

Learning Emotion from 100 Observations: Unexpected Robustness of Deep Learning under Strong Data Limitations

Sven Buechel*

Friedrich-Schiller-Universität Jena
sven.buechel@uni-jena.de

João Sedoc

New York University
jsedoc@stern.nyu.edu

H. Andrew Schwartz

Stony Brook University
has@cs.stonybrook.edu

Lyle Ungar

University of Pennsylvania
ungar@cis.upenn.edu

Abstract

One of the major downsides of Deep Learning is its supposed need for vast amounts of training data. As such, these techniques appear ill-suited for NLP areas where annotated data is limited, such as less-resourced languages or emotion analysis, with its many nuanced and hard-to-acquire annotation formats. We conduct a questionnaire study indicating that indeed the vast majority of researchers in emotion analysis deems neural models inferior to traditional machine learning when training data is limited. In stark contrast to those survey results, we provide empirical evidence for English, Polish, and Portuguese that commonly used neural architectures can be trained on surprisingly few observations, outperforming n -gram based ridge regression on only 100 data points. Our analysis suggests that high-quality, pre-trained word embeddings are a main factor for achieving those results.

1 Introduction

Deep Learning (DL) has radically changed the rules of the game in NLP by boosting performance figures in almost all application areas. Yet in contrast to more conventional techniques, such as n -gram based linear models, *neural* methodologies seem to rely on vast amounts of training data, as is obvious in areas such as machine translation (Vaswani et al., 2017) or representation learning for individual words (Mikolov et al., 2013; Pennington et al., 2014) or contextualized word sequences (Devlin et al., 2019; Yang et al., 2019; Joshi et al., 2020).

With this profile, DL seems ill-suited for many prediction tasks in sentiment and subjectivity analysis (Balahur et al., 2014). For the widely studied problem of polarity prediction (distinguishing only between *positive* and *negative* emotion), training data is relatively abundant especially for the social media domain (Rosenthal et al., 2017). However, in recent years, there has been a growing interest in more nuanced and informative annotation formats for affective states (Bostan and Klinger, 2018; De Bruyne et al., 2019). Such annotation schemes often follow distinct psychological theories such as the dimensional approach to emotion representation (Bradley and Lang, 1994) or basic emotions (Ekman, 1992). Yet, annotating for more complex representations of affective states seems to be significantly harder in terms of both time consumption and inter-annotator agreement (IAA) (Strapparava and Mihalcea, 2007). Adding even more complexity, computational work following this trend often uses numerical scores as target variables making, emotion analysis a regression, rather than a classification problem (Buechel and Hahn, 2016; Mohammad et al., 2018). What makes this situation even worse is that, first, we currently have a situation where there is no community-wide consensus on how emotion should be represented. That is, different ways of annotating emotion (see, e.g., Table 2) compete with each other, leading to decreased inter-operability of language resources and provoking additional data sparsity (Buechel and Hahn, 2018b). And, second, especially large-scale annotated corpora are almost exclusively available for English, leaving most of the world’s languages with little or no gold data at all.

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

* Work partially conducted at the University of Pennsylvania.

Data Requirements for Deep Learning in Emotion Analysis

We're interested in your beliefs about what training size is necessary for deep learning techniques. Your response will be used for academic research and stored anonymously. Thank you very much!

Consider the task of fine-grained emotion analysis, using what you believe to be the best deep learning architecture (e.g., RNN, CNN, Self-Attention) with input from pre-trained word embeddings (e.g., word2vec, GloVe; NOT contextualized embeddings like ELMO or BERT):

1. How many training examples do you believe are necessary for deep learning to provide results in line with traditional discriminative learning (e.g. SVM, penalized linear regression, random forests)?
2. How many observations do you believe are necessary for deep learning approaches to provide a clear benefit over traditional discriminative learning?

Thank you for completing the survey! Do you have any additional comments regarding this questionnaire?

Figure 1: Survey on expected data requirements of deep learning.

For the social media domain, this lack of gold data can be partly countered by (pre-)training with distant supervision using signals such as emojis or hashtags as a surrogate for manual annotation (Mohammad and Kiritchenko, 2015; Felbo et al., 2017; Abdul-Mageed and Ungar, 2017). Yet, this procedure is less viable for target domains other than social media, as well as for predicting other subjective phenomena such as empathy, uncertainty, or personality (Khanpour et al., 2017; Rubin, 2007; Liu et al., 2017). Besides pre-training the entirety of the model with distant supervision, an alternative strategy is pre-training word representations, only. This approach is feasible for a wide range of languages, including otherwise less-resourced ones, since raw text is much more readily available than gold data, e.g., through Wikipedia (Grave et al., 2018). Very recently, *contextualized* word representations generated by pre-trained language models have established themselves as a powerful alternative (Peters et al., 2018; Devlin et al., 2019).

In summary, deep learning supposedly depends on vast amounts of annotated data—and this seems particularly troublesome for the field of emotion analysis because such phenomena are intrinsically hard to annotate. However, we suspect that this *gold* data dependency may, for emotion analysis at least, be less severe than anticipated because large, pre-trained embedding models already seem to encode word-level emotion quite well (Du and Zhang, 2016; Li et al., 2017; Buechel and Hahn, 2018c), possibly allowing to fit sentence-level DL architectures on rather small datasets. If that was true, it would be the *reputation* of DL rather than its actual characteristics which prevent its wider application for emotion analysis in low-resource environments.

Contribution. We start by quantifying the expectations of the research community regarding the data requirements of DL. To this end, we first conduct a questionnaire study among NLP researchers in the field of emotion analysis finding that the median respondent expects DL to be viable only from 10,000 training examples onward. Next, we perform a series of experiments on English, Polish, and Portuguese emotion corpora. In contrast to the survey results, we show that commonly used architectures can be fitted on as little as 100 data points and still outperform supposedly more robust *n*-gram based approaches. We believe these findings potentially open up DL to many low-resource areas and, by extension, wider cross-lingual or cross-domain applications.

2 Survey

We conducted a short questionnaire study asking the research community about their beliefs regarding data requirements of deep learning in the context of emotion analysis. We included two questions, one asking for the number of training examples “necessary for deep learning to provide results *in line* with traditional discriminative learning” (question 1), the other asking for the number of examples necessary for “deep learning approaches to provide a *clear benefit* over traditional discriminative learning” (question 2). The full text of the questionnaire is given in Figure 1. Participants were instructed to focus on *non*-contextualized word representations (in line with our latter experiments; the use of contextualized word embeddings is left for future work) being used as input to the, from their view, most suitable model

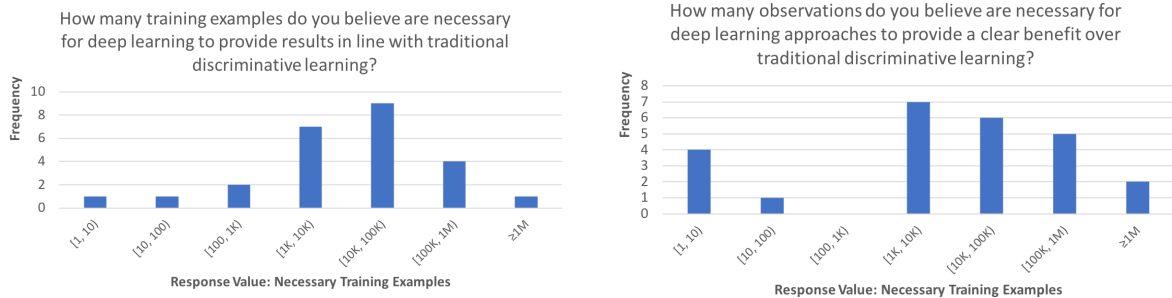


Figure 2: Responses to survey questions 1 (left) and 2 (right).

Corpus	Language	Size	Annotation	Emb. Alg.	Dims.	Emb. Data	Emb. Data Size
SE07	English	1000	BE6 [1, 100]	word2vec	300	Google News	100B tokens
WASSA	English	≤ 2252	BE4 [0, 1]	word2vec	400	Twitter	400M tweets
ANPST	Polish	718	VAD [1, 9]	FastText	300	Wikipedia	4B tokens
MAS	Portuguese	192	VAD [1, 9] + BE5 [1, 5]	FastText	300	Wikipedia	4B tokens

Table 1: Annotated corpora and embedding models used for experiments; with language, number of instances, annotation format, embedding algorithm, embedding dimensions, and dataset (size) embeddings were trained on.

architecture.

To recruit participants, we queried the ACL Anthology for papers from between 2016 and 2018 using the keyword “emotion”, we collected all email addresses in the author field of all retrieved PDFs (166 papers in total). Invitations to participate in the survey were sent to the resulting 391 email addresses on February 28, 2019. We received 26 responses within four weeks (6% response rate). One response (stating in the optional comment field that no numeric answer could be given) was excluded. Figure 2 shows the distribution of the 25 remaining responses on a logarithmic scale.

As can be seen, the responses to both questions clearly support the intuition that deep learning is thought of as being dependent on large amounts of training data by the scientific community. In both cases, the median response was 10,000. Perhaps surprisingly, a total of 5 participants stated that fewer than 100 instances are necessary for deep learning to show clear improvements over more traditional methods (compared to 20 who believed the opposite), whereas only 2 believed that less than 100 instances are enough to show results “in line with” traditional methods. Inspecting the individual responses, we found that this discrepancy stems from a minority of participants (4 of out 25) who indicated that traditional learning performs worse than DL on small datasets but may catch up as dataset sizes grow.

Another interesting, most likely related outcome is that the responses to question 2 show a bi-modal distribution: While a minority of 5 participants believed that DL approaches are superior below 100 observations, the vast majority of participants (20) states that 1,000 or more instances are necessary for that. Yet, no one responded with a number between 100 and 1,000.

While we do not validate the claim of this minority, the remainder of the paper provides strong evidence that the majority of the participants largely overestimated the data requirements of deep learning.

3 Data

For the following study, we selected four small (< 3000 instances) and typologically diverse datasets described below. Pre-trained, publicly available word2vec (Mikolov et al., 2013) and FastText vectors (Bojanowski et al., 2017) of matching language and target domain were used as model input. Table 1 summarizes the employed data. Illustrative examples of the particular styles and annotation formats of those corpora are provided in Table 2.

SE07: The test set of SemEval 2007 Task 14 (Strapparava and Mihalcea, 2007) comprises 1000 English news headlines that are annotated according to six Basic Emotions, joy, anger, sadness, fear, disgust,

Corpus	Text	Val	Aro	Dom	Joy	Ang	Sad	Fea	Dis	Sur
SE07	<i>Inter Milan set Serie A win record</i>	-	-	-	50	2	0	0	0	9
	<i>TBS to pay \$2M fine for ad campaign bomb scare</i>	-	-	-	11	25	28	45	32	43
WASSA	<i>@TheRevAl please tell us why 'protesting' injustice requires #burning #beating and #looting terrible optics #tousaintromain is true leader!</i>	-	-	-	-	.73	-	-	-	-
	<i>@TauDeltaPhiDK THANK YOU FOR MY OBAMA CUT OUT!!!!!! I am elated that he's back home 😊</i>	-	-	-	.83	-	-	-	-	-
ANPST	<i>Decyzje podjęte w przeszłość i kształtują naszą teraźniejszość. 'Decisions made in the past shape our present.'</i>	4.9	4.1	5.8	-	-	-	-	-	-
	<i>Dopóki walczysz i podejmujesz starania, jesteś zwycięzcą. 'As long as you fight and keep trying, you are a winner.'</i>	7.3	5.3	7.4	-	-	-	-	-	-
MAS	<i>A praia é espetacular. 'The beach is spectacular.'</i>	8.0	4.0	6.7	4.2	1.0	1.0	1.0	1.0	-
	<i>A tinta é azul. 'The ink is blue.'</i>	5.0	3.9	5.5	1.3	1.0	1.0	1.0	1.0	-

Table 2: Exemplary entries from our four datasets illustrating differences in linguistic characteristics and emotion annotation scheme. Emotion variables: valence, arousal, dominance, joy, anger, sadness, fear, disgust, and surprise. English translations for ANPST and MAS were provided by the respective dataset creators.

and surprise on a $[0, 100]$ -scale (BE6 annotation format). The news headlines are quite short and rather objective, being written by professional journalists. However, they may still elicit strong emotional reactions in readers as illustrated in Table 2. For this corpus, we used the word2vec embeddings trained on Google News.¹

WASSA: The English Twitter dataset of the WASSA 2017 shared task (Mohammad and Bravo-Marquez, 2017b) contains four subsets, one for each of the first four basic emotions, annotated on a $[0, 1]$ scale (BE4 format). Their sizes vary between 1533 and 2252 samples (union of the respective train, dev, and test set). Being a Twitter corpus, these data display features typical for the social media domain, e.g., extensive use of colloquialism, emojis, and explicit language, as well as platform-specific phenomena such as hashtags ($\#$) and user mentions ($@$). Note that different from other corpora, here each individual instance is annotated according to only one emotion variable, whereas in SE07, ANPST, and MAS every instance is annotated for all variables covered by the respective dataset. Here, we used Twitter word2vec embeddings by Godin et al. (2015).

ANPST: The Affective Norms for Polish Short Texts (Imbir, 2017) is a dataset designed as stimulus in psychological experiments. It is annotated according to valence, arousal, and dominance on a $[1, 9]$ -scale (VAD). ANPST comprises sentences of various genres (such as proverbs, jokes, literature quotes, or newswire material) from a wide range of sources (Imbir, 2017). The resulting selection of raw data seems often quite complex and ambiguous in terms of the elicited emotion (see examples in Table 2). We used the FastText embeddings by Grave et al. (2018) trained on the Polish Wikipedia.

MAS: Like ANPST, the Minho Affective Sentences (Pinheiro et al., 2017) is a dataset designed by psychologists, also being annotated according to valence, arousal, and dominance on a $[1, 9]$ -scale (VAD). Yet, additionally, MAS is also annotated according to the first five Basic Emotions (omitting ‘surprise’) on a $[1, 5]$ -scale (BE5). It consists of very short situation descriptions in the third person in European Portuguese (Pinheiro et al., 2017). Those sentences were purposefully constructed by psychologists to be simple in language and familiar in content for a large proportion of the population. This was done to make the dataset more widely applicable as experimental stimulus, yet also resulted in a slightly ar-

¹code.google.com/archive/p/word2vec/

Model	Description
Ridge _{ngram}	n -gram features with $n \in \{1, 2, 3\}$; feature normalization; automatically chosen regularization coefficient from $\{10^{-4}, 10^{-3}, \dots, 10^4\}$
Ridge _{BV}	<i>bag of vectors</i> -features; regularization coefficient chosen as in ‘Ridge _{ngram} ’
FFN	<i>bag of vectors</i> -features; two dense layers (256 and 128 units)
CNN	one conv. layer (filter size 3, 128 channels), max-pooling layer with .5 dropout; dense layer (128 units)
GRU	recurrent layer (128 units, unidirectional); last timestep receives .5 vertical dropout and is fed into a dense layer (128 units)
LSTM	identical to ‘GRU’
CNN-LSTM	conv. layer as in ‘CNN’; max-pooling layer (pool size 2, stride size 1) with .5 dropout; LSTM identical to ‘GRU’

Table 3: Model-specific design choices.

tificial style. MAS is the smallest dataset considered, having only 192 instances. We used the FastText embeddings by Grave et al. (2018), trained on the Portuguese Wikipedia.

4 Methods

We provide two distinct linear baseline models which both rely on Ridge regression, an ℓ^2 -regularized version of linear regression. The first one, Ridge_{ngram}, is based on n -gram features where we use $n \in \{1, 2, 3\}$. The second one, Ridge_{BV} uses *bag-of-vectors* features, i.e., the pointwise mean of the embeddings of the words in a text. Regarding the deep learning approaches, we compare Feed-Forward Networks (FFN), Gated Recurrent Unit Networks (GRU), Long Short-Term Memory Networks (LSTM), Convolutional Neural Networks (CNN), as well as a combination of the latter two (CNN-LSTM) (Cho et al., 2014; Hochreiter and Schmidhuber, 1997; Kalchbrenner et al., 2014).

Since holding out a dev set from the already limited training data does not seem feasible for some of the datasets (see Table 1), we decided to instead use constant hyperparameter settings across all corpora. We also keep most hyperparameters constant between models. Hence, hyperparameter choices followed well-established recommendations described in the next paragraph.

The input to the DL models is based on pre-trained word vectors. ReLu activation was used everywhere except in recurrent layers. Dropout is used for regularization with a probability of .2 for embedding layers and .5 for dense layers following the recommendations by Srivastava et al. (2014). We use .5 dropout also on other types of layers where it would conventionally be considered too high (e.g. on max pooling layers). Our models are trained for 200 epochs using the Adam optimizer (Kingma and Ba, 2015) with a fixed learning rate of .001 and a batch size of 32. Word embeddings were not updated during training. Since, in compliance with our gold data, we treat emotion analysis as a regression problem (Buechel and Hahn, 2016), the output layers of our models consist of an affine transformation, i.e., a dense layer without non-linearity. To reduce the risk of overfitting on such small data sets, we used relatively simple models both in terms of the number of layers and units in them (mostly 2 and 128, respectively). An overview of our models and details about their *individual* hyperparameter settings are provided in Table 3. `Keras.io` and `scikit-learn.org` (Pedregosa et al., 2011) were used for the implementation.

	SE07	WASSA	ANPST	MAS	Mean
Ridge _{ngram}	.53	.67	.32	.16	.42
Ridge _{BV}	.62	.64	.52	.62	.60
CNN-LSTM	.66	.69	.50	.63	.62
CNN	.67	.70	.47	.61	.62
FFN	.67	.69	.50	.65	.63
LSTM	.65	.73	.52	.65	.64
GRU	.67	.73	.54	.66	.65

Table 4: Comparative results of the 10×10 -cross-validation in Pearson’s r ; averaged over all variables of the respective annotation format.

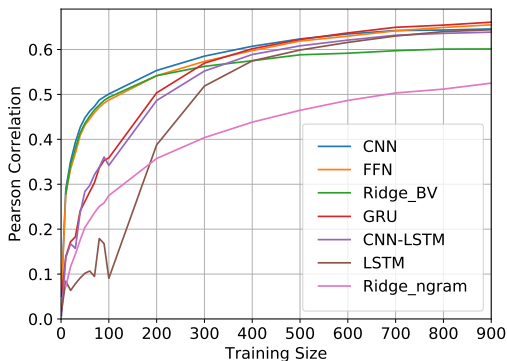


Figure 3: Comparison of model performance vs. training size on the SE07 dataset in Pearson’s r .

5 Experimental Results

5.1 Repeated Cross-Validation

Given our small datasets, conventional 10-fold cross-validation (CV) would lead to very small test splits (only 19 instances in the case of MAS) thus causing high variance between the individual splits and, ultimately, even regarding the average of all 10 runs. Therefore, we *repeat 10-fold CV ten times* (10×10 -CV) with different data splits, then averaging the results (Dietterich, 1998). Performance is measured as Pearson correlation r between predicted and human gold ratings. To further increase reliability, identical data splits were used for each of the approaches under comparison. Results are given in Table 4.

All DL approaches (FFN, CNN, GRU, LSTM, CNN-LSTM) yield a satisfying performance of $r > .6$ on average over all corpora, despite the small data sizes. Each one of them clearly outperforms Ridge_{ngram}, representing more conventional learning techniques, on every single dataset. This stands in sharp contrast to our survey results where the median respondent indicated that DL would need at least 10,000 instances to provide a clear benefit over conventional techniques—the datasets we employed are between 4 and 50 times smaller.

Overall, the GRU performs best. However, differences between the DL models are quite small on average. (We emphasize that our primary concern is to compare deep vs. conventional learning techniques under data limitations whereas comparisons within the group of DL architectures are secondary.) Perhaps surprisingly, Ridge_{BV}, which takes a middle ground between DL and conventional approaches, also performs very competitively. Given its low computational cost, our results indicate that this model may constitute an excellent baseline.

Observe that Ridge_{BV} and FFN rely solely on lexical information—their inputs are computed by mere averaging of word embeddings whereas CNN, LSTM, GRU, and CNN-LSTM learn their own composition functions from gold data. Still, the former two both display satisfying performance. This suggests that the quality of the pre-trained embeddings may be a key factor for their strong results.

	FFN	CNN	GRU	LSTM	CNN-LSTM	Mean
Learned	.24	.23	.38	.26	.30	.28
Tuned	.59	.55	.59	.59	.57	.58
Frozen	.63	.62	.65	.64	.62	.63

Table 5: Comparison of embedding training strategies (average Pearson’s r over all datasets).

5.2 Embedding Training Strategies

To further examine this conjecture, we repeated the above experiment two more times, altering the training strategy of the embeddings (only applicable to DL models). Instead of using pre-trained vectors without updating them (**Frozen**), we looked at embeddings which were either randomly initialized and updated (**Learned**) or pre-trained and updated (**Tuned**). As can be seen from Table 5, both strategies involving pre-trained vectors (frozen and tuned) outperform learned word embeddings by a large margin (about 30%-points on average). Frozen embeddings yield the highest performance, even outperforming fine-tuned vectors (5%-point margin on average), a possible reason being that the large increase in the number of parameters leads to overfitting.

5.3 Training Size vs. Model Performance

We will now continue to explore the unexpected behavior of DL architectures by further limiting the available training data. For each number $N \in \{1, 10, 20, \dots, 100, 200, \dots, 900\}$, we randomly sampled N instances from the SE07 corpus for training and tested on the held-out data. This procedure was repeated 100 times for each of the training data sizes before averaging the results. Each of the models was evaluated with identical data splits. The outcome of this experiment is depicted in Figure 3. As can be seen, recurrent models suffer only a moderate loss of performance down to a third of the original training data (about 300 observations). The CNN, FFN, and Ridge_{BV} models remain stable even longer—their performance only begins to decline rapidly at about 100 instances. In contrast, Ridge_{ngram} declines more steadily yet its overall performance is much lower as well. Most notably, all DL models but the LSTM *always* performed better than the conventional Ridge_{ngram} baseline no matter how little training data was used.

	Joy	Anger	Sadness	Fear	Disgust	Surprise	Mean
WINNER	.23	.32	.41	.45	.13	.17	.28
IAA	.60	.50	.68	.64	.45	.36	.54
BECK	.59	.65	.70	.74	.54	.47	.62
GRU	.60	.70	.75	.77	.61	.53	.66

Table 6: Comparison of previously reported results, human performance (IAA), and our proposed GRU model on the SE07 dataset in Pearson’s r .

Official Rank	Team/System	Joy	Anger	Sadness	Fear	Mean
1	Prayas	.762	.765	.732	.732	.747
2	IMS	.726	.767	.690	.705	.722
3	SeeNet	.698	.745	.715	.676	.708
–	Our Work	.658	.668	.724	.717	.692
4	UWaterloo	.699	.703	.693	.643	.685

Table 7: Comparison against official WASSA 2017 shared task results (in Pearson’s r).

5.4 Comparison against Previous Work

The above experiments have shown that our DL models perform robustly under strong data limitations, beating a conventional baseline in the vast majority of cases. Yet, perhaps this was achieved by designing overly simple network architectures, thus trading an excessive amount of performance for robustness in low-data scenarios. To rule out this possibility, we will now move forward and compare our findings against previous work.

SemEval 2007 Affective Text First, we compare our best performing model, the GRU, against previously reported results for the SE07 corpus. Table 6 provides the performance of the winning system of the original shared task (WINNER; Chaumartin (2007)), the inter-annotator agreement (IAA) as given by the organizers (Strapparava and Mihalcea, 2007), the performance by Beck (2017), the highest one reported for this dataset so far (BECK), as well as the results for our GRU from the 10×10 -CV set-up.

As can be seen, the GRU established a new state-of-the-art surpassing the previous one by about 4%-points on average over all emotion categories. The difference is statistically significant (two-tailed one-sample t -test comparing the results of the 10 cross-validation runs against the reported performance by Beck (2017); $p < .001$). Our GRU also outperforms IAA, as already did BECK. This may sound improbable at first glance. However, Strapparava and Mihalcea (2007) employ a rather weak notion of human performance which is—broadly speaking—based on the reliability of a single human rater.² Interestingly, the GRU shows particularly large improvements over human performance for categories where the IAA is low (anger, disgust, and surprise).

WASSA 2017 Shared Task Data Table 7 displays the official results of the four best systems (out of 21 submissions) of the WASSA 2017 shared task (Mohammad and Bravo-Marquez, 2017b) as well as the performance our GRU achieved. For this experiment, we deviated from the above 10×10 -CV set-up but instead used the official train-dev-test split for comparability. As for all experiments in this paper, hyperparameters were kept constant and were not adjusted to this dataset. Consequently, train and dev sets were combined for training. Training and testing were repeated ten times with different random seeds but otherwise identical configuration following the recommendation by Reimers and Gurevych (2018). Table 7 shows our *average* performance over those ten runs.

As can be seen, our GRU performs very competitively and would have been ranked fourth place, outperforming 18 out of 21 submissions. The difference to the next lower-performing system (UWaterloo) is statistically significant (two-tailed one-sample t -test comparing our ten runs against their official results; $p < .001$).

6 Conclusion

Annotating emotion is necessarily subjective thus making gold data in this area particularly rare. As such, applying DL may seem ill-advised since supposedly large amounts of training data are required. But is this really the case? We started our investigation by conducting a survey among researchers in emotion analysis. 80% of the respondents believed that DL is superior to traditional machine learning techniques only when at least 1,000 training examples are available. Half of the participants even believed that 10,000 or more examples are necessary. Putting this popular notion to the test, we provided the first examination of neural emotion analysis under severe data constraints, featuring five distinct neural architectures and three typologically diverse languages. In stark contrast to the survey results, we found that *all* architectures could be fitted on datasets comprising as little as 200 observations, CNNs and FFNs even being robust on 100 observations. A subsequent analysis indicated that high-quality, pre-trained word embeddings are a key factor in achieving those results. In the future, we would like to extend this work to *contextualized* word representations, e.g., by ELMo or BERT (Peters et al., 2018; Devlin et al., 2019).

²Instead, other approaches to IAA computation for numerical values, such as split-half or inter-study reliability, constitute a more challenging comparison since they are based on the reliability of *many* raters, not one (Mohammad and Bravo-Marquez, 2017a; Buechel and Hahn, 2018a).

Acknowledgements

We thank the anonymous reviewers for their helpful suggestions and comments. Sven Buechel thanks his doctoral advisor Udo Hahn, JULIE Lab, for funding his research visit at the University of Pennsylvania.

References

- Muhammad Abdul-Mageed and Lyle Ungar. 2017. EmoNet: Fine-grained emotion detection with gated recurrent neural networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 718–728.
- Alexandra Balahur, Rada Mihalcea, and Andrés Montoyo. 2014. Computational approaches to subjectivity and sentiment analysis: Present and envisaged methods and applications. *Computer Speech & Language*, 28(1):1–6.
- Daniel Beck. 2017. Modelling representation noise in emotion analysis using gaussian processes. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 140–145.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomáš Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5(1):135–146.
- Laura-Ana-Maria Bostan and Roman Klinger. 2018. An analysis of annotated corpora for emotion classification in text. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2104–2119.
- Margaret M. Bradley and Peter J. Lang. 1994. Measuring emotion: The Self-Assessment Manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry*, 25(1):49–59.
- Sven Buechel and Udo Hahn. 2016. Emotion analysis as a regression problem: Dimensional models and their implications on emotion representation and metrical evaluation. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1114–1122.
- Sven Buechel and Udo Hahn. 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.
- Sven Buechel and Udo Hahn. 2018b. Representation mapping: A novel approach to generate high-quality multi-lingual emotion lexicons. In *Proceedings of the 11th International Conference on Language Resources and Evaluation*, pages 184–191.
- Sven Buechel and Udo Hahn. 2018c. 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.
- François-Régis Chaumartin. 2007. UPAR7: A knowledge-based system for headline sentiment tagging. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, pages 422–425.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, 8th Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111.
- Luna De Bruyne, Pepa Atanasova, and Isabelle Augenstein. 2019. Joint emotion label space modelling for affect lexica. *arXiv:1911.08782 [cs.CL]*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Thomas G Dietterich. 1998. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7):1895–1923.
- Steven Du and Xi Zhang. 2016. Aicyber’s system for IALP 2016 Shared Task: Character-enhanced word vectors and boosted neural networks. In *Proceedings of the 2016 International Conference on Asian Language Processing*, pages 161–163.
- Paul Ekman. 1992. An argument for basic emotions. *Cognition & Emotion*, 6(3-4):169–200.

- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1615–1625.
- Frédéric Godin, Baptist Vandersmissen, Wesley De Neve, and Rik Van de Walle. 2015. Multimedia Lab @ ACL WNUT NER Shared Task: Named entity recognition for Twitter microposts using distributed word representations. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 146–153.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the 11th International Conference on Language Resources and Evaluation*, pages 3483–3487.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Kamil K Imbir. 2017. The affective norms for polish short texts (ANPST) database properties and impact of participants’ population and sex on affective ratings. *Frontiers in Psychology*, 8:855.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 655–665.
- Hamed Khanpour, Cornelia Caragea, and Prakhar Biyani. 2017. Identifying empathetic messages in online health communities. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 246–251.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations*.
- Minglei Li, Qin Lu, Yunfei Long, and Lin Gui. 2017. Inferring affective meanings of words from word embedding. *IEEE Transactions on Affective Computing*, 8(4):443–456.
- Fei Liu, Julien Perez, and Scott Nowson. 2017. A language-independent and compositional model for personality trait recognition from short texts. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 754–764.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Saif Mohammad and Felipe Bravo-Marquez. 2017a. Emotion intensities in tweets. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics*, pages 65–77.
- Saif Mohammad and Felipe Bravo-Marquez. 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.
- Saif M Mohammad and Svetlana Kiritchenko. 2015. Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 31(2):301–326.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. SemEval-2018 Task 1: Affect in tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 1–17.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, and others. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.

- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237.
- Ana P. Pinheiro, Marcelo Dias, João Pedrosa, and Ana P. Soares. 2017. Minho Affective Sentences (MAS): Probing the roles of sex, mood, and empathy in affective ratings of verbal stimuli. *Behavior Research Methods*, 49(2):698–716.
- Nils Reimers and Iryna Gurevych. 2018. Why comparing single performance scores does not allow to draw conclusions about machine learning approaches. *arXiv:1803.09578 [cs.LG]*.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 Task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation*, pages 502–518.
- Victoria L. Rubin. 2007. Stating with certainty or stating with doubt: Intercoder reliability results for manual annotation of epistemically modalized statements. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 141–144.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Carlo Strapparava and Rada Mihalcea. 2007. SemEval-2007 Task 14: Affective text. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, pages 70–74.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5753–5763. Curran Associates, Inc.