

PVG at WASSA 2021: A Multi-Input, Multi-Task, Transformer-Based Architecture for Empathy and Distress Prediction

Atharva Kulkarni¹, Sunanda Somwase¹, Shivam Rajput¹, and Manisha Marathe²

^{1,2}Department of Computer Engineering, PVG's College of Engineering and Technology, affiliated to Savitribai Phule Pune University, India.

¹{atharva.j.kulkarni1998, sunandasomwase, shivamraje43}@gmail.com

² mvm.comp@pvgcoet.ac.in

Abstract

Active research pertaining to the affective phenomenon of empathy and distress is invaluable for improving human-machine interaction. Predicting intensities of such complex emotions from textual data is difficult, as these constructs are deeply rooted in the psychological theory. Consequently, for better prediction, it becomes imperative to take into account ancillary factors such as the psychological test scores, demographic features, underlying latent primitive emotions, along with the text's undertone and its psychological complexity. This paper proffers team PVG's solution to the WASSA 2021 Shared Task on Predicting Empathy and Emotion in Reaction to News Stories. Leveraging the textual data, demographic features, psychological test score, and the intrinsic interdependencies of primitive emotions and empathy, we propose a multi-input, multi-task framework for the task of empathy score prediction. Here, the empathy score prediction is considered the primary task, while emotion and empathy classification are considered secondary auxiliary tasks. For the distress score prediction task, the system is further boosted by the addition of lexical features. Our submission ranked 1st based on the average correlation (0.545) as well as the distress correlation (0.574), and 2nd for the empathy Pearson correlation (0.517).

1 Introduction

In recent years, substantial progress has been made in the NLP domain, with sentiment analysis and emotion identification at its core. The advent of attention-based models and complex deep learning architectures has led to substantial headways in sentiment/ emotion classification and their intensity prediction. However, the studies addressing the prediction of affective phenomena of empathy and distress have been relatively limited. Factors such as lack of large-scale quality labeled datasets,

the weak notion of the constructs themselves, and inter-disciplinary dependencies have hindered the progress. The WASSA 2021 Shared Task on Predicting Empathy and Emotion in Reaction to News Stories (Tafreshi et al., 2021) provides a quality, gold-standard dataset of the empathic reactions to news stories to predict Batson's empathic concern and personal distress scores.

Empathy, as defined by Davis (1983), is considered as "reactions of one individual to the observed experiences of another." It is more succinctly summarized by Levenson and Ruef (1992) in three key components, "(a) knowing what another person is feeling (cognitive), (b) feeling what another person is feeling (emotional), and (c) responding compassionately to another person's distress (behavioral)." Distress, on the other hand, as delineated by Dowling (2018), is "a strong aversive and self-oriented response to the suffering of others, accompanied by the desire to withdraw from a situation in order to protect oneself from excessive negative feelings." Empathy and distress are multifaceted interactional processes that are not always self-evident and often depend on the text's undertone. Moreover, along with the textual data, multiple psychological and demographic features also play a vital role in determining these complex emotions. Evidence by Fabi et al. (2019) suggests that empathy and distress are not independent of the basic emotions (happiness, sadness, disgust, fear, surprise, and anger) the subject feels during a given scenario. This appositeness of the primitive emotions with empathy and distress can be aptly exploited using a multi-task learning approach.

Multi-task learning has led to successes in many applications of NLP such as machine translation (McCann et al., 2017), speech recognition (Arik et al., 2017), representation learning (Hashimoto et al., 2017), semantic parsing (Peng et al., 2017), and information retrieval (Liu et al., 2015) to name

a few. NLP literature (Standley et al., 2020) suggests that under specific circumstances and with well-crafted tasks, multi-task learning frameworks can often aid models to achieve state-of-the-art performances. Standley et al. (2020) further asserts that seemingly related tasks can often have similar underlying dynamics. With the same intuition, Deep et al. (2020) designed a multi-task learning model for sentiment classification and their corresponding intensity predictions. Building on these findings, we propose a multi-input, multi-task, transformer-based architecture for the prediction of empathy and distress scores. The multi-input nature of the framework aggregates information from textual, categorical, and numeric data to generate robust representations for the regression task at hand. Exploiting the latent interdependencies between primitive emotions and empathy/ distress, we formulate the multi-task learning problem as a combination of classification and regression. The model simultaneously classifies the text into its correct basic emotion, detects if it exhibits high empathy/ distress, and accordingly predicts its appropriate empathy/ distress intensity score according to Batson’s scale (Batson et al., 1987). This multi-input, multi-task learning paradigm is further bolstered with the addition of NRC Emotion Intensity Lexicons (Mohammad, 2018b); NRC Valence, Arousal, and Dominance Lexicons (Mohammad, 2018a); and relevant features from Empath (Fast et al., 2016). Moreover, our proposed models have less than 110k trainable parameters and are still able to achieve relatively high Pearson’s correlation of 0.517 and 0.574 for empathy and distress, respectively, and 0.545 for average correlation, outperforming other teams.

2 Related Work

Over the last few years, earnest endeavours have been made in the NLP community to analyze empathy and distress. Earlier work in empathy mostly addressed the presence or absence of empathy in spoken dialogue (Gibson et al., 2015; Alam et al., 2016; Fung et al., 2016; Pérez-Rosas et al., 2017; Alam et al., 2018). For text-based empathy prediction, Buechel et al. (2018) laid a firm foundation for predicting Batson’s empathic concern and personal distress scores in reaction to news articles. They present the first publicly available gold-standard dataset for text-based empathy and distress prediction. Sharma et al. (2020) contem-

Dataset	Datapoints
Train	1860
Development	270
Test	525

Table 1: Data distribution.

plated a computational approach for understanding empathy in text-based health support. They developed a multi-task RoBERTa-based bi-encoder paradigm for identifying empathy in conversations and extracting rationales underlying its predictions. Wagner (2020) analysed the linguistic undertones for empathy present in avid fiction readers.

Computational work done for predicting distress is relatively modest. Shapira et al. (2020) analysed textual data to examine associations between linguistic features and clients distress during psychotherapy. They combined linguistic features like positive and negative emotion words with psychological measures like Outcome Questionnaire-45 (Lambert et al., 2004) and Outcome Rating Scale (Miller, 2003). Zhou and Jurgens (2020) studied the affiliation between distress, condolence, and empathy in online support groups using nested regression models.

3 Data Description

The WASSA 2021 Shared Task (Tafreshi et al., 2021) provides an extended dataset to the one compiled by Buechel et al. (2018). The dataset has a total of 14 features spanning textual, categorical, and numeric data types. Essays represent the subject’s empathic reactions to news stories he/she has read. The demographic features of gender, race, education, and the essay’s gold-standard emotion label cover the categorical input features. The numeric features include the subject’s age and income, followed by personality traits scores (conscientiousness, openness, extraversion, agreeableness, stability) and interpersonal reactivity index (IRI) scores (fantasy, perspective taking, empathetic concern, personal distress). The train-development-test split of the dataset is illustrated in Table 1.

4 Proposed Methodology

In this section, we posit two multi-task learning frameworks for empathy and distress score prediction. They are elucidated in detail as follows:

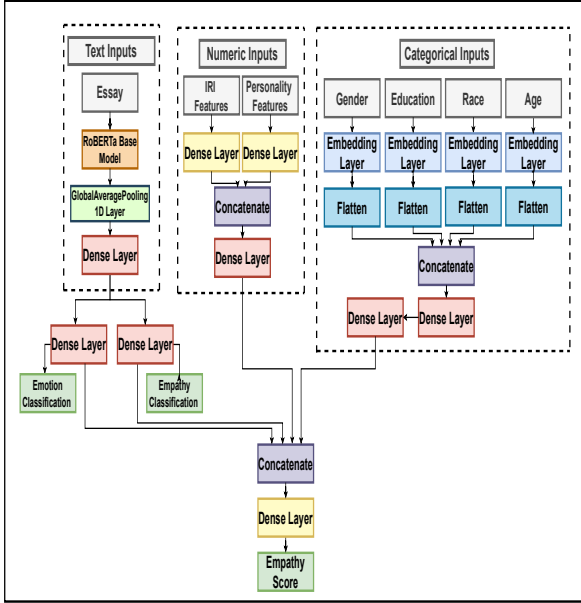


Figure 1: System architecture for empathy score prediction. The Dense layers in red have an ℓ^2 kernel regularization applied to them.

4.1 Empathy Score Prediction

Figure 1 depicts the system architecture for empathy score prediction. We formulate the task of empathy score prediction as a multi-input, multi-task learning problem. The proposed multi-task learning framework aims to leverage the empathy bin¹ and the text’s emotion to predict its empathy score. Here, the empathy score prediction is treated as the primary task, whereas emotion and empathy classification are considered secondary auxiliary tasks. The multi-input, multi-task nature of the framework efficiently fuses the diverse set of information (textual, categorical, and numeric) provided in the dataset to generate robust representations for empathy score prediction.

For the task of empathy and emotion classification, we make use of the pre-trained RoBERTa base model (Liu et al., 2019). The contextualized representations generated by RoBERTa help extract the context-aware information and capture the undertone of the text better than standard deep learning models. For each word in the essay, we extract the default pre-trained embeddings from the last hidden layer of RoBERTa. The 768-dimensional word

¹empathy bin is a feature given in the training dataset, where its value is 1 if empathy score greater than or equal to 4.0, and 0 if empathy score is less than 4.0. Thus, an essay exhibits high empathy if its empathy bin is 1 and exhibits low empathy if its empathy bin is 0. The same analogy is true for distress bin.

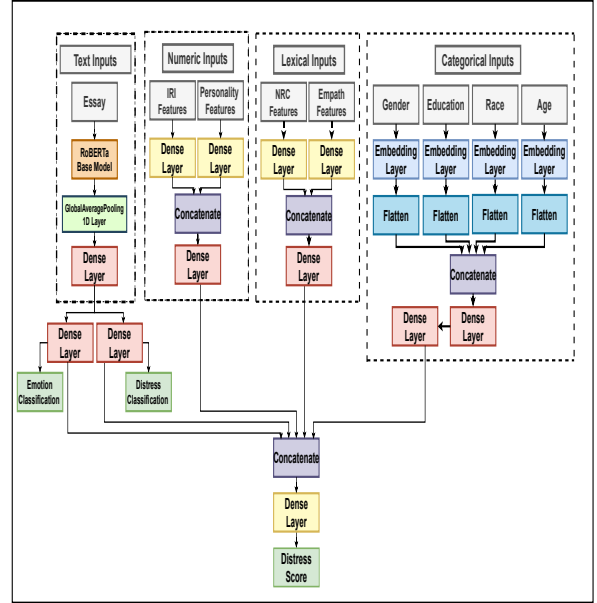


Figure 2: System architecture for distress score prediction. The Dense layers in red have an ℓ^2 kernel regularization applied to them.

embeddings are averaged to generate essay-level representations, followed by a hidden layer (128 units) for dimensionality reduction. The model further branches into two parallel fully-connected layers (16 units each), which form the task-specific layers for empathy and emotion classification, respectively. Let T_1 and T_2 denote the generated task-specific representations for empathy and emotion classification, respectively, of dimension d_1 . Finally, a classification layer (binary classification for empathy and multi-class classification for emotion) is added to each of the task-specific layers.

To incorporate the demographic information encoded in the categorical variables, we make use of entity embeddings (Guo and Berkahn, 2016). Formally, entity embeddings are domain-specific multi-dimensional representations of categorical variables, automatically learned by a neural network when trained on a particular task. For each of the demographic features (gender, education, race, and age²), 3-dimensional entity embeddings are generated. All the resultant embeddings are flattened, concatenated, and passed through two fully-connected layers (32 and 16 units, respectively), generating a layer having relevant information from each categorical input. Let this representation be denoted as C of dimension d_2 .

The numeric inputs of the personality and the

²In the dataset age is given as a numeric input. We split it into intervals of below 25, 26-40, 41-60, and 61 and above.

	Correlation with empathy score	Correlation with distress score
NRC features	sadness (0.19), fear (0.12), arousal (0.11), joy (0.11), valence (0.11), dominance (0.10)	fear (0.22), sadness (0.20), disgust (0.15), arousal (0.15), anger (0.15), dominance (0.12)
Empath features	domestic_work (0.15), death (0.14), home (0.12), sadness (0.12), suffering (0.11), help (0.11), party (0.11), family (0.10), shame (0.10), celebration (0.10), leisure (0.9), negative_emotion (0.9), air_travel (0.8), violence (0.8), pain (0.7)	suffering (0.16), death (0.15), torment (0.13), hate (0.11), negative_emotion (0.10), sadness (0.10), aggression (0.09), fight (0.09), help (0.09), pain (0.08), kill (0.08), horror (0.08), violence (0.08), war (0.08), ugliness (0.7)

Table 2: Selected NRC and Empath features and their pearson’s correlation with training data’s empathy and distress score.

Interpersonal Reactivity Index (IRI) scores are incorporated in the model by passing them individually through a single hidden layer (8 units). The results are concatenated and further passed on to a fully-connected layer (32 units) to generate their combined representations. Let this representation be denoted as N of dimension d_3 .

The task-specific layers for empathy and emotion classification and the representations generated by the final hidden layer of the entity embeddings and the numeric psychological score inputs are concatenated as given in equation 1 to generate the final representation F_1 for empathy score prediction.

$$F_1 = [T_1; T_2; C; N] \in \mathbb{R}^{d_1+d_2+d_3} \quad (1)$$

It is further passed to another hidden layer (16 units). Thus, this layer contains the compressed information from different knowledge views of the input data. A final regression layer is added for the empathy score prediction task. The multi-input, multi-task model is trained end-to-end with the objective loss calculated as the sum of the loss for each of the three tasks.

4.2 Distress Score Prediction

Figure 2 depicts the system architecture for distress score prediction. The distress score prediction model is the same as that of the empathy score prediction model, but with the addition of hand-crafted lexical features. We use 6 NRC Emotion Intensity Lexicons (Mohammad, 2018b); 2 NRC Valence, Arousal, and Dominance Lexicons (Mohammad, 2018a); and 15 features from Empath (Fast et al., 2016) as stated in Table 2. These features are chosen as they exhibit a high Pearson correlation with the training data. For each essay, the respective NRC and Empath score is calculated as the sum of each word’s score in the essay. The NRC lexicons and Empath features are passed to a single hidden layer (8 and 16 units for NRC and

Empath, respectively.) independently before concatenation. The resultant representation is further passed through another hidden layer (48 units). Let this representation be denoted by L of dimension d_4 . It is concatenated with the final layer representations from other inputs to generate the final representation F_2 for distress score prediction, as given by equation 2.

$$F_2 = [T_1; T_2; C; N; L] \in \mathbb{R}^{d_1+d_2+d_3+d_4} \quad (2)$$

5 Experimental Setup

5.1 Data Preparation

We applied standard text cleaning steps for each essay in the dataset, such as removing the punctuations, special characters, digits, single characters, multiple spaces, and accented words. The essays are further normalized by removing wordplay, replacing acronyms with full forms, and expanding contractions³.

Each essay is tokenized and padded to a maximum length of 200 tokens. Longer essays are truncated. Each Empath feature is converted into its percentage value. For distress score prediction, the numeric features, NRC lexicon scores, and Empath features are standardized by removing the mean and scaling to unit variance before being passed to the model.

5.2 Parameter Setting and Training Environment

Given the small amount of data, the weights of the RoBERTa layers were frozen and not updated during the training. The multi-task model is trained end-to-end using the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 1×10^{-3} and a batch size of 32. We used Hyperbolic Tangent

³<https://pypi.org/project/pycontractions/>

Model Name	Empathy		Distress	
	Train	Dev	Train	Dev
Plain RoBERTa model	0.56	0.47	0.56	0.41
Plain multi-input model	0.64	0.50	0.64	0.49
Multi-input, multi-task model	0.74	0.57	0.66	0.53
Multi-input, multi-task model+ NRC + Empath	0.64	0.54	0.66	0.56

Table 3: Performance comparison of various models as per Pearson’s correlation ($p < 0.05$).

(tanh) activation for all the hidden layers as it performed better than ReLU activation and its other variants. The model is trained for 200 epochs with early stopping applied if the validation loss does not improve after 20 epochs. Furthermore, the learning rate is reduced by a factor of 0.2 if validation loss does not decline after ten successive epochs. The model with the best validation loss is selected. ℓ^2 kernel regularizer of 5×10^{-4} applied to certain hidden layers as shown in Figures 1 and 2. A dropout of probability 0.2 is applied after the average pooling of contextual RoBERTa embeddings. Our code is available at our GitHub repository.⁴

6 Results and Discussion

The performance comparison of the various models on the validation data is reported in Table 3. As illustrated in the Table 3, we compare the performances of a RoBERTa based regression model, a multi-input model (text + categorical + numeric); a multi-input (text + categorical + numeric), multi-task model; and a multi-input (text + categorical + numeric), multi-task model with added lexical features (NRC + Empath). The systems submitted by our team are highlighted in Table 3. It is evident from Table 3 that the addition of categorical and numeric input features leads to an appreciable improvement in the models performance. This further attests that the demographic features and the psychological scores contribute valuable information for predicting scores of complex emotions like empathy and distress. The performance is further improved by adopting a multi-task approach, thus, reinstating our belief in the interdependencies between the primitive emotions and the complex emotions of empathy and distress. The addition of lexical features, however, leads to an improvement for distress prediction but a decrease in performance

⁴<https://github.com/mr-atharva-kulkarni/WASSA-2021-Shared-Task>

for empathy prediction. This may be explained by the theories in the affective neuroscience literature, wherein empathy is considered a neocortex emotion, describing it as an emotional, cognitive, and behavioral process. Thus, ones empathic quotient is reflected in how one expresses ones feelings and the undertone of it, rather than the mere words one uses. Moreover, as stated by Sedoc et al. (2020) in their work, there exist no clear set of lexicons that can accurately distinguish empathy from self-focused distress. Another reason for the decrease in the performance for empathy score prediction may be attributed to the fact that the Empath features for empathy, as reported in Table 2 lack some interpretability from a human perspective. Empath features such as *domestic work*, *party*, *celebration*, *leisure*, and *air travel* are not innately empathic categories and perhaps show high correlation due to corpus-based topical bias. The Empath features for distress, on the other hand, seem quite relevant and thus, might explain the increase in performance for distress prediction. We encourage further research in this direction.

7 Conclusion and Future Work

In this work, we propose a multi-input, multi-task, transformer-based architecture to predict Batson’s empathic concern and personal distress scores. Leveraging the dependency between the basic emotions and empathy/ distress, as well as incorporating textual, categorical, and numeric data, our proposed model generates robust representations for the regression task at hand. The addition of certain lexical features further improves the model’s performance for distress score prediction. Our submission ranked 1st based on average correlation (0.545) as well as distress correlation (0.574), and 2nd for empathy Pearson correlation (0.517) on the test data. As for the future work of this research, a weighted loss scheme could be employed to enhance the results. From a psychological and linguistic standpoint, features such as part-of-speech tags, syntax parse tree, and the text’s subjectivity and polarity scores could also be exploited.

References

- F. Alam, M. Danieli, and G. Riccardi. 2016. [Can we detect speakers’ empathy?: A real-life case study](#). In *2016 7th IEEE International Conference on Cognitive Infocommunications (CogInfoCom)*, pages 000059–000064.

- Firoj Alam, Morena Danieli, and Giuseppe Riccardi. 2018. [Annotating and modeling empathy in spoken conversations](#). *Comput. Speech Lang.*, 50(C):4061.
- Sercan Ö. Arik, Mike Chrzanowski, Adam Coates, Gregory Diamos, Andrew Gibiansky, Yongguo Kang, Xian Li, John Miller, Andrew Ng, Jonathan Raiman, Shubho Sengupta, and Mohammad Shoeybi. 2017. [Deep voice: Real-time neural text-to-speech](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 195–204, International Convention Centre, Sydney, Australia. PMLR.
- C Daniel Batson, Jim Fultz, and Patricia A Schoenrade. 1987. [Distress and empathy: Two qualitatively distinct vicarious emotions with different motivational consequences](#). *Journal of personality*, 55(1):19–39.
- Sven Buechel, Anneke Buffone, Barry Slaff, Lyle Ungar, and João Sedoc. 2018. [Modeling empathy and distress in reaction to news stories](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4758–4765, Brussels, Belgium. Association for Computational Linguistics.
- Mark H Davis. 1983. [The effects of dispositional empathy on emotional reactions and helping: A multidimensional approach](#). *Journal of personality*, 51(2):167–184.
- Kumar Shikhar Deep, Md Shad Akhtar, Asif Ekbal, and Pushpak Bhattacharyya. 2020. [Related tasks can share! a multi-task framework for affective language](#). *arXiv preprint arXiv:2002.02154*.
- Trisha Dowling. 2018. [Compassion does not fatigue!](#) *The Canadian Veterinary Journal*, 59(7):749.
- Sarah Fabi, Lydia Anna Weber, and Hartmut Leuthold. 2019. [Empathic concern and personal distress depend on situational but not dispositional factors](#). *PLoS ONE*, 14(11):e0225102.
- Ethan Fast, Binbin Chen, and Michael S. Bernstein. 2016. [Empath: Understanding topic signals in large-scale text](#). In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, page 46474657, New York, NY, USA. Association for Computing Machinery.
- Pascale Fung, Anik Dey, Farhad Bin Siddique, Ruixi Lin, Yang Yang, Yan Wan, and Ho Yin Ricky Chan. 2016. [Zara the Supergirl: An empathetic personality recognition system](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 87–91, San Diego, California. Association for Computational Linguistics.
- James Gibson, Nikolaos Malandrakis, Francisco Romero, David C Atkins, and Shrikanth S Narayanan. 2015. [Predicting therapist empathy in motivational interviews using language features inspired by psycholinguistic norms](#). In *Sixteenth annual conference of the international speech communication association*.
- Cheng Guo and Felix Berkhahn. 2016. [Entity embeddings of categorical variables](#).
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2017. [A joint many-task model: Growing a neural network for multiple NLP tasks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1923–1933, Copenhagen, Denmark. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Michael J Lambert, Ann T Gregersen, and Gary M Burlingame. 2004. [The outcome questionnaire-45. The use of psychological testing for treatment planning and outcomes assessment: Instruments for adults](#).
- Robert W Levenson and Anna M Ruef. 1992. [Empathy: a physiological substrate](#). *Journal of personality and social psychology*, 63(2):234.
- Xiaodong Liu, Jianfeng Gao, Xiaodong He, Li Deng, Kevin Duh, and Ye-yi Wang. 2015. [Representation learning using multi-task deep neural networks for semantic classification and information retrieval](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 912–921, Denver, Colorado. 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](#). *arXiv preprint arXiv:1907.11692*.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. [Learned in translation: Contextualized word vectors](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 62976308, Red Hook, NY, USA. Curran Associates Inc.
- Scott D Miller. 2003. [The outcome rating scale: A preliminary study of the reliability, validity, and feasibility of a brief visual analog measure](#). *Journal of brief Therapy*, 2(2):91–100.
- Saif Mohammad. 2018a. [Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 174–184, Melbourne, Australia. Association for Computational Linguistics.

- Saif Mohammad. 2018b. [Word affect intensities](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Hao Peng, Sam Thomson, and Noah A. Smith. 2017. [Deep multitask learning for semantic dependency parsing](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2037–2048, Vancouver, Canada. Association for Computational Linguistics.
- Verónica Pérez-Rosas, Rada Mihalcea, Kenneth Resnicow, Satinder Singh, and Lawrence An. 2017. [Understanding and predicting empathic behavior in counseling therapy](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1435, Vancouver, Canada. Association for Computational Linguistics.
- João Sedoc, Sven Buechel, Yehonathan Nachmany, Anneke Buffone, and Lyle Ungar. 2020. [Learning word ratings for empathy and distress from document-level user responses](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1664–1673, Marseille, France. European Language Resources Association.
- Natalie Shapira, Gal Lazarus, Yoav Goldberg, Eva Gilboa-Schechtman, Rivka Tuval-Mashiach, Daniel Juravski, and Dana Atzil-Slonim. 2020. [Using computerized text analysis to examine associations between linguistic features and clients distress during psychotherapy](#). *Journal of counseling psychology*.
- Ashish Sharma, Adam Miner, David Atkins, and Tim Althoff. 2020. [A computational approach to understanding empathy expressed in text-based mental health support](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5263–5276, Online. Association for Computational Linguistics.
- Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. 2020. [Which tasks should be learned together in multi-task learning?](#) In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9120–9132. PMLR.
- Shabnam Tafreshi, Orphée De Clercq, Valentin Barriere, Sven Buechel, João Sedoc, and Alexandra Balahur. 2021. [WASSA2021 Shared Task: Predicting Empathy and Emotion in Reaction to News Stories](#). In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Association for Computational Linguistics.
- Rachel Wagner. 2020. [It’s not what you say, it’s how you say it: a linguistic analysis of empathy in fiction readers](#). Ph.D. thesis, Rutgers University-Camden Graduate School.
- Naitian Zhou and David Jurgens. 2020. [Condolence and empathy in online communities](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 609–626, Online. Association for Computational Linguistics.