO9000	Model	Caption	SF
		[Peddler, Seller, Beer Drinker, Vintner, Dad, Cereal Bowl]	SAT	a group of people standing around a table.	0.1
			NOC	A group of people standing around a table with a large white plate of food..	0.34
			DNOC	a group of people sitting around a table with food.	0.23


	YOLO9000	Model	Caption	SF
10) 	[Entrepreneur, Background, Laptop, Camp Chair, Settler]	SAT	<i>A man sitting at a table with a laptop.</i>	0.42
		NOC	<i>A man in a kitchen with a large display of food..</i>	0.44
		DNOC	<i>A man in a suit and tv standing in front of a tv.</i>	0.62

Table 2: Captions generated by each of the pre-trained models for images from the Egoshots dataset. The SF metric is used in order to evaluate the captions by taking into consideration the objects detected by YOLO9000. As observed in image 3, due to a poor object detector failing to detect all objects correctly, SF penalize the caption only for its incorrect cosine similarity and skips object diversity ($SF=s_i$, Assumption 1 broken).

A.4 METRIC LIMITATIONS

We must note some limitations of the metric, which should be complemented/extended to (1) account for verbs and other syntactic elements of the caption; (2) rate a caption in terms of the quality of the interpretation, taking into account the count of objects of the same type in the image with respect to those present in the caption. Particular models for counting such as (Paul Cohen et al., 2017) is a particular example on how to enhance the label-less dataset annotation pipeline proposed here.

Metrics should be evaluated in more targeted application use cases, e.g. the usefulness of such captions for targeted users such as the blind, in concrete applications such as navigation settings (Weiss et al., 2019).

A.5 ACKNOWLEDGEMENTS

We thank Valerio Iebba, Gemma Roig and David Filliat for practical and insightful feedback. We thank all Eindhoven friends for being part of the dataset.