

LISTEN CAREFULLY AND TELL: AN AUDIO CAPTIONING SYSTEM BASED ON RESIDUAL LEARNING AND GAMMATONE AUDIO REPRESENTATION

Sergi Perez-Castanos, Javier Naranjo-Alcazar¹, Pedro Zuccarello, Maximo Cobos¹

¹ Universitat de València, Burjassot, Spain, {janal2@alumni.uv.es}, {maximo.cobos}@uv.es

ABSTRACT

Automated audio captioning is machine listening task whose goal is to describe an audio using free text. An automated audio captioning system has to be implemented as it accepts an audio as input and outputs as textual description, that is, the caption of the signal. This task can be useful in many applications such as automatic content description or machine-to-machine interaction. In this work, an automatic audio captioning based on residual learning on the encoder phase is proposed. The encoder phase is implemented via different Residual Networks configurations. The decoder phase (create the caption) is run using recurrent layers plus attention mechanism. The audio representation chosen has been Gammatone. Results show that the framework proposed in this work surpass the baseline system in challenge results.

Index Terms— Audio captioning, Residual learning, Attention, Encoder-Decoder, Gammatone

1. INTRODUCTION

Audio captioning is a novel machine listening task that was first presented in [1]. Audio captioning can be understood as a intermodal translation. Its goal is to create an autonomous and smart description on an audio signal. The state-of-the-art solution employs an encoder-decoder architecture [2]. The encoder block embeds the audio representation (i.e. log Mel-Spectrogram representation) into a lower dimensionality feature map while the decoder creates a sequence of words from that new representation, that is, the smart caption. This caption must be as close as possible as human performance which means that captions must be structured according to the language in which it is being described. This task differs from other classic machine listening tasks such as audio tagging or sound event detection. Audio captioning is not intended to assign labels to audio or to calculate onset and offset times.

First approximations to autonomous captioning where done in image domain [3, 4] followed by autonomous video captioning [5, 6]. In [3], this problem was addressed for the first time. Captioning, whether in the image or audio domain, can be interpreted as an artificial intelligence problem that connects two fields. In the case of the image domain, computer vision and language processing techniques must be merged. In audio domain, the representation of the audio, the processing of this representation (similar to computer vision techniques) and the language processing must be taken into account. The intuition on which this work is based is in the work done in sentence translation where a sentence must be translated from an initial language to a target language. The emergence of recurrent neural networks (RNNs) led to the creation of simpler solutions (without reordering or individual processing of each word) and maintained state-of-the-art performance [7, 8]. In this type of

problem, an RNN encodes the input sentence into a fixed-size codification and a RNN decoder generates the translated sentence in the target language. As image captioning takes as input a fixed size image, in [3] it is decided to replace the encoder block by a convolutional neural network (CNNs). This type of network has shown good results in extracting descriptive information about the images. Thus, a network of this nature will be used to encode the images for the decoder, which in this case is a recurrent network.

Autonomous captioning frameworks (regardless the data domain, e.g. image or audio) can be divided in two blocks as mentioned before: encoder and decoder. Going a little more in detail in each of the parts can be stated:

- **Encoder:** it processes the input data e.g. RGB image and creates more sophisticated high level representations from the input. This block is the one that has produced the most different solutions. Some state-of-the-art works propose the encoder to be composed of convolutional layers [3, 5], while others by recurring layers [8, 7, 1].
- **Decoder:** it takes the encoder's output and creates the final caption. This block is usually implemented using a Recurrent Neural Network (RNN). The final layer is usually a fully connected one [8, 7] with the number of units equal to number of possible words [1].

As it can be observed, an autonomous captioning system can be reformulated as a system that, given an entry \mathbf{X} , the system is able to obtain the most relevant characteristics that allow it to title the entry through a series of meaningful words lexicon.

The aim of this work is to propose a solution based on an encoder-decoder structure where the encoder corresponds to a convolutional neural network and the decoder to a recurrent neural network. To achieve more accurate results the decoder implements an attention mechanism. An analysis of different state-of-the-art residual networks as possible encoders is carried out. In addition, the representation of the audio is done using the Gammatone filter bank.

The rest of the paper is organized as follows: Section 2 explains the proposed method in this paper. Section 3 describes the dataset and the training procedure. Section 4 shows the results obtained by all the experiments compared to the baseline. Finally, Section 5 concludes this work.

2. METHOD

The system proposed in this work is based on [3, 9]. The main idea on [9] is that the final step of the captioning system (decoder) learn where to look in the representation in order to predict the next sequence word. This mechanism is known as attention and forces the system to look for the relevant part in the encoding feature map. A




Figure 1: Proposed system representation. Audio is first represented using Gammatone filter bank. Later, a richer representation of the audio is obtained by means of a residual encoder. The final caption is obtained by using a recurrent decoder with the attention mechanism.

full representation of the proposed framework can be seen in Figure 1.

2.1. Audio representation

The input of this framework has been decided to be the Gammatone audio representation [10, 11]. The number of frequency bins is set to 64 as well as the baseline system. All post-processing of the captions is done with the same procedure as the baseline (see subsection 4.1). This work implementation can be found in the following link¹.

2.2. Residual Convolutional Encoder

In this work, the encoder block correspond to a Convolutional Neural Network (CNN). Residual networks have been the choice. In particular, residual CNNs have been the choice. Residual CNNs were first introduced in [12]. Residual layers are designed in order to approximate a residual function as $\mathcal{F}(\mathbf{X}) := \mathcal{H}(\mathbf{X}) - \mathbf{X}$, where \mathcal{H} represents the feature to be fit by a set of stacked layers and \mathbf{X} represents the input to the first of such stacked layers. Therefore, \mathcal{H} can be expressed as $\mathcal{H}(\mathbf{X}) = \mathcal{F}(\mathbf{X}) + \mathbf{X}$. The reason why this kind of CNN have been chosen corresponds to the idea that the network training might be easier if optimizing a residual mapping instead of an unreferenced one, as in a classical CNN [13]. This residual learning can be easily implemented by adding a shortcut connection that would perform as identity mapping, that is, adding the input \mathbf{X} to the output of the residual block $\mathcal{F}(\mathbf{X})$. In this work, different state-of-the-art Residual networks have been implemented as encoder block [12, 14].

The *Resnet50*, *Resnet101* and *Resnet152* residual architectures were proposed in the original paper where the residual learning was first presented [12]. The idea behind them is the same and they are implemented with the same convolutional blocks placed in the same sequential way. That is, they have the same kernel size or number of filters. The difference between them is how many times each convolutional block is repeated. Thus *Resnet50* is a 50-layer residual network, *Resnet101* is a 101-layer network, and *Resnet152* is a 152-layer network.

Since the introduction of this type of network, a multitude of combinations or studies have been proposed to improve the behaviour of the residual network, either in reducing the number of layers or the training time by maintaining or improving the results obtained. One solution that aims to maintain the benefits of residual learning but with less deep networks is the Wide Residual Networks

[14]. In this submission, the *Resnet101* network (*Wide101*) has been tested but implementing the modifications presented in [14].

All implementations have been done with the Deep Learning Pytorch library [15] without changing the original kernel sizes. It is important to point out that the networks are trained from scratch. Although the architectures are common choice in the state of the art and are available pre-trained, in this case it has been decided not to perform any fine-tuning technique that would involve the use of external data to those proposed in the dataset.

2.3. Recurrent decoder

The decoder is implemented via a Recurrent Neural Network (RNN). As explained before, this block task is to look at the encoder’s output and generate the final caption word by word. The decoder is implemented with a LSTM layer. The final layer of the decoder corresponds to a Dense layer of 4637 (4635 words + $\langle \text{eos} \rangle$ + $\langle \text{sos} \rangle$) units that predicts the word output probability according to the LSTM state, the context vector (obtained by the attention mechanism explained above) and the previous word [9, 16]. The context vector is obtained by an attention mechanism that allows the decoder to analyze different parts of the encoder’s output as it has to predict the next word. In practice, this is a smart weighted average across all encoder’s outputs feature maps. The higher weights will indicated relevancy in that features. This averaged feature map is used to generate the next word, this is known as gate mechanism (for more insight see [9]). The attention mechanism takes into account the sequence created at the specific moment before predicting the next word and looks for the next part of the encoder’s output to be more relevant. This is done by a DNN with a softmax activation at the last layer in order to assign the relevancy of each part of the encoder’s output.

2.4. Beam Search

Different approaches are available in order to implement the decoding stage. The simplest decoding approach would be to consider the highest probability word in each iteration as the correct. Although this approach is valid, it is not the optimal decoding process. At each iteration, the best possible word depends on the former selected words but it is possible that other combinations, considering not only the first best word but the k best words, might be a better solution. As a result, k sequences are obtained instead of only one and then, the best one is considered as the decoder output. This idea is called Beam Search. This method is used in the testing subset (see subsection 3.1.)

¹https://github.com/sergipc22/dcase20_task6/tree/develop

System	BLEU ₁	BLEU ₂	BLEU ₃	BLEU ₄	METEOR	ROUGE _L	Cider	SPICE	SPIDEr
Baseline	0.389	0.136	0.055	0.015	0.084	0.262	0.074	0.033	0.054
Enc-50-DecAtt	0.353	0.132	0.054	0.023	0.089	0.248	0.078	0.037	0.058
Enc-101-DecAtt	0.357	0.138	0.059	0.026	0.091	0.249	0.088	0.038	0.063
Enc-152-DecAtt	0.348	0.129	0.054	0.026	0.090	0.246	0.079	0.038	0.058
Enc-Wide101-DecAtt	0.345	0.129	0.054	0.023	0.090	0.245	0.079	0.038	0.058

Table 1: Results with different Residual configurations in development stage. Metrics obtained from evaluation subset. Teaching forcing method is used for this metrics.

In this case, beam width is 3, that is $k = 3$. At first iteration, the 3 best words are considered and at second iteration, each one considers their 3 best options as a second word. At this point, there are 9 possible paths with their corresponding 9 combined scores. The 6 worst paths are pruned, keeping the other 3 at the end of the iteration. This procedure is repeated iteratively until each path reaches the $\langle \text{eos} \rangle$ token. When there are 3 completed sentences, their combined scores are calculated and the sentence with the best one is considered the output of the decoder.

3. EXPERIMENTAL DETAILS

3.1. Clotho dataset

In order to validate automated audio captioning systems the Clotho dataset will be used [2]. This is the first captioning dataset manually labeled using only audio data information. All the information on how it has been labeled and the postprocessing procedure can be found in [17]. This dataset is made up of audios from 15 seconds long to 30 seconds long. Each audio has 5 captions that can vary from 8 to 20 words. There are a total of 4981 audio samples and therefore 24905 captions.

The total number of possible words in the dataset is 4365. In turn, the dataset is divided into 3 parts: development, evaluation and testing. Some of the details of this separation are that only an audio sample can appear in one of the partitions and not in two. There is no word that appears only in a split. The appearance of the words is proportional to the partition percentage: 60 % development (training), 20 % evaluation (validate) and 20 % testing (challenge classification).

During the development phase of the contest, the development and evaluation partitions have been released. The development split consists of 2893 audio clips and 14465 captions. This partition is used to train the system. The evaluation split consists of 1045 audio clips and 5225 captions. This partition is used to validate the training and check its generalization. This block is used to choose the model that will be used to predict once the testing partition is released. All partitions take into account the consideration of the proportional appearance of words in each split as explained before.

3.2. Training procedure

The system is optimized during training using Adam optimizer [18]. The batch size is set to 16 samples and the number of epochs to 300. The loss function is set to cross-entropy loss. The learning rate start with a value of 10^{-4} and before every weight update, the 2-norm of the gradients clipped using as a threshold the value of 2. The systems have only been trained with the development subset. Captions are calculated with teacher forcing method [19].

4. RESULTS

4.1. Baseline system

The starting point of this task consists in an encoder-decoder structure formed by a Recurrent Neural Network (RNN) in both blocks. The encoder corresponds to a RNN of 3 bi-directional GRU layers. All GRU layers have the same number of features, 256. The decoder is implemented with just one GRU layer of 256 features and a classification layer of 4637 (length necessary to represent all possible words in one-hot encoding).

The input is a log-Mel Spectrogram audio representation with 64 Mel filters. The temporal bins are obtained using a 46 ms window length with 50% overlap. Representation are padded with zeros at the beginning when the audio length is less than the maximum possible in the batch.

4.2. Metrics

There are a series of metrics that allow evaluating the performance of an autonomous captioning system. These metrics are: BLEU_n [20], ROUGE_L [21], METEOR [22], CIDER_r [23], SPICE [24] and SPIDEr [25]. BLEU_n calculates precision for n -grams, for that purpose, the number of matchings between the target caption and the predicted caption. BLEU_n penalizes if the predicted caption is shorter than the target. BLEU_n does not take into account the context. The subscript n represent the grams. METEOR is considered a recall-based metric. The harmonic mean of precision and recall of segments of the captions between the predicted and the target ones is computed. It employs alignment between predicted and target words. This alignment is computed over segments of the caption, the number of chunks needed is the least possible. ROUGE_L is a Longest Common Subsequence (LCS) based metric. The F-measure using LCS between the predicted and the target caption is computed. In this case, more weight is given to the recall [1]. CIDER computes a weighted sum of the cosine similarity between both captions (predicted and target). Finally, SPICE is a more sophisticated metric that aims at capturing human judgments over the predicted caption better than the presented metrics. SPICE computation is divided into 3 steps. The first one consists in transforming both predicted and target captions into an intermediate representation that encodes semantic propositional content. This metrics focuses on semantic meaning. After this parsing a F-score calculation is run. Finally, a gameability step is implemented in order to measure how well the predicted caption has recover objects, attributes and the relation between them. Take into account that this metric is an image-based metric. SPIDEr can be understood as a combination of SPICE and CIDER metrics.

System	BLEU ₁	BLEU ₂	BLEU ₃	BLEU ₄	METEOR	ROUGE _L	Cider	SPICE	SPIDeR
Baseline	0.344	0.082	0.023	0.000	0.066	0.234	0.022	0.013	0.018
Enc-50-DecAtt	0.464	0.260	0.157	0.092	0.135	0.308	0.195	0.083	0.139
Enc-101-DecAtt	0.469	0.265	0.162	0.096	0.136	0.310	0.214	0.086	0.150
Enc-152-DecAtt	0.466	0.261	0.156	0.091	0.137	0.310	0.207	0.086	0.147
Enc-Wide101-DecAtt	0.464	0.259	0.154	0.086	0.137	0.310	0.205	0.087	0.146

Table 2: Results with different Residual configurations plus Beam Search method in challenge results. Metrics obtained from testing subset.

4.3. Performance Analysis

Table 1 shows the results obtained for the different implementations submitted to this task. As it can be appreciated, the study module in this task has been the encoder. For this, different state-of-the-art architectures incorporating residual learning have been implemented.

Table 1 presents all the metrics obtained on the evaluation subset. As it can be noticed, not all the metrics show an improvement of the baseline system. If the focus is on the SPIDeR metric, (since it is the one that is going to be used to rank the systems in the task) it can be noticed that the network with *Resnet101* encoder shows a better behavior. The system obtains a value of 0.063, improving the result presented in the baseline. As it can be observed, all other settings show the same value in the SPIDeR metric. As it can be observed, with this subset of the dataset, the behavior of the rest of the implemented encoders is almost the same for the SPICE and CIDEr metrics. This can be due to the fact that the deeper networks to *Resnet101* are suffering from overfitting. On the other hand, *Resnet50* is too tiny to extract relevant features from the input Gammatone spectrogram.

4.4. Evaluation examples

In this section some examples of how the proposed system generates captions from the provided audio are shown. The captions shown correspond to samples of the testing partition of the Clotho dataset.

Encoder	Predicted Caption
Enc-50	a car engine revs up and then drives away
Enc-101	a car engine is revving up and is driving away
Enc-152	a motor is running and people are talking in the background
Enc-Wide101	cars drive by on a busy road as people are talking

Table 3: Predicted captions for test_0002.wav

Encoder	Predicted Caption
Enc-50	people talk to each other while cars drive by
Enc-101	a group of people are talking in the background
Enc-152	a crowd of people are talking in the background and gets louder as time goes on
Enc-Wide101	people talk to each other while the vehicles drive by in the background

Table 4: Predicted captions for test_0014.wav

Encoder	Predicted Caption
Enc-50	water is running from a faucet and is filling up something
Enc-101	water is running and dripping out of a faucet
Enc-152	water is running and a person is eating a shower
Enc-Wide101	water is being poured from a pitcher into a cup and glass again

Table 5: Predicted captions for test_0028.wav

As it can be seen, the captions have an acceptable grammatical structure. Some of the visible differences may be the decision between the use of the present or the present continuous as can be seen in Table 3. Another difference may be the decision when defining a concept as, for example, “a group of people”, “people” or “a crowd of people” as can be seen in Table 4. Table 5 is another example and how the system still generates nonsense “eating a shower”.

4.5. Evaluation discussion

The systems with encoder *Resnet50*, *Resnet101*, *Resnet152* and *Wide101* have obtained a SPIDeR metric value of 0.139, 0.150, 0.147 and 0.146 respectively (see Table 2). As it can be observed it has been a huge improvement to the one shown in the development phase (see Table 1). The best system has ranked in the position 13 of 31 systems. As it can be appreciated, this system (*Resnet101*) shows the best performance in all metrics. If the rest of the encoders are analyzed, it can be seen that *Resnet50* shows a better behavior in the BLUE_n metrics and *Resnet152* and *Wide101* in the rest of the metrics. The systems ranked between the first and eleventh positions have implemented different data augmentation techniques. The system that ranks (without data augmentation) above those presented in this paper, has the difference of using embeddings when representing the words and implementing a CRNN as an encoder. The use of embeddings has also been used by the winning system. It has not been our case since it was not intended to add more degrees of freedom to the system, for example, the Word2Vec model and only analyze the behavior of the encoder-decoder.

5. CONCLUSION

Audio captioning is a very novel task in the field of machine listening. Automated captioning is a problem that has been getting the attention of the image research community for a few years now. Thanks to the recent release of a dataset specially designed for audio captioning and the proposal of this task, the first novel solutions will be proposed. In this work, a state-of-the-art image captioning network (decoder-encoder structure) is implemented in the problem of audio captioning by making a study in the encoder block that is in charge of extracting the information from the audio. It has been decided to change the state-of-the-art representation based on Mel filters and to use the Gammatone filter bank that has shown good results in other machine listening tasks. The results presented in this paper show a considerable improvement of the baseline system in the challenge results, enhancing its performance in the SPIDeR metric. In addition, the systems have shown competitive behaviour in the test phase in accordance with the limitations imposed on their design, such as the non-use of data augmentation techniques or word embeddings.

6. REFERENCES

- [1] K. Drossos, S. Adavanne, and T. Virtanen, “Automated audio captioning with recurrent neural networks,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, New Paltz, New York, U.S.A., Oct. 2017. [Online]. Available: <https://arxiv.org/abs/1706.10006>
- [2] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: An audio captioning dataset,” in *45th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Barcelona, Spain, May 2020. [Online]. Available: <https://arxiv.org/abs/1910.09387>
- [3] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, “Show and tell: A neural image caption generator,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3156–3164.
- [4] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo, “Image captioning with semantic attention,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4651–4659.
- [5] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko, “Translating videos to natural language using deep recurrent neural networks,” *arXiv preprint arXiv:1412.4729*, 2014.
- [6] A. Shin, K. Ohnishi, and T. Harada, “Beyond caption to narrative: Video captioning with multiple sentences,” in *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2016, pp. 3364–3368.
- [7] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [8] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” *arXiv preprint arXiv:1406.1078*, 2014.
- [9] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *International conference on machine learning*, 2015, pp. 2048–2057.
- [10] S. Tabibi, A. Kegel, W. K. Lai, and N. Dillier, “Investigating the use of a gammatone filterbank for a cochlear implant coding strategy,” *Journal of neuroscience methods*, vol. 277, pp. 63–74, 2017.
- [11] Z. Zhang, S. Xu, S. Cao, and S. Zhang, “Deep convolutional neural network with mixup for environmental sound classification,” in *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*. Springer, 2018, pp. 356–367.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [13] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [14] S. Zagoruyko and N. Komodakis, “Wide residual networks,” *arXiv preprint arXiv:1605.07146*, 2016.
- [15] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.
- [16] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, “How to construct deep recurrent neural networks: Proceedings of the second international conference on learning representations (iclr 2014),” in *2nd International Conference on Learning Representations, ICLR 2014*, 2014.
- [17] S. Lipping, K. Drossos, and T. Virtanen, “Crowdsourcing a dataset of audio captions,” in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Nov. 2019. [Online]. Available: <https://arxiv.org/abs/1907.09238>
- [18] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [19] R. J. Williams and D. Zipser, “A learning algorithm for continually running fully recurrent neural networks,” *Neural computation*, vol. 1, no. 2, pp. 270–280, 1989.
- [20] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th annual meeting on association for computational linguistics*. Association for Computational Linguistics, 2002, pp. 311–318.
- [21] C.-Y. Lin, “ROUGE: A package for automatic evaluation of summaries,” in *Text Summarization Branches Out*. Barcelona, Spain: Association for Computational Linguistics, July 2004, pp. 74–81. [Online]. Available: <https://www.aclweb.org/anthology/W04-1013>
- [22] A. Lavie and A. Agarwal, “METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments,” in *Proceedings of the Second Workshop on Statistical Machine Translation*. Prague, Czech Republic: Association for Computational Linguistics, June 2007, pp. 228–231. [Online]. Available: <https://www.aclweb.org/anthology/W07-0734>
- [23] R. Vedantam, C. L. Zitnick, and D. Parikh, “Cider: Consensus-based image description evaluation,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 4566–4575.
- [24] P. Anderson, B. Fernando, M. Johnson, and S. Gould, “Spice: Semantic propositional image caption evaluation,” in *European Conference on Computer Vision*. Springer, 2016, pp. 382–398.
- [25] S. Liu, Z. Zhu, N. Ye, S. Guadarrama, and K. Murphy, “Improved image captioning via policy gradient optimization of spider,” 10 2017, pp. 873–881.