

Ring That Bell: A Corpus and Method for Multimodal Metaphor Detection in Videos

Khalid Alnajjar^{1,2}, Mika Hämmäläinen¹ and Shuo Zhang²

¹University of Helsinki, Finland

²Bose Corporation, USA

firstname.lastname@{helsinki.fi or bose.com}

Abstract

We present the first openly available multimodal metaphor annotated corpus. The corpus consists of videos including audio and subtitles that have been annotated by experts. Furthermore, we present a method for detecting metaphors in the new dataset based on the textual content of the videos. The method achieves a high F1-score (62%) for metaphorical labels. We also experiment with other modalities and multimodal methods; however, these methods did not out-perform the text-based model. In our error analysis, we do identify that there are cases where video could help in disambiguating metaphors, however, the visual cues are too subtle for our model to capture. The data is available on Zenodo.

1 Introduction

Figurative language is a challenging topic for computational modeling as the meaning of a figurative expression is non-compositional and typically very context dependent (see Roberts and Kreuz 1994). Metaphor is one of the most important figures of language; it is constantly used in every day language (Steen et al., 2010a) to draw comparisons or to express something difficult and foreign in more familiar terms. Metaphors can be conventional (Traugott, 1985) and they are often found in idioms, but at the same time metaphors are used to create something new (see Kantokorpi et al. 1990).

Given its ubiquitous presence, understanding metaphors is integral in achieving true natural language understanding (NLU) in the real world. Without their successful interpretation, our models are bound to make mistakes whenever anything is expressed in an indirect or creative fashion. Metaphors are often very contextual and their successful detection and interpretation requires a wide range of contextual cues that would be captured in audio (e.g., prosody) and video (e.g., gestures and actions). Therefore, we believe a multimodal

dataset is a great contribution to metaphor research within and outside of the field of NLP.

Two important parts of a metaphor are a tenor and a vehicle (see Richards 1936). For example, in the metaphor *life is a journey*, *life* is the tenor and *journey* is the vehicle. How metaphors essentially operate is that a vehicle is used to give some of its attributes to the tenor. In the case above, *journeys* are long and full of adventure, which means that these properties are attributed to *life* in an indirect fashion. The meaning of a metaphor is never literal nor compositional, but rather calls for interpretation on the level of pragmatics (see Rosales Sequeiros 2016).

Meanwhile, multimodality is becoming increasingly important for many tasks (see Castellucci et al. 2020; Mogadala et al. 2020; Declerk et al. 2020). We believe the availability of multimodal datasets for a variety of NLP tasks is lacking, and we hope to contribute to the community with our multimodal metaphor dataset.

In this paper, we present the first fully open expert annotated multimodal dataset for metaphor detection¹. In addition, we experiment with unimodal and multimodal methods for metaphor detection. Our results indicate that the text-based model achieved the best performance. We discuss the results of our experiments and conduct an extensive error analysis to shed light on what was learned successfully by the model and its shortcomings.

Using CC BY licensed videos in our corpus has been the primary design principle of our data collection so that we can release our corpus without any restrictions in its entirety. This, we believe, is more useful for research purposes than a corpus consisting of short video clips to compile with copyright laws such as the fair use law in the US.

¹<https://doi.org/10.5281/zenodo.7217991>

2 Related Work

Metaphors have, thus far, been computationally detected using only text. In this section, we describe some of the recent approaches for textual metaphor detection, the corpora used to achieve that and some of the multimodal research conducted on NLP tasks other than metaphor detection. There are several takes on metaphor interpretation (Xiao et al., 2016; Rai et al., 2019; Bar et al., 2020) and generation (Hämäläinen, 2018; Terai and Sugyo, 2019; Zheng et al., 2019), but we do not describe them in detail as interpretation is a very different problem.

There are two corpora currently used for metaphor detection, VU Amsterdam (VUA) Metaphor Corpus (Steen et al., 2010b) and Corpus of Non-Native Written English Annotated for Metaphor (Beigman Klebanov et al., 2018). Unlike our corpus, both of these datasets contain textual modality only.

For textual metaphor detection, Gao et al. (2018) has used a bi-directional LSTM (long short-term memory) based model with ELMo embeddings. Similarly, Liu et al. (2020) have used a bi-LSTM model with BERT and XLNet for the same task. Not unlike the previous approaches, Dankers et al. (2020) has also applied bi-LSTM models comparing ELMo and GloVe embeddings to BERT embeddings with global and hierarchical attention models. Traditional machine learning methods, Logistic Regression, Linear SVC (Support Vector Classification) and Random Forest Classifier, have been used recently with feature engineering to detect metaphors (Wan et al., 2020). In DeepMet, proposed by Su et al. (2020), a siamese neural network have been utilized, where textual RoBERTa (Liu et al., 2019) embeddings are computed from the context, the token in question and its part-of-speech and fine-grained part-of-speech. DeepMet was the best performing solution for detecting textual metaphors in the VUA dataset, based on a recent shared task (Leong et al., 2020).

There are several recent works on multimodal detection of a variety of linguistic phenomena. For example, SVMs (Support Vector Machines) with word embeddings and feature extraction have been used for multimodal sarcasm detection (Castro et al., 2019; Alnajjar and Hämäläinen, 2021). Mittal et al. (2020) uses GloVe embeddings, features extracted from audio and facial recognition system output to predict emotion in a multimodal dataset.

These multimodal features are fused using a memory fusion network (MFN) (Zadeh et al., 2018). Similarly, Li et al. (2021) detect emotion in a multimodal dataset by modeling the problem from the point of view of the quantum theory. While the field has seen increasing research on multimodal NLP (Tsai et al., 2019; Mai et al., 2020; Sahu and Vechtomova, 2021), no data or model has been proposed for multimodal metaphor detection.

3 Our Metaphor Corpus

In this section, we present our video, audio and textual corpus of manually annotated metaphorical language. Our selection of the video clips includes only CC-BY licensed videos on YouTube that have human authored closed captions in English. The content of the videos presents mainly real people talking, which rules out animations and video game streams. The availability of human authored closed captioning is important as it speeds up our annotation time and provides us with subtitles that are already aligned with video and audio. The CC-BY license was an important selection criterion because it makes it possible for us to release the dataset openly.

We used the filters provided by YouTube to limit our search to videos that were marked as CC-BY and had closed captioning. However, the YouTube filter does not distinguish between automatically generated closed captioning and a human authored one. Fortunately, it is relatively easy to tell these two apart from each other. Automated closed captioning tends to appear one word at a time, whereas human authored closed captioning is visualized more like traditional subtitles. These criteria greatly reduced the number of eligible videos to include in our corpus. Apart from these criteria, we also filtered videos with sensitive and offensive languages. No further restrictions have been explicitly placed on the genres or types of videos, as we do not want to introduce biases for which types of contents are more likely to contain metaphors. Therefore, the availability of the metaphors naturally occurring in the corpus is the result of the ubiquity of the metaphor in everyday language use. All Youtube queries were conducted in incognito mode to avoid biased YouTube suggestions based on our viewing habits.

Figure 1 shows real examples from our corpus where video can be useful in detecting metaphors. On the left, the woman wearing a gray shirt is

sentence
that you can use to really up your <v>game</v>
because while a <t>quick fix</t> can be <v>appetizing</v> and appealing
<t r="domain name">That</t>'s <v>the street address</v> for your website
you're ready to <v>give it a shot</v>

Table 1: Example of the annotations for the metaphor detection corpus.



Figure 1: Metaphors made visible in the video through gestures.

talking about *sprinkling keywords* and showing a sprinkling gesture. On the right, the woman wearing the wine red shirt says *ring that bell* and shows a bell ringing gesture.

Our corpus consists of 27 YouTube videos with a total duration of 3 hours, 53 minutes and 47 seconds of video. For comparison, a recently released multimodal dataset for sarcasm detection (Castro et al., 2019) has the duration of 3 hours, 40 minutes and 47 seconds. The videos belong mostly to a start-up domain and many of them deal with issues of online visibility for a start-up company. This domain was a consequence of our selection criteria for videos. It turns out that YouTube has plenty of high-quality human close-captioned videos released under the CC-BY license that relate to this particular domain.

Our corpus provides linguistics researchers with the ability to study the use of metaphor in a multimodal setting, something that has gained attention in their field of science as well (Müller and Cienki, 2009). This can, indeed, foster a wider interdisciplinary collaboration leading to a deeper understanding of the phenomenon.

3.1 Annotation

Two expert annotators went through the video files and annotated metaphors by surrounding them with *v* tags for vehicles and *t* tags for tenors. The use of experts is motivated by the fact that previous research has found that non-expert annotators struggle with metaphors (Hämäläinen and Alnajjar, 2019).

The annotators followed a simple procedure in

annotating the data:

- Is the meaning literal?
- If the meaning of the word is abstract, is it a dictionary meaning?
- Does the potential metaphor express pragmatic insincerity?
- If the answer to all of the questions is no, annotate it as a metaphor.

In other words, if the meaning of a word or a phrase is not literal, it is annotated as a metaphor. However, just the mere fact of a word being used in an abstract way is not enough to mark it as metaphorical. For example, in the sentence *it is tied to revenue*, “tied” is not tagged as a metaphor just because it is used in a more abstract sense than the typical concrete sense of tying one’s shoes, for example. If the abstract meaning of a word appears in a dictionary, the word is not considered metaphorical. However, conventional metaphors that consist of multiple words, and are thus idioms, are tagged as metaphors. We do not make a distinction between metaphors and similes.

Pragmatic insincerity (see Grice 1975) is a phenomenon related to sarcasm as one of its preconditions (see Kumon-Nakamura et al. 1995). There is a certain overlap between metaphors and sarcastic expressions in the sense that both use words in their non-literal meaning. In order to ensure that we do not mix these two notions with each other, it is important to avoid annotating pragmatically insincere expressions as metaphorical.

Table 1 shows an example of annotations. The annotations were done directly in the subtitles. The utterances are time stamped and aligned with the video. In the table, tenors are indicated with <t> and vehicles with <v>. For deictic tenors, an *r* attribute is provided to resolve the deixis by indicating the actual tenor that has appeared earlier in the conversation. In the examples, *game* is used metaphorically to talk about marketing, *quick fix* is

called *appetizing* as though it was something edible and *domain name* is contrasted to a physical *street address* by direct comparison. *Give it a shot* is a conventional metaphor.

All in all, after multiple annotation iterations, the dataset consists of 304 vehicles and 67 tenors. This totals to 371 metaphorical expressions. They vary in length: the shortest tenor is one word, such as *it*, while the longest tenor is several words *the discovery of those five noble gases to illuminate like that*. The same goes for vehicles where their length varies from one word such as *dive* to multiple words: *the history of the internet itself*. On a token level, we have 672 vehicle tokens and 113 tenor tokens, so altogether 785 metaphorical tokens.

In total, 6% of the expressions in the corpus are metaphorical. While this percentage might appear low, it is natural and more representative of the real usage of metaphors in typical conversations which makes this corpus suitable for building metaphor detection models applicable for real-world scenarios.

Around 55% of the vehicles are conventional metaphors and 45% are novel metaphors. However, it is fairly common that same words appear in the corpus in a metaphorical and non-metaphorical sense. In our corpus, there are two videos that deal with actual cooking, in which many food-related metaphors appear non-metaphorically, such as *sprinkle those in*, said metaphorically about keywords and *a little sprinkle*, said non-metaphorically about sugar. Another example is the use of *house* non-metaphorically as in *come pick it up at my house* and metaphorically as in *think of hosting as your house*, where a metaphorical connection is drawn between *hosting* and a *house*.

3.2 Data preparation

As YouTube serves files in several different formats such as *webm*, *mkv* and *mp4* the first step is to use FFmpeg² to convert all videos into mp4 format. We also use the same tool to clip the video files into sentence-length clips based on the time stamps in the subtitles and extract their audio into wav files. This process yielded 6,565 video and audio clips that are aligned with text.

We split the dataset randomly so that 70% of sentences that contain metaphors and 70% of sentences that don't contain any metaphors are used for train-

²<https://ffmpeg.org/>

ing, 15 % of both types of sentences for validation and 15% of both for testing. This way we ensure that both metaphorical and non-metaphorical sentences are divided proportionally with the same ratios. These splits are used for all the models.

4 Metaphor Detection

We experiment with uni- and multi-modal models for metaphor detection. In this section, we describe the preprocessing steps applied and the experimental setups conducted.

4.1 Preprocessing

For each modality, we make use of the latest advances in neural network models to capture important features that have achieved state-of-the-art results in various NLP tasks. As metaphor detection has been conducted solely based on text, we follow the DeepMet approach by Su et al. (2020) and process the entire textual content using spaCy (Honnibal et al., 2020) to tokenize it and acquire Universal Dependencies style syntactic trees (Nivre et al., 2020) and Penn Treebank parts-of-speech tags (Santorini, 1990). Similarly to the original approach, all of our textual models predict metaphors at the token level given the context surrounding it and its POS tags as input.

We resample the audio to 16kHz. Audio features are extracted using *Wav2Vec2FeatureExtractor* provided by the Transformers Python library (Wolf et al., 2020).

Video features are obtained by taking equally-distributed 16 frames from a clip and then resize them into 128x171, followed by normalization and center cropping to 112x112.

4.2 Textual model

We train two text-only models, both follow the architecture and approach of DeepMet where we obtain textual embeddings using RoBERTa (Liu et al., 2019) and feed them into two transformer encoding layers which are then combined by applying global average pooling and concatenation. A dense fully-connected layer takes in the combined output of both encoders and predicts whether the token is metaphorical (c.f., Su et al. 2020 for more details).

In our first textual model, we train the model using our corpus, whereas in the second one we train it using VUA corpus (with a learning rate of

0.00001, akin the original paper) and later fine-tune it using our corpus.

4.3 Audio model

We extend and fine-tune Facebook’s pretrained multilingual XLSR-Wav2Vec2 large model (Baevski et al., 2020). The model is trained on Multilingual LibriSpeech (Pratap et al., 2020), CommonVoice (Ardila et al., 2020) and Babel (Roach et al., 1996) for speech recognition. We employ this model to encode speech into vector representations from raw audio.

We replace the classification layer of the original model with a dense fully-connected layer that produces two outputs, one for each label. Unlike the textual model, here we classify whether the entire spoken expression contains a metaphor or not (i.e., not on a word level).

4.4 Video model

For our video unimodal model, we incorporate a pretrained model for human action detection. The model is based on the 18 layer deep R(2+1)D network (Tran et al., 2018) and it is trained on the Kinetics-400 (Zisserman et al., 2017) dataset. The intuition behind using this model is that it was able to detect actions (e.g., playing organ), gestures (e.g., pointing) and movements (e.g., waving). Realizing such information is crucial in understanding the context, and would provide further cues for detecting metaphors.

Similar to the audio model, we substitute the original classification layer with a fully connected layer and fine-tune the pretrained model to predict whether a scene is metaphorical or not.

4.5 Multimodal metaphor detection

We test out three multimodal metaphor detection models; 1) text and audio, 2) text and video and 3) text, audio and video. The textual model is the fine-tuned model using the VUA corpus and our textual corpus.

In all of the models, the final classification layer of their sub-models are removed. Unimodal models are combined by concatenating the weights of their last layer, which are then fed to a classification layer.

4.6 Common configuration

All of the models described above share common configurations, unless we explicitly indicate otherwise. Prior to the last classification layer of all of

our mono- and multimodal models, we introduce a dropout layer (Srivastava et al., 2014) (with a probability of 20%) to accelerate training, and reduce internal covariate shift and overfitting.

We use the cross entropy loss function along with Adam optimizer (Kingma and Ba, 2014; Loshchilov and Hutter, 2019) to update the weights and train the models. All the fine-tuned models are trained with a learning rate of 0.0001 and for 3 full epochs.

5 Results

In this section, we follow the evaluation metrics commonly used for the metaphor detection task by reporting the precision, recall and F1 scores for the metaphorical label.

Regarding the textual models, we report three sets of results, which are for the models trained on: 1) VUA corpus, 2) our corpus and 3) both the VUA and our corpus. All the models predict metaphoricality on the token level. To ensure that our implementation of the DeepMet approach is correct, we tested the first model on the VUA test dataset of the metaphor detection shared task and achieved an F1-score of 0.68 and 0.73 on all POS and verb subsets of the data, respectively. These results are relatively close to the results reported by the authors.

Table 2 shows the classification results of all three models on the test set. The test set contained 90 metaphorical tokens and 6,961 non-metaphorical tokens. The results indicate that the textual model trained solely on the VUA dataset performed poorly on our test set. In comparison, training the model using our metaphor corpus only resulted in a great increase of correct predictions. Nonetheless, combining both corpora by fine-tuning the first model with our corpus produced the winning model, which managed to spot 76% of the metaphorical tokens correctly.

We believe that the huge differences between the first and second textual models, despite the larger size of VUA’s training dataset, are due to the differences in domains. The VUA corpus contains academic texts, conversation, fiction, and news texts, whereas our corpus is dominated by conversations on the web and start-ups. It is evident that by exposing the model to general domains (i.e., VUA’s corpus) and, thereafter, concentrating it on the start-up domain, the model was able to identify the highest number of metaphorical usages.

Trained on	Precision	Recall	F1-score
VUA	0.04	0.33	0.07
Ours	0.38	0.63	0.47
VUA + Ours	0.53	0.76	0.62

Table 2: Classification results of the textual monomodal models on the test set of our corpus, for the metaphorical label.

Results from the other models (unimodal or multimodal) that involving audio and video showed that adding these modalities actually did not help improving the model - rather, they are detrimental to the model performance on metaphor detection. We extend two possible explanations for this failure. First, it is possible that because the visual and audio cues of metaphor are subtle, these models failed to learn from such a small amount of annotated data.

Second, it is unclear that the specific models we are using for audio and video modalities encode the information relevant for the metaphor detection task. For instance, whereas it is impossible to completely disentangle what exactly the Wav2Vec model is encoding, we can conjecture that it encodes information about phoneme identity considering it is optimized for the speech recognition task. Therefore, it may not be entirely surprising that the Wav2Vec encoding is not useful for the metaphor detection task because it is adding redundant or irrelevant information to the model. It is our future work (or the future work for the community who utilizes this dataset) to refine our understanding of the multimodal encoding for the metaphor detection task (for instance, employing a model that more directly encodes information about speech prosody from the audio).

5.1 Error analysis

When looking at the results of the text only model, we can see that the model identifies metaphors correctly as metaphors more often than not. There are some metaphorical tokens in metaphors consisting of multiple words that get classified wrong, for example, in *You could think of hosting as your house*, the tenor *hosting* and the determinant *your* of the metaphorical word *house* are not identified as metaphorical, while *house* is correctly identified. Another example is the conventional metaphor *toot their own horn*, where all other words except for *own* are correctly identified as metaphorical.

There are also a fewer number of cases where all

words get identified wrongly as non-metaphorical, for example, the model did not predict any metaphorical tokens in *It's where you live*, while in reality *it* is the tenor and *where you live* is the vehicle. Also, individual tenors where the vehicle comes later get often not recognized such as in *Yes, malware you could think of like*, where *malware* is the tenor for a vehicle that appears later in the dialog.

When the tenor and the vehicle co-exist nearby, the model can get all metaphorical tokens right such as in *It's kinda like real estate right?* where both the tenor *it* and the vehicle *real estate* are correctly identified. Also many tenorless expressions are fully recognized correctly as metaphorical, such as *Spreadin' the love*.

There were plenty of cases (61) where the model predicted a metaphor tag for a token while there was no metaphor. Curiously, prepositions were often tagged metaphorical, such as *to* in *ring that bell to see these episodes first*. The actual metaphorical part *ring that bell* ends before the preposition *to* that has a non-metaphorical meaning *in order to*.

We can also see that the model was indeed fooled by cooking terms that were used both metaphorically and non-metaphorically. In *Yeah a little sprinkle*, both *a* and *sprinkle* were classified as metaphors, while the context was about sprinkling sugar. Another similar case was *there's five noble gases that illuminate*, where *noble gases* and *illuminate* were erroneously classified to be metaphorical. This was clearly due to the tenor in the corpus: *the discovery of those five noble gases to illuminate like that* that contained similar words. It is evident that the model relies on word similarities more than reaching to a higher pragmatic representation of the phenomenon, however, this is not an unexpected behavior from a machine learning model.

There are also cases where the model detects a metaphor, that could theoretically be a metaphor, but is not because of the way it was used in the corpus. For example, the model predicts *Give it a go* as metaphorical in the expression *button, "Give it a go."*, where people are talking about a button with a particular text rather than using the expression metaphorically. Another such an example is *flying in (money flying)*. Such an expression might be used metaphorically, but in this case this was a note for the hearing impaired as money was actually flying on the video.

6 Discussion and Conclusions

In this work, we have only focused on metaphor as a strictly linguistic phenomenon and we have built a multimodal dataset where these linguistic metaphors have been tagged in terms of tenors and vehicles. However, it is apparent that metaphor is a phenomenon that occurs on a higher level of our cognitive capacities than mere language. There are several cases in our corpus, where we can evidence the existence of a metaphor but it is never expressed verbally. For example in Figure 2, *money flying* cannot be a metaphor when inspected purely from the point of view of language and its relation to the video when money is actually flying in the scene. However, it is a metaphor on a higher level in the sense that the entire scene where money was flying was to indicate someone becoming rich. In other words, stating a fact that is happening is not metaphorical if the fact is literally taking place, however the fact itself might be metaphorical.




Figure 2: Money actually flying on the video.

At the same time, as evidenced by our error analysis, there are certainly cases where video modality could help in disambiguating whether something is said metaphorically or not. For instance, talking about *sprinkling* in a kitchen environment (see Figure 3) is a very strong sign that the word is potentially non-metaphorical. Integrating these weak cues into a multimodal system is, however, not an easy task given that the current methods for video processing are limited in their coverage.

Therefore, in the future, it would be useful to annotate metaphors also in the other modalities. Money flying can be a visual metaphor, and so can a sound effect, and they can exist independently from each other in different modalities. Perhaps the reason why our multimodal attempts failed was that metaphor can be independent of the other modalities. Producing such a dataset where these modal specific metaphors are also annotated for video and



Figure 3: *Sprinkling* used in a kitchen in reference to sugar.

audio is definitely a huge undertaking that requires research in its own right.

It is clear that our model can detect metaphors correctly, but also the mistakes it makes highlight that despite using a large RoBERTa model, the meaning representation the model has cannot reach to such a nuanced level as to confidently detect metaphors. Metaphor is a figurative device that cannot be explained by semantics, but rather requires pragmatic inspection. It is not clear based on our research and other contemporary approaches whether the current word or sentence embedding models are sufficient to navigate in the depths of pragmatics and subjective interpretation in any other way than learning some irrelevant co-occurring phenomena from a biased corpus. At the same time there is no such thing as an unbiased corpus, either, given that bias (and mostly heuristics causing it) is a fundamental part of our cognition as human beings.

In this paper, we have presented a new open and multimodal dataset for metaphor detection. Because we have focused strictly on CC-BY licensed videos, we can make the entire dataset available on Zenodo. In our current work, we have not taken the context widely into account when predicting metaphoricity, but rather resorted to a very local context. The fact that the videos can be published in full length makes it possible for any future work to explore different ways of including contextual cues freely.

Acknowledgments

This work was partially financed by the Society of Swedish Literature in Finland with funding from Enhancing Conversational AI with Computational Creativity, and by the Ella and Georg Ehrnrooth Foundation for Modelling Conversational Artificial Intelligence with Intent and Creativity.

References

- Khalid Alnajjar and Mika Härmäläinen. 2021. ¡Qué maravilla! multimodal sarcasm detection in Spanish: a dataset and a baseline. *arXiv preprint arXiv:2105.05542*.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. **Common voice: A massively-multilingual speech corpus**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4218–4222, Marseille, France. European Language Resources Association.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33.
- Kfir Bar, Nachum Dershowitz, and Lena Dankin. 2020. Automatic metaphor interpretation using word embeddings. *arXiv preprint arXiv:2010.02665*.
- Beata Beigman Klebanov, Chee Wee (Ben) Leong, and Michael Flor. 2018. **A corpus of non-native written English annotated for metaphor**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 86–91, New Orleans, Louisiana. Association for Computational Linguistics.
- Giuseppe Castellucci, Simone Filice, Soujanya Poria, Erik Cambria, and Lucia Specia, editors. 2020. *Proceedings of the First International Workshop on Natural Language Processing Beyond Text*. Association for Computational Linguistics, Online.
- Santiago Castro, Devamanyu Hazarika, Verónica Pérez-Rosas, Roger Zimmermann, Rada Mihalcea, and Soujanya Poria. 2019. Towards multimodal sarcasm detection (an _obviously_ perfect paper). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Florence, Italy. Association for Computational Linguistics.
- Verna Dankers, Karan Malhotra, Gaurav Kudva, Volodymyr Medentsiy, and Ekaterina Shutova. 2020. **Being neighbourly: Neural metaphor identification in discourse**. In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 227–234, Online. Association for Computational Linguistics.
- Thierry Declerk, Itziar Gonzalez-Dios, and German Rigau, editors. 2020. *Proceedings of the LREC 2020 Workshop on Multimodal Wordnets (MMW2020)*. The European Language Resources Association (ELRA), Marseille, France.
- Ge Gao, Eunsol Choi, Yejin Choi, and Luke Zettlemoyer. 2018. **Neural metaphor detection in context**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 607–613, Brussels, Belgium. Association for Computational Linguistics.
- Herbert P Grice. 1975. Logic and conversation. In *Speech acts*, pages 41–58. Brill.
- Mika Härmäläinen. 2018. Harnessing nlg to create finnish poetry automatically. In *Proceedings of the ninth international conference on computational creativity*. Association for Computational Creativity (ACC).
- Mika Härmäläinen and Khalid Alnajjar. 2019. **Let’s FACE it. Finnish poetry generation with aesthetics and framing**. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 290–300, Tokyo, Japan. Association for Computational Linguistics.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. **spaCy: Industrial-strength Natural Language Processing in Python**.
- Mervi Kantokorpi, Lyytikäinen Pirjo, and Viikari Auli. 1990. *Runousopin perusteet*. Gaudeamus.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Sachi Kumon-Nakamura, Sam Glucksberg, and Mary Brown. 1995. How about another piece of pie: The allusional pretense theory of discourse irony. *Journal of Experimental Psychology: General*, 124(1):3.
- Chee Wee (Ben) Leong, Beata Beigman Klebanov, Chris Hamill, Egon Stemle, Rutuja Ubale, and Xinyang Chen. 2020. **A report on the 2020 VUA and TOEFL metaphor detection shared task**. In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 18–29, Online. Association for Computational Linguistics.
- Qiuchi Li, Dimitris Gkoumas, Christina Lioma, and Massimo Melucci. 2021. Quantum-inspired multimodal fusion for video sentiment analysis. *Information Fusion*, 65:58–71.
- Jerry Liu, Nathan O’Hara, Alexander Rubin, Rachel Draelos, and Cynthia Rudin. 2020. **Metaphor detection using contextual word embeddings from transformers**. In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 250–255, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *International Conference on Learning Representations*.

- Sijie Mai, Haifeng Hu, and Songlong Xing. 2020. Modality to modality translation: An adversarial representation learning and graph fusion network for multimodal fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 164–172.
- Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. 2020. M3er: Multiplicative multimodal emotion recognition using facial, textual, and speech cues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1359–1367.
- Aditya Mogadala, Sandro Pezzelle, Dietrich Klakow, Marie-Francine Moens, and Zeynep Akata, editors. 2020. *Proceedings of the Second Workshop on Beyond Vision and Language: inTEgrating Real-world kNowledge (LANTERN)*. Association for Computational Linguistics, Barcelona, Spain.
- Cornelia Müller and Alan Cienki. 2009. *Chapter 13. Words, gestures, and beyond: Forms of multimodal metaphor in the use of spoken language*, pages 297–328. De Gruyter Mouton.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajic, Christopher D Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. Universal dependencies v2: An evergrowing multilingual treebank collection. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4034–4043.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. MIs: A large-scale multilingual dataset for speech research. *ArXiv*, abs/2012.03411.
- Sunny Rai, Shampa Chakraverty, Devendra K Tayal, Divyanshu Sharma, and Ayush Garg. 2019. Understanding metaphors using emotions. *New Generation Computing*, 37(1):5–27.
- Ivor Armstrong Richards. 1936. *The Philosophy of Rhetoric*. Oxford University Press, London, United Kingdom.
- Peter Roach, Simon Arnfield, W Barry, J Baltova, Marian Boldea, Adrian Fourcin, W Gonet, Ryszard Gubrynowicz, E Hallum, Lori Lamel, et al. 1996. Babel: An eastern european multi-language database. In *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP'96*, volume 3, pages 1892–1893. IEEE.
- Richard M Roberts and Roger J Kreuz. 1994. Why do people use figurative language? *Psychological science*, 5(3):159–163.
- Xose Rosales Sequeiros. 2016. Metaphor: Pragmatics, relevance and cognition. *English Studies*, 97(6):656–677.
- Gaurav Sahu and Olga Vechtomova. 2021. [Adaptive fusion techniques for multimodal data](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3156–3166, Online. Association for Computational Linguistics.
- Beatrice Santorini. 1990. Part-of-speech tagging guidelines for the penn treebank project (3rd revision). *Technical Reports (CIS)*, page 570.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15(56):1929–1958.
- GJ Steen, AG Dorst, JB Herrmann, AA Kaal, and T Krennmayr. 2010a. Metaphor in usage. *Cognitive Linguistics*, 21(4):757–788.
- GJ Steen, AG Dorst, JB Herrmann, AA Kaal, T Krennmayr, and T Pasma. 2010b. A method for linguistic metaphor identification. from mip to mipvu. *Converging Evidence in Language and Communication Research*, (14).
- Chuangdong Su, Fumiyo Fukumoto, Xiaoxi Huang, Jiyi Li, Rongbo Wang, and Zhiqun Chen. 2020. [DeepMet: A reading comprehension paradigm for token-level metaphor detection](#). In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 30–39, Online. Association for Computational Linguistics.
- Asuka Terai and Taiki Sugyo. 2019. Construction of a corpus-based metaphor generation support system built on japanese literature. In *2019 IEEE 11th International Workshop on Computational Intelligence and Applications (IWCIA)*, pages 1–6. IEEE.
- Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. 2018. [A closer look at spatiotemporal convolutions for action recognition](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6450–6459.
- Elizabeth Closs Traugott. 1985. ‘conventional’ and ‘dead’ metaphors revisited. *The ubiquity of metaphor: Metaphor in language and thought*, pages 17–56.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access.
- Mingyu Wan, Kathleen Ahrens, Emmanuele Chersoni, Menghan Jiang, Qi Su, Rong Xiang, and Chu-Ren Huang. 2020. [Using conceptual norms for metaphor detection](#). In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 104–109, Online. Association for Computational Linguistics.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ping Xiao, Khalid Alnajjar, Mark Granroth-Wilding, Kat Agres, and Hannu Toivonen. 2016. Meta4meaning: Automatic metaphor interpretation using corpus-derived word associations. In *Proceedings of the Seventh International Conference on Computational Creativity*. Sony CSL Paris.
- Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Memory fusion network for multi-view sequential learning. *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Danning Zheng, Ruihua Song, Tianran Hu, Hao Fu, and Jin Zhou. 2019. “love is as complex as math”: Metaphor generation system for social chatbot. In *Workshop on Chinese Lexical Semantics*, pages 337–347. Springer.
- Andrew Zisserman, Joao Carreira, Karen Simonyan, Will Kay, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, and Mustafa Suleyman. 2017. The kinetics human action video dataset.