

Learning Word Ratings for Empathy and Distress from Document-Level User Responses

João Sedoc^{*♣}, Sven Buechel^{*†♣}, Yehonathan Nachmany[◇], Anneke Buffone[◇], and Lyle Ungar[◇]

[♣]Johns Hopkins University, [♣]Friedrich-Schiller-Universität Jena, [◇]University of Pennsylvania

jsedoc@jhu.edu, sven.buechel@uni-jena.de, yoninachmany@gmail.com, buffonea@sas.upenn.edu, ungar@cis.upenn.edu

Abstract

Despite the excellent performance of black box approaches to modeling sentiment and emotion, lexica (sets of informative words and associated weights) that characterize different emotions are indispensable to the NLP community because they allow for interpretable and robust predictions. Emotion analysis of text is increasing in popularity in NLP; however, manually creating lexica for psychological constructs such as empathy has proven difficult. This paper automatically creates empathy word ratings from document-level ratings. The underlying problem of learning word ratings from higher-level supervision has to date only been addressed in an *ad hoc* fashion and has not used deep learning methods. We systematically compare a number of approaches to learning word ratings from higher-level supervision against a Mixed-Level Feed Forward Network (MLFFN), which we find performs best, and use the MLFFN to create the first-ever empathy lexicon. We then use Signed Spectral Clustering to gain insights into the resulting words. The empathy and distress lexica are publicly available at: <http://www.wwpdb.org/lexica.html>.

Keywords: lexicon creation, empathy, distress

1. Introduction

Deep learning, applied to ever larger datasets, has led to large improvements in performance in sentiment and emotion analysis. In light of this development, lexica, lists of words and associated weights for a particular affective variable, which used to be a key component for feature extraction (Mohammad and Bravo-Marquez, 2017a), may seem obsolete. However, this is far from the truth.

Lexica can be used as features to improve performance for sentence-level emotion prediction even in advanced neural architectures (Mohammad and Bravo-Marquez, 2017b; De Bruyne et al., 2019). Word ratings are also often used to refine pre-trained embedding models for specific tasks (Yu et al., 2017; Khosla et al., 2018). But much more importantly, word ratings are relatively cheap to acquire and have been found to be robust across domains and even languages, regarding their translational equivalents (Leveau et al., 2012; Warriner et al., 2013). This gives lexica a pivotal role for processing under-resourced languages. Perhaps most importantly, using lexica allows for interpretable models since the resulting document-level predictions can be easily broken down to the words within it. This gives lexica an important role for building justifiable AI and addressing related ethical challenges (Clos et al., 2017). Interpretability is also crucial for NLP use in other academic disciplines such as psychology, social science, discourse linguistics, and the digital humanities, where understanding the nature of “constructs” (as psychologists call them) such as emotions is far more important than making accu-

rate predictions (Schwartz et al., 2013; Eichstaedt et al., 2015; Pennebaker, 2011; Liu et al., 2016).

While lexica for many kinds of emotion already exist (see Section 2.1.), there is no such resource for empathy despite its growing popularity in the NLP community (Khanpour et al., 2017; Buechel et al., 2018). Hand-curated lexica for empathy are difficult to create in part because there is no clear set of words that can accurately distinguish empathy from self-focused distress. The gold standard for discerning these is an emotion rating scale by Batson et al. (1987). This scale is a collection of emotion words (e.g., compassionate, tender, warm) that could serve as a rudimentary lexicon, but it contains many words that are rarely used (e.g., “perturbed”), and many words that can take on meanings that are far from empathy (e.g., “warm”, “tender”). These word-based scales have shown good reliability for self-report of these emotional states, but would make poor guides for a proper lexicon of empathy.

In this paper, we construct the first empathy lexicon. Specifically, we learn ratings for *two kinds of empathy*—empathic concern (feeling for someone) and personal distress (suffering with someone)—for words given existing document-level ratings from the recently published Empathic Reactions dataset (Buechel et al., 2018). We first train a model to predict document-level empathy in a regular supervised set-up and then “invert” the resulting model to derive word ratings. We conclude with an in-depth analysis of the resulting resource.

2. Related Work

2.1. Lexica for Psychological Quantities

The notion of describing (part of) a word’s meaning, such as the emotion typically associated with it, in terms of numerical ratings has a long tradition in psychology, dating back at least to Osgood et al. (1957). Today, many sets of word ratings exist, covering numerous constructs and languages, particularly relating to sentiment and emo-

* These authors contributed equally to this work. Sven Buechel designed the MLFFN model (Section 3.). João Sedoc, along with Yoni Nachmany, was responsible for conducting experiments (Section 4.) and clustering analysis (Section 5.2.). Experimental design and supervision of the implementation of the algorithms were done jointly by both first authors.

†Work partially conducted as a visiting researcher at the University of Pennsylvania.

tion. Early work in NLP was mostly focused on positive-vs.-negative resources such as SentiWordNet and VADER (Baccianella et al., 2010; Hutto and Gilbert, 2014). In contrast, resources from psychologists tend to focus on *valence* and *arousal* (or other representations of affective states (Ekman, 1992)). In particular, this includes the Affective Norms for English Words (ANEW; (Bradley and Lang, 1999)) which have been adopted to many languages (Redondo et al., 2007; Montefinese et al., 2014), and their extension by Warriner et al. (2013). Such lexica have recently become popular in NLP (Wang et al., 2016; Sedoc et al., 2017b; Mohammad, 2018; Buechel and Hahn, 2018a). Lexica also exist for many other constructs, including concrete/abstractness, familiarity, imageability and humor (Brysbaert et al., 2014; Yee, 2017; Engelthaler and Hills, 2018). Yet, noticeably, an empathy lexicon is missing.

Psychologists use such lexica either for content analysis, most noticeably using the Linguistic Inquiry and Word Count (LIWC) lexica (Tausczik and Pennebaker, 2010), or as controlled stimuli for experiments, e.g., on language processing and memory (Hofmann et al., 2009; Monnier and Syssau, 2008). Applications of lexica in NLP have been discussed in Section 1.

Whereas most lexica are created manually, there is an extensive body of work on learning such ratings automatically (see Kulkarni et al. (2019) for a survey). Early work focused on deriving scores through linguistic patterns or statistical association with a small set of seed words (Hatzivassiloglou and McKeown, 1997; Turney and Littman, 2003). More recent approaches almost always rely on word embeddings (Hamilton et al., 2016; Li et al., 2017; Buechel and Hahn, 2018b). This line of work is predominantly based on *word-level* supervision. In contrast, we learn word ratings from document-level ratings.

2.2. Empathy and Distress in Psychology

Empathic emotions, “reactions of one individual to the observed experiences of another” (Davis, 1983), often in response to their need or suffering, is complex and controversial, with luminary scientists both arguing for the benefits of empathy (De Waal, 2009) and “against empathy” (Bloom, 2016). Empathy has been linked to a multitude of positive outcomes, from volunteering (Batson et al., 1997), to charitable giving (Pavey et al., 2012), and even longevity (Poulin et al., 2013), but it can also cause the empathizing person increased stress (Buffone et al., 2017) and emotional pain (Chikovani et al., 2015). In this paper, we build lexica for two distinct types of *state* (momentary) empathy, empathic concern and personal distress, which are based on the subscales of Interpersonal Reactivity Index (IRI) questionnaire (Davis, 1980). The scale creator defines these as follows: **Empathic Concern** assesses “other-oriented” feelings of sympathy and concern for unfortunate others and **Personal Distress** measures “self-oriented” feelings of personal anxiety and unease in tense interpersonal settings. For conciseness, we will use the terms “empathy” and “distress” to refer to this pair of constructs throughout the paper.

2.3. Empathy and Distress in AI

Most previous work in language-centered AI for empathy has been conducted with a focus on speech and especially spoken dialogue. Conversational agents, psychological interventions, and call center applications have been addressed particularly often (McQuiggan and Lester, 2007; Fung et al., 2016; Pérez-Rosas et al., 2017; Alam et al., 2017). In contrast, studies addressing empathy in written language are surprisingly rare. Abdul-Mageed et al. (2017), in contrast, focus on *trait* empathy, a temporally more stable personal attribute. In particular, they studied the detection of “pathogenic empathy”, marked by self-focused distress, a potentially detrimental form of empathy associated with health risks, in social media language using a wide array of features, including *n*-grams and demographic information.

Khanpour et al. (2017) present a corpus of messages from online health communities which has binary empathy annotations on the sentence-level. They report an .78 F-Score using a CNN-LSTM. The corpus, however, is not publicly available. In contrast, Buechel et al. (2018) recently presented the first publicly available gold standard dataset supported by proper psychological theories. The dataset consists of responses to news articles and scales for empathy and distress between 1 and 7. They collected empathy and distress ratings from the writer of an informal message using a sophisticated annotation methodology borrowed from psychology. In this contribution, we build upon their work by using their document-level ratings to predict word labels.

2.4. Lexicon Learning from Document Labels

Few studies address learning word ratings based on document-level supervision. However, those studies (described in detail below) focus on their particular application rather than addressing the underlying, abstract learning problem (formalized in Section 3.). As a result, previously proposed methods have not been quantitatively compared. In an early study, Mihalcea and Liu (2006) computed the *happiness factor* of a word type as the ratio of documents labeled “happy” to all blog posts it occurs in. Labels were given by the blog users. The resulting lexicon was used to estimate user happiness over the course of an average 24-hour day as well as a seven-day week. Rill et al. (2012) independently came up with a very similar approach for identifying the evaluative meaning of adjectives and adjective phrases (*absolutely fantastic* vs. *just awful*) based on a corpus of online product reviews. Since the individual reviews come with a one-to-five star rating, the evaluative meaning of an adjective or phrase was computed as the average rating of all reviews it occurs in (Mean Star Rating, see Section 3.). This approach was later adopted by Ruppenhofer et al. (2014) who found that it works quite well for classifying quality and intelligence adjectives into intensity classes (*excellent* vs. *mediocre* and *brilliant* vs. *dim*, respectively). Another related approach was proposed by Mohammad (2012), who used hashtags in Twitter posts as distant supervision labels of emotion categories, e.g., *#sadness*. Word ratings were then computed based on pointwise mutual information between word types and emotion

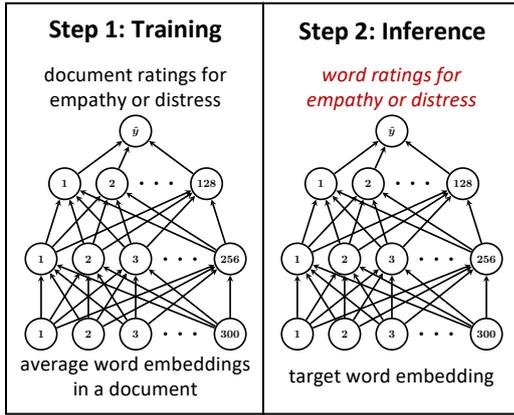


Figure 1: Schematic illustration of the MLFFN: **Step 1** Train a model to predict a document rating from the embedding of the document text; **Step 2** "Invert" the trained model to compute a word rating for each word embedding.

labels.

The above methods all derive word labels directly using relatively simple statistical operations. From this group, we selected the Mean Star Rating approach for experimental comparison (Section 4.), as it expects numerical document labels, in line with the later employed empathy gold standard (Section 5.).

Note that these contributions are distinct from pattern-based approaches, e.g., presented by Hatzivassiloglou and McKeown (1997), who distinguish positive and negative words based on their usage pattern with particular conjunctions: "A *and* B" implies that A and B have the same polarity whereas "A *but* B" implies opposing polarity. Such approaches are not considered here because they base lexicon learning on linguistic usage patterns instead of document-level supervision and hence rely on large quantities of raw text.

In another line of work, (Sap et al., 2014) address the task of modeling user age and gender in social media. They showed that by training a linear model with Bag-of-Words (BoW) unigram features, the resulting feature weights can effectively be interpreted as word-level ratings. In a later study Preoțiuc-Pietro et al. (2016) employed the same method to create a valence and arousal lexicon based on annotated Facebook posts. This is the second baseline method we used in our evaluation; Technical details are given in Section 3. (Regression Weights).

In a recent study, Wang and Xia (2017) present a three step approach to infer word polarity. Based on a Twitter corpus with hashtag-derived polarity labels, they (1) apply the method of Mohammad (2012) to generate a first set of word labels (see above). Those ratings are used (2) to train sentiment-aware word embeddings. The embeddings are then used (3) as input to a classifier which is trained on a set of seed words to predict the final word ratings. In essence, this is a semi-supervised approach because the last step requires word-level gold data and does not address the problem at hand.

3. Methods

This section formalizes the learning problem we address, describes the three baseline methods we compare against and the Mixed-Level Feed Forward Network, and concludes with a brief discussion. Signed Spectral Clustering, which we use for qualitative interpretation of our resulting empathy lexicon, is described in Section 5..

Problem Statement. We address the problem of learning word ratings for an arbitrary lexical characteristic based on gold labels of the same characteristic, but for a higher linguistic level (see Figure 1). For example, how can one learn *word-level* polarity ratings based on *document-level* polarity gold labels? More formally, let $\mathcal{W} = \{w_1, w_2, \dots, w_n\}$ denote a set of words with corresponding gold labels $Y^w = (y_1^w, y_2^w, \dots, y_n^w)$. Let $\mathcal{D} = \{d_1, d_2, \dots, d_m\}$ be a set of higher level linguistic units with corresponding gold labels $Y^d = (y_1^d, y_2^d, \dots, y_m^d)$. Those linguistic units can be anything from phrases over paragraphs to whole books, yet for conciseness we will refer to those units as *documents*. Our problem is to predict Y^w given \mathcal{W} , \mathcal{D} , and Y^d . To the best of our knowledge, this is the first contribution ever formulating this as an abstract learning problem (rather than looking at concrete applications in isolation) and studying it in a systematic manner—the baseline methods have so far not been compared against each other. We now proceed by introducing methods for solving this problem.

MEAN STAR RATING. Following Rill et al. (2012), we predict y_i^w by averaging the gold labels of documents d_j in which the word w_i occurs. We denote the set of documents containing w_i as $D(w_i)$. Hence this baseline method can be described as follows:

$$\hat{y}_i^w = \frac{1}{|D(w_i)|} \sum_{d_j \in D(w_i)} y_j^d. \quad (1)$$

MEAN BINARY RATING. As previously mentioned, Mihalcea and Liu (2006) created lexica for happiness and sadness from binary labels. To apply this method to numerical document labels (as present in the Empathic Reactions dataset; see Section 5.) we first apply a median split (documents labels below (above) the median are recorded as 0 (1)). Subsequently, we calculate Mean Binary Rating using the same equation as for Mean Star Rating (Equation 1) thus showing the resemblances between Mihalcea and Liu (2006) and Rill et al. (2012).

REGRESSION WEIGHTS. Following (Sap et al., 2014), this baseline method learns word ratings by fitting a linear regression model with Bag-of-Words (BoW) features.

First, consider a linear regression model for predicting document ratings Y^d . In general, such a model is given by

$$\hat{y}_i^d = a_0 + \sum_{j \in \text{features}} a_j \cdot x_j,$$

where a_0 denotes the intercept and a_j and x_j represent weight and value for feature j , respectively.

Using a BoW approach, *relative frequency* of a word in a document is often used as features. Except for the intercepts, the linear model can conversely be interpreted as computing the weighted average of *all weight terms*

a_j , the relative term frequency then being the weighting factor. With this interpretation in mind, a linear BoW model aligns perfectly with a lexicon-based approach to achieve document-level prediction, with feature weights corresponding to word ratings (see Sap et al. (2014) for a more detailed explanation). Hence the above equation can be rewritten as

$$\hat{y}_i^d = a_0 + \sum_{w_j \in \mathcal{W}} \hat{y}_j^w \cdot \text{rf}(w_j, d_i),$$

where $\text{rf}(w_j, d_i)$ denotes the relative frequency of word w_j in document d_i . Thus, by fitting the model to predict document ratings, we learn word ratings, simultaneously. In practice, ridge regression is used for fitting the model parameters (word ratings) thus introducing ℓ_2 penalization to avoid overfitting.

MIXED-LEVEL FEED FORWARD NETWORK. We learn a Feed Forward Network (FFN; illustrated in Figure 1) on the document-level using a neural BoW approach with an external, pre-trained embedding model. By training the FFN on this task, it implicitly learns to map points of the embedding space to gold labels, which we then exploit for predicting word level ratings.

In general, a Feed Forward Network consists of an input layer $a^{(0)} \in \mathbb{R}^{\text{dim}}$ followed by multiple hidden layers with activation,

$$a^{(l+1)} = \sigma(W^{(l+1)}a^{(l)} + b^{(l+1)}),$$

where $W^{(l+1)}$, $b^{(l+1)}$ denote weights and biases of layer $l + 1$, respectively, and σ is a nonlinear function. Since we predict numerical values (document-level ratings), the activation on the output layer $a^{(\text{out})}$, where out is the number of non-input layers, is given by the affine transformation

$$\hat{y}^d = a^{(\text{out})} = W^{(\text{out})}a^{(\text{out}-1)} + b^{(\text{out})}.$$

For fitting the model parameters, consider a pre-trained embedding model such that $\text{vec}(\omega) \in \mathbb{R}^{\text{dim}}$ denotes the vector representation of a word ω . This would be either the learned representation of ω or a zero vector of length dim if ω is not in the embedding model. We can now train the model to predict the document gold ratings Y^d using a gradient descent-based method. For a document d_i , the embedding centroid of tokens present in d_i is used as input $a^{(0)}(d_i)$. That is,

$$a^{(0)}(d_i) = \frac{1}{\text{len}(d_i)} \sum_{\omega \in d_i} \text{vec}(\omega),$$

where $\text{len}(d_i)$ is the number of tokens in d_i . Embeddings are not updated during training.

Until this point, the described model is quite similar to deep averaging networks (DAN) proposed by Iyyer et al. (2015) in that it is a Feed Forward Network that predicts document labels from embedding centroid features. What differs is that we used the model to predict *word* labels, once it is fit to predict the document labels Y^d . Conceptually, by fitting the model parameters, the FFN learns to map points of the (pre-trained) embedding space to points in the label space of Y^d . But using the same embedding model, we can also

represents words $w_i \in \mathcal{W}$, the ones we want to predict labels for, within the same feature space. Moreover, note that per our problem definition, word and document labels populate the same label space. Hence, we can predict y_i^w by feeding $\text{vec}(w_i)$ into the FFN without any further adjustments. Since the FFN can predict both word *and* document labels, we call this model Mixed-Level Feed Forward Network (MLFFN).¹

Hyperparameters and Implementation. The implementation of Mean Star Rating and Mean Binary Rating is straightforward and requires no further details. For Regression Weights, we used the same setup as Sap et al. (2014), as implemented in the Differential Language Analysis Toolkit (DLATK, (Schwartz et al., 2017)).

For MLFFN, we built on the implementation² and hyperparameter choices Buechel et al. (2018) used for the Empathic Reactions dataset. Thus, MLFFN has two hidden layers (256 and 128 units, respectively) with ReLU activation. The model was trained using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 10^{-3} and a batch size of 32. We trained for a maximum of 200 epochs, and applied early stopping if the performance on the validation set did not improve for 20 consecutive epochs. We applied dropout with probabilities of .2 and .5 on input and dense layers, respectively. Moreover ℓ^2 regularization of .001 was applied to the weights of dense layers. Keras (Chollet and others, 2015) was used for implementation.

Discussion of Model Properties. Mean Star Rating, Binary Star Rating, and Regressions Weights learn exclusively from the available document-level gold data. In contrast, one of the major advantages of the MLFFN is that it builds on pre-trained word embeddings, thus implicitly leveraging vast amounts of unlabeled text data. For our experiments we use publicly available embeddings which are trained on hundreds of billions of tokens. MLFFN is also more flexible than Regression Weights, since it can learn nonlinear dependencies between relative word frequencies of a document and its gold label.

Another major advantage of the MLFFN model relates to the set of words that gold labels can be predicted for. Whereas Mean Star Rating, Mean Binary Rating, and Regression Weights are conceptually limited to words which occur in the gold data, MLFFN can predict ratings for any word for which embeddings are known. In practice, this implies that with our approach empathy ratings for *millions* of word types can be induced.

4. Experiments

We next conduct a systematic comparison of the above approaches. The best evaluation strategy would require having both document-level and word-level ratings for empathy. One could then train the models on the former and test the performance in predicting the later, possibly using resampling to get a distribution of scores. However, this

¹For the MLFFN, our method of using individual words to derive ratings is mathematically equivalent to SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017).

²Implementation available along with the Empathic Reactions dataset; see Footnote 4.

Method	Intr.			Extr.
	V	A	D	
Mean Binary Rating	.31	.18	.11	.07
Regression Weights	.36	.22	.13	.09
Mean Star Rating	.39	.22	.14	.09
MLFFN	.64	.45	.50	.18

Table 1: Evaluation of lexicon learning methods in Pearson correlation. Left column group: intrinsic evaluation on emotion corpus and lexicon (VAD = Valence-Arousal-Dominance); Right column group: Extrinsic evaluation on Twitter corpus annotated with person-level empathy.

option is not available, since the difficulty of acquiring empathy word ratings is exactly the point of this paper.³

We adopt two alternative approaches: First, in place of empathy, we rely on other affective variables, namely, valence, arousal, and dominance (VAD), for which both document and word ratings are available. The assumption here is that performance results for VAD are transferable to empathy. Second, we use the Empathic Reactions⁴ dataset to create one lexicon for each method under consideration. We then use it to predict *trait*-level empathy ratings, thus testing the generalizability of the resulting lexica to other domains as well as from state to trait empathy (see Section 2.).

4.1. Intrinsic Evaluation with Emotion Data

We use the following gold standards for evaluation: Document-level supervision is provided by EmoBank (Buechel and Hahn, 2017)⁵, a large-scale corpus manually annotated with emotion according to the psychological Valence-Arousal-Dominance scheme. EmoBank contains ten thousand sentences with multiple genres and has annotations from both writer and reader emotion. Word-level supervision to test against comes from the well-known affective norms (psychological valence, arousal, and dominance) dataset collected by Warriner et al. (2013) containing 13,915 English word types.

We fit all four models on EmoBank and evaluate against the word ratings by Warriner et al. (2013) using 10-fold cross-validation. For word embeddings we used off-the-shelf Fasttext subword embeddings (Mikolov et al., 2018).⁶ The embeddings are trained with subword information on Common Crawl (600B tokens). Performance will be mea-

³Another seemingly obvious evaluation strategy would be to predict document-level ratings from derived word-level lexica using the empathic reactions dataset in a cross-validation setup. However, we found that this approach has two major drawbacks. First, rather than validating the resulting word ratings directly, this strategy constitutes a down-stream (extrinsic) evaluation, similar to what we propose in Section 4.2.. Second, we found empirically that this approach lacks statistical power to distinguish between methods due to an insufficient number of examples.

⁴https://github.com/wwbp/empathic_reactions

⁵<https://github.com/JULIELab/EmoBank>

⁶<https://dl.fbaipublicfiles.com/fasttext/vectors-english/crawl-300d-2M-subword.zip>.

sured in terms of Pearson correlation between predicted and gold labels. As shown in Table 1, the MLFFN by far outperforms Mean Binary Rating, Mean Star Rating and Regression Weights, the latter three being roughly equal. This is most probably, because the MLFFN builds on top of a pre-trained embedding model, thus leverage vast amounts of unlabeled data in addition to the document-level supervision.

4.2. Extrinsic Evaluation

To evaluate the lexica created using each of the underlying methods, we applied them to measure *trait*-level empathy dataset. We validate our empathy lexica by showing that they predict personal-level empathy traits on another dataset which collected trait-level empathy questionnaires and users’ Facebook posts. Abdul-Mageed et al. (2017) used this dataset to predict person trait-based pathogenic empathy. Here, instead we aggregate the empathy survey results of Facebook users recruited via Qualtrics. We filtered users to include those who had posted at least five times in the last 30 days and had at least 100 lifetime posts. The survey included an integrated app grabbing participants’ Facebook posts. In total there are 2,405 users with 1,835,884 Facebook posts after filtering non-English posts (see Abdul-Mageed et al. (2017) for further dataset details). The lexica were employed in a very simple fashion: For each user, we computed the weighted average of the empathy scores of the the words they used across all tweets. Relative frequency was used as weighting feature. As shown in Table 1, the performance is generally much poorer, indicating the change of domain, the more difficult task of inferring *trait*-level empathy from state-level ratings, and the overall reduced performance of purely lexicon-based approaches.⁷ Nevertheless, our findings are consistent with intrinsic results. The MLFFN widely outperforms the other approaches (having about twice as high performance figures), the latter being roughly similar in performance.

5. The Empathy Dictionary

The final empathy dictionary consists of the predictions of the MLFFN from the last experiment, which we then adjusted using log min-max rescaling into the interval [1, 7] for consistency with Buechel et al. (2018). We restricted ourselves to words which appear in the Empathic Reaction dataset and did not make use of the ability of the MLFFN to predict ratings for all word of the embedding model (left for future work). This was done to ensure interpretability of word ratings relative to their usage in the corpus (achieved via clustering analysis in Section 5.2.).

5.1. Dataset Description

Our final lexicon consists of 9,356 word types (lower-cased, non-lemmatized, including named entities and spelling errors) each with associated empathy and distress ratings. For illustration Table 2, lists the highest and lowest ranking

⁷While 0.18 may seem low, our results are similar to those from Abdul-Mageed et al. (2017) who use a regression model with LDA topics trained on Facebook posts.

		Empathy	Distress
High Empathy	lukemia	6.90	5.09
	lakota	6.70	4.88
	healing	6.60	4.90
Low Empathy	joke	1.10	1.52
	worrying	1.10	3.93
	wacky	1.10	1.49
High Distress	inhumane	4.07	6.55
	dehumanizes	5.46	6.40
	mistreating	4.85	6.31
Low Distress	somewhere	1.82	1.05
	dunno	1.31	1.05
	guessing	1.38	1.06

Table 2: Illustrative examples from final lexicon: highest/lowest ranking words for empathy and distress.

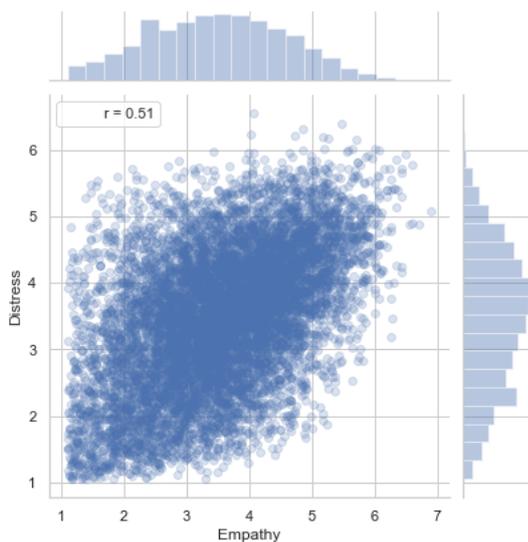


Figure 2: Label distribution in our empathy lexicon.

words for each construct (empathy and distress). High-empathy words contain many named entities that experience or cause suffering making a reader feel empathic (e.g., *lukemia* or *lakota*). This is likely because the Empathic Reactions corpus used news stories to evoke empathy in subjects who then referred to those named entities for expressing their feeling (Section 5.2. provides an estimate of the total number of named entities in the lexicon). Low-empathy words, on the other hand, are often ones used for ridiculing, hence expressing a lack of empathy (*joke*, *wacky*). High-Distress words contain predominantly adjectives, nouns, and participles which can be used to characterize abusive behaviour (*inhumane*, *mistreating*) thus causing personal distress in readers when taking the perspective of the affected entity. Interestingly, low-distress words do not seem to display any clear pattern, making us suspect that personal distress should be addressed in terms of a unipolar rather than a bipolar scale.

Uni- and bi-variate distribution of empathy and distress scores is displayed in Figure 2. As can be seen, both sets of labels are fairly close to a normal distribution. Empathy and

distress word-level ratings display only a moderate Pearson correlation of $r \approx .51$ which confirms that both are distinct constructs as already indicated by the qualitative analysis above. It is also highly consistent with Empathic Reactions where Buechel et al. (2018) found $r \approx .45$ between both sets of *document-level* ratings.

5.2. Clustering Analysis

To assess the face validity of the lexica, we partition the lexica into groups (clusters) of words that are semantically similar and simultaneously have similar ratings. Straight-forward clustering does not take the ratings into account, and is less interpretable. We use the Signed Spectral Clustering (SSC) algorithm to cluster words that are similar semantically and in their ratings (Sedoc et al., 2017b). Weighted edges are added between words such that words of similar empathy have positive connections and those of differing empathy are negative (see Sedoc et al. (2017a) for precise mathematical formulation). SSC minimizes the cumulative edge weights cut within clusters versus between clusters, while simultaneously minimizing the negative edge weights within the clusters, thus pulling words of similar empathy or distress into the same clusters and pushing those that differ away. We follow the method used by Sedoc et al. (2017b). In terms of linguistic analysis, the resulting clusters can help us describe the *language of empathy* by providing us with semantic groups of words which are high or low on either of the two empathy scales, i.e., allowing us to answer questions such as “what kind of words do people use when they feel empathic?”

As seen in Table 3, the clusters of words for high and low empathy and for high distress are strikingly well illustrated. There are many clusters of topics around situations where people feel empathy. Furthermore, there are lists of different negative emotions. The lists that are all places and people names are less useful obviously for psychological analysis. However, these lists are places where bad things happen, and people to whom bad things happen, which is useful for predictive models. Usable lexica must be interpretable, SSC allows us not only to give words and ratings, but also, groups of high magnitudes. These allow domain experts then to analyze and possibly modify the lexica.

5.3. Named Entities

Motivated by the observation that groups of named entities (NEs) play a (perhaps surprisingly) prominent role in our lexica, we used the clusters to derive an estimate of the number of NE entries in the resulting dataset. The authors manually labeled clusters which belong to either of the following classes of NEs: person (names), organization (including geopolitical entity), date or time, number (including units of measurement), and punctuation. If a cluster predominantly contains entries of one NE type, then all words within this cluster were counted as belonging to this category. The categorization was done independently for the empathy and distress lexica. The clusters tend to be consistent regarding NE types, as illustrated in Table 3. However, both false positives (FP; non-NE entries in NE clusters) as well as false negatives (FN; NE entries in non-NE clusters) do occur. This leads to slightly different es-

High Empathy	<ul style="list-style-type: none"> • grieve, grieving, loss, prayers, grief, heartbroken, losses, depression, condolences, widowed • wounds, wounded, scars, heal, blisters, trauma, wound, heals, bleeding, fasciitis • duckworth, salama, mansour, santiago, gilbert, fernandez, braves, vaughn, colonialism, crowe • minneapolis, neighborhoods, detroit, chicago, charlotte, cincinatti, brisbane, angeles, atlanta, drescher
Low Empathy	<ul style="list-style-type: none"> • fool, clueless, dumbass, idiotic, lazy, stupidity, morons, idiot, idiots, dumb • bother, slightest, anything, else, nobody, anybody, any, nothing, anyone, himself • loser, bs, moron, dingus, maniac, buffoon, ffs, loon, crap, psycho • wacky, bizarre, odd, creepy, weird, unnerving, masochistic, freaks, unusual, strange
High Distress	<ul style="list-style-type: none"> • homicide, killings, murdered, massacre, murdering, homicides, genocide, murder, murderers, killed • brutalized, assaulted, raped, bullied, tormented, harassed, detained, molested, reprimanded, beaten • horrific, witnessed, retched, wretched, atrocious, awful, horrid, foul, shoddy, unpleasant • horrifying, terrifying, harrowing, overdoses, suicides, deaths, suicide, gruesome, devastating, tragedy
Low Distress	<ul style="list-style-type: none"> • dunno, guessing, guess, gues, probably, assuming, maybe, clue, bet, assume • wont, knowlegde, alot, doesnt, isnt, wasnt, ahve, dont, didnt, exempt • sort, lot, bunch, sorts, type, whatever, plenty, depending, types, range • intact, stays, rememeber, keeping, keeps, always, kept, vague, rember, stay

Table 3: Signed Spectral Clustering results for qualitative analysis.

estimates of the number of NEs in the empathy and distress lexicon, respectively.⁸

The approximate NE counts are presented in Table 4. The estimates are largely consistent between the empathy and distress lexica, with only “Organization” displaying more pronounced differences. In total, the results indicate the ratio of NE entries in our lexica is roughly 6%.

Named Entity Class	Empathy	Distress
Person	136	140
Organization	238	165
Date	52	54
Number	128	117
Punctuation	48	30
Sum	602	506

Table 4: Cluster based estimate of the number of lexicon entries belonging to various classes of named entities.

6. Conclusion

This contribution reported on the creation of the first-ever lexica for empathy and distress using a newly proposed model: The Mixed-Level Feed-Forward Network (MLFFN) successfully learns word ratings from document-level ratings by backing out word ratings from a trained neural net, performing substantially better than methods that others have used for lexicon creation. Signed Spectral Clustering was applied to the resulting lexical scores to gain insights into the language of empathy. We look forward to further validating the lexica by using them in predictive models and psychological experiments, and to exploring the extent to which using the SHAP (SHapley Additive exPlanations) (Lundberg and Lee, 2017) calculations of feature importance for CNNs, RNNs, or Transformers

⁸Although both dictionaries comprise the same list of word types, the clusters are split differently leading to deviating FP and FN rates.

improve lexicon quality over the simple neural nets which we used.

7. Acknowledgements

We would like to thank Daphne Ippolito, Reno Kriz, and the anonymous reviewers for their helpful feedback. This work was partially supported by João Sedoc’s Microsoft Dissertation Grant. Sven Buechel was partially funded by the German Federal Ministry for Economic Affairs and Energy (funding line “Big Data in der makroökonomischen Analyse”; Fachlos 2; GZ 23305/003#002). He would also like to thank his doctoral advisor Udo Hahn, JULIE Lab, for supporting his research visit at the University of Pennsylvania. This research is supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the BETTER Program contract #2019-19051600005. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

8. Bibliographical References

- Abdul-Mageed, M., Buffone, A., Peng, H., Eichstaedt, J. C., and Ungar, L. H. (2017). Recognizing pathogenic empathy in social media. In *Eleventh International AAAI Conference on Web and Social Media*, pages 448–451, Montreal, Canada, May. Association for the Advancement of Artificial Intelligence (AAAI).
- Alam, F., Danieli, M., and Riccardi, G. (2017). Annotating and modeling empathy in spoken conversations. *Computer Speech & Language*, 50:40–61.
- Baccianella, S., Esuli, A., and Sebastiani, F. (2010). SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta,

- Malta, May. European Language Resources Association (ELRA).
- Batson, C. D., Fultz, J., and Schoenrade, P. A. (1987). Distress and empathy: Two qualitatively distinct vicarious emotions with different motivational consequences. *Journal of personality*, 55(1):19–39.
- Batson, C. D., Polycarpou, M. P., Harmon-Jones, E., Imhoff, H. J., Mitchener, E. C., Bednar, L. L., Klein, T. R., and Highberger, L. (1997). Empathy and attitudes: Can feeling for a member of a stigmatized group improve feelings toward the group? *Journal of personality and social psychology*, 72(1):105.
- Bloom, P. (2016). *Against Empathy: The Case for Rational Compassion*. Ecco.
- Bradley, M. M. and Lang, P. J. (1999). Affective Norms for English Words (ANEW): Stimuli, instruction manual and affective ratings. Technical Report C-1, The Center for Research in Psychophysiology, University of Florida, Gainesville, Florida, USA.
- Brysbaert, M., Warriner, A. B., and Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3):904–911.
- Buechel, S. and Hahn, U. (2017). EmoBank: Studying the impact of annotation perspective and representation format on dimensional emotion analysis. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 578–585, Valencia, Spain, April. Association for Computational Linguistics.
- Buechel, S. and Hahn, U. (2018a). Emotion representation mapping for automatic lexicon construction (mostly) performs on human level. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2892–2904, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Buechel, S. and Hahn, U. (2018b). Word emotion induction for multiple languages as a deep multi-task learning problem. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1907–1918, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Buechel, S., Buffone, A., Slaff, B., Ungar, L., and Sedoc, J. (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, October–November. Association for Computational Linguistics.
- Buffone, A. E., Poulin, M., DeLury, S., Ministerio, L., Morrisson, C., and Scalco, M. (2017). Don’t walk in her shoes! Different forms of perspective taking affect stress physiology. *Journal of Experimental Social Psychology*, 72:161 – 168.
- Chikovani, G., Babuadze, L., Iashvili, N., Gvalia, T., and Surguladze, S. (2015). Empathy costs: Negative emotional bias in high empathisers. *Psychiatry Research*, 229(1-2):340–346.
- Chollet, F. et al. (2015). Keras. <https://keras.io>.
- Clos, J., Wiratunga, N., and Massie, S. (2017). Towards explainable text classification by jointly learning lexicon and modifier terms. In *IJCAI-17 Workshop on Explainable AI (XAI)*, Melbourne, Australia, August. <http://home.earthlink.net/~dwaha/research/meetings/ijcail17-xai>.
- Davis, M. H. (1980). A multidimensional approach to individual differences in empathy. *JSAS Catalog of Selected Documents in Psychology*, 10:85.
- Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, 44:113–126.
- De Bruyne, L., Atanasova, P., and Augenstein, I. (2019). Joint emotion label space modelling for affect lexica. <http://arxiv.org/abs/1911.08782>.
- De Waal, F. (2009). *The Age of Empathy: Nature’s Lessons for a Kinder Society*. Crown.
- Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., et al. (2015). Psychological language on twitter predicts county-level heart disease mortality. *Psychological Science*, 26(2):159–169.
- Ekman, P. (1992). An argument for basic emotions. *Cognition & Emotion*, 6(3-4):169–200.
- Engelthaler, T. and Hills, T. T. (2018). Humor norms for 4,997 English words. *Behavior Research Methods*, 50:1116–1124.
- Fung, P., Dey, A., Siddique, F. B., Lin, R., Yang, Y., Wan, Y., and Chan, H. Y. R. (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, June. Association for Computational Linguistics.
- Hamilton, W. L., Clark, K., Leskovec, J., and Jurafsky, D. (2016). Inducing domain-specific sentiment lexicons from unlabeled corpora. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 595–605, Austin, Texas, November. Association for Computational Linguistics.
- Hatzivassiloglou, V. and McKeown, K. R. (1997). Predicting the semantic orientation of adjectives. In *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pages 174–181, Madrid, Spain, July. Association for Computational Linguistics.
- Hofmann, M. J., Kuchinke, L., Tamm, S., Vö, M. L. H., and Jacobs, A. M. (2009). Affective processing within 1/10th of a second: High arousal is necessary for early facilitative processing of negative but not positive words. *Cognitive, Affective, & Behavioral Neuroscience*, 9(4):389–397, December.
- Hutto, C. J. and Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International AAAI Conference on Weblogs and Social Media*, pages 216–225, Ann Arbor,

- USA, June. Association for the Advancement of Artificial Intelligence (AAAI).
- Iyyer, M., Manjunatha, V., Boyd-Graber, J., and Daumé III, H. (2015). Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1681–1691, Beijing, China, July. Association for Computational Linguistics.
- Khanpour, H., Caragea, C., and Biyani, P. (2017). Identifying empathetic messages in online health communities. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 246–251, Taipei, Taiwan, November. Asian Federation of Natural Language Processing.
- Khosla, S., Chhaya, N., and Chawla, K. (2018). Aff2Vec: Affect-enriched distributional word representations. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2204–2218, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Kingma, D. and Ba, J. (2015). Adam: A method for stochastic optimization. In *Third International Conference on Learning Representations*, San Diego, USA, May. <https://arxiv.org/abs/1412.6980>.
- Kulkarni, P. V., Nagori, M. B., and Kshirsagar, V. P. (2019). An in-depth survey of techniques employed in construction of emotional lexicon. In *Information and Communication Technology for Intelligent Systems. Proceedings of ICTS 2018, Volume 1*, pages 609–620, Ahmedabad, India, April. Springer.
- Leveau, N., Jhean-Larose, S., Denhière, G., and Nguyen, B.-L. (2012). Validating an interlingual metanorm for emotional analysis of texts. *Behavior Research Methods*, 44(4):1007–1014.
- Li, M., Lu, Q., Long, Y., and Gui, L. (2017). Inferring affective meanings of words from word embedding. *IEEE Transactions on Affective Computing*, 8(4):443–456.
- Liu, L., Preotiuc-Pietro, D., Samani, Z. R., Moghaddam, M. E., and Ungar, L. (2016). Analyzing personality through social media profile picture choice. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media*, pages 211–220, Cologne, Germany, May. Association for the Advancement of Artificial Intelligence (AAAI).
- Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4765–4774, Long Beach, USA, December. Curran Associates.
- McQuiggan, S. W. and Lester, J. C. (2007). Modeling and evaluating empathy in embodied companion agents. *International Journal of Human-Computer Studies*, 65(4):348–360.
- Mihalcea, R. and Liu, H. (2006). A corpus-based approach to finding happiness. In *Computational Approaches to Analyzing Weblogs, Papers from the 2006 AAAI Spring Symposium, Technical Report SS-06-03*, pages 139–144, Stanford, USA, March. Association for the Advancement of Artificial Intelligence (AAAI).
- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., and Joulin, A. (2018). Advances in pre-training distributed word representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Mohammad, S. and Bravo-Marquez, F. (2017a). Emotion intensities in tweets. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 65–77, Vancouver, Canada, August. Association for Computational Linguistics.
- Mohammad, S. and Bravo-Marquez, F. (2017b). WASSA-2017 shared task on emotion intensity. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 34–49, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Mohammad, S. (2012). #Emotional tweets. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 246–255, Montréal, Canada, 7-8 June. Association for Computational Linguistics.
- Mohammad, S. (2018). 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, July. Association for Computational Linguistics.
- Monnier, C. and Syssau, A. (2008). Semantic contribution to verbal short-term memory: Are pleasant words easier to remember than neutral words in serial recall and serial recognition? *Memory & Cognition*, 36(1):35–42, January.
- Montefinese, M., Ambrosini, E., Fairfield, B., and Mammarella, N. (2014). The adaptation of the Affective Norms for English Words (ANEW) for Italian. *Behavior Research Methods*, 46(3):887–903.
- Osgood, C., Suci, G. J., and Tannenbaum, P. (1957). *The Measurement of Meaning*. University of Illinois.
- Pavey, L., Greitemeyer, T., and Sparks, P. (2012). “I help because I want to, not because you tell me to”: Empathy increases autonomously motivated helping. *Personality and Social Psychology Bulletin*, 38(5):681–689.
- Pennebaker, J. (2011). *The Secret Life of Pronouns: What Our Words Say About Us*. Bloomsbury Press.
- Pérez-Rosas, V., Mihalcea, R., Resnicow, K., Singh, S., and An, L. (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, July. Association for Computational Linguistics.
- Poulin, M. J., Brown, S. L., Dillard, A. J., and Smith, D. M. (2013). Giving to others and the association between

- stress and mortality. *American Journal of Public Health*, 103(9):1649–1655.
- Preoțiuc-Pietro, D., Schwartz, H. A., Park, G., Eichstaedt, J., Kern, M., Ungar, L., and Shulman, E. (2016). Modelling valence and arousal in Facebook posts. In *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 9–15, San Diego, California, June. Association for Computational Linguistics.
- Redondo, J., Fraga, I., Padrón, I., and Comesaña, M. (2007). The Spanish adaptation of ANEW (Affective Norms for English Words). *Behavior Research Methods*, 39(3):600–605.
- Rill, S., Drescher, J., Reinel, D., Scheidt, J., Schuetz, O., Wogenstein, F., and Simon, D. (2012). A generic approach to generate opinion lists of phrases for opinion mining applications. In *Proceedings of the First International Workshop on Issues of Sentiment Discovery and Opinion Mining*, Beijing, China, August. Association for Computing Machinery.
- Ruppenhofer, J., Wiegand, M., and Brandes, J. (2014). Comparing methods for deriving intensity scores for adjectives. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 117–122, Gothenburg, Sweden, April. Association for Computational Linguistics.
- Sap, M., Park, G., Eichstaedt, J., Kern, M., Stillwell, D., Kosinski, M., Ungar, L., and Schwartz, H. A. (2014). Developing age and gender predictive lexica over social media. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1146–1151, Doha, Qatar, October. Association for Computational Linguistics.
- Schwartz, H. A., Eichstaedt, J. C., Dziurzynski, L., Kern, M. L., Blanco, E., Kosinski, M., Stillwell, D., Seligman, M. E., and Ungar, L. H. (2013). Toward personality insights from language exploration in social media. In *Analyzing Microtext: Papers from the 2013 AAAI Spring Symposium*, pages 72–79, Stanford, USA, March. Association for the Advancement of Artificial Intelligence (AAAI).
- Schwartz, H. A., Giorgi, S., Sap, M., Crutchley, P., Ungar, L., and Eichstaedt, J. (2017). DLATK: Differential language analysis ToolKit. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 55–60, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Sedoc, J., Gallier, J., Foster, D., and Ungar, L. (2017a). Semantic word clusters using signed spectral clustering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 939–949, Vancouver, Canada, July. Association for Computational Linguistics.
- Sedoc, J., Preoțiuc-Pietro, D., and Ungar, L. (2017b). Predicting emotional word ratings using distributional representations and signed clustering. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 564–571, Valencia, Spain, April. Association for Computational Linguistics.
- Tausczik, Y. R. and Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1):24–54.
- Turney, P. D. and Littman, M. L. (2003). Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21(4):315–346.
- Wang, L. and Xia, R. (2017). Sentiment lexicon construction with representation learning based on hierarchical sentiment supervision. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 502–510, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Wang, J., Yu, L.-C., Lai, K. R., and Zhang, X. (2016). Dimensional sentiment analysis using a regional CNN-LSTM model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 225–230, Berlin, Germany, August. Association for Computational Linguistics.
- Warriner, A. B., Kuperman, V., and Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4):1191–1207.
- Yee, L. T. S. (2017). Valence, arousal, familiarity, concreteness, and imageability ratings for 292 two-character Chinese nouns in Cantonese speakers in Hong Kong. *PLoS one*, 12(3):e0174569.
- Yu, L.-C., Wang, J., Lai, K. R., and Zhang, X. (2017). Refining word embeddings for sentiment analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 534–539, Copenhagen, Denmark, September. Association for Computational Linguistics.