

Enriching Ontology with Temporal Commonsense for Low-Resource Audio Tagging

Zhiling Zhang
Shanghai Jiao Tong University
blmoistawinde@sjtu.edu.cn

Zelin Zhou
Shanghai Jiao Tong University
ze-lin@sjtu.edu.cn

Haifeng Tang
CMB Credit Card Center
thfeng@cmbchina.com

Guangwei Li[†]
Shanghai Jiao Tong University
ligw20@sjtu.edu.cn

Mengyue Wu^{†*}
Shanghai Jiao Tong University
mengyuewu@sjtu.edu.cn

Kenny Q. Zhu[†]
Shanghai Jiao Tong University
kzhu@cs.sjtu.edu.cn

ABSTRACT

Audio tagging aims at predicting sound events occurred in a recording. Traditional models require enormous laborious annotations, otherwise performance degeneration will be the norm. Therefore, we investigate robust audio tagging models in low-resource scenarios with the enhancement of knowledge graphs. Besides existing ontological knowledge, we further propose a semi-automatic approach that can construct temporal knowledge graphs on diverse domain-specific label sets. Moreover, we leverage a variant of relation-aware graph neural network, D-GCN, to combine the strength of the two knowledge types. Experiments on AudioSet and SONYC urban sound tagging datasets suggest the effectiveness of the introduced temporal knowledge, and the advantage of the combined KGs with D-GCN over single knowledge source.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning.**

KEYWORDS

Audio Tagging, Low-Resource, Graph Neural Network, Knowledge Graph

ACM Reference Format:

Zhiling Zhang, Zelin Zhou, Haifeng Tang, Guangwei Li[†], Mengyue Wu^{†*}, and Kenny Q. Zhu. 2021. Enriching Ontology with Temporal Commonsense for Low-Resource Audio Tagging. In *ACM International Conference on Information and Knowledge Management, November 1–5, 2021, Virtual Event*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/1122445.1122456>

^{*}Corresponding author, supported by National Natural Science Foundation of China (No.61901265), Shanghai Pujiang Program (No.19PJ1406300), and Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102).

[†]Corresponding author, and is partially supported by SJTU-CMB Credit Card Center Joint Research Grant and SJTU Medical-Engineering Cross Disciplinary Scheme.

[†] authors are with MoE Key Lab of Artificial Intelligence and X-LANCE Lab.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).
CIKM '21, November 1–5, 2021, Virtual Event

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Audio tagging is the task to label the sound recordings with representative tags like sound events. It can be used in many applications like music tagging [6], sound retrieval [5] and urban sound planning [1]. However, the manual labeling required by a well-performing tagging model can be heavy and expensive. Compared to annotating for images classification, annotators have to spend substantially more time on finishing the whole recording before tagging. Consequently, even the largest dataset for audio tagging, AudioSet [7], is orders of magnitude smaller than image classification datasets [2, 19], and many other common datasets are even smaller, like SONYC [1] with only 2,351 training recordings. Therefore, when we want to tailor a new dataset for particular needs, the data size tends to be limited, and thus robust methods in such low-resource scenarios are largely desired.

Structured knowledge has been shown effective in the low-resource setting of many tasks [3, 9, 13, 17, 22], while the exploration of knowledge is still limited for audio tagging. These works mainly focus on ontological knowledge [8, 16, 18]. For example, AudioSet involves *IsA* relations between sound events, such as *Speech IsA Human voice*. Models enhanced with such knowledge can effectively capture the similarity between tags of similar category. However, ontological knowledge failed to cover other relations between different categories, like the frequent *Co-occurrence* of *Car* sounds and *Engine* sounds. Such relationships are still under-explored.

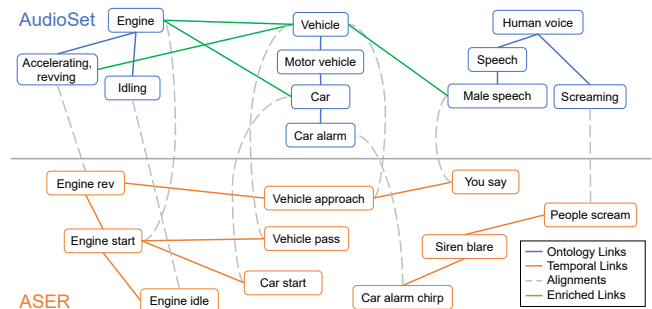


Figure 1: Aligning ASER to AudioSet. We establish an enriched link between two AudioSet events if their corresponding ASER events have temporal links.

To extend the dimensions of available knowledge for audio analysis, we propose to construct temporal knowledge graph on top of pre-defined audio tags with the knowledge transferred from ASER [21], a large-scale commonsense knowledge graph about

the relations between events extracted from textual corpus. Here we pay special attention to temporal relations like *Precedence* and *Conjunction*, which we will collectively refer to as “temporal links”. They can provide hints whether 2 events might co-occur in a short period, as is the case for multi-label audio tagging. This kind of knowledge is similar to the sound event co-occurrence graph mined from audio dataset [16, 20] to some extent, but they come from different sources, i.e., text vs. audio. Algorithms to mine sound event co-occurrence from audio data rely on dataset-specific hyperparameter tuning, and large-scale, accurate, multi-label annotations, which can be hard to obtain due to missing labels problems [4], or from single-label or low-resource datasets.

The knowledge construction requires the alignments between ASER events and domain-specific label sets, usually accompanied with ontologies, such as AudioSet (Figure 1). We will then add an enriched link between two Audioset events if any pair of their aligned events in ASER have temporal links. There are many challenges in establishing the alignments. First, there are mismatches between the event representations. For example, events are usually named with representative noun phrases or verb phrases in AudioSet, while ASER provides a finer representation of short clauses. Therefore, one AudioSet event can be aligned to multiple related ASER events, like *Vehicle* to *Vehicle approach* and *Vehicle pass*. Moreover, the alignments should be decided according to acoustic relatedness rather than lexical or semantical similarity alone. For instance, we should not treat *I see engine* as related to the sound event of *Engine* despite the shared object, as the former event doesn’t make a sound. In contrast, we should link *Male Speech* to *You say* although they have no words in common.

In this work, we propose a semi-automatic approach to align the events in ASER and those in the target label set. We use heuristics to automatically extract audible events in ASER that are synonymous to those events in AudioSet, with the help of Word Sense Disambiguation (WSD) [12] and WordNet [14]. For AudioSet, candidate alignments are also verified manually to ensure the quality of final temporal KG. The resulting temporal KGs can be directly incorporated into existing Graph Convolutional Network (GCN)-based models. We conduct experiments on two audio tagging datasets in different domains: the open domain AudioSet [7] and SONYC urban sound tagging dataset [1]. Results suggest the usefulness of temporal knowledge in low resource settings. Nonetheless, simple combination of temporal and ontological knowledge sources in one graph doesn’t provide performance gain. We hypothesize that a single graph can’t model the heterogeneity of the two knowledge types properly. Inspired by R-GCN [15], we further propose D-GCN that applies relation-specific transformation for two knowledge types.

In summary, our contributions are: (1) We propose a semi-automatic approach to construct temporal commonsense knowledge graph to enrich different audio tagging ontologies like AudioSet and SONYC. (2) We propose to use a variant of R-GCN of two relation types, D-GCN, to leverage the heterogeneity of ontological and temporal knowledge. (3) Experiments on AudioSet and SONYC dataset in low-resource settings show that temporal KG-enhanced model outperforms backbone model without KG, and combining ontology and temporal KG with D-GCN can provide further improvement.¹

¹Code and data: https://github.com/blmoistawinde/dgcn_tagging

2 METHOD

2.1 Knowledge Resources

AudioSet [7] is a hierarchically structured ontologies comprised of 632 audio events. It is by far the largest ontology, covering most common sound events. *SONYC* [1] is a two-level taxonomy consists of 8 coarse level tags, 23 fine level tags about urban sounds. This taxonomy is smaller due to its specific target of detecting urban noises. *ASER* [21] is a large-scale eventuality knowledge graph extracted from textual corpus. Each eventuality is represented as a short clause containing lemmatized words of subject, verb, object, etc. In this work, we will use its core version with 27,565,673 event nodes and 8,834,257 relation edges of 15 types. We will focus on temporal relations types like *Precedence* and *Conjunction*.

2.2 Knowledge Construction

In this section, we will introduce how to perform the alignment between an audio event label set and ASER, and construct the temporal KG to enrich the existing ontologies (Figure 2). Without loss of generality, we will mainly describe the procedure with AudioSet as target, and explain the difference for SONYC when necessary.

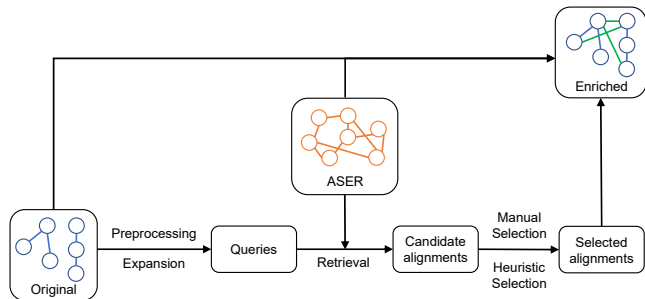


Figure 2: Illustration of alignment pipeline.

2.2.1 Preprocessing and Expansion. First, we conduct preprocessing to deal with the morphological differences between two type of representations. The words in ASER events are already lemmatized, while the verbs in AudioSet tags are not. Therefore, we enforce lemmatization on AudioSet tags, and also do lowercasing, remove the parentheses and stopwords.

After that, since there are lexical variations in the expressions of similar audio events in ASER, we propose to expand a single AudioSet tag into multiple equivalent queries to improve the recall of alignment candidates retrieval (§2.2.2).

First, some of the tags already contain parallel concepts or synonyms, like “Roaring cats (lions, tigers)”, and we will split it into multiple queries with one concept each, like “roaring cats”, “lions” and “tigers”. Moreover, for those tags without provided synonyms, we can also extract their synonyms from the lexical database WordNet [14]. To leverage the knowledge from WordNet, we applied the Lesk algorithm [12] for Word Sense Disambiguation (WSD) to link each word to its corresponding Wordnet synset. There are also ambiguous events in AudioSet with multiple father events. For example, the event “Hiss” is a child event of “Cat”, “Snake” and “Steam”, and the acoustic property of the event might be different under different father events. Hence, we produce different queries by pairing such event with each of its father events, so the expanded

queries for the above example would be “cat hiss”, “snake hiss” and “steam hiss”.

2.2.2 Alignment Candidates Retrieval. Since ASER contains numerous events, we need to retrieve a small number of likely candidates before alignment selection. At first, we filter out some unlikely candidates by excluding events with noisy patterns like duplicate verbs (“I say say”), and infrequent ones (with frequency < 5). We then retrieve the top 10 events for each query with ElasticSearch. First, we adopt pure text-matching, and give higher weight to the matching of verbs, as verbs are the key components of events. Moreover, since we noticed that the above text matching approach may sometimes retrieve infrequent events that is either too rare or too specific, we also supplemented the above method with another weighting scheme that adds additional weights according to event frequency. Consequently, retrieval results will contain more general and frequent events, which will have more linked relations to leverage. The results from both weighting schemes are combined to balance accuracy and frequency. Finally, the retrieved ASER events with the multiple queries of the same AudioSet event are aggregated as the candidates for alignment. As a result, each event has 31.3 candidates in average with a minimum of 2 and maximum of 190.

2.2.3 Selection. Now we need to select events that are precisely related to a sound event in terms of acoustic property from the retrieved candidates.

For the alignments to AudioSet, we manually annotate them ourselves to ensure the reliability. We need to decide for each candidate eventuality of an AudioSet event, whether they are related, unrelated or that the relation is ambiguous. The event name, description, father/child events, corresponding queries and example videos containing that event are shown to aid the decision. Among all the annotated alignments, about 31.96% are considered related, 13.52% ambiguous and 54.51% unrelated. To ensure the precision of the results, we only use the “related” alignments later.

For the alignments to SONYC, we completely automate the selection process. We observe that most candidates for certain specific event labels (like “Reverse Beeper”) are already of acceptable quality. Therefore, we preserve all the candidates of specific labels except “other/unknown” labels and under-specified labels like “machinery impact”.

2.2.4 Construction. The relations between ASER events are transferred to their corresponding AudioSet events through the alignments, and their relations will be aggregated. For example, let’s assume that an AudioSet event a_1 is aligned to events e_{11}, e_{12} , and another event a_2 is aligned to e_{21} . In ASER, $e_{11}-e_{21}$ has the relation (‘Co_Occurrence’, ‘Conjunction’), and $e_{12}-e_{21}$ has the relation (‘Co_Occurrence’, ‘Precedence’). Then a_1-a_2 will have the aggregated relation (‘Co_Occurrence’, ‘Conjunction’, ‘Precedence’). The resulting KG will inherit various types of relations from ASER, and we will only use *Conjunction* and *Precedence* in later experiments as other relations are either not temporal or too sparse.

2.3 Double GCN (D-GCN)

Typically, CNN-based models [10] are used for audio tagging. To leverage KGs, we roughly follow [20] to add a GCN component to learn the representation of labels. Then the audio representations

from the CNN encoder are dot-producted with the learned label representations, followed by a sigmoid layer to get the prediction.

Either the ontology or temporal KG alone can be incorporated into GCN-based models for audio tagging. However, a single GCN may not be suitable to handle the heterogeneity of two knowledge types when we want to combine their strength in one model. We thus draw inspiration from R-GCN [15], which introduced relation-specific transformations.

$$h_i^{l+1} = \sigma\left(\sum_{r \in \mathcal{R}} \sum_{j \in N_i^r} \frac{1}{c_{ij,r}} W_r^l h_j^l\right) \quad (1)$$

where h_i^l are the features of node i at the l -th layer, \mathcal{R} is the set of relations, N_i^r is the indices of the neighbors of i via edges with relation r (self-loop included), $c_{ij,r}$ is a normalization constant, and W_r^l is the specific weight matrix for relation r at the l -th layer. The original R-GCN also introduce basis function decomposition and block decomposition for W_r^l as a regularization to reduce parameters and prevent overfitting on rare relations.

Different from the original R-GCN, which aims to tackle the completion of knowledge bases with over 1,000 relation types, we only need to model 2 kinds of relations with relatively balanced numbers. Therefore, we don’t apply the regularizations and further add relation specific bias terms, to improve the model’s expressiveness on each relation. We will refer to this variant as D-GCN.

3 EXPERIMENTS

We experiment our methods on datasets with corresponding ontologies to verify the effectiveness of either ontological or temporal knowledge alone, and their combination.

AudioSet is a large-scale multi-label audio tagging dataset collected from Youtube videos with the annotations of 527 categories out of the 632 tags defined in the ontology. Most recordings are processed into 10 seconds single-channel 16kHz, 16-bit wave format. Due to the changes of videos, it is not possible to recover the whole dataset. We downloaded 19,400 (87.5%), 1,851,420 (90.7%), and 17,756 (87.2%) recordings for the balanced train, full train, and evaluation set, respectively. To simulate the low-resource scenario, we randomly sample 1% of the unbalanced set as the training set, which has 18,514 samples, where 134 classes have no more than 5 samples, and 10 classes have no training sample. We also sample 5% and 10% sets for comparison. We use the commonly used mAP, mAUC as the evaluation metrics. SONYC is the multi-label Urban Sound Tagging dataset used in DCASE 2019 Task5 (D19T5). It contains 2,351 train recordings and 443 validate recordings, which can be considered relatively low-resource. All recordings are 10 seconds single-channel 44.1kHz, 16-bit wave format. We use the official metrics, micro AUPRC and macro AUPRC.

3.1 Experimental Setup

Audio Features. For SONYC, we adopt a similar setting as [11], all audios are re-sampled to 32 kHz and 64-Mel-bin log-Mel spectrograms are used to represent the audios. The window size is 1024 samples, the hop size of 500 samples, and cut-off frequencies of 50 Hz to 14 kHz. For AudioSet, our setting is similar to [10], all audios are re-sampled to 16 kHz and represented as 64-Mel-bin log-Mel

Methods	Balance		Unbalance (100%)	
	mAP	mAUC	mAP	mAUC
CNN14	0.2441	0.8930	0.4090	0.9669
AT-GCN	0.2510	0.9278	0.4095	0.9664
GCN(ASER)	0.2500	0.9283	0.3994	0.9660
GCN(AudioSet)	0.2543	0.9420	0.4063	0.9665
GCN(ASER+AudioSet)	0.2490	0.9277	0.3999	0.9690
D-GCN	0.2554	0.9377	0.4109	0.9648

Table 1: Results on AudioSet evaluation set with models trained on balanced and unbalanced set.

Methods	1%	5%	10%
CNN14	0.1118	0.2343	0.2770
AT-GCN	0.1243	0.2331	0.2785
GCN(ASER)	0.1252	0.2269	0.2735
GCN(AudioSet)	0.1280	0.2336	0.2747
GCN(AudioSet+ASER)	0.1214	0.2283	0.2741
D-GCN	0.1283	0.2387	0.2799
D-GCN rel. improvement	14.73%	1.88%	1.05%

Table 2: mAP for each model trained on different portion of the unbalanced set, and the relative improvement(%) of D-GCN over CNN14 backbone.

spectrograms. The window size is 512 samples, the hop size of 160 samples, and cut-off frequencies of 50 Hz to 8 kHz.

Models and Baselines. We adopt standard CNN models as our baseline, that is, the CNN9 model use in [11] for SONYC and the CNN14 (16kHz) in [10] for AudioSet. We also introduce the SOTA method AT-GCN [20] as baseline, which is based on co-occurrence graph mined from the whole AudioSet annotation, and uses tuned, dataset-specific hyperparameter for edge thresholding and smoothing. It is thus not directly applicable to SONYC. Our models include GCN(ASER), GCN(AudioSet), GCN(ASER+AudioSet), which refers to single GCN with temporal knowledge, AudioSet ontology, and their combination. D-GCN denotes double-GCN with 2 types of knowledge. We replace AudioSet with SONYC’s ontology (OT) for experiments on SONYC dataset. We use batch size of 32 for all models and the learning rate is 1e-3 for all models except D-GCN using 3e-4.

3.2 Results

3.2.1 AudioSet. Table 1 shows the performance of models trained on the official balanced and unbalanced set, while Table 2 shows the results on sampled subsets. We can see that all GCN-based models significantly outperform the baseline CNN14 on balanced and low-resource (1%) set, suggesting the usefulness of the knowledge sources including the newly proposed temporal knowledge in low-resource scenarios. As the size of training data grows, the advantage of GCN models ceases to exist, expect for AT-GCN, possibly due to its knowledge of the co-occurrences on the whole training set, which matches more with the larger training data. D-GCN performs consistently better than single GCN with one KG or the simple addition of both KGs, showing the effectiveness of the separate relation modeling, and it also outperforms CNN14 by mAP in all settings despite the diminishing gain.

Methods	Fine-level		Coarse-level	
	Mi AUPRC	Ma AUPRC	Mi AUPRC	Ma AUPRC
CNN9	0.675	0.493	0.808	0.580
GCN(ASER)	0.703	0.459	0.822	0.548
GCN(OT)	0.680	0.494	0.821	0.596
GCN(ASER+OT)	0.706	0.492	0.823	0.616
D-GCN	0.709	0.516	0.820	0.647

Table 3: Results on SONYC validate set, OT: SONYC ontology, Mi: Micro, Ma: Macro.

To study the reason for the effectiveness of GCN models in low-resource scenario (1% set) and their degeneration in large-data settings, we divide the classes into groups according to the numbers of training samples, and calculate D-GCN’s improvement over baseline on these groups. From Figure 3, we can see that D-GCN can benefit classes with extremely few samples ([0, 5]), and the gain is the highest on classes with moderate number of samples, but not on the most prevalent classes. We may conclude that the prior knowledge in KG can effectively help the model learn the dependency between labels especially for the few-shot ones. However, as we have more resources, the large backbone model may be capable of learning such relations without KG, which explains why the advantage of GCN-based models would shrink.




Figure 3: Absolute improvement of D-GCN over CNN14 on classes with different number of training samples.

3.2.2 SONYC. The results on the SONYC dataset is shown in Table 3. Similar to AudioSet (1%), all GCN models significantly outperform baseline by the main metric Micro AUPRC. The temporal knowledge of ASER seems to be more useful here compared to ontology, as the ontology for SONYC is more sparse, and the labels for each level are predicted separately, so that they don’t co-occur. D-GCN again gives consistently best or competitive performance on both level, suggesting the generalizability of this method on effectively combining the strength of two knowledge types.

4 CONCLUSION

We investigated the efficacy of KG-enhanced models for low-resource audio tagging. We proposed a semi-automatic procedure to build temporal knowledge graph in multiple domains to enrich existing ontologies. We further proposed D-GCN to effectively combine knowledge of two distinct types. Results on AudioSet and SONYC showed that GCN-based model with the introduced temporal knowledge can significantly outperform baseline without prior knowledge especially in low-resource settings, and D-GCN model with combined knowledge can provide further improvement over models with single type of knowledge.

REFERENCES

- [1] Juan P Bello, Claudio Silva, Oded Nov, R Luke Dubois, Anish Arora, Justin Salamon, Charles Mydlarz, and Harish Doraiswamy. 2019. Sonyc: A system for monitoring, analyzing, and mitigating urban noise pollution. *Commun. ACM* 62, 2 (2019), 68–77.
- [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [3] Yanlin Feng, Xinyue Chen, Bill Yuchen Lin, Peifeng Wang, Jun Yan, and Xiang Ren. 2020. Scalable multi-hop relational reasoning for knowledge-aware question answering. *arXiv preprint arXiv:2005.00646* (2020).
- [4] Eduardo Fonseca, Shawn Hershey, Manoj Plakal, Daniel PW Ellis, Aren Jansen, and R Channing Moore. 2020. Addressing missing labels in large-scale sound event recognition using a teacher-student framework with loss masking. *IEEE Signal Processing Letters* 27 (2020), 1235–1239.
- [5] Frederic Font, Gerard Roma, and Xavier Serra. 2018. Sound sharing and retrieval. In *Computational analysis of sound scenes and events*. Springer, 279–301.
- [6] Zhouyu Fu, Guojun Lu, Kai Ming Ting, and Dengsheng Zhang. 2010. A survey of audio-based music classification and annotation. *IEEE transactions on multimedia* 13, 2 (2010), 303–319.
- [7] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 776–780.
- [8] Arindam Jati, Naveen Kumar, Ruxin Chen, and Panayiotis Georgiou. 2019. Hierarchy-aware loss function on a tree structured label space for audio event detection. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6–10.
- [9] Haozhe Ji, Pei Ke, Shaohan Huang, Furu Wei, Xiaoyan Zhu, and Minlie Huang. 2020. Language generation with multi-hop reasoning on commonsense knowledge graph. *arXiv preprint arXiv:2009.11692* (2020).
- [10] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D Plumbley. 2020. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28 (2020), 2880–2894.
- [11] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yong Xu, Wenwu Wang, and Mark D Plumbley. 2019. Cross-task learning for audio tagging, sound event detection and spatial localization: DCASE 2019 baseline systems. *arXiv preprint arXiv:1904.03476* (2019).
- [12] Michael Lesk. 1986. Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In *Proceedings of the 5th annual international conference on Systems documentation*. 24–26.
- [13] Hongsong Li, Kenny Q Zhu, and Haixun Wang. 2013. Data-driven metaphor recognition and explanation. *Transactions of the Association for Computational Linguistics* 1 (2013), 379–390.
- [14] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.
- [15] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European semantic web conference*. Springer, 593–607.
- [16] Harsh Shrivastava, Yfang Yin, Rajiv Ratn Shah, and Roger Zimmermann. 2020. MT-GCN for multi-label audio tagging with noisy labels. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 136–140.
- [17] Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Unsupervised commonsense question answering with self-talk. *arXiv preprint arXiv:2004.05483* (2020).
- [18] Yiwei Sun and Shabnam Ghaffarzadegan. 2020. An ontology-aware framework for audio event classification. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 321–325.
- [19] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: The new data in multimedia research. *Commun. ACM* 59, 2 (2016), 64–73.
- [20] Helin Wang, Yuexian Zou, Dading Chong, and Wenwu Wang. 2020. Modeling Label Dependencies for Audio Tagging With Graph Convolutional Network. *IEEE Signal Processing Letters* 27 (2020), 1560–1564.
- [21] Hongming Zhang, Xin Liu, Haojie Pan, Yangqiu Song, and Cane Wing-Ki Leung. 2020. ASER: A large-scale eventuality knowledge graph. In *Proceedings of The Web Conference 2020*. 201–211.
- [22] Zhixian Zhang, Kenny Q Zhu, Haixun Wang, and Hongsong Li. 2013. Automatic extraction of top-k lists from the web. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE, 1057–1068.