

# GWU NLP Lab at SemEval-2019 Task 3 : *EmoContext*: Effective Contextual Information in Models for Emotion Detection in Sentence-level in a Multigenre Corpus

**Shabnam Tafreshi**

George Washington University  
Department of Computer Science  
Washington, DC  
shabnamt@gwu.edu

**Mona Diab**

AWS AI  
George Washington University  
Department of Computer Science  
Washington, DC  
mtdiab@gwu.edu

## Abstract

In this paper we present an emotion classifier models that submitted to the SemEval-2019 Task 3 : *EmoContext*. The task objective is to classify emotion (*i.e.* happy, sad, angry) in a 3-turn conversational data set. We formulate the task as a classification problem and introduce a Gated Recurrent Neural Network (GRU) model with attention layer, which is bootstrapped with contextual information and trained with a multigenre corpus. We utilize different word embeddings to empirically select the most suited one to represent our features. We train the model with a multigenre emotion corpus to leverage using all available training sets to bootstrap the results. We achieved overall %56.05 f1-score and placed 144.

## 1 Introduction

In recent studies, deep learning models have achieved top performances in emotion detection and classification. Access to large amount of data has contributed to these high results. Numerous efforts have been dedicated to build emotion classification models, and successful results have been reported. In this work, we combine several popular emotional data sets in different genres, plus the one given for this task to train the emotion model we developed. We introduce a multigenre training mechanism, our intuition to combine different genres are a) to augment more training data, b) to generalize detection of emotion. We utilize Portable textual information such as subjectivity, sentiment, and presence of emotion words, because emotional sentences are subjective and affectual states like sentiment are strong indicator for presence of emotion.

The rest of this paper is structured as followings: section 2 introduce our neural net model, in section 3 we explain the experimental setup and data

that is been used for training and development sets, section 4 discuss the results and analyze the errors, section 5 describe related works, section 6 conclude our study and discuss future direction.

## 2 Model Description

Gates Recurrent Neural Network (GRU) (Cho et al., 2014; Chung et al., 2015) and attention layer are used in sequential NLP problems and successful results are reported in different studies. Figure 1 shows the diagram of our model. <sup>1</sup>

GRU- has been widely used in the literature to model sequential problems. RNN applies the same set of weights recursively as follow:

$$h_t = f(W_{x_t} + U h_{t-1} + b) \quad (1)$$

GRU is very similar to LSTM with the following equations:

$$r_t = \sigma(W_{x_t}^r + U^r h_{t-1} + b^r) \quad (2)$$

$$z_t = \sigma(W_{x_t}^z + U^z h_{t-1} + b^z) \quad (3)$$

$$\hat{h}_t = \tanh(W_{x_t} + r_t \times U^{\hat{h}} h_{t-1} + b^{\hat{h}}) \quad (4)$$

$$h_t = z_t \times h_{t-1} + (1 - z_t) \times \hat{h}_t \quad (5)$$

GRU has two gates, a reset gate  $r_t$ , and an update gate  $z_t$ . Intuitively, the reset gate determines how to combine the new input with the previous memory, and the update gate defines how much of the previous memory to keep around. We use Keras<sup>2</sup> GRNN implementation to setup our experiments. We note that GRU units are a concatenation of GRU layers in each task.

**Attention layer** - GRUs update their hidden state  $h(t)$  as they process a sequence and the final

<sup>1</sup>Data and system will be released upon the request.

<sup>2</sup><https://keras.io/>

hidden state holds the summation of all other history information. Attention layer (Bahdanau et al., 2014) modifies this process such that representation of each hidden state is an output in each GRU unit to analyze whether this is an important feature for prediction.

**Model Architecture** - our model has an embedding layer of 300 dimensions using *fasttext* embedding, and 1024 dimensions using ELMo (Peters et al., 2018) embedding. GRU layer has 70 hidden unites. We have 3 perceptron layers with size 300. Last layer is a *softmax* layer to predict emotion tags. Textual information layers (explained in section 2.1) are concatenated with GRU layer as auxiliary layer. We utilize a dropout (Graves et al., 2013) layer after the first perceptron layer for regularization.

## 2.1 Textual Information

**Sentiment and objective Information (SOI)**-relativity of subjectivity and sentiment with emotion are well studied in the literature. To craft these features we use SentiwordNet (Baccianella et al., 2010), we create sentiment and subjective score per word in each sentences. SentiwordNet is the result of the automatic annotation of all the synsets of WORDNET according to the notions of *positivity*, *negativity*, and *neutrality*. Each synset  $s$  in WORDNET is associated to three numerical scores Pos(s), Neg(s), and Obj(s) which indicate how positive, negative, and objective (i.e., neutral) the terms contained in the synset are. Different senses of the same term may thus have different opinion-related properties. These scores are presented per sentence and their lengths are equal to the length of each sentence. In case that the score is not available, we used a fixed score 0.001.

**Emotion Lexicon feature (emo)**- presence of emotion words is the first flag for a sentence to be emotional. We use NRC Emotion Lexicon (Mohammad and Turney, 2013) with 8 emotion tags (*e.i.* joy, trust, anticipation, surprise, anger, fear, sadness, disgust). We demonstrate the presence of emotion words as an 8 dimension feature, presenting all 8 emotion categories of the NRC lexicon. Each feature represent one emotion category, where 0.001<sup>3</sup> indicates of absent of

<sup>3</sup>empirically we observed that 0 is not a good initial value in neural net.

the emotion and 1 indicates the presence of the emotion. The advantage of this feature is their portability in transferring emotion learning across genres.

## 2.2 Word Embedding

Using different word embedding or end to end models where word representation learned from local context create different results in emotion detection. We noted that pre-trained word embeddings need to be tuned with local context during our experiments or it causes the model to not converge. We experimented with different word embedding methods such as *word2vec*, *GloVe* (Pennington et al., 2014), *fasttext* (Mikolov et al., 2018), and *ELMo*. Among these methods *fasttext* and *ELMo* create better results.

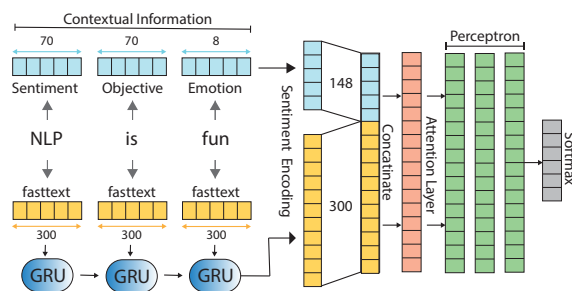


Figure 1: GRU-Attention neural net architecture. In this model framework, context information are features generated from SentiWordNet and emotion lexicon. We use *fasttext* to show the embedding layer (we use ELMo too, but we do not show it in here). Features are presented to GRU and attention layer and the output of attention layer is sent to 3 perceptron layer. Last layer is a softmax layer to predict emotion labels. Model without contextual info, exclude the contextual info input, which we do not show in the architecture.

## 3 Experimental Setup

We split *MULTI* dataset into 80%,10%,10% for train, dev, and test, respectively. We use AIT and EmoContext (data for this task) split as it is given by SemEval 2018 and semEval 2019. We describe these data sets in details in the next section. All experiments are implemented using Keras<sup>4</sup> and Tensorflow<sup>5</sup> in the back-end.

### 3.1 Data

We used three different emotion corpora in our experiments. Our corpora are as follows: a)

<sup>4</sup><https://keras.io/>

<sup>5</sup><https://www.tensorflow.org/>

A multigenre corpus created by (Tafreshi and Diab, 2018) with following genres: *emotional blog posts*, collected by (Aman and Szpakowicz, 2007), *headlines* data set from SemEval 2007-task 14 (Strapparava and Mihalcea, 2007), *movie review* data set (Pang and Lee, 2005) originally collected from Rotten tomatoes <sup>6</sup> for sentiment analysis and it is among the benchmark sets for this task. We refer to this multigenre set as (MULTI), b) SemEval-2018 Affect in Tweets data set (Mohammad et al., 2018) (AIT) with most popular emotion tags: *anger*, *fear*, *joy*, and *sadness*, c) the data set that is given for this task, which is 3-turn conversation data. From these data sets we only used the emotion tags *happy*, *sad*, and *angry*. We used tag *no-emotion* from MULTI data set as *others* tag. Data statistics are shown in figures 2, 3, 4 .

**Data pre-processing** - we tokenize all the data. For tweets we replace all the URLs, image URLs, hashtags, @users with specific anchors. Based on the popularity of each emoticon per each emotion tag, we replace them with the corresponding emotion tag. We normalized all the repeated characters, finally caps words are replaced with lower case but marked as caps words.

### 3.2 Training the Models

We have input size of 70 for sentence length, sentiment, and objective features and emotion lexicon feature has size 8. All these features are explained in section 2.1 and are concatenated with GRU layer as auxiliary (input) layer. Attention comes next after GRU and have size 70. We select dropout of size 0.2. We select 30 epochs in each experiment, however, training is stopped earlier if 2 consecutive larger loss values are seen on evaluation of dev set. We use Adam (Kingma and Ba, 2014) optimizer with a learning rate 0.001. We use dropout with rates 0.2. The loss function is a categorical-cross-entropy function. We use a mini batch (Cotter et al., 2011) of size 32. All hyper-parameter values are selected empirically. We run each experiment 5 times with random initialization and report the mean score over these 5 runs. In section 4 we describe how we choose the hyper-parameters values.

Data set	#train	#dev	#test	total
MULTI	13776	1722	1722	17220
AIT	6839	887	4072	11798
EmoContext	30160	2755	5510	38425

Table 1: Data statistics illustrating the distributions of the train, dev, and test sets across different data sets.

**baseline-** in each sentence we tagged every emotional word using NRC emotion lexicon (Mohammad and Turney, 2013), if any emotion has majority occurrence we pick that emotion tag as sentence emotion tag, when all emotion tags happen only once we randomly choose among them, when there is no emotional word we tag the sentence as *others*. We only use the portion of the emotion lexicon that covers the tags in the task (*i.e.* happy, sad, and angry).

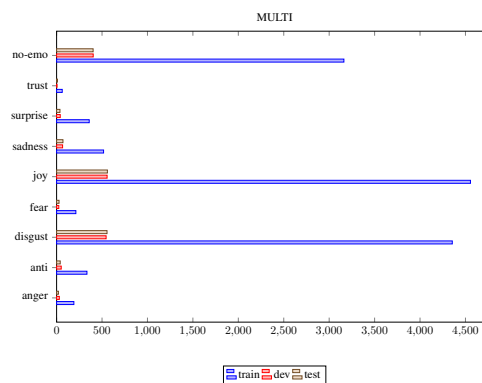


Figure 2: MULTI data set - train, dev, test data statistic

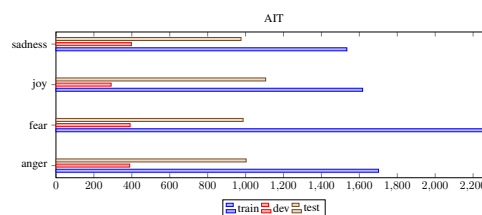


Figure 3: AIT data set - train, dev, test data statistic

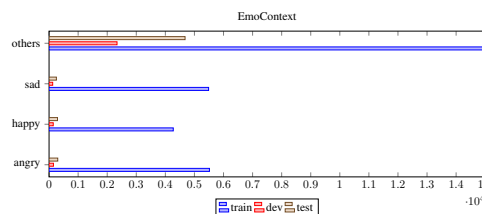


Figure 4: EmoContext data set - train, dev, test data statistic

<sup>6</sup><https://www.rottentomatoes.com/>

Methods/Data set	EmoContext				
	pr.	re.	f.	acc.	sp./#epo.
Baseline	-	-	46.20	46.20	n.a.
GRU-att-fasttext	88.12	81.24	83.44	80.84	103/14
GRU-att-fasttext+F	88.27	84.47	85.27	82.07	321/8
GRU-att-ELMo	88.50	82.65	83.05	82.65	310/20
GRU-att-ELMo+F	<b>88.61</b>	<b>84.34</b>	<b>85.54</b>	<b>83.62</b>	960/28
Context Results(emotion only)	54.28	57.93	56.04	-	960/28

Table 2: Results on the *EmoContext* test sets. We report the mean score over 5 runs. Standard deviations in score are around 0.8. The experiments are demonstrating different embedding (*i.e.* ELMo and fasttext), with features (F), which are *emo* and *SOI* explained in section 2.1

Emotion tags/Data set	EmoContext		
	pr.	re.	f.
happy	45.37	53.52	49.11
sad	57.92	55.60	56.73
angry	61.02	64.09	62.52

Table 3: Context results of each emotion tag.

## 4 Results and Analysis

The results indicates the impact of contextual information using different embeddings, which are different in feature representation. Results of class *happy* without contextual features has %44.16 by GRU-att-ELMo model, and %49.38 by GRU-att-ELMo+F.

We achieved the best results combining ELMo with contextual information, and achieve %85.54 f-score overall, including class *others*. In this task we achieved %56.04 f-score overall for emotion classes, which indicates our model needs to improve the identification of emotion. Table 3 shows our model performance on each emotion tag. The results show a low performance of the model for emotion tag *happy*, which is due to our data being out of domain. Most of the confusion and errors are happened among the emotion categories, which suggest further investigation and improvement. We achieved %90.48, %60.10, %60.19, %49.38 f-score for class *others*, *angry*, *sad*, and *happy* respectfully.

Processing ELMo and attention is computationally very expensive, among our models GRU-att-ELMo+F has the longest training time and GRU-att-fasttext has the fastest training time. Results are shown in table 2 and table refemoresultss

## 5 Related Works

In semEval 2018 task-1, *Affect in Tweets* (Mohammad et al., 2018), 6 team reported results on sub-task E-c (emotion classification), mainly using neural net architectures, features and

resources, and emotion lexicons. Among these works (Baziotis et al., 2018) proposed a Bi-LSTM architecture equipped with a multi-layer self attention mechanism, (Meisheri and Dey, 2018) their model learned the representation of each tweet using mixture of different embedding. in *WASSA 2017 Shared Task on Emotion Intensity* (Mohammad and Bravo-Marquez, 2017), among the proposed approaches, we can recognize teams who used different word embeddings: GloVe or word2vec (He et al., 2017; Duppada and Hiray, 2017) and exploit a neural net architecture such as LSTM (Goel et al., 2017; Akhtar et al., 2017), LSTM-CNN combinations (Köper et al., 2017; Zhang et al., 2017) and bi-directional versions (He et al., 2017) to predict emotion intensity. Similar approach is developed by (Gupta et al., 2017) using sentiment and LSTM architecture. Proper word embedding for emotion task is key, choosing the most efficient distance between vectors is crucial, the following studies explore solution sparsity related properties possibly including uniqueness (Shen and Mousavi, 2018; Mousavi and Shen, 2017) .

## 6 Conclusion and Future Direction

We combined several data sets with different annotation scheme and different genres and train an emotional deep model to classify emotion. Our results indicate that semantic and syntactic contextual features are beneficial to complex and state-of-the-art deep models for emotion detection and classification. We show that our model is able to classify non-emotion (*others*) with high accuracy. In future we want to improve our model to be able to distinguish between emotion classes in a more sufficient way. It is possible that hierarchical bi-directional GRU model can be beneficial, since these models compute history and future sequence while training the model.

## References

- Md S Akhtar, Palaash Sawant, Asif Ekbal, JD Pawar, and Pushpak Bhattacharyya. 2017. Iitp at emoint-2017: Measuring intensity of emotions using sentence embeddings and optimized features. In *Association for Computational Linguistics*.
- Saima Aman and Stan Szpakowicz. 2007. Identifying expressions of emotion in text. In *International*

- Conference on Text, Speech and Dialogue*, pages 196–205. Springer.
- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In *Lrec*, pages 2200–2204.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Christos Baziotis, Nikos Athanasiou, Alexandra Chronopoulou, Athanasia Kolovou, Georgios Paraskevopoulos, Nikolaos Ellinas, Shrikanth Narayanan, and Alexandros Potamianos. 2018. Ntua-slp at semeval-2018 task 1: Predicting affective content in tweets with deep attentive rnns and transfer learning. *arXiv preprint arXiv:1804.06658*.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2015. Gated feedback recurrent neural networks. In *International Conference on Machine Learning*, pages 2067–2075.
- Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. 2011. Better mini-batch algorithms via accelerated gradient methods. In *Advances in neural information processing systems*, pages 1647–1655.
- Venkatesh Duppada and Sushant Hiray. 2017. Seernet at emoint-2017: Tweet emotion intensity estimator. *arXiv preprint arXiv:1708.06185*.
- Pranav Goel, Devang Kulshreshtha, Prayas Jain, and Kaushal Kumar Shukla. 2017. Prayas at emoint 2017: An ensemble of deep neural architectures for emotion intensity prediction in tweets. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 58–65.
- Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *Acoustics, speech and signal processing (icassp), 2013 IEEE international conference on*, pages 6645–6649. IEEE.
- Umang Gupta, Ankush Chatterjee, Radhakrishnan Srikanth, and Puneet Agrawal. 2017. A sentiment-and-semantics-based approach for emotion detection in textual conversations. *arXiv preprint arXiv:1707.06996*.
- Yuanye He, Liang-Chih Yu, K Robert Lai, and Weiyi Liu. 2017. Yzu-nlp at emoint-2017: Determining emotion intensity using a bi-directional lstm-cnn model. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 238–242.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Maximilian Köper, Evgeny Kim, and Roman Klinger. 2017. Ims at emoint-2017: emotion intensity prediction with affective norms, automatically extended resources and deep learning. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 50–57.
- Hardik Meisheri and Lipika Dey. 2018. Tcs research at semeval-2018 task 1: Learning robust representations using multi-attention architecture. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 291–299.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhresch, and Armand Joulin. 2018. Advances in pre-training distributed word representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- 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.
- Saif M Mohammad and Felipe Bravo-Marquez. 2017. Wassa-2017 shared task on emotion intensity. *arXiv preprint arXiv:1708.03700*.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, pages 436–465.
- Seyedahmad Mousavi and Jinglai Shen. 2017. Solution uniqueness of convex piecewise affine functions based optimization with applications to constrained ell<sub>1</sub> minimization. *arXiv preprint arXiv:1711.05882*.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 115–124. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.

- Jinglai Shen and Seyedahmad Mousavi. 2018. Least sparsity of p-norm based optimization problems with  $p_i \geq 1$ . *SIAM Journal on Optimization*, 28(3):2721–2751.
- 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. Association for Computational Linguistics.
- Shabnam Tafreshi and Mona T Diab. 2018. Sentence and clause level emotion annotation, detection, and classification in a multi-genre corpus. In *LREC*.
- You Zhang, Hang Yuan, Jin Wang, and Xuejie Zhang. 2017. Ynu-hpcc at emoint-2017: Using a cnn-lstm model for sentiment intensity prediction. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 200–204.