



Enhance Temporal Relations in Audio Captioning with Sound Event Detection

Zeyu Xie, Xuenan Xu, Mengyue Wu[†], Kai Yu[†]

MoE Key Lab of Artificial Intelligence X-LANCE Lab,
Department of Computer Science and Engineering AI Institute,
Shanghai Jiao Tong University, Shanghai, China
{zeyu.xie, wsntxxn, mengyuewu, kai.yu}@sjtu.edu.cn

Abstract

Automated audio captioning aims at generating natural language descriptions for given audio clips, not only detecting and classifying sounds, but also summarizing the relationships between audio events. Recent research advances in audio captioning have introduced additional guidance to improve the accuracy of audio events in generated sentences. However, temporal relations between audio events have received little attention while revealing complex relations is a key component in summarizing audio content. Therefore, this paper aims to better capture temporal relationships in caption generation with sound event detection (SED), a task that locates events' timestamps. We investigate the best approach to integrate temporal information in a captioning model and propose a temporal tag system to transform the timestamps into comprehensible relations. Results evaluated by the proposed temporal metrics suggest that great improvement is achieved in terms of temporal relation generation.

Index Terms: Audio captioning, Sound Event Detection, Temporal-enhanced model

1. Introduction

Increasing amount of research has shed light on machine perception of audio events, for instance label-wise classification and detection. Recently automated audio captioning (AAC) [1] has gathered much attention due to its resemblance to human perception, which involves not only detecting and classifying sounds, but also summarizing the relationship between different audio events [2]. Over the last few years, AAC has witnessed remarkable advances in recent works. The utilization of pre-trained audio classification and language generation models improve the captioning performance significantly [3, 4]. The incorporation of semantic guidance (e.g., keywords [5, 6, 7], sound tags [8] or similar captions [9]) and new loss functions [10, 11, 12] are also hot topics. While previous work endeavors to better detect audio events and improve caption quality, little attention is paid to summarizing relations between different sound events in a caption. The current captioning model rarely outputs sentences involving temporal conjunctions like “before”, “after” and “followed by” that suggest the sequential relations between events. A statistical examination on a well-performing AAC model [3] indicates that only 11.1% generated captions include precise temporal relations.

Different from vision-based captioning where a plethora of spatial attributes can be extracted, audio events' relations are mainly focused on their time specificity as shown in Figure 1.

Whether two audio events occur sequentially or simultaneously

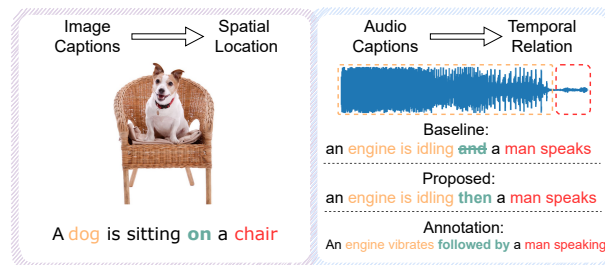


Figure 1: Expressions of relationships in image versus audio. Image pays more attention to spatial relations while audio focus on temporal relations.

is important to understand the audio content correctly, which is as critical as whether two objects in an image are adjacent, stacked, or overlaid.

Sound event detection (SED), a task to detect on- and off-sets of each sound event, on the other hand, provides extensive information on the temporal location of each event. Previous works integrated SED outputs by direct concatenation to improve the overall quality and accuracy of generated captions [8, 13]. However, whether such straightforward fusion methods can help a captioning model learn about temporal relations between events remains unexplored. SED output contains information about the occurring probability of hundreds of sound events in each frame. These redundant low-level features are difficult to align with the temporal conjunction words in a caption, making it difficult for the captioning model to leverage SED outputs. In this work, we first directly integrate SED outputs by concatenation (cat-prob-AAC) and attention (attn-prob-AAC), to investigate the performance of direct SED integration methods. The results demonstrate that such approaches bring little improvement in temporal relationship description accuracy.

Therefore, it is necessary to distill high-level, comprehensible temporal information from SED outputs, for a better alignment with audio caption content to mimic humans' temporal information processing procedure. Inspired by this, we first analyse the current AAC data and propose a 4-scale temporal relation tagging system (i.e. simultaneous, sequential) based on human annotations. A clear matching mechanism is further proposed to infer the temporal relations from SED outputs and align with the temporal tags. Based on this, we propose a temporal tag-guided captioning system (temp-tag-AAC), which takes temporal tag guidance inferred from SED output that represents the complexity of temporal information to facilitate the model to generate captions with accurate temporal expressions.

[†]Mengyue Wu and Kai Yu are the corresponding authors.

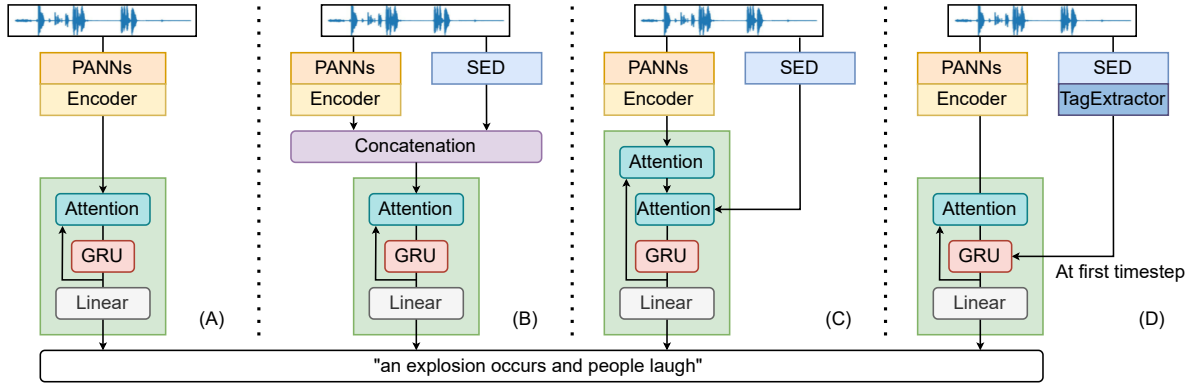


Figure 2: An overview of different AAC models. (A) Baseline AAC model: the decoder generates captions solely based on audio embeddings; (B) Cat-prob-AAC: audio embeddings and SED outputs are concatenated and used as the input to the decoder; (C) Attn-prob-AAC: an attention mechanism is used to integrate SED outputs and decoder hidden states; (D) Temp-tag-AAC: mimicking human judgment, tags are extracted and used as the input at the first timestep instead of $\langle \text{BOS} \rangle$.

To measure the quality of generated captions in terms of temporal relationship descriptions, we propose ACC_{temp} and F1_{temp} . Evaluated by these temporal-focused metrics and commonly-adopted captioning metrics (e.g., BLEU) indicate that temp-tag-AAC significantly outperforms the baseline model and the direct SED integration approach, especially in temporal relationship description accuracy. Our contributions are summarized as follows:

1. Innovative utilization of SED to enhance the temporal information in AAC, with a temporal tag to better imitate humans' inference on temporal relations.
2. Metrics that are specifically designed to measure a system's capability in describing sound events' temporal relations.
3. Validation shows that the proposed temp-tag-AAC leverages SED outputs to significantly improves the accuracy of temporal expression as well as the caption quality.

2. Temporal-Enhanced Captioning System

This section illustrates our temporal-enhanced captioning system shown in Figure 2, includes: 1) the baseline model for audio captioning; 2) the SED model that predicts the probability of events; 3) two direct approaches for integrating probability as temporal information; 4) proposed temp-tag-AAC approach.

2.1. Baseline Approach

The baseline framework follows an encoder-decoder architecture which achieves competitive performance in DCASE challenges [14].

Audio Encoder PANNs [15] CNN14, a pre-trained convolutional neural network, is adopted to extract the feature from the input audio \mathcal{A} . We use a bidirectional gated recurrent unit (GRU) network as the audio encoder to transform the feature into an embedding sequence $\mathbf{e}^A \in \mathbb{R}^{T \times D}$. The combination takes advantage of the pre-trained large model while setting some parameters trainable for adaptation to the target captioning task.

$$\mathbf{e}^A = \text{Encoder}(\text{PANNs}(\mathcal{A})) \quad (1)$$

Text Decoder We use a unidirectional GRU as the text decoder to predict the caption word by word. At each timestep n , a context vector \mathbf{c} is calculated by attention mechanism [16], given \mathbf{e}^A and the previous hidden state \mathbf{h}_{n-1} :

$$\alpha_{n,t} = \frac{\exp(\text{score}(\mathbf{h}_{n-1}, \mathbf{e}_t^A))}{\sum_{t=1}^T \exp(\text{score}(\mathbf{h}_{n-1}, \mathbf{e}_t^A))} \quad (2)$$

$$\mathbf{c} = \text{ATTN}(\mathbf{h}_{n-1}, \mathbf{e}^A) = \sum_{t=1}^T \alpha_{n,t} \mathbf{e}_t$$

Then the text decoder predicts the next word based on previously generated words $w_{0:n}$ and \mathbf{c} . At the first timestep, w_0 is a special " $\langle \text{BOS} \rangle$ " token denoting the beginning of a sentence.

2.2. SED Architecture

To ensure the reliability of the SED results, we use a separately-trained SED model. It adopts a convolutional recurrent neural network architecture with 8 convolutional layers attached by a BiGRU. The convolution layers take a structure similar to the CNN10 in PANNs, with the difference that we use a downsampling ratio of 4 on the temporal axis. Compared with other SED models provided in PANNs which typically utilize a downsampling ratio of 32, we keep a relatively high temporal resolution for more accurate SED.

Given an audio clip, the SED model outputs the predicted probability $\tilde{\mathbf{e}}^S \in \mathbb{R}^{\tilde{T} \times M}$, where \tilde{T} and M denote the sequence length and the number of sound event categories respectively. Due to the higher resolution of the SED model, $\tilde{T} > T$. The probability is temporally aligned to the audio embedding to obtain $\mathbf{e}^S \in \mathbb{R}^{T \times M}$ by pooling on every $\frac{\tilde{T}}{T}$ segments along the temporal axis.

2.3. Direct SED Integration

Cat-prob-AAC The probability is concatenated onto audio embedding, resulting in $\mathbf{e}^{A_{new}} \in \mathbb{R}^{T \times (D+M)}$, which is used as the input to the decoder instead of the original \mathbf{e}^A .

$$\mathbf{e}^{A_{new}} = \text{CONCAT}(\mathbf{e}^A, \mathbf{e}^S) \quad (3)$$

Table 1: *Temporal Tags Extracted from Text (Captions) and Audio (SED Results), c.w. = Conjunction Words.*

Temporal Tag	0	1	2	3
Caption	No c.w.	Simultaneous c.w.	Sequential c.w.	More Complex Relations
SED	Only 1 Event	Simultaneous Events	Sequential Events	More Complex Events

Attn-prob-AAC Another attention is used to integrate the probability and context vector \mathbf{c} obtained from Equation (2). The result is used as the input to the GRU instead of \mathbf{c} .

$$\mathbf{c}^{new} = \text{ATTN}(\mathbf{c}, \mathbf{e}^S) \quad (4)$$

2.4. Temp-tag-AAC

In our proposed temp-tag-AAC system, we transform the SED outputs into quantized temporal tags to make it easier for the model to learn the correspondence between SED outputs and captions.

We use double threshold post-processing [17] with a low threshold of 0.25 and a high threshold of 0.75 to obtain the on- and off-sets of detected sound events from probability $\tilde{\mathbf{e}}^S$. To infer relation between two different audio events, we compare the *overlap* of them and the *duration* of the shorter event. If *overlap* is less than half of *duration*, these two events are considered to occur sequentially; otherwise, they are considered to occur simultaneously. Based on the relations in audio clip obtained above, a 4-scale temporal tag representing the complexity of temporal information is extracted according to Table 1. The process is shown as in Algorithm 1. During **inference**, the temporal tag **inferred from the SED outputs** is used as w_0 fed to the decoder as temporal guidance, replacing the original $\langle \text{BOS} \rangle$.

To help the model learn the correspondence between the temporal tag and the temporal descriptions in captions better, the **ground truth** tag is fed to the decoder during **training**. The ground truth tag is extracted from the annotations based on the occurrence of conjunction words according to Table 1. We manually collect these conjunction words by analyzing the existing AAC datasets, such as “while”, “and”, *etc.* indicating “simultaneously,” and “follow”, “then”, *etc.* indicating “sequentially”.

Algorithm 1: Infer temporal tags from SED results.

Data: Predicted probability $\tilde{\mathbf{e}}^S$
Result: Temporal tag

```

1 Relations  $\leftarrow$  empty list;
2  $\{E : E_{on}, E_{off}\} \leftarrow \text{DOUBLE\_THRES}(\tilde{\mathbf{e}}^S)$ ;
3 for every pair  $A, B$  in  $\{E\}$  where  $A_{on} < B_{on}$  do
4   | Overlap  $\leftarrow A_{off} - B_{on}$ ;
5   | Duration  $\leftarrow \text{MIN}(A_{off} - A_{on}, B_{off} - A_{on})$ ;
6   | if Overlap  $< 0.5 \times \textit{Duration}$  then
7     | | Relations.insert(“sequential”);
8   | else
9     | | Relations.insert(“simultaneous”);
10  | end
11 end
12 Tag  $\leftarrow$  Query Table 1 using Relations;
13 Return Tag;
```

3. Experimental Setup

3.1. Datasets

AudioSet [18] is a large-scale weakly-annotated sound event dataset, where sound events appearing in each audio clip are annotated, which consists of 527 categories. AudioSet also provides a small-scale strongly-annotated subset [19] which contains additional on- and off-sets of present events.

AudioCaps [20] is the current largest AAC dataset, containing 50k+ audio clips collected from AudioSet. According to the extraction method mentioned in Section 2.4, 13487, 29399, 5438 and 8472 captions in AudioCaps annotations belong to the 4 scales respectively. The latter two more complex scenarios account for $\approx \frac{1}{4}$ of the total.

Clotho [21] is another AAC dataset, containing 5k+ audio clips. The distribution of ground truth tag numbers in Clotho annotations is 8246, 18077, 926, 2396 respectively. The latter two scales account for $\approx \frac{1}{10}$, which is much more imbalanced than AudioCaps. As a matter of fact, Clotho derives from Freesound, where the audio clips often contain only the indicated sound with minimal background noise [22, p. 51][23].

3.2. Hyper-parameters

The SED model is first pre-trained on the weakly-annotated AudioSet, and then fine-tuned on the strongly-annotated AudioSet subset [19]. It achieves a d' of 2.37 on the strongly-annotated AudioSet evaluation set, compared with 1.39 in [19], indicating that it provides reliable results for caption generation.

The training of audio captioning models, including the baseline model and other three approaches, follows the setup in [14]. Models are trained for 25 epochs. Cross-entropy loss is used along with label smoothing ($\alpha = 0.1$). We use a linear warm-up and an exponential decay strategy to schedule the learning rate, whose maximum value is 5×10^{-4} . Scheduled sampling is used with the proportion of teacher forcing decreasing linearly from 1 to 0.7. Beam search with a size of 3 is adopted during inference.

3.3. Metrics

Generated captions are evaluated by our proposed temporal metrics and commonly-adopted metrics in AAC task.

Temporal Metrics To better evaluate whether a generated caption includes temporal relations, we take the time-related conjunction words as a clue. Captions can be classified upon whether there exists sequential conjunction words or not. The conjunction words include “follow, followed, then, after”, which are only used to suggest temporal relations between sound events. For example, “Door closed *then* a man talking” is regarded as a positive example (with temporal output) since “then” appears. We exclude simultaneous conjunction words (e.g. “and, with, as, while”) because they might carry semantic conjunction function and does not always signify temporal relations. Whether these words express temporal relations cannot be recognized automatically and accurately. Naturally, the temporal evaluation can be regarded as a binary classification eval-

Table 2: Results of system performance. “FENSE_{p=0}” indicates not penalizing grammatical errors in FENSE.

Dataset	System	BLEU ₄	ROUGE _L	CIDEr	METEOR	SPICE	FENSE _{p=0}	ACC _{temp}	F1 _{temp}
AudioCaps	Baseline	26.3	49.4	72.6	23.9	17.6	62.3	37.3	23.4
	Cat-prob-AAC	25.8	49.4	71.2	24.1	17.5	62.6	39.2	27.1
	Attn-prob-AAC	26.3	49.1	72.1	23.9	17.2	62.3	40.7	30.4
	Temp-tag-AAC (Ours)	28.8	49.4	68.2	25.5	18.5	61.6	66.5	75.2
Clotho	Baseline	16.8	38.3	40.8	17.5	12.1	48.4	69.2	17.4
	Cat-prob-AAC	16.2	38.0	38.8	17.3	11.9	48.6	69.7	17.7
	Attn-prob-AAC	16.7	37.7	38.7	17.1	11.9	48.2	69.7	16.4
	Temp-tag-AAC (Ours)	13.3	35.3	33.8	15.9	10.9	47.2	59.0	46.5

uation: to determine whether there are sequential conjunction words or not in a caption. We therefore use the binary classification evaluation metrics ACC_{temp} and F1_{temp} to measure the accuracy of temporal relation description. Within 5 reference captions, the maximum value contains the most detailed information and is taken as label to ensure the metric correctness.

Overall Quality Evaluation Metrics We also adopt common audio captioning metrics to evaluate the overall quality of generated captions, including BLEU [24], ROUGE_L [25], METEOR [26], CIDEr [27], SPICE [28] and FENSE [29]. For FENSE we do not penalize grammatical errors to focus on evaluating the accuracy of captions’ semantic information.

4. Results and Analysis

4.1. Temporal Relation Enhancement

Comparing temp-tag-AAC with the baseline model, our tag mechanism greatly improves the accuracy of temporal expressions on both datasets (shown in Table 2). For AudioCaps, both ACC_{temp} and F1_{temp} are significantly improved, suggesting the effectiveness of our method in enhancing temporal relations. Due to the imbalanced categories of Clotho, F1_{temp} is more reliable compared with ACC_{temp}, which also indicates a better capability to generate temporal-rich captions.

Without guidance, the baseline model tends to output general conjunction words that do not represent specific relations (“and” and “while” are typical examples), resulting in loss of attention to the temporal relations between sound events. Temp-tag-AAC restricts the output by inputting a tag which guides the model to use conjunction words to describe temporal relations. Typical examples are shown in Figure 3: by incorporating the temporal tags, temp-tag-AAC successfully expresses the temporal relations while the baseline model simply uses “and”.

4.2. Overall Quality of Generated Sentences

The overall quality is evaluated by commonly-adopted metrics and shown in columns 3 to 8 of Table 2. On AudioCaps, temp-tag-AAC outperforms the baseline model on some metrics, but falls behind on others, indicating that our method is comparable to the baseline. However, on Clotho, the quality of the caption sentences decreases, though the accuracy of temporal relations still sees an increase. The performance drop is attributed to the data discrepancies between AudioSet and Clotho. As stated in Section 3.1, Clotho audio samples exhibit vastly different characteristics from those in AudioSet. As a result, the SED model trained on Audioset tends to output complex temporal tags (i.e., “3”) for Clotho data when only one sound event is present. The captioning model trained with such tags is prompted to gener-

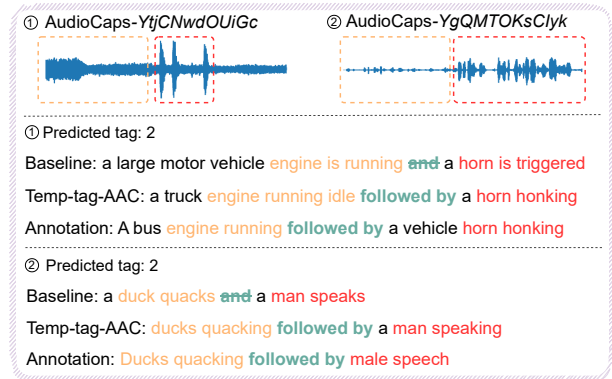


Figure 3: Output examples generated by baseline systems and tag guided approach.

ate sentences with complex conjunctions for single-event audios, which undermines its ability in generating reference-alike captions. The declined quality on Clotho indicates that adaptive SED deserves further exploration for generalization purpose.

4.3. Comparison Between Different Approaches

Comparing three different methods of integrating temporal information, we can conclude that direct integration by concatenation or attention only slightly improves the temporal description accuracy, but are far less effective than temp-tag-AAC. This validates our intuition that human-like quantized prompts are more conducive to learning the correspondence between temporal information and conjunctions than direct outputs of SED.

5. Conclusions

This paper aims to improve the performance of expressing temporal information in AAC task. We demonstrate that direct integration of SED outputs provides little help in improving the temporal relation description accuracy. To overcome such challenge, we propose temp-tag-AAC which mimics human judgment by introducing 4-scale tags to guide the model to utilize temporal information. Binary classification metrics ACC_{temp} and F1_{temp} are proposed to measure the accuracy of the temporal relation description. Experimental results show that temp-tag-AAC significantly improves the temporal relation description accuracy. With the guidance from the temporal tag, temp-tag-AAC uses conjunctions to express the temporal relations between sound events. It is also comparable with the baseline in terms of the overall semantic quality of generated captions.

6. References

- [1] K. Drossos, S. Adavanne, and T. Virtanen, "Automated audio captioning with recurrent neural networks," in *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2017, pp. 374–378.
- [2] M. Wu, H. Dinkel, and K. Yu, "Audio caption: Listen and tell," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 830–834.
- [3] X. Xu, H. Dinkel, M. Wu, Z. Xie, and K. Yu, "Investigating local and global information for automated audio captioning with transfer learning," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 905–909.
- [4] X. Mei, X. Liu, Q. Huang, M. D. Plumbley, and W. Wang, "Audio captioning transformer," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Barcelona, Spain, November 2021, pp. 211–215.
- [5] Y. Koizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, "The ntt dcase2020 challenge task 6 system: Automated audio captioning with keywords and sentence length estimation," DCASE2020 Challenge, Tech. Rep., June 2020.
- [6] Z. Ye, H. Wang, D. Yang, and Y. Zou, "Improving the performance of automated audio captioning via integrating the acoustic and semantic information," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Barcelona, Spain, 2021, pp. 40–44.
- [7] A. Ö. Eren and M. Sert, "Audio captioning based on combined audio and semantic embeddings," in *2020 IEEE International Symposium on Multimedia (ISM)*. IEEE, 2020, pp. 41–48.
- [8] F. Gontier, R. Serizel, and C. Cerisara, "Automated audio captioning by fine-tuning bart with audioset tags," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Barcelona, Spain, November 2021, pp. 170–174.
- [9] Y. Koizumi, Y. Ohishi, D. Niizumi, D. Takeuchi, and M. Yasuda, "Audio captioning using pre-trained large-scale language model guided by audio-based similar caption retrieval," *arXiv preprint arXiv:2012.07331*, 2020.
- [10] E. Cakır, K. Drossos, and T. Virtanen, "Multi-task regularization based on infrequent classes for audio captioning," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Tokyo, Japan, November 2020, pp. 6–10.
- [11] X. Xu, H. Dinkel, M. Wu, and K. Yu, "Audio caption in a car setting with a sentence-level loss," in *Proceedings of the International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.
- [12] X. Liu, Q. Huang, X. Mei, T. Ko, H. Tang, M. D. Plumbley, and W. Wang, "Cl4ac: A contrastive loss for audio captioning," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Barcelona, Spain, November 2021, pp. 196–200.
- [13] A. Ö. Eren and M. Sert, "Audio captioning using sound event detection," *arXiv preprint arXiv:2110.01210*, 2021.
- [14] X. Xu, Z. Xie, M. Wu, and K. Yu, "The SJTU system for DCASE2022 challenge task 6: Audio captioning with audio-text retrieval pre-training," DCASE2022 Challenge, Tech. Rep., 2022.
- [15] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "Panns: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [16] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [17] H. Dinkel, M. Wu, and K. Yu, "Towards duration robust weakly supervised sound event detection," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 29, pp. 887–900, 2021.
- [18] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [19] S. Hershey, D. P. Ellis, E. Fonseca, A. Jansen, C. Liu, R. C. Moore, and M. Plakal, "The benefit of temporally-strong labels in audio event classification," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 366–370.
- [20] C. D. Kim, B. Kim, H. Lee, and G. Kim, "Audiocaps: Generating captions for audios in the wild," in *Proceedings of Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019, pp. 119–132.
- [21] K. Drossos, S. Lipping, and T. Virtanen, "Clotho: An audio captioning dataset," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 736–740.
- [22] N. Turpault, "Analyse des problématiques liées à la reconnaissance de sons ambiants en environnement réel," Theses, Université de Lorraine, May 2021. [Online]. Available: <https://hal.inria.fr/tel-03304880>
- [23] I. Martin and A. Mesaros, "Diversity and bias in audio captioning datasets," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, 2021, pp. 90–94.
- [24] P. Kishore, R. Salim, W. Todd, and Z. Wei-Jing, "BLEU: a Method for Automatic Evaluation of Machine Translation," in *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL)*, 2002, pp. 311–318.
- [25] L. Chin-Yew, "Rouge: A package for automatic evaluation of summaries," in *Proceedings of the workshop on text summarization branches out*, no. 1, 2004, pp. 25–26.
- [26] A. Lavie and A. Agarwal, "METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments," in *Proceedings of the Second Workshop on Statistical Machine Translation*, no. June, 2007, pp. 228–23.
- [27] R. Vedantam, C. L. Zitnick, and D. Parikh, "CIDEr: Consensus-based image description evaluation," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 07-12-June, 2015, pp. 4566–4575.
- [28] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "SPICE: Semantic propositional image caption evaluation," in *Proceedings of European Conference on Computer Vision (ECCV)*. Springer, 2016, pp. 382–398.
- [29] Z. Zhou, Z. Zhang, X. Xu, Z. Xie, M. Wu, and K. Q. Zhu, "Can audio captions be evaluated with image caption metrics?" in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 981–985.