

Distant supervision for emotion detection using Facebook reactions

Chris Pool

Anchormen, Groningen
The Netherlands
c.pool@anchormen.nl

Malvina Nissim

CLCG, University of Groningen
The Netherlands
m.nissim@rug.nl

Abstract

We exploit the Facebook reaction feature in a distant supervised fashion to train a support vector machine classifier for emotion detection, using several feature combinations and combining different Facebook pages. We test our models on existing benchmarks for emotion detection and show that *employing only information that is derived completely automatically*, thus without relying on any handcrafted lexicon as it’s usually done, we can achieve competitive results. The results also show that there is large room for improvement, especially by gearing the collection of Facebook pages, with a view to the target domain.

1 Introduction

In the spirit of the brevity of social media’s messages and reactions, people have got used to express feelings minimally and symbolically, as with hashtags on Twitter and Instagram. On Facebook, people tend to be more wordy, but posts normally receive more simple “likes” than longer comments. Since February 2016, Facebook users can express specific emotions in response to a post thanks to the newly introduced *reaction feature* (see Section 2), so that now a post can be wordlessly marked with an expression of say “joy” or “surprise” rather than a generic “like”.

It has been observed that this new feature helps Facebook to know much more about their users and exploit this information for targeted advertising (Stinson, 2016), but interest in people’s opinions and how they feel isn’t limited to commercial reasons, as it invests social monitoring, too, including health care and education (Mohammad, 2016). However, emotions and opinions are not always expressed this explicitly, so that there is high interest in developing systems towards their automatic detection. Creating manually annotated datasets large enough to train supervised models is not only costly, but also—especially in the case of opinions and emotions—difficult, due to the intrinsic subjectivity of the task (Strapparava and Mihalcea, 2008; Kim et al., 2010). Therefore, research has focused on unsupervised methods enriched with information derived from lexica, which are manually created (Kim et al., 2010; Chaffar and Inkpen, 2011). Since Go et al. (2009) have shown that happy and sad emoticons can be successfully used as signals for sentiment labels, *distant supervision*, i.e. using some reasonably safe signals as proxies for automatically labelling training data (Mintz et al., 2009), has been used also for emotion recognition, for example exploiting both emoticons and Twitter hashtags (Purver and Battersby, 2012), but mainly towards creating emotion lexica. Mohammad and Kiritchenko (2015) use hashtags, experimenting also with highly fine-grained emotion sets (up to almost 600 emotion labels), to create the large *Hashtag Emotion Lexicon*. Emoticons are used as proxies also by Hallsmar and Palm (2016), who use distributed vector representations to find which words are interchangeable with emoticons but also which emoticons are used in a similar context.

We take advantage of distant supervision by using Facebook reactions as proxies for emotion labels, which to the best of our knowledge hasn’t been done yet, and we train a set of Support Vector Machine models for emotion recognition. Our models, differently from existing ones, exploit information which

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

is *acquired entirely automatically*, and achieve competitive or even state-of-the-art results for some of the emotion labels on existing, standard evaluation datasets. For explanatory purposes, related work is discussed further and more in detail when we describe the benchmarks for evaluation (Section 3) and when we compare our models to existing ones (Section 5). We also explore and discuss how choosing different sets of Facebook pages as training data provides an intrinsic domain-adaptation method.

2 Facebook reactions as labels

For years, on Facebook people could leave comments to posts, and also “like” them, by using a thumbs-up feature to explicitly express a generic, rather underspecified, approval. A “like” could thus mean “I like what you said”, but also “I like that you bring up such topic (though I find the content of the article you linked annoying)”.

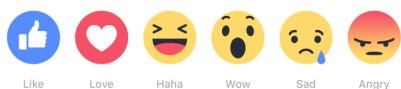


Figure 1: Facebook reactions

In February 2016, after a short trial, Facebook made a more explicit *reaction* feature available world-wide. Rather than allowing for the underspecified “like” as the only wordless response to a post, a set of six more specific reactions was introduced, as shown in Figure 1: Like, Love, Haha, Wow, Sad and Angry. We use such reactions as proxies for emotion labels associated to posts.

We collected Facebook posts and their corresponding reactions from public pages using the Facebook API, which we accessed via the Facebook-sdk python library¹. We chose different pages (and therefore domains and stances), aiming at a balanced and varied dataset, but we did so mainly based on intuition (see Section 4) and with an eye to the nature of the datasets available for evaluation (see Section 5). The choice of which pages to select posts from is far from trivial, and we believe this is actually an interesting aspect of our approach, as by using different Facebook pages one can intrinsically tackle the domain-adaptation problem (See Section 6 for further discussion on this). The final collection of Facebook pages for the experiments described in this paper is as follows: FoxNews, CNN, ESPN, New York Times, Time magazine, Huffington Post Weird News, The Guardian, Cartoon Network, Cooking Light, Home Cooking Adventure, Justin Bieber, Nickelodeon, Spongebob, Disney.

For each page, we downloaded the latest 1000 posts, or the maximum available if there are fewer, from February 2016, retrieving the counts of reactions for each post. The output is a JSON file containing a list of dictionaries with a timestamp, the post and a reaction vector with frequency values, which indicate how many users used that reaction in response to the post (Figure 2). The resulting emotion vectors must then be turned into an emotion label.³

In the context of this experiment, we made the simple decision of associating to each post the emotion with the highest count, ignoring like as it is the default and most generic reaction people tend to use. Therefore, for example, to the first post in Figure 2, we would associate the label sad, as it has the highest score (284) among the meaningful emotions we consider, though it also has non-zero scores for other emotions. At this stage, we didn’t perform any other entropy-based selection of posts, to be investigated in future work.

```
[
  {
    "created_time": "2016-06-19T01:40:00+0000",
    "message": "Walt Disney World representatives said they plan to put up fencing and signs at all resorts and waterways.",
    "reactions": [5073, 4483, 60, 22, 54, 284, 170, 0]
  },
  [
    {
      "created_time": "2016-06-19T01:00:00+0000",
      "message": "Charlene and Joseph Handrik face more than 550 counts of animal cruelty.",
      "reactions": [2256, 1011, 16, 6, 123, 409, 691, 0]
    }
  ],
]
```

Figure 2: Sample of resulting JSON file. The order of values/reactions is total, like, love, haha, wow, sad, angry, thankful.³

¹<https://pypi.python.org/pypi/facebook-sdk>

³Note that thankful was only available during specific time spans related to certain events, as Mother’s Day in May 2016.

3 Emotion datasets

Three datasets annotated with emotions are commonly used for the development and evaluation of emotion detection systems, namely the *Affective Text* dataset, the *Fairy Tales* dataset, and the *ISEAR* dataset. In order to compare our performance to state-of-the-art results, we have used them as well. In this Section, in addition to a description of each dataset, we provide an overview of the emotions used, their distribution, and how we mapped them to those we obtained from Facebook posts in Section 3.4. A summary is provided in Table 1, which also shows, in the bottom row, what role each dataset has in our experiments: apart from the development portion of the Affective Text, which we used to develop our models (Section 4), all three have been used as benchmarks for our evaluation.

3.1 Affective Text dataset

Task 14 at SemEval 2007 (Strapparava and Mihalcea, 2007) was concerned with the classification of emotions and valence in news headlines. The headlines were collected from several news websites including Google news, The New York Times, BBC News and CNN. The used emotion labels were *Anger*, *Disgust*, *Fear*, *Joy*, *Sadness*, *Surprise*, in line with the six basic emotions of Ekman’s standard model (Ekman, 1992). Valence was to be determined as positive or negative. Classification of emotion and valence were treated as separate tasks. Emotion labels were not considered as mutually exclusive, and each emotion was assigned a score from 0 to 100. Training/developing data amounted to 250 annotated headlines (*Affective development*), while systems were evaluated on another 1000 (*Affective test*). Evaluation was done using two different methods: a fine-grained evaluation using Pearson’s r to measure the correlation between the system scores and the gold standard; and a coarse-grained method where each emotion score was converted to a binary label, and precision, recall, and f-score were computed to assess performance. As it is done in most works that use this dataset (Kim et al., 2010; Chaffar and Inkpen, 2011; Calvo and Mac Kim, 2013), we also treat this as a classification problem (coarse-grained). This dataset has been extensively used for the evaluation of various unsupervised methods (Strapparava and Mihalcea, 2008), but also for testing different supervised learning techniques and feature portability (Mohammad, 2012).

3.2 Fairy Tales dataset

This is a dataset collected by Alm (2008), where about 1,000 sentences from fairy tales (by B. Potter, H.C. Andersen and Grimm) were annotated with the same six emotions of the Affective Text dataset, though with different names: *Angry*, *Disgusted*, *Fearful*, *Happy*, *Sad*, and *Surprised*. In most works that use this dataset (Kim et al., 2010; Chaffar and Inkpen, 2011; Calvo and Mac Kim, 2013), only sentences where all annotators agreed are used, and the labels *angry* and *disgusted* are merged. We adopt the same choices.

3.3 ISEAR

The ISEAR (International Survey on Emotion Antecedents and Reactions (Scherer and Wallbott, 1994; Scherer, 1997)) is a dataset created in the context of a psychology project of the 1990s, by collecting questionnaires answered by people with different cultural backgrounds. The main aim of this project was to gather insights in cross-cultural aspects of emotional reactions. Student respondents, both psychologists and non-psychologists, were asked to report situations in which they had experienced all of seven major emotions (*joy*, *fear*, *anger*, *sadness*, *disgust*, *shame* and *guilt*). In each case, the questions covered the way they had appraised a given situation and how they reacted. The final dataset contains reports by approximately 3000 respondents from all over the world, for a total of 7665 sentences labelled with an emotion, making this the largest dataset out of the three we use.

3.4 Overview of datasets and emotions

We summarise datasets and emotion distribution from two viewpoints. First, because there are different sets of emotions labels in the datasets and Facebook data, we need to provide a mapping and derive a subset of emotions that we are going to use for the experiments. This is shown in Table 1, where in

the “Mapped” column we report the final emotions we use in this paper: anger, joy, sadness, surprise. All labels in each dataset are mapped to these final emotions, which are therefore the labels we use for training and testing our models.

Second, the distribution of the emotions for each dataset is different, as can be seen in Figure 3.

In Figure 4 we also provide the distribution of the emotions anger, joy, sadness, surprise per Facebook page, in terms of number of posts (recall that we assign to a post the label corresponding to the majority emotion associated to it, see Section 2).

Table 1: Emotion labels in existing datasets, Facebook, and resulting mapping for the experiments in this work. The last row indicates which role each dataset has in our experiments.

Affective Text	Fairy tales	ISEAR	Facebook	Mapped
Anger	Angry-Disgusted	Anger	Angry	anger
Disgust	Angry-Disgusted	Disgust		anger
Fear	Fearful	Fear		
Joy	Happy	Joy	Haha, Love	joy
Sadness	Sad	Sadness	Sad	sadness
Surprise	Suprised		Wow	surprise
		Shame		
		Guilt		
development/test	test	test	train	

We can observe that for example pages about news tend to have more sadness and anger posts, while pages about cooking and tv-shows have a high percentage of joy posts. We will use this information to find the best set of pages for a given target domain (see Section 5).

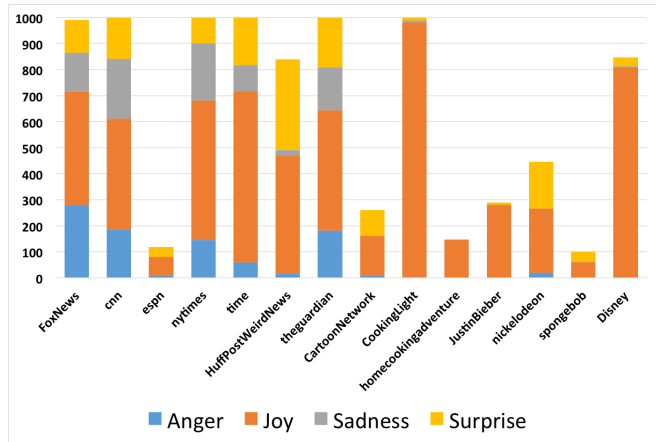
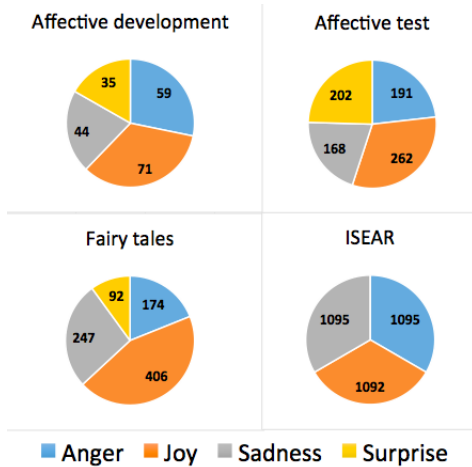


Figure 3: Emotion distribution in the datasets

Figure 4: Emotion distribution per Facebook page

4 Model

There are two main decisions to be taken in developing our model: (i) which Facebook pages to select as training data, and (ii) which features to use to train the model, which we discuss below. Specifically, we first set on a subset of pages and then experiment with features. Further exploration of the interaction between choice of pages and choice of features is left to future work, and partly discussed in Section 6. For development, we use a small portion of the Affective data set described in Section 3.1, that is the portion that had been released as development set for SemEval’s 2007 Task 14 (Strapparava and Mihalcea, 2007), which contains 250 annotated sentences (*Affective development*, Section 3.1). All results reported in this section are on this dataset. The test set of Task 14 as well as the other two datasets described in Section 3 will be used to evaluate the final models (Section 4).

4.1 Selecting Facebook pages

Although page selection is a crucial ingredient of this approach, which we believe calls for further and deeper, dedicated investigation, for the experiments described here we took a rather simple approach. First, we selected the pages that would provide training data based on intuition and availability, then chose different combinations according to results of a basic model run on development data, and eventually tested feature combinations, still on the development set.

For the sake of simplicity and transparency, we first trained an SVM with a simple bag-of-words model and default parameters as per the Scikit-learn implementation (Pedregosa et al., 2011) on different combinations of pages. Based on results of the attempted combinations as well as on the distribution of emotions in the development dataset (Figure 3), we selected a *best model (B-M)*, namely the combined set of *Time*, *The Guardian* and *Disney*, which yields the highest results on development data. *Time* and *The Guardian* perform well on most emotions but *Disney* helps to boost the performance for the *Joy* class.

4.2 Features

In selecting appropriate features, we mainly relied on previous work and intuition. We experimented with different combinations, and all tests were still done on *Affective development*, using the pages for the best model (B-M) described above as training data. Results are in Table 2. Future work will further explore the simultaneous selection of features and page combinations.

Standard textual features We use a set of basic text-based features to capture the emotion class. These include a tf-idf bag-of-words feature, word (2-3) and character (2-5) ngrams, and features related to the presence of negation words, and to the usage of punctuation.

Affect Lexicons This feature is used in all unsupervised models as a source of information, and we mainly include it to assess its contribution, but eventually do not use it in our final model.

We used the NRC10 Lexicon because it performed best in the experiments by (Mohammad, 2012), which is built around the emotions *anger*, *anticipation*, *disgust*, *fear*, *joy*, *sadness*, and *surprise*, and the valence values *positive* and *negative*. For each word in the lexicon, a boolean value indicating presence or absence is associated to each emotion. For a whole sentence, a global score per emotion can be obtained by summing the vectors for all content words of that sentence included in the lexicon, and used as feature.

Word Embeddings As additional feature, we also included Word Embeddings, namely distributed representations of words in a vector space, which have been exceptionally successful in boosting performance in a plethora of NLP tasks. We use three different embeddings:

- *Google embeddings*: pre-trained embeddings trained on Google News and obtained with the skip-gram architecture described in (Mikolov et al., 2013). This model contains 300-dimensional vectors for 3 million words and phrases.
- *Facebook embeddings*: embeddings that we trained on our scraped Facebook pages for a total of 20,000 sentences. Using the *gensim* library (Řehůřek and Sojka, 2010), we trained the embeddings with the following parameters: window size of 5, learning rate of 0.01 and dimensionality of 100. We filtered out words with frequency lower than 2 occurrences.
- *Retrofitted embeddings*: Retrofitting (Faruqui et al., 2015) has been shown as a simple but efficient way of informing trained embeddings with additional information derived from some lexical resource, rather than including it directly at the training stage, as it's done for example to create sense-aware (Iacobacci et al., 2015) or sentiment-aware (Tang et al., 2014) embeddings.⁴ In this work, we retrofit general embeddings to include information about emotions, so that emotion-similar words can get closer in space. Both the Google as well as our Facebook embeddings were retrofitted with

⁴Training emotion-aware embeddings is a strategy that we plan to explore in future work.

lexical information obtained from the NRC10 Lexicon mentioned above, which provides emotion-similarity for each token. Note that differently from the previous two types of embeddings, the retrofitted ones do rely on handcrafted information in the form of a lexical resource.

4.3 Results on development set

We report precision, recall, and f-score on the development set. The average f-score is reported as *micro-average*, to better account for the skewed distribution of the classes as well as in accordance to what is usually reported for this task (Mohammad and Kiritchenko, 2015).

Table 2: Results on the development set (*Affective development*). *avg f* is the micro-averaged f-score.

Feature	anger	joy	sadness	surprise	avg f
	prec,rec,f	prec,rec,f	prec,rec,f	prec,rec,f	
Tf-idf	0.57,0.22,0.32	0.44,0.51,0.47	0.41,0.25, 0.31	0.22,0.49,0.30	0.368
Lexicon	0.28,0.08,0.13	0.43,0.37,0.40	0.31,0.30, 0.30	0.20,0.51,0.29	0.297
Token n-grams(2,5)	0.00,0.00,0.00	1.00,0.01,0.03	0.00,0.00, 0.00	0.17,1.00,0.29	0.172
Character n-grams(2,5)	0.50,0.03,0.06	0.39,0.73,0.51	0.38,0.07, 0.12	0.17,0.31,0.22	0.325
All features	0.40,0.03,0.06	0.35,0.97,0.52	0.62,0.11, 0.19	1.00,0.03,0.06	0.368
Google (G) embeddings	0.41,0.49,0.45	0.56,0.46,0.51	0.48,0.57, 0.52	0.22,0.17,0.19	0.445
Facebook (FB) embeddings	0.33,0.15,0.21	0.31,0.45,0.37	0.23,0.11,0.15	0.20,0.31,0.24	0.273
Retrofitted G-embeddings	0.36,0.20,0.26	0.42,0.48,0.45	0.30,0.25, 0.27	0.20,0.34,0.26	0.330
Retrofitted FB-embeddings	0.07,0.02,0.03	0.34,0.86,0.49	0.36,0.09,0.15	0.17,0.03,0.05	0.321
Tf-idf + G-emb	0.42,0.46,0.44	0.45,0.49,0.47	0.49,0.41, 0.44	0.29,0.26,0.27	0.426
All features + G-emb	0.63,0.29,0.40	0.43,0.83,0.56	0.46,0.27, 0.34	0.33,0.17,0.23	0.450
All features – Lexicon + G-emb	0.62,0.34,0.44	0.43,0.85,0.57	0.57,0.30, 0.39	0.36,0.14,0.20	0.469

From Table 2 we draw three main observations. First, a simple tf-idf bag-of-words mode works already very well, to the point that the other textual and lexicon-based features don’t seem to contribute to the overall f-score (0.368), although there is a rather substantial variation of scores per class. Second, Google embeddings perform a lot better than Facebook embeddings, and this is likely due to the size of the corpus used for training. Retrofitting doesn’t seem to help at all for the Google embeddings, but it does boost the Facebook embeddings, leading to think that with little data, more accurate task-related information is helping, but corpus size matters most. Third, in combination with embeddings, all features work better than just using tf-idf, but removing the Lexicon feature, which is the only one based on hand-crafted resources, yields even better results. Then our best model (**B-M**) on development data relies *entirely on automatically obtained information*, both in terms of training data as well as features.

5 Results

In Table 3 we report the results of our model on the three datasets standardly used for the evaluation of emotion classification, which we have described in Section 3.

Our **B-M** model relies on subsets of Facebook pages for training, which were chosen according to their performance on the development set as well as on the observation of emotions distribution on different pages and in the different datasets, as described in Section 4. The feature set we use is our best on the development set, namely all the features plus Google-based embeddings, but excluding the lexicon. This makes our approach completely independent of any manual annotation or handcrafted resource. Our model’s performance is compared to the following systems, for which results are reported in the referred literature. Please note that no other existing model was re-implemented, and results are those reported in the respective papers.

Kim et al. (2010) experiment with four different unsupervised techniques that rely on lexicon-derived information. In Table 3 we report the scores for their best average performing approach, namely a

CNMF-based categorical classification. They made the decision not to deal with `surprise` because this emotion is not present in the ISEAR dataset.

Strapparava and Mihalcea (2008) experiment with several models based on a core LSA model and, in their best performing model (`LSA-all emotion words`) whose results we report in Table 3, also use information from lexical resources both in their general (WordNet (Fellbaum, 1998)) and emotion-aware (WordNet Affect (Strapparava et al., 2004)) form.

Danisman and Alpkocak (2008) adopt a supervised approach, training a model using the ISEAR dataset and testing it on the Affective text dataset. They only report results per category in terms of f-score, without further specification of how precision and recall contribute.

We have mentioned that the selection of Facebook pages is relevant and can be also thought of as a tool for domain adaptation in accordance with the characteristics of the target domains/datasets (see also Section 2 and Figures 3–4). Although we believe that such an interesting aspect will require deeper investigation (see also Section 6), we preliminary test this assumption by developing and comparing two more models: a model that uses a combination of pages that we expect will perform best on the Fairy Tales dataset (**FT-M**), and a model that uses a combination of pages that should perform best on the ISEAR dataset (**ISE-M**). The feature set is kept the same for all three models.

FT-M The sentences in the Fairy Tales dataset are quite different compared to the news headlines in the development set. Looking at the distribution in this dataset, as can be seen in Figure 3, `Joy` is the most frequent class. We selected the pages `HuffPostWeirdNews`, `ESPN` and `CNN` for this model especially looking at the performance for the emotions that are most frequent in this dataset.

ISE-M As described in Section 3.3, the sentences in the ISEAR collection are also different compared to the two other datasets. Looking at the distribution in Figure 3 and according to performance on relevant emotions (we took into account the absence of `Surprise` in this dataset), we selected the pages `Time`, `The Guardian` and `CookingLight` for this model.

In Table 3 we report results for all of the models mentioned above. We indicate averages only for our models, since not all approaches deal with the same sets of emotions and we cannot easily compute them. We discuss results both in terms of how our models fair with respect to other systems as reported the literature, as well as how they compare to one another with a view to the selection of Facebook pages.

Compared to other systems, our models are globally competitive, given that **B-M** is entirely unsupervised. Overall, the unsupervised but heavily lexicon-based best model of (Kim et al., 2010) performs well on all emotions, excluding surprise, which they do not address (thus also making their classification task slightly easier). Differently from existing systems, our models appear rather balanced in terms of performance on the different emotions as well as in precision and recall, and are able to deal well with the variance of the datasets.

On the Affective Text dataset, we have the highest precision for all emotions but `joy`, though on this emotion our models have very good recall. The highest recall for all emotions for this dataset is reported in (Strapparava and Mihalcea, 2008), together with extremely low precision. Such skewed performance for all emotions can only be explained if different emotion-specific models were trained rather than a single multiclass model, but this is not described as such in the paper. The authors state that their models are completely unsupervised, which is true in terms of training data, but they nevertheless augment them with information derived from hand-crafted resources.

On the Fairy Tales dataset, (Kim et al., 2010) Chaffar and Inkpen (2011) also used the Fairy tales dataset to evaluate a supervised model using features like bag-of-words, N-grams and lexical emotion features, but report cross-validated results using accuracy only, and are therefore harder to compare.

On the ISEAR dataset, which is the largest, our models perform best for all emotions but `anger`, for which however we achieve the highest precision with all our models.

From the perspective of comparing our models, we do not observe any real correlation between our actual best performances and the models designed to best perform on a given dataset. For example, **B-M**

Table 3: Results on test datasets according to Precision, Recall and F-score.

	Affective test			Fairy Tales			ISEAR		
	P	R	F	P	R	F	P	R	F
	anger								
B-M	0.50	0.35	0.41	0.33	0.04	0.07	0.72	0.06	0.11
FT-M	0.51	0.30	0.38	0.27	0.02	0.04	0.57	0.10	0.17
ISE-M	0.48	0.35	0.40	0.36	0.05	0.08	0.74	0.06	0.11
(Strapparava and Mihalcea, 2008)	0.06	0.88	0.12						
(Kim et al., 2010)	0.29	0.26	0.28	0.77	0.56	0.65	0.41	0.99	0.58
(Danisman and Alpkocak, 2008)			0.24						
	joy								
B-M	0.39	0.85	0.54	0.49	0.77	0.60	0.41	0.79	0.53
FT-M	0.41	0.77	0.54	0.49	0.69	0.58	0.42	0.63	0.50
ISE-M	0.39	0.82	0.53	0.48	0.81	0.60	0.40	0.83	0.54
(Strapparava and Mihalcea, 2008)	0.19	0.90	0.31						
(Kim et al., 2010)	0.77	0.58	0.65	0.80	0.76	0.78	0.39	0.01	0.01
(Danisman and Alpkocak, 2008)			0.50						
	sadness								
B-M	0.51	0.21	0.30	0.43	0.39	0.41	0.50	0.39	0.44
FT-M	0.53	0.28	0.37	0.50	0.24	0.33	0.79	0.28	0.41
ISE-M	0.49	0.21	0.29	0.43	0.34	0.38	0.51	0.38	0.44
(Strapparava and Mihalcea, 2008)	0.12	0.87	0.22						
(Kim et al., 2010)	0.50	0.45	0.48	0.71	0.82	0.77	0.37	0.01	0.25
(Danisman and Alpkocak, 2008)			0.37						
	surprise								
B-M	0.20	0.05	0.08	0.12	0.04	0.06			
FT-M	0.25	0.17	0.20	0.14	0.33	0.19			
ISE-M	0.27	0.08	0.12	0.17	0.04	0.07			
(Strapparava and Mihalcea, 2008)	0.08	0.95	0.14						
(Kim et al., 2010)									
(Danisman and Alpkocak, 2008)									
	AVERAGE (micro f-score)								
B-M	0.409			0.459			0.411		
FT-M	0.412			0.408			0.336		
ISE-M	0.405			0.460			0.422		

was expected to perform best on the Affective Text, but it is outperformed by **FT-M** in the precision of detecting anger and sadness, and overall for the detection of surprise. Generally, by looking at averages, it seems that our best performing model across datasets is **ISE-M**. However, the extremely large variance among scores for the same emotion on the three datasets, highlights the differences among such datasets and the need to better tailor training data to different domains. The large discrepancy in detecting different emotions in the same dataset also deserves further investigation. We discuss such issues further in the next section, with a view to future work.

6 Discussion, conclusions and future work

We have explored the potential of using Facebook reactions in a distant supervised setting to perform emotion classification. The evaluation on standard benchmarks shows that models trained as such, especially when enhanced with continuous vector representations, can achieve competitive results without relying on any handcrafted resource. An interesting aspect of our approach is the view to domain adap-

tation via the selection of Facebook pages to be used as training data.

We believe that this approach has a lot of potential, and we see the following directions for improvement. Feature-wise, we want to train emotion-aware embeddings, in the vein of work by Tang et al. (2014), and Iacobacci et al. (2015). Retrofitting FB-embeddings trained on a larger corpus might also be successful, but would rely on an external lexicon.

The largest room for yielding not only better results but also interesting insights on extensions of this approach lies in the choice of training instances, both in terms of Facebook pages to get posts from, as well as in which posts to select from the given pages. For the latter, one could for example only select posts that have a certain length, ignore posts that are only quotes or captions to images, or expand posts by including content from linked html pages, which might provide larger and better contexts (Plank et al., 2014). Additionally, and most importantly, one could use an entropy-based measure to select only posts that have a strong emotion rather than just considering the majority emotion as training label. For the former, namely the choice of Facebook pages, which we believe deserves the most investigation, one could explore several avenues, especially in relation to *stance*-based issues (Mohammad et al., 2016). In our dataset, for example, a post about Chile beating Colombia in a football match during the Copa America had very contradictory reactions, depending on which side readers would cheer for. Similarly, the very same political event, for example, would get very different reactions from readers if it was posted on Fox News or The Late Night Show, as the target audience is likely to feel very differently about the same issue. This also brings up theoretical issues related more generally to the definition of the emotion detection task, as it's strongly dependent on personal traits of the audience. Also, in this work, pages initially selected on availability and intuition were further grouped into sets to make training data according to performance on development data, and label distribution. Another criterion to be exploited would be *vocabulary overlap* between the pages and the datasets.

Lastly, we could develop single models for each emotion, treating the problem as a multi-label task. This would even better reflect the ambiguity and subjectivity intrinsic to assigning emotions to text, where content could be at same time joyful or sad, depending on the reader.

Acknowledgements

In addition to the anonymous reviewers, we want to thank Lucia Passaro and Barbara Plank for insightful discussions, and for providing comments on draft versions of this paper.

References

- Ebba Cecilia Ovesdotter Alm. 2008. *Affect in text and speech*. ProQuest.
- Rafael A Calvo and Sunghwan Mac Kim. 2013. Emotions in text: dimensional and categorical models. *Computational Intelligence*, 29(3):527–543.
- Soumaya Chaffar and Diana Inkpen. 2011. Using a heterogeneous dataset for emotion analysis in text. In *Advances in Artificial Intelligence*, pages 62–67. Springer.
- Taner Danisman and Adil Alpkocak. 2008. Feeler: Emotion classification of text using vector space model. In *AISB 2008 Convention Communication, Interaction and Social Intelligence*, volume 1, page 53.
- Paul Ekman. 1992. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200.
- Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard Hovy, and Noah A. Smith. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1606–1615, Denver, Colorado, May–June. Association for Computational Linguistics.
- Christiane Fellbaum. 1998. *WordNet*. Wiley Online Library.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *CS224N Project Report, Stanford*, 1:12.

- Fredrik Hallsmar and Jonas Palm. 2016. Multi-class sentiment classification on twitter using an emoji training heuristic. Technical report, KTH/Skolan för datavetenskap och kommunikation (CSC). University essay.
- Ignacio Iacobacci, Mohammad Taher Pilehvar, and Roberto Navigli. 2015. Sensembled: learning sense embeddings for word and relational similarity. In *Proceedings of ACL*, pages 95–105.
- Sunghwan Mac Kim, Alessandro Valitutti, and Rafael A Calvo. 2010. Evaluation of unsupervised emotion models to textual affect recognition. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 62–70. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2 - Volume 2*, ACL '09, pages 1003–1011, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Saif M Mohammad and Svetlana Kiritchenko. 2015. Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 31(2):301–326.
- Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. A dataset for detecting stance in tweets. In *Proceedings of 10th edition of the the Language Resources and Evaluation Conference (LREC)*, Portorož, Slovenia.
- Saif Mohammad. 2012. Portable features for classifying emotional text. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 587–591, Montréal, Canada, June. Association for Computational Linguistics.
- Saif M. Mohammad. 2016. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In Herb Meiselman, editor, *Emotion Measurement*. Elsevier.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Barbara Plank, Dirk Hovy, Ryan T McDonald, and Anders Søgaard. 2014. Adapting taggers to twitter with not-so-distant supervision. In *COLING*, pages 1783–1792.
- Matthew Purver and Stuart Battersby. 2012. Experimenting with distant supervision for emotion classification. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 482–491. Association for Computational Linguistics.
- Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May. ELRA. <http://is.muni.cz/publication/884893/en>.
- Klaus R Scherer and Harald G Wallbott. 1994. Evidence for universality and cultural variation of differential emotion response patterning. *Journal of personality and social psychology*, 66(2):310.
- Klaus R Scherer. 1997. The role of culture in emotion-antecedent appraisal. *Journal of personality and social psychology*, 73(5):902.
- Liz Stinson. 2016. Facebook reactions, the totally redesigned like button, is here. *Wired*. <http://www.wired.com/2016/02/facebook-reactions-totally-redesigned-like-button/>.
- 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.
- Carlo Strapparava and Rada Mihalcea. 2008. Learning to identify emotions in text. In *Proceedings of the 2008 ACM symposium on Applied computing*, pages 1556–1560. ACM.
- Carlo Strapparava, Alessandro Valitutti, et al. 2004. Wordnet affect: an affective extension of wordnet. In *LREC*, volume 4, pages 1083–1086.
- Duyu Tang, Furu Wei, Nan Yang, Ming Zhou, Ting Liu, and Bing Qin. 2014. Learning sentiment-specific word embedding for Twitter sentiment classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1555–1565.