

AUTOMATED AUDIO CAPTIONING WITH RECURRENT NEURAL NETWORKS

*Konstantinos Drossos**, *Sharath Adavanne**, and *Tuomas Virtanen*

Tampere University of Technology, Tampere, Finland

ABSTRACT

We present the first approach to automated audio captioning. We employ an encoder-decoder scheme with an alignment model in between. The input to the encoder is a sequence of log mel-band energies calculated from an audio file, while the output is a sequence of words, i.e. a caption. The encoder is a multi-layered, bi-directional gated recurrent unit (GRU) and the decoder a multi-layered GRU with a classification layer connected to the last GRU of the decoder. The classification layer and the alignment model are fully connected layers with shared weights between timesteps. The proposed method is evaluated using data drawn from a commercial sound effects library, ProSound Effects. The resulting captions were rated through metrics utilized in machine translation and image captioning fields. Results from metrics show that the proposed method can predict words appearing in the original caption, but not always correctly ordered.

Index Terms— audio captioning, recurrent neural networks, RNN, gated recurrent unit, GRU, attention mechanism

1. INTRODUCTION

The automated audio captioning problem can be defined as the task of automatically generating a textual description (i.e. caption) for an audio signal, where the caption is as close as possible to a human-assigned one for the same audio signal. This is different from the sound event detection (SED) and audio tagging [1, 2], because the audio captioning method does not predict sound events and their start and end times (as in SED) nor assigns labels to an audio file or parts of it (as in audio tagging). For example, a method for audio captioning must be able to generate descriptions like “large flock off and flying away”, “footsteps down on wooden stairs, two men slow approach”, and “domestic clock striking five”¹.

*Equally contributing authors in this paper.

The research leading to these results has received funding from the European Research Council under the European Unions H2020 Framework Programme through ERC Grant Agreement 637422 EVERYSOUND.

Part of the computations leading to these results were performed on a TITAN-X GPU donated by NVIDIA. The authors also wish to acknowledge CSC-IT Center for Science, Finland, for computational resources.

The authors would like to thank S. Venugopalan (U.T.A., U.S.A.), S. Mimilakis (Fraunhofer IDMT, Germany), and A. Mesaros (TUT, Finland) for the fruitful conversations and valuable feedback.

¹These captions are from PSE Library [3]

In other domains, the automated captioning task has previously been explored with different types of multimedia content. In particular, automated image captioning [4, 5] can be considered as the first attempt to automatically create descriptions for multimedia data, followed by automated video captioning [6, 7]. All the above-cited works of automated captioning employ a scheme which is also seen in machine translation field [8, 9] and in the zero-shot translation system by Google [10, 11]. In fact, some image captioning works were inspired by or based on existing works in machine translation, e.g. [4]. In this scheme, the methods can be divided into two parts, the encoder and the decoder. The encoder processes the input data and creates a higher level and/or rich representation of them [4]. The decoder takes as input the output of the encoder and generates the final output sequence, i.e. translated version of the input for the machine translation or caption(s) for the automated captioning task. Some works implement the encoder as a multi-layered convolutional neural network (CNN) [4, 6] and others as multi-layered RNN [8, 9]. The decoder is usually implemented as an RNN (multi-layer or single-layer) with a fully connected or maxout [12] layer on top [8, 9, 4].

In this work, we present the first attempt for automated audio captioning. Given an audio file, we seek to generate a textual description that can describe that audio file and is as close as possible to human-assigned captions for the same audio file. To do so, we employ an encoder that processes extracted features from the audio file and feeds it to a decoder, which in turn generates the caption for the audio file word-by-word. The network is trained using audio and captions from a commercial sound effects library, the ProSound effects (PSE) library [3], and we evaluate the produced captions with the metrics employed in Microsoft COCO Caption evaluation [13], namely BLEU [14], ROUGE_L [15], METEOR [16], and CIDEr [17]. The rest of the paper is organized as follows. The proposed method is presented in Section 2. The evaluation procedure, obtained results and discussions are presented in Section 3. The conclusions and the future expansions for the proposed audio captioning task is discussed in Section 4.

2. PROPOSED METHOD

Our proposed method takes as an input an audio file of standard CD quality (i.e. 44.1 kHz sampling frequency and 16 bits

sample width) and creates a textual description (i.e. caption) for it. At first a matrix of features, $\mathbf{X} \in \mathbb{R}^{T \times N_{\text{feats}}}$, is extracted from the audio file. Each row of the matrix \mathbf{X} contains the extracted features from the input audio file in frame t . The matrix \mathbf{X} is used as an input to a neural network, which outputs a matrix $\mathbf{Y} \in \mathbb{R}^{I \times N_{\text{words}}}$. Each row of the output matrix \mathbf{Y} contains the probability distribution over unique words. Then, we pick the most probable word according to each row of the output matrix, creating a sequence of words. The resulting sequence of words is the caption of the input audio file.

For feature extraction, the input audio file is segmented using Hamming window of 46 ms (2048 samples) and 50% overlap. From each resulting frame we extract $N_{\text{feats}} = 64$ log mel-band energies. With the extracted audio features, we create the matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$, where $\mathbf{x}_t \in \mathbb{R}^{64}$ is a vector containing the log mel-band energies in frame t , and T is the number of frames that the input audio data is divided into.

The neural network consists of an encoder, a soft alignment model, and a decoder. The encoder is a three-layered bi-directional GRU with tanh output activations and a residual connection between the second and the third layer. All but the last layer have 64 cells and the last one 128. This results to 128 and 256 cells due to bi-directionality (2×64 and 2×128), respectively. The output \mathbf{h}_t^l of the bi-directional GRU layer l is, according to [9],

$$\mathbf{h}_t^l = [\vec{\mathbf{h}}_t^l; \overleftarrow{\mathbf{h}}_t^l]^T, \quad (1)$$

where $\vec{\mathbf{h}}_t^l$ is the hidden output of the forward GRU layer l , $\overleftarrow{\mathbf{h}}_t^l$ of the backward one. The encoder takes as input the matrix \mathbf{X} and produces as an output the matrix $\mathbf{H}^3 = [\mathbf{h}_1^3, \dots, \mathbf{h}_T^3]$, with $\mathbf{h}_t^3 \in \mathbb{R}^{256}$.

The soft alignment model is a fully-connected layer with softmax activation. It takes as input the output of the encoder (i.e. \mathbf{H}^3) and the hidden recurrent output \mathbf{h}'_{i-1} of the decoder, and produces as an output the i -th input to the decoder, \mathbf{c}_i . The soft alignment model helps the decoder focus on different parts of the output of the encoder (i.e. weight differently each \mathbf{h}_t^3) for his i -th prediction, while incorporating information from the previous states and predictions of the decoder (due to \mathbf{h}'_{i-1}). For that reason, and similar to [9], we interpret the alignment model as a soft attention mechanism and the resulting weighted sum, \mathbf{c}_i , as a context vector. The context vector \mathbf{c}_i is the weighted sum of the output sequence of the encoder and, according to [9], calculated as

$$\mathbf{c}_i = \sum_{t=1}^T \text{FullyConnected}(\mathbf{h}_t^3, \mathbf{h}'_{i-1}) \times \mathbf{h}_t^3, \quad (2)$$

where the *FullyConnected* is the fully-connected layer of the alignment model with softmax activation which outputs a scalar in the range of (0, 1), \mathbf{h}'_{i-1} is the hidden recurrent output of first layer of the decoder for the $i - 1$ prediction,

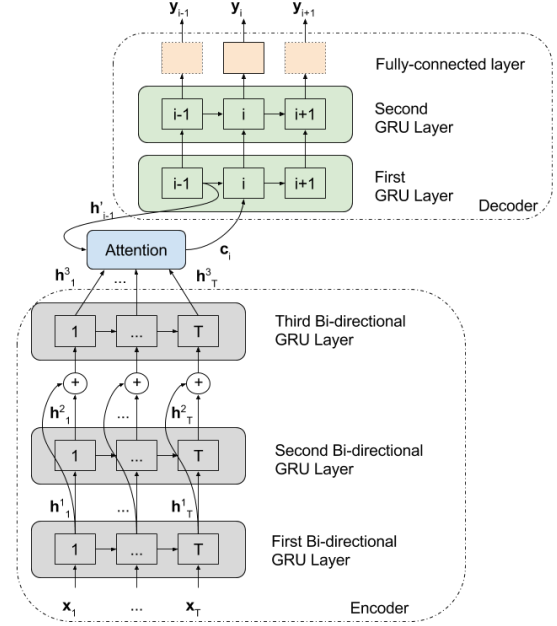


Fig. 1. Illustration of the presented neural network.

and \mathbf{h}_0 is a vector of zeros. The soft alignment model has shared weights between calculations of the different \mathbf{c}_i .

The decoder is a two-layered GRU with tanh activation functions, followed by a fully-connected layer with softmax activation. The first GRU layer of the decoder has 128 cells and the second 256. It takes as input the outputs of the soft alignment model and produces the matrix $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_I]$, where $\mathbf{y}_i \in \mathbb{R}^{N_{\text{words}}}$ is a vector containing the probability distribution over N_{words} unique words. From each \mathbf{y}_i , we select the word with the highest probability. This results at a sequence of I words. This sequence of I words is the predicted caption for the input audio file.

An illustration of the neural network is in Figure 1. The amount of the total parameters of the network is 928 638. The encoder, the attention mechanism, and the decoder are jointly trained using Adam optimizer with the default parameters mentioned in the original paper [18], and categorical-crossentropy loss. Dropout rate is 0.5 and 0.25 for the input and recurrent connections, respectively, in the GRUs of the encoder and the decoder, and the batch size is 64. All the above mentioned hyper-parameters, including the amount of layers for the encoder and decoder, are set after a grid search, using validation data. The initializations of the weights of the GRU inputs and the fully-connected layers are performed using Glorot & Bengio uniform initialization [19]. The weights of the recurrent connections of the GRU layers are initialized using orthogonal initialization. The code is developed using the keras [20] framework with Theano backend [21]. The attention mechanism is implemented by the authors according to [9].

3. EVALUATION

3.1. Dataset pre-processing

In order to evaluate our method, we employ the audio data and their corresponding captions from PSE library [3]. This library consists of 175,000 sound recordings distributed over 501 categories in total, e.g. ambience, audience, foley, nature, wildlife, sports, transportation and weather. The audio recordings in the PSE library are inconsistent in terms of sampling rate, number of channels, and sample width. In order to have consistency in the data we used recordings with sampling rate of at least 44.1 kHz, maximum two channels, and sample width of at least 16 bits. This resulted in 153 220 recordings.

All 153 220 recordings were processed to have a sampling rate of 44.1 kHz, sample width of 16 bits, and single channel. Each of the recordings is associated with one textual description (i.e. caption). The captions in most cases are not proper sentences but rather a set of keywords. Within a caption some keywords repeat more than once and in different forms (e.g. plural and singular). The repeating words in the same form were removed (i.e. from “bird bird” we removed only one word resulting in “bird”). Removal of punctuations was performed in order to reduced the final amount of unique tokens/words that have to be predicted. The words in captions have also a plenty of typographic/orthographic errors. This results in a very high number of unique words. In order to overcome this, we removed all words which are not in the US or UK English dictionary using Enchant python package [22] and GNU-Aspell dictionaries [23]. This reduced the unique words from 19,638 to 11,478 and the total number of words from 1.3 M to 1.16 M. This also resulted in 601 captions with zero words that were excluded from our experiments. This resulted in 152 619 recordings, each having caption length from one to 60 words.

We use a randomly chosen subset of approximately 10% of the 152 619 cleaned library as our data set. This subset is further randomly split to approximately 60% training, 20% validation, and 20% testing data. This corresponds to 8757, 2975 and 2960 recordings, respectively, and amounts to 14,692 recordings. Some captions of this 10% subset are identical between different recordings. In order to ensure that our proposed method does not memorize captions or similar acoustic conditions (implied by the similarity of the captions), we create the data test split with recordings that have captions which do not occur in either training or validation split. In the data splits there might be cases where the captions in the test split contain words not appearing in the training and validation splits. This would result in trying to predict unknown words (i.e. words not appearing in the training split). Given also that in the set of 14,692 recordings 75.36% of words appeared less than five times, the probability of having words in the test split that are not appearing in the training one, increased. We tackle this by creating 1000 random sets of splits

(i.e. training, validation, and testing splits is one set) fulfilling the restriction of not having identical captions in the three splits. From these 1000 sets, we choose the one that had the maximum amount of common words, with the restriction that these words also occurred at least five times in each of the splits. The chosen set of splits contained a total of 71000 words with $N_{\text{words}} = 633$ unique words. Finally, we chose a caption length of maximum 10 words followed by an end of sentence (<EOS>) token as the caption for each of the recordings. This results to $I = 11$. Captions longer than 10 words were trimmed, and shorter were padded with <EOS> token. The length was motivated from the dataset, where approximately 90% of the recordings (13220 out of 14692) had captions equal to or less than 10 words. We utilize only the first thirty seconds of the audio of each recording. Recordings shorter than this were padded with zeros. This results to $T = 1289$ frames of audio.

3.2. Evaluation metrics and procedure

The proposed method was evaluated according to the following metrics used in the machine translation and image captioning fields. These metrics were calculated for 10 different training and testing runs, i.e. for each run the parameters of the neural network were re-initialized according to the initialization functions and the neural network was re-trained on the training split.

BLEU [14] is a precision-based metric. It calculates a weighted geometric mean of a modified precision of n -grams between predicted and ground truth captions. Due to the calculation of the modified precision that favors short predicted captions, BLEU uses a brevity penalty in the calculation of the geometric mean. This penalty penalizes predicted captions that are shorter than the ground truth ones. Typical lengths for n -grams are one to four, resulting in $\text{BLEU}_{1/2/3/4}$, respectively [14, 13]. METEOR [16] calculates a harmonic mean of precision and recall of segments of the captions between the predicted and ground truth captions. The recall is weighted significantly more than precision and thus METEOR is considered a recall-based metric [13]. It employs alignment between the words of the predicted and ground truth captions and matches exact words, stems of words, synonyms, and paraphrases. The alignment is computed over segments of the captions (chunks) between the ground truth and predicted captions while also minimizing a number of chunks needed [16, 13]. CIDEr [17] calculates a weighted sum of the cosine similarity between the predicted and ground truth captions for n -grams of length n with $n \in [1, 4]$. The cosine similarity is calculated using Term Frequency Inverse Document Frequency (TF-IDF) weighting for each n -gram [17, 13]. The CIDEr that is employed to MS COCO Caption evaluation (and is used in this work) is a modified version of CIDEr that is more robust to gaming and is called CIDEr-D. Gaming refers to the fact that sentences which

Table 1. Obtained results for the evaluation metrics

Metric	Case		
	Random words	Random input data	Proposed method
BLEU ₁	0.003 ±0.000	0.006 ±0.001	0.191 ±0.004
BLEU ₂	0.000 ±0.000	0.002 ±0.000	0.129 ±0.003
BLEU ₃	0.000 ±0.000	0.000 ±0.000	0.106 ±0.003
BLEU ₄	0.000 ±0.000	0.000 ±0.000	0.094 ±0.003
ROUGE _L	0.004 ±0.000	0.008 ±0.002	0.149 ±0.002
METEOR	0.004 ±0.000	0.004 ±0.001	0.092 ±0.002
CIDEr	0.005 ±0.001	0.026 ±0.012	0.526 ±0.012

scored high with an automated metric, tend to be evaluated poorly by humans [13]. ROUGE_L [15] is a Longest Common Subsequence (LCS) based metric. It calculates an F-measure using LCS between the predicted and ground truth caption. The F-measure is oriented towards recall using a value for the $\beta = 1.2$ in the F-measure calculation [15, 13].

Since the current paper presents the very first method for audio captioning, there are no previous results to compare the presented ones. For that reason, we employed two additional evaluation cases. The predicted captions for the first case are generated by randomly picking one to ten words from the set of unique words in the dataset. This case will show the values for the metrics obtained by random predictions of words. Given the lack of previous studies on automated audio captioning, we use these values of the metrics as indicative lower threshold for the performance of our proposed method. The captions of the second case are generated from testing the network with random input data instead of extracted audio features, along with the ground truth captions. The random input data is a matrix with same dimensions as \mathbf{X} and with values randomly drawn from the distribution of \mathbf{X} for the data in the training split. We use random input data case in order to investigate if the neural network has learned to process the input data in a useful manner and discover relationships between patterns of the input data and the patterns of the desired output.

3.3. Results and discussion

Table 1 presents the obtained results for the three evaluation cases. The results reported are the mean and standard deviation of the 10 runs. As can be seen, for the proposed method the precision-based metric BLEU drops with the increase of the n -gram length. Comparing it to the recall-based metric METEOR, it can be seen that the proposed method has better precision than recall. Additionally, according to the fragment penalization that is performed in METEOR and the trend of values of BLEU scores, it can be inferred that the proposed method tends to correctly produce words that appear in the original caption but not in the right order. This is also supported by the high value of CIDEr (namely 0.526), the value of ROUGE_L (namely 0.149), and the fact that most of the pro-

duced captions (56.59%) are shorter than the original ones. The case with the random words indicate that the results obtained with our method are well above of a random guess of a sequence of words. The values of the metrics for the random input data case indicate that the proposed neural network has learned to process the input data in a useful manner and discover the relationships between the input audio features and output caption.

The captions of the employed data set (PSE library) are not syntactically structured. By checking the predicted captions, we also observed that our method tends to generate captions that can describe the sound but in different words. For example, our method predicted “field dark” for a recording with an original caption of “electric hum deep”. This fact brings forward the well-known ambiguity for sound recognition when no other context is been given. This ambiguity was greatly employed in the initial days of audio effects where sound events were represented with similar sounds, e.g. the sound of crinkled thin plastic used to represent the sound of fire. Below are three example predictions of our method that try to demonstrate the above facts and not the performance of the method itself. **GX** is the ground truth and **PX** is the corresponding predicted caption. We present more exhaustive examples in the corresponding on-line demo².

G1: goose cu flock calls with distant and near end

P1: cu gull calls medium other birds md birds in

G2: footsteps wood down interior

P2: footsteps belt

G3: door close slide

P3: squeak squeak

4. CONCLUSION AND FUTURE WORK

In the present work, we proposed the first method for automated audio captioning. The method is based on an encoder-decoder scheme with an attentional layer in between them. The method was evaluated using a commercial dataset of recordings, each of which is associated with a textual description (caption) within the dataset. The evaluation of the proposed method was performed with machine translation metrics commonly used in image and video captioning tasks. The results showed that our method can identify the information in an audio recording and produce a set of words that can describe the recording to a certain extent, but still far from intelligible caption.

Recommendations for future work on the problem of audio captioning include the usage of methods that compensate for audio and captions with different lengths, utilization of language model, and the employment of a different data set with more than one caption per recording and captions that exhibit increased syntactic structure.

²<http://arg.cs.tut.fi/demo/captioning/>

5. REFERENCES

- [1] S. Adavanne, P. Pertilä, and T. Virtanen, “Sound event detection using spatial features and convolutional recurrent neural network,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017.
- [2] G. Wichern, M. Yamada, H. Thornburg, M. Sugiyama, and A. Spanias, “Automatic audio tagging using covariate shift adaptation,” in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, March 2010, pp. 253–256.
- [3] ProSoundEffects, “Master Library 2.0,” <http://www.prosoundeffects.com/blog/master-library-2-0-nab/>, accessed March 2017, 2015.
- [4] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, “Show and tell: A neural image caption generator,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 3156–3164.
- [5] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo, “Image captioning with semantic attention,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 4651–4659.
- [6] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko, “Translating videos to natural language using deep recurrent neural networks,” in *North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, May - June 2015.
- [7] 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)*, Sept 2016, pp. 3364–3368.
- [8] K. Cho et al, “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Oct 2014, pp. 1724–1734.
- [9] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [10] Y. Wu et al, “Google’s neural machine translation system: Bridging the gap between human and machine translation,” arXiv:1609.08144, Oct 2016, arXiv.org e-print.
- [11] M. Johnson et al, “Google’s multilingual neural machine translation system: Enabling zero-shot translation,” arXiv:1611.04558, Nov 2016, arXiv.org e-print.
- [12] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio, “Maxout networks,” in *International Conference on Machine Learning (ICML)*, 2013.
- [13] X. Chen et al, “Microsoft COCO captions: Data collection and evaluation server,” *CoRR*, vol. abs/1504.00325, 2015.
- [14] 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 (ACL ’02)*, 2002, pp. 311–318.
- [15] C.-Y. Lin, “Rouge: a package for automatic evaluation of summaries,” in *Proceedings of the Workshop on Text Summarization Branches Out (WAS 2004)*, 2004, pp. 25–26.
- [16] 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 (StatMT ’07)*, 2007, pp. 228–231.
- [17] 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)*, June 2015, pp. 4566–4575.
- [18] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference for Learning Representations*, May 2015.
- [19] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS10)*, May 2010.
- [20] F. Chollet, “Keras,” <https://github.com/fchollet/keras>, 2015.
- [21] Theano Development Team, “Theano: A Python framework for fast computation of mathematical expressions,” *arXiv e-prints*, vol. abs/1605.02688, May 2016.
- [22] R. Kelly, “PyEnchant,” <https://pypi.python.org/pypi/pyenchant>, 2016.
- [23] K. Atkinson, “GNU aspell,” <https://github.com/GNUAspell/aspell>, 2016.