

DSC IIT-ISM at SemEval-2020 Task 8: Bi-Fusion Techniques for Deep Meme Emotion Analysis

Pradyumna Gupta*, Himanshu Gupta* and Aman Sinha

Indian Institute of Technology (Indian School of Mines) Dhanbad, India
{pradyumna.gupta22, hg24091996, amansinha091}@gmail.com

Abstract

Memes have become an ubiquitous social media entity and the processing and analysis of such multimodal data is currently an active area of research. This paper presents our work on the Memotion Analysis shared task of SemEval 2020, which involves the sentiment and humor analysis of memes. We propose a system which uses different bimodal fusion techniques to leverage the inter-modal dependency for sentiment and humor classification tasks. Out of all our experiments, the best system improved the baseline with macro F1 scores of 0.357 on Sentiment Classification (Task A), 0.510 on Humor Classification (Task B) and 0.312 on Scales of Semantic Classes (Task C).

1 Introduction

Social media is becoming increasingly abundant in multimedia content with internet memes being one of the major types of such content circulated online. Internet memes, or simply, memes are mostly derived from news, popular TV shows and movies making them capable of conveying ideas that the masses can understand readily. Today memes are shared in numerous online communities and are not only limited to one language. It is therefore imperative to study this emerging form of mass communication. Although they are mostly for amusement and satire, memes are also being used to propagate controversial, hateful and offensive ideologies. Since it is not feasible to manually inspect such content it is important to build systems that can process memes and segregate them into appropriate categories.

The SemEval 2020 shared task on Memotion Analysis (Sharma et al., 2020) aims to bring attention to the study of sentiment and humor conveyed through memes. The challenge involves 3 tasks:

- Task A involves classifying the sentiment of an internet meme as positive, negative or neutral.
- Task B is the multilabel classification of a meme into 4 categories viz. sarcastic, humorous, offensive and motivational.
- Task C is to quantify, on a scale of 0-3 (0 = not, 1 = slightly, 2 = mildly, 3 = very), the extent of sarcasm, humor and offence expressed in a meme.

Here, we consider Task B as a set of 4 independent binary classification subtasks while Task C comprises 3 subtasks each being a multiclass classification problem.

This paper describes our work on all three tasks. Here, a meme has two modalities - textual and visual cues. We use a combination of state-of-the-art model architectures and adapt them for processing multimodal memes through transfer learning while also training models from scratch for comparison. We explore various modality fusion techniques and also the inter-task dependency. Our systems[†] score higher than the baseline results released on all the 3 tasks. We investigate our model performances from relative and absolute perspectives and also explore some possible improvements. We make our code publicly available.

* Equal contribution.

† Submissions have been made under the username of *hg* on CodaLab.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

2 Related Work

Discussing the sentiments of memes has been a relatively new research area. Although interest in the research related to meme generation, circulation and perception has been increasing in the last couple of years where studies of Vyalla and Udandarao (2020), Wang and Wen (2015) have shown pioneering work, studies exclusively analysing meme emotion are relatively scarce.

With the explosion of social media in recent years and human generated content on them like tweets, memes etc., deep learning techniques have captured the attention of researchers due to their ability to significantly outperform traditional methods in sentiment analysis tasks. Studies based on discussing textual segments (Zhang et al., 2018) and based on other channels like images and videos (Poria et al., 2017) have been in progress. Visual sentiment and emotion analysis, although having received less attention than text based analysis, has involved studies of Borth et al. (2013), You et al. (2017) and Sinha et al. (2019) where features extracted from images using deep learning have been used for predictions.

Apart from the unimodal methods we have seen so far, the multimodal methods involving simultaneous processing of textual and visual information are of special interest, as studies by Guillaumin et al. (2010) and Zahavy et al. (2016) have reported improvement in performance by combining modalities. Multimodal processing of features using hierarchical methods (Majumder et al., 2018; Hu et al., 2018) and attention based methods (Huang et al., 2019; Moon et al., 2018) have shown to be effective fusion techniques and have been used by teams in previous iterations of SemEval as well. Multitask learning for processing multimodal information has also been studied by works like Akhtar et al. (2019), where the authors have tried to capture inter-dependence between multimodal tasks like emotion recognition and sentiment analysis. In addition to above, processing of multimodal information specifically for memes have also been studied by some papers which have greatly inspired our work, including papers by French (2017), which discusses correlation between the implied semantic meaning of image-based memes and the textual content of discussions in social media and also Prajwal et al. (2019), which have developed models to generate rich face emotion captions from memes.

3 Data

The dataset (Sharma et al., 2020) consists of almost 7K meme images along with their extracted text. Each sample has a sentiment annotation into one of 3 classes- negative (0), neutral (1) or positive (2). There are another 3 semantic classes namely humor, sarcasm and offence with their fine grained labels as not (0), slightly (1), mildly (2) and very (3). Also there is a binary column denoting whether a meme is motivational. Another set of around 2K samples form the test set. We split the dataset into train and validation sets of around 6K and 1K samples respectively. From the data distribution in Table 1, we can see that all the labels in the dataset have considerable class imbalance.

Label	Train Set				Validation Set			
	0	1	2	3	0	1	2	3
Sentiment	519	1894	3530	-	112	307	630	-
Humor	1400	2060	1952	531	251	392	286	120
Sarcasm	1304	2983	1323	333	240	524	224	61
Offence	2289	2203	1259	192	424	389	207	29
Motivation	3844	2099	-	-	681	368	-	-

Table 1: Dataset label distributions for the train and validation sets.

4 Methodology

A meme typically consists of an image paired with a catchphrase on it. The image is generally derived from some popular source and is reused as a template to represent an underlying theme. The text on the

other hand is usually original and specific, written by the creator, and it is what makes the meme unique. Together they convey some sentiment and emotion which we study here. In this task we focus on three main approaches:

- Textual, based on the text part of the meme.
- Visual, based on the image of the meme.
- Combined or fused, unifying the textual and visual inputs in a single end-to-end model.

A more detailed explanation of our approaches can be found in the following subsections.

4.1 Textual Features

Similar to textual content on social media from other sources like tweets and comments, the textual part of a meme also contains important cues that can help extract the sentiment and emotion of the meme. Since the textual part of a meme gives it some specific context, it can sometimes be more detrimental in analysing the sentiment or emotion behind a meme. Therefore, as the first approach in our work we analyse some of the text based models to solve the three tasks.

BiLSTM: Long Short-Term Memory (LSTM), (Hochreiter and Schmidhuber, 1997) is a widely known recurrent neural network (RNN) architecture. We use Bidirectional LSTM (Schuster and Paliwal, 1997) models for our experiments. A Bidirectional LSTM (BiLSTM) layer processes the text in the forward as well as backward direction and hence is known to give better context understanding.

RoBERTa: Bidirectional Encoder Representations from Transformers, or BERT (Devlin et al., 2018) has been a revolutionary language model for a while now which is based on the self-supervised pretraining technique and has shown to generalize extremely well to downstream tasks. The recently released RoBERTa - A Robustly Optimized BERT Pretraining Approach (Liu et al., 2019), builds on BERT's language masking strategy and modifies key hyperparameters in BERT, including removing BERT's next-sentence pretraining objective, and training with much larger mini-batches and learning rates. RoBERTa has also been trained on a much larger text corpus than BERT, for a longer amount of time. This allows RoBERTa representations to generalize even better to downstream tasks compared to BERT. Here we use Roberta-Base to encode the textual part of the meme.

4.2 Visual Features

It is well known that the human brain prefers and is faster at processing visuals as compared to reading text. This is mainly because the entire image is perceived at once while text has to be processed linearly. On seeing a meme, the image is the part one perceives first and so it can provide a generic context on top of which the text part acts. We extract the visual representation of images in the form of abstract as well as specific features.

AlexNet: AlexNet (Krizhevsky et al., 2012) is a widely known convolutional neural network (CNN) which consists of convolutional and fully connected layers and is one of the first networks to use rectified linear unit as its activation function. We do not use a pretrained model and train from scratch on the dataset.

ResNet: ResNet (He et al., 2015), short for Residual Networks is a classic neural network used as a backbone for many computer vision tasks. ResNet makes use of the skip connection to add the output from an earlier layer to a later layer thereby helping it mitigate the vanishing gradient problem. For the sake of experiments reported in this paper we use ResNet with its ImageNet (Deng et al., 2009) pretrained weights.

Facial Expressions: On deeper exploratory analysis of the given dataset we found that most of the meme images contain faces of one or more people in them. The facial expression of these people might contain useful information about the overall sentiments expressed by the meme. Prajwal et al. (2019) have already implemented an intuitive end-to-end model for generating rich face emotion captions from meme images using a ResNet to generate emotion features which are used by a captioning model to generate caption sequences. Therefore, as a part of our visual approach, for capturing facial expressions in our

dataset we try using their pretrained ResNet block for generating predictions of the current task. However, because of poor performance of this model on the validation set we exclude it from our experiments on the test set.

Face Emotion Captions: Unlike our earlier approach of processing facial emotions in meme images, this time we try incorporating facial expressions as a third channel of information in addition to the image and meme overlay text. This is done by building our textual models on top of the text captions generated by the Prajwal et al. (2019) model and then treating these textual models as a third input in our fusion approaches, described later. It is also worthwhile to mention here that for the memes that did not contain faces, random captions were generated by the model. However again due to poor performance of these models on the validation set we exclude this from our experiments on the test set.

4.3 Multimodal Approach

The above methods consider both modalities individually for representation learning. However, effective computational processing of internet memes needs a combined approach. After extracting the text and image features they need to be combined to get a complete representation of a meme. Figure 1 gives a high level illustration of the modality fusion techniques that we use.

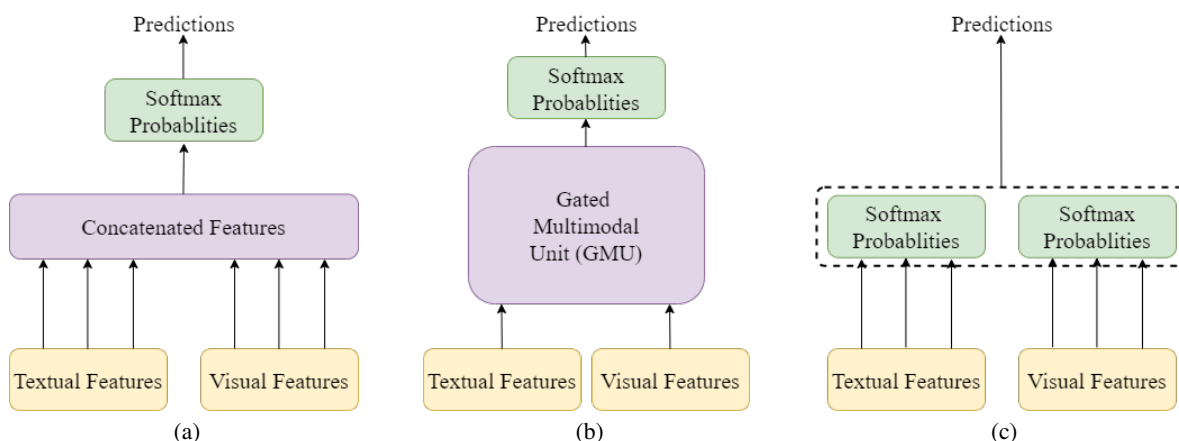


Figure 1: Modality fusion strategies- (a) Early Fusion, (b) GMU and (c) Late Fusion.

Early Fusion: This method involves simply concatenating the textual and visual feature vectors extracted by above mentioned models to produce a joint representation before passing it through a softmax classifier. We do this for the features generated by BiLSTM & AlexNet, and also for RoBERTa & ResNet.

GMU: The Gated Multimodal Unit (Ovalle et al., 2017) is a building block to be used as an internal unit in a neural network architecture whose purpose is to find an intermediate representation based on a combination of data from different modalities. The GMU learns to decide how modalities influence the activation of the unit using multiplicative gates. We use GMU particularly to combine the representations produced by RoBERTa and ResNet.

Late Fusion: In this method, instead of combining text and visual models at the time of training, we fuse the models at inference time by using their individual softmax prediction scores. During our experiments we assign equal weights to text and visual models to take effectively an average of the prediction scores of the models to make the final predictions. We do this to combine the predictions of RoBERTa with ResNet.

4.4 Multitask Learning

In all the above approaches we have to train separate models for each subtask of the three tasks thereby requiring a total of 8 models to be trained per submission. Using multitask learning (MTL) enables us to train a model with multiple outputs which can be used for multiple classification tasks simultaneously. In MTL we have the hidden layers of the network shared among the tasks and also keep some individual

layers for each output. MTL is known to leverage useful information contained in multiple related tasks to help improve the generalization performance of all the tasks (Zhang and Yang, 2017). To this end, we use multitask models with early fusion of features from BiLSTM & AlexNet for Task A and the 4 subtasks of Task B together. We also use a similar model for solving Task A and the subtasks of Task C together.

5 Experiments

Using the techniques mentioned in the previous section we train several models for our submissions. In this section, we provide the training parameters for reproducibility¹ and also their results on the test data.

5.1 Training Details

Except for the MTL model, we train the following models separately for Task A and each of the subtasks of Task B & Task C. We use Keras with Tensorflow backend (Chollet and others, 2015) on Google Colab GPU to implement and train our models. Unless otherwise stated, all models use the default Adam optimizer with the Cross Entropy loss and trained till the validation loss stopped improving. We use a batch size of 32.

BiLSTM: As a preprocessing step, we remove the punctuation and symbols from the text data, convert to lowercase and tokenize to generate sequences with truncated length of 64. We train models with an input layer having embedding size 32, followed by a bidirectional LSTM layer with recurrent dropout and then a softmax dense layer. We keep the number of LSTM units to be 32 as higher values didn't improve performance. Also, we apply a dropout of 0.5 in between the layers. To deal with the class imbalance, we use appropriate class weights for loss during training.

RoBERTa: For implementation details of RoBERTa-Base readers can refer to Liu et al. (2019). Since transformer based models can handle long sequences we take the input sequence length for RoBERTa equal to the length of the longest text which was 199. Apart from preprocessing steps like lower casing, punctuation and symbol removal, we use oversampling to make the target class ratio 1 for handling class imbalance as the model tends to get biased towards one class on the given dataset. We fine tune the pretrained model on the given dataset using the output corresponding to the CLS token of the last layer followed by a batch normalization layer, a dense layer of hidden size 256 and finally a softmax layer. The model is trained with a learning rate of $3e-4$.

AlexNet: As mentioned in the original paper we resize the images to 224 x 224 with 3 channels as input to the model. The filter sizes of the convolutional layers follow [96, 128, 128, 256, 256]. After flattening we apply 2 dense layers of size 1024 each before the final softmax layer. We apply a dropout of 0.4 after each layer and use class weights to handle imbalance.

ResNet: Before passing the meme images to the model all the images are resized to 256 x 256, rescaled and all the three RGB channels of the images are retained. We use the pretrained ImageNet weights to initialize the model and the outputs are followed by dense layers of hidden sizes [768, 256] and a softmax layer. A learning rate of $3e-4$ is used to train the model.

BiLSTM+AlexNet: After using the same process and architecture as mentioned in their individual training details, we replace the softmax layer in the BiLSTM and the AlexNet model with a dense layer of 64 units to get the text and image vector representations. Then we pass both through a shared dense layer of size 64 and concatenate (Early Fusion) the outputs to form a joint vector of length 128. Then we keep another dense layer of 128 units before the softmax prediction layer. Here also we employ class weight balancing.

RoBERTa+ResNet: For the pretrained models we have explored three ways of fusion of information, as mentioned earlier, viz Early Fusion, GMU and Late Fusion. For Early Fusion (Figure 1(a)), we batch normalize the outputs of the models, concatenate them and finally pass them through a softmax layer. For GMU (Figure 1(b)), we take a bimodal GMU unit which is connected to the two models as input and has a softmax layer connected at its output. Finally, for Late Fusion (Figure 1(c)), we take the predictions of the two earlier separately trained models and average their prediction probabilities to get the final predictions. Since both the pretrained architectures are relatively deeper neural networks we use SGD optimizer while

¹Link to our code- <https://github.com/dscitism/SemEval-2020-Task-8>

training the Early Fusion and GMU models to prevent gradient vanishing and oversampling to prevent the model from getting biased.

MTL: For the multitask learning approach we use the BiLSTM+AlexNet architecture, with slight modifications, and train from scratch. We train 2 models here, one for Task A with Task B and another for Task A with Task C. To handle the increased complexity of the model objective we add another bidirectional LSTM layer and double the vector size for each modality to 128. We also increase the size of the joint vector to 256 which is fed into separate 128 unit dense layers for each model task followed by individual softmax layers. For training the model we keep the loss contribution or weights equal for all outputs. We also try using the Task B semantic labels for auxiliary tasks with Task A as the main objective but it doesn't give any improvements.

Model	Modality	Fusion Strategy	Task A	Task B	Task C
Baseline Results	-	-	0.218	0.500	0.301
BiLSTM	Text	-	0.319	0.502	0.290
RoBERTa	Text	-	0.314	0.494	0.276
AlexNet	Image	-	0.309	0.490	0.302
ResNet	Image	-	0.324	0.508	0.312
BiLSTM+AlexNet	Text & Image	Early Fusion	0.323	0.495	0.291
BiLSTM+AlexNet (MTL)	Text & Image	Early Fusion	0.322	0.495	0.267
RoBERTa+ResNet	Text & Image	Early Fusion	0.357	0.510	0.306
RoBERTa+ResNet	Text & Image	GMU	0.346	0.503	0.303
RoBERTa+ResNet	Text & Image	Late Fusion	0.321	0.508	0.312

Table 2: Model details and their macro F1 scores on the test set. For Task B and Task C, the macro F1 scores for each subtask were averaged.

5.2 Results

The evaluation metric used here is the macro F1 score for Task A. For Task B and Task C, the macro F1 scores for each subtask were averaged. Table 2 has the model test set scores from our experiments as well as the baseline results released by the competition organisers. From the table we can see that the RoBERTa+ResNet model with Early Fusion has the best performance for Task A and Task B while the same model using the Late Fusion strategy has the best score for Task C. For the BiLSTM+AlexNet model, using MTL did not give any improvements. We find that among the 3 fusion strategies, Early Fusion scores best for Task A and Task B, Late Fusion gives the best performance on Task C while GMU gives intermediate results. We chose the BiLSTM+AlexNet model for our contest submission as it gave the best results during the development phase. However, after final evaluation the RoBERTa+ResNet models were found to be better.

6 Discussion

In order to justify the results and to discuss some observations we have about the dataset and memes in general, we present a detailed analysis in this section. We investigate why the state-of-the-art pretrained models could not give the superior results that is expected from them and highlight the areas of potential improvement.

Text Arrangement: Apart from deep learning models we also try the TF-IDF model with classical machine learning classifiers like logistic regression. Although the deep learning models perform better, the validation set score difference is not large. On examining the data we find examples in which the text is scattered around the canvas and since an OCR is used for preparing the dataset, it ended up giving jumbled phrases in an order which is different from what is intended by the maker of the meme. In one

such example², the image had an overlay phrase “*CHALLENGE ACCEPTED!*” and under that was the dialogue “*Friend: You can’t honestly watch How I Met Your Mother again for like the 4th time... Me:*”. In this case, the phrase “challenge accepted” will appear at the beginning of the text when read by an OCR whereas in reality it is intended to be perceived at the end of the dialogue when read by a human. Such samples undermine the advantage that sequential models have over frequency based models.

Language Use: For the unimodal BiLSTM model, apart from training the embedding layer parameters from scratch we also use the pretrained GloVe embeddings (Pennington et al., 2014). Although the GloVe embeddings are trained on a much larger corpus than the dataset here, we found that it did not improve the performance of the text model. This can be attributed to the difference in the vocabulary usage of memes over standard English and the intentional misspellings used for common words. This can be seen in a meme³ in the dataset which had the Russian President Vladimir Putin with a smirk on his face and the text reading “*NOT UKRAINE ... MYKRAINE*” to imply “*My Ukraine*”. Moreover, the incorrect grammar used in a lot of memes may also be the reason why RoBERTa, despite being pretrained on a huge English corpus, failed to generalize well over the meme language and scored below BiLSTM for all the three tasks. Using character level embedding approaches to the language model can be helpful in dealing with such cases.

Image Manipulation: For the image models also, in spite of having an edge over AlexNet, the pretrained deep models, Inception (Szegedy et al., 2014) and ResNet, which we try, didn’t outperform it by a large margin. We believe this is because there is very little similarity between the ImageNet dataset and the meme images used here. Memes are usually images that have been edited and also have some text overlay which can act as noise for the network thereby hindering representation learning. A lot of memes in the dataset were found to have considerable editing or using multiple images combined into one along with text annotations on top⁴. To deal with the text overlay in memes, some techniques in image masking can be used as a preprocessing step to replace the text with its background.

Overall, the winning macro-averaged F1 scores in the competition for Task A, Task B and Task C were less than 0.36, 0.52 and 0.33 respectively which shows there is ample room for improvement. Out of our three fusion strategies, simple concatenation (Early Fusion) turned out to be the best. But when we see a meme both text and image rely on each other to jointly convey an idea. We strongly feel that there is a need for a better fusion strategy or learning technique that can generate a unified and coherent representation, which is analogous to how a human brain processes memes. For the image captioning task, there has been some work done regarding inter-modal correspondences between language and visual data (Karpathy and Li, 2014) to align the two modalities. But in image captioning, the image and the text essentially have the same semantic information which is not the case for memes.

7 Conclusion

In this paper, we explore the challenges in meme sentiment and humor analysis using deep learning. We compare the performance of large pretrained models with the simpler models that we train from scratch. Through our detailed analysis we show that the unconventional presence of both the modalities in memes reduces the effectiveness of domain adaptation or transfer learning techniques and discuss some ways to improve our models as future work.

References

- Md. Shad Akhtar, Dushyant Singh Chauhan, Deepanway Ghosal, Soujanya Poria, Asif Ekbal, and Pushpak Bhat-tacharyya. 2019. Multi-task learning for multi-modal emotion recognition and sentiment analysis. *CoRR*, abs/1905.05812.
- Damian Borth, Rongrong Ji, Tao Chen, Thomas Breuel, and Shih-Fu Chang. 2013. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *Proceedings of the 21st ACM International Conference on Multimedia*, MM ’13, page 223–232, New York, NY, USA. Association for Computing Machinery.

²https://pikdo.info/p/robinsparklesmemes/1984316872681767883_11298693004?hl=en

³<https://knowyourmeme.com/photos/931101-vladimir-putin>

⁴https://www.reddit.com/r/raimimemes/comments/9d1sme/he_was_beautiful/

- François Chollet et al. 2015. Keras. <https://keras.io>.
- Jia Deng, Wenjun Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR 2009*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Jean H. French. 2017. Image-based memes as sentiment predictors. *2017 International Conference on Information Society (i-Society)*, pages 80–85.
- M. Guillaumin, J. Verbeek, and C. Schmid. 2010. Multimodal semi-supervised learning for image classification. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 902–909.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *CoRR*, abs/1512.03385.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November.
- Di Hu, Feiping Nie, and Xuelong Li. 2018. Dense multimodal fusion for hierarchically joint representation. *CoRR*, abs/1810.03414.
- Po-Yao Huang, Xiaojun Chang, and Alexander G. Hauptmann. 2019. Multi-head attention with diversity for learning grounded multilingual multimodal representations. In *EMNLP/IJCNLP*.
- Andrej Karpathy and Fei-Fei Li. 2014. Deep visual-semantic alignments for generating image descriptions. *CoRR*, abs/1412.2306.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc.
- 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. *CoRR*, abs/1907.11692.
- Navonil Majumder, Devamanyu Hazarika, Alexander F. Gelbukh, Erik Cambria, and Soujanya Poria. 2018. Multimodal sentiment analysis using hierarchical fusion with context modeling. *CoRR*, abs/1806.06228.
- Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. 2018. Multimodal named entity disambiguation for noisy social media posts. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2000–2008, Melbourne, Australia, July. Association for Computational Linguistics.
- John Edison Arevalo Ovalle, Tamar Solorio, Manuel Montes y Gómez, and Fabio A. González. 2017. Gated multimodal units for information fusion. *ArXiv*, abs/1702.01992.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. Context-dependent sentiment analysis in user-generated videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 873–883, Vancouver, Canada, July. Association for Computational Linguistics.
- K R Prajwal, C V Jawahar, and Ponnurangam Kumaraguru. 2019. Towards increased accessibility of meme images with the help of rich face emotion captions. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, page 202–210, New York, NY, USA. Association for Computing Machinery.
- Mike Schuster and Kuldip K. Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.*, 45:2673–2681.
- Chhavi Sharma, Deepesh Bhageria, William Paka, Scott, Srinivas P Y K L, Amitava Das, Tanmoy Chakraborty, Viswanath Pulabaigari, and Björn Gambäck. 2020. SemEval-2020 Task 8: Memotion Analysis-The Visuo-Lingual Metaphor! In *Proceedings of the 14th International Workshop on Semantic Evaluation (SemEval-2020)*, Barcelona, Spain, Sep. Association for Computational Linguistics.

- Aman Sinha, Parth Patekar, and Radhika Mamidi. 2019. Unsupervised approach for monitoring satire on social media. In *Proceedings of the 11th Forum for Information Retrieval Evaluation, FIRE '19*, page 36–41, New York, NY, USA. Association for Computing Machinery.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2014. Going deeper with convolutions. *CoRR*, abs/1409.4842.
- Suryatej Reddy Vyalla and Vishaal Udandarao. 2020. Memeify: A large-scale meme generation system. In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD, CoDS COMAD 2020*, page 307–311, New York, NY, USA. Association for Computing Machinery.
- William Yang Wang and Miaomiao Wen. 2015. I can has cheezburger? a nonparanormal approach to combining textual and visual information for predicting and generating popular meme descriptions. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 355–365, Denver, Colorado, May–June. Association for Computational Linguistics.
- Quanzeng You, Hailin Jin, and Jiebo Luo. 2017. Visual sentiment analysis by attending on local image regions. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17*, page 231–237. AAAI Press.
- Tom Zahavy, Alessandro Magnani, Abhinandan Krishnan, and Shie Mannor. 2016. Is a picture worth a thousand words? A deep multi-modal fusion architecture for product classification in e-commerce. *CoRR*, abs/1611.09534.
- Yu Zhang and Qiang Yang. 2017. A survey on multi-task learning. *CoRR*, abs/1707.08114.
- Lei Zhang, Shuai Wang, and Bing Liu. 2018. Deep learning for sentiment analysis : A survey. *CoRR*, abs/1801.07883.