

EmoTag – Towards an Emotion-Based Analysis of Emojis

Abu Awal Md Shoeb Shahab Raji Gerard de Melo

Department of Computer Science

Rutgers University, New Brunswick, NJ, USA

{abu.shoeb, shahab.raji, gerard.demelo}@rutgers.edu

Abstract

Despite being a fairly recent phenomenon, emojis have quickly become ubiquitous. Besides their extensive use in social media, they are now also invoked in customer surveys and feedback forms. Hence, there is a need for techniques to understand their sentiment and emotion. In this work, we provide a method to quantify the emotional association of basic emotions such as anger, fear, joy, and sadness for a set of emojis. We collect and process a unique corpus of 20 million emoji-centric tweets, such that we can capture rich emoji semantics using a comparably small dataset. We evaluate the induced emotion profiles of emojis with regard to their ability to predict word affect intensities as well as sentiment scores.

1 Introduction

In recent years, information technology has profoundly altered the way humans communicate. A substantial proportion of the global population has adopted the use of social media platforms (such as Twitter, Facebook, and Instagram) and messaging technology (such as Facebook Messenger, WeChat, and WhatsApp) to interact and voice their opinion. The unique properties and expressive capabilities afforded by computer and mobile device-mediated communication has led to quite distinct forms of expression in comparison with classic email etiquette, let alone traditional written correspondence.

Meanwhile, for any meaningful analysis of social interactions or expression of opinions, it is critical to extract and understand the sentiment and the affect of the source. There are numerous studies investigating the connection between

words or sentences and the affects they convey. However, emojis are a particularly prominent feature of modern online interaction. Thus, this paper introduces a new basis for studying this new modality with regard to conveyed affective associations. Emojis have become widespread in social media, and are variously used to carry emotional and contextual information pertaining to the content of social media posts. There have been studies exploring the relationship between hashtags and tweets (Ferragina et al., 2015), and between emojis and tweets (Campero et al., 2017). Additional research has aimed at conducting sentiment analysis based on emojis and hashtags (Novak et al., 2015). A number of other works study the connection between words and emotions, resulting in datasets such as *EmoLex* (Mohammad and Turney, 2013). Most of these studies relied upon a crowdsourcing approach to compile the data and lexicons and to capture relationships among linguistic and paralinguistic elements (Kulahcioglu and de Melo, 2019).

However, previous work has neglected to focus on the emotional aspects of emojis. For instance, we may ultimately be interested in devising a system that jointly assesses the affect conveyed by a tweet based not only on the words, but also in part on the emojis occurring within it. In some cases, an emoji may reinforce the emotion conveyed by the text. In other cases, it may reveal an additional dimension of affect. In some cases, it may also point in the opposite direction, e.g., by helping to discern sarcasm, which otherwise might be hard to ascertain in certain contexts. Currently, there are no readily available resources to understand the direct relationship between emojis and emotions.

We address this gap by harvesting an emoji-centric collection of tweets. From this, we create the *EmoTag* resource. The name alludes to its usefulness in exploiting emoji for emotional tag-

ging. The resource is based on a series of co-occurrence statistics that allow us to quantify the emotional associations of individual emojis. We subsequently assess these connections in a series of experiments and case studies.

2 Background

Emotion and Communication. Facial expression has been an important aspect of communication that predates the emergence of mankind. Chevalier-Skolnikoff, in ascending order of phylogenetic complexity, draws connections between the degree of evolution of the brain and the spectrum of facial expression observed for a species (Chevalier-Skolnikoff, 1973). Charles Darwin’s well-known volume on the expression of emotions (Darwin, 1872) analysed the connection between emotions and their expression. He remarked for instance, that for both animals and humans, anger coincides with eye muscle contractions and teeth exposure, and commented on the fact that humans lift their eyebrows in moments of surprise. His work then goes on to study the role of such forms of facial expression in conveying to others how an animal feels, studying primates as well as human infants and adults.

In light of this, humans continue to rely extensively on such nonverbal cues even in oral forms of linguistic communication. Although a person’s emotion and mood can to some extent be conveyed by means of suitable content words (e.g., “I am happy to hear that!”) or interjections (“Wow!”), face-to-face communication has important properties that written communication tends to lack (Bordia, 1997). These include facial expressions of the aforementioned sort, but also gesture and intonation. In certain circumstances, e.g. certain problem-solving settings, face-to-face communication may hence prove more efficient and effective (Bordia, 1997).

Accordingly, since the beginning of writing, humans have resorted to surrogate mechanisms to convey emotive signals, attempting to push the boundaries and overcome some of the inherent restrictions of plain written language as a medium. Examples include illustrative embellishments and ornaments, calligraphy, a judicious use of color, and various typographic instruments. For instance, it has been shown that the choice of font may radically alter the affective perception of a text (Juni and Gross, 2008; Kulahcioglu and de Melo, 2018).

Emoticons and Emoji. While emoticons such as “:-)” and Japanese 顔文字 (kaomaji) such as “(^_^)”, both based on regular characters, have been in use for several decades, emojis originated in Japan in the 1990s and have only recently spread globally. Despite the lexicographic similarity between the two words *emoji* and *emotion*, etymologically, the former stems from the Japanese words 絵 (e, *picture*) and 文字 (moji, *character*). Emoji characters, similar to earlier dingbat characters, are pictorial and colorful.

Their principal use has indeed been to convey emotion, particularly via facial expression emojis. In 2015, Oxford Dictionaries declared the *Face with Tears of Joy* emoji its Word of the Year 2015. Kaye et al. (2017) explained how emojis may aid the interlocutor in disambiguating utterances that would otherwise remain ambiguous. Emojis may also be useful as a more instantaneously and widely recognized form of communicating degrees of satisfaction. Kay et al. go as far as suggesting them for consideration as possible alternatives to regular Likert scales (Kaye et al., 2017).

Historically, the spread of emojis has been driven in large part by their adoption in popular messaging and social media platforms, which led, among things, to their inclusion in Shift JIS, and, subsequently, the Unicode standard. Nowadays, they are ubiquitous in social media and chat applications, but increasingly also in emails and other digital correspondence.

3 Related Work

Emoticons. Early studies focused on the use of emoticons in social media. Go et al. (2009) proposed a form of distant supervision by using emoticons as noisy labels for Twitter sentiment classification. Davidov et al. (2010) adopted a fairly similar approach by handpicking smileys and hashtags as tweet labels and relying on a supervised method for sentiment analysis of tweets.

Emoji Semantics. A prominent work on emojis is the DeepMoji project (Campero et al., 2017) from MIT. It provided a model that recommends emojis given a natural language sentence as input. The deep learning model was trained on a collection of 1.2B tweets to learn the sentiment, emotions, and the use of sarcasm in short text.

Barbieri et al. (2016) proposed a method to learn vector space embeddings of emojis using the

standard word2vec skip-gram approach, applied to a large collection of tweets. In contrast, [Eisner et al. \(2016\)](#) attempted to learn vector embeddings of emojis based on their short descriptions in the Unicode standard.

Emoji Associations. The first paper that thoroughly investigated the sentiment of emojis ([Novak et al., 2015](#)) proposed a sentiment ranking of 715 emojis on a corpus of 70,000 tweets. This work provides a basis for future research on the logographic usage of emojis in social media.

[Zhou and Wang \(2017\)](#) trained a natural language conversation model that accounts for the underlying emotion of utterances by exploiting the existence of emojis as a signal.

[Rakhmetullina et al. \(2018\)](#) proposed a method to classify emojis with regard to their sentiment and emotion. Their corpus consists of 500 labeled tweets, and they categorize emojis by assigning them labels for 8 emotions. For this, they applied a distant supervision technique for a reliable mapping based on manually annotated data.

4 EmoTag

Given the prominence of emojis in human communication, our work seeks to study relevant associations of emojis. We begin by assembling a dataset for this purpose (Section 4.1), and subsequently induce a series of lexicons that reveal potential connections (Section 4.2), including between words and emojis, as well as between emojis and emotions.

4.1 Data Collection

In assembling a collection of social media postings containing both emojis and hashtags with tweets, one strategy would be to rely on available datasets and filter them so as to retain only those entries that contain both emojis and hashtags. However, this approach results in a comparably small number of postings. Despite the overall surge in popularity of emojis, only a fraction of all postings includes emojis.

Instead, we proceeded to compile a new dataset of about 20.8 million tweets by specifically searching for postings that contain emojis. For the set of target emojis, our goal was to focus on emojis associated with emotions, as opposed to generic symbols from domains such as transportation or household appliances. To this end, we relied on a set of most frequently used 620 emojis from [No-](#)

[vak et al. \(2015\)](#) and from [Emoji Tracker¹](#), a website that monitors the use of emojis on Twitter in realtime.

Using our set of frequent emojis as search terms, we retrieved tweets that specifically contain one or more of these target emojis. The number of tweets is evenly distributed across different emojis. While tweets can be in any language, we only collected tweets labeled as being in English. In total, we obtained a set of 20.8 million tweets over a span of one year. In addition to the volume that such a large time span provides, collecting the data for every day of the year aids in mitigating the effect of potential biases in the data. All collected tweets contain at least one emoji.

Note that only a fraction of all tweets have hashtags. Specifically, within our collected data, we found that only 10-15% of our tweets with emojis also include hashtags. To clean up the data, we removed usernames (marked with @-symbol), tweets consisting only of hashtags and emojis but no text, tweets that only contain a short time stamp such as “6AM” or simply a URL (with or without the “http://” prefix), as well as all duplicate tweets.

4.2 Lexicon Induction

Based on the corpus, EmoTag is constructed as a series of lexicons.

4.2.1 Co-occurring Emojis

We first collect a series of co-occurrence based lexicons. Each entry in such a lexicon is the representation of pairwise count of desired unigram tokens. These resources can be useful for the community, but also allow us to conduct analyses of the data.

In our tweet collection, there are roughly 36K tweets per emoji, and these have a uniform distribution across the collection time period.

Inspecting the results, we observe that the overall top-ranked pair of co-occurring emojis in our dataset is U+1F61D 🤪 and U+1F61C 🤩. These showed up together 42K times, which is fairly frequent in comparison with other pairs. Note that U+1F61D 🤪 is the “face with stuck-out tongue and tightly-closed eyes” emoji, while U+1F61C 🤩 is the “face with stuck-out tongue and winking eye” emoji.

Another emoji, U+1F602 😂, the “face with tears of joy” one, appears to be the most common emojis to co-occur saliently with others. It appears

¹<http://emojitracker.com/>

Unicode	Emoji	Sentiment Score	Description
U+1F649	🙈	0.333	Hear-no-evil monkey
U+1F648	🙉	0.432	See-no-evil monkey
U+1F649	🙈	0.333	Hear-no-evil monkey
U+1F64A	🙊	0.459	Speak-no-evil monkey
U+1F62D	😭	-0.093	Loudly crying face
U+1F602	😄	0.221	Face with tears of joy
U+1F620	😡	-0.299	Angry face
U+1F608	👹	0.265	Smiling face with horns
U+1F620	😡	-0.299	Angry face
U+1F629	😓	-0.368	Weary face

Table 1: Similarities and Contrasts of Co-occurring Emojis

with a broad range of other emojis with a relatively high frequency. Three other popular emojis that co-occurred with U+1F602 😄 include U+1F62D 😭 (“loudly crying face”), U+1F648 🙈 (“see-no-evil monkey”), and U+1F629 😓 (“weary face”).

Somewhat different from the previous cases, the fourth pair in Table 1 involves the emoji U+1F62D 😭, i.e., a crying face, and U+1F602 😄, i.e., a face with tears of joy. This is unusual in the sense that these two emojis possess opposite sentiment polarities. According to Novak et al. (2015), the sentiment value of U+1F62D 😭 is -0.093, whereas the sentiment value of U+1F602 😄 is 0.221, i.e., a positive sentiment. This suggests that people tend to conflate the two due to their similar appearance, as both involve tears. Another possibility is that people may be using one of the two sarcastically. As shown in the table, similar observations can also be made for certain other pairs of emojis.

Our results also show a correlation between U+1F60D 😍 and U+1F629 😓. The two are paired up around 2,500 times, illustrating another connection between a positive and a negative sentiment emoji. U+1F629 😓 is the “weary face” emoji, whereas U+1F60D 😍 is the “smiling face with heart-shaped eyes” one. This appears to stem from tweets that express positive sentiment about a target entity, but also negative sentiment about the current situation.

4.2.2 Emoji–Words Lexicon

Another lexicon that we produce aims to provide co-occurring words for a given emoji, or vice versa, emojis for a given word. Table 2 shows an excerpt of the emoji–word lexicon, grouped by words. For example, the word “miss” co-

Word	Emoji	Description
miss	U+1F62D 😭	Loudly crying face
	U+2764 ❤️	Heavy black heart
	U+1F622 😓	Crying face
happy	U+1F389 🎉	Party popper
	U+2764 ❤️	Heavy black heart
	U+1F618 😘	Face throwing a kiss
love	U+1F60D 😍	Smiling face with heart-shaped eyes
	U+1F618 😘	Face throwing a kiss
	U+2764 ❤️	Heavy black heart

Table 2: Co-occurring Emojis and Words

occurs with a wide range of emojis, but the top co-occurring emojis are U+1F62D 😭, U+2764 ❤️, and U+1F622 😓. These emojis are likely to be used when someone misses someone or something. Similarly, the words “happy” and “love” appear with numerous emojis that carry happy and positive sentiment.

4.2.3 Emoji–Hashtags Lexicon

This lexicon provides a collection of hashtags along with the emojis that they co-occur with. The resource also includes the corresponding co-occurrence frequencies between emojis and hashtags. According to our findings, the emoji U+1F637 🤒 (“face with medical mask”) co-occurs with the hashtags #sick, #flu, #yuck, #cold, #insomnia, and #dying, which all are clearly semantically relevant for this emoji.

4.3 Interpretable Emoji-Based Word Vectors

Interpretability and explainability are widely regarded as highly desirable attributes of AI-driven decision making (Xian et al., 2019). Dense word vectors such as those produced by word2vec (Mikolov et al., 2013) are ubiquitous in NLP (de Melo, 2017). However, it is often remarked that they lack interpretability, in the sense that individual values in such vectors do not carry any easily interpretable inherent significance. Previous work has proposed interpretable word vectors consisting of one or more sentiment polarity scores for a word (Dong and de Melo, 2018; Dong and de Melo, 2018). Given that emojis represent a wide spectrum of aspects considered relevant in human communication, we study to what extent emojis can serve as a means of inducing word vectors endowed with interpretability.

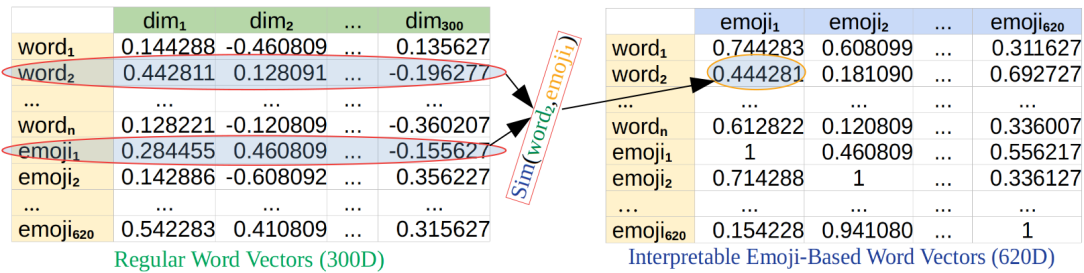


Figure 1: Inducing Interpretable Word Vectors via Emojis

This can be achieved by assigning every word a 620-dimensional word vector, in which each dimension reflects the association of that word with one out of 620 emojis. Since we use a list of the 620 most frequent emojis, the dimensionality of a vector becomes 620. An obvious method would be to adopt just simple frequency counts as the values in these vectors, i.e., for a given word, the entries in its word vector would simply reflect the number of times that word co-occurred with a given emoji.

However, we can improve over this by relying on the word2vec Skip-Gram with Negative Sampling algorithm (Mikolov et al., 2013) as an intermediate representation, as illustrated in Fig. 1. We first train such a word2vec model on the EmoTag corpus. Then a cosine similarity score is calculated between all words and emojis. This yields a semantic relatedness score in $[0, 1]$ for any word–emoji pair. Thus, we can view the score as reflecting to what extent a word correlates with an emoji. We use these correlation coefficients to form a word vector $\mathbf{v}_w \in [0, 1]^d$ for every word w , such that each of the $d = 620$ dimensions reflects the correlation with a particular emoji. This is the final EmoTag word vector representation that we use in all experiments.

5 Evaluation

In the following, we evaluate EmoTag for machine learning-driven emotion analysis of tweets and show how it can be used to reveal the sentiment and emotion of individual emojis.

The first study aims at evaluating the usefulness of our interpretable EmoTag word vectors in a downstream task, exploiting them in a machine learning-driven system that seeks to identify the emotion intensity of tweets.

Subsequently, we use our data to compute sentiment polarity scores for emojis, comparing these against existing human annotations of emoji senti-

ment.

Finally, we develop the first resource providing emotion scores for emojis. We evaluate these by showing how they can be used to automatically induce emotion scores for words.

5.1 Emotion Intensity Prediction with Interpretable Emoji-Based Word Vectors

We begin by evaluating the interpretable emoji-based word vectors, assessing to what extent they are able to keep up with regular word vectors in a downstream task relating to emotions.

Benchmark. In particular, we consider the EmoInt Shared Task from WASSA (Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis) 2017 (Mohammad and Bravo-Marquez, 2017a), which involves determining the intensity or degree of emotion felt by a speaker when a tweet and a target emotion are given. Tweets were provided for four different emotion categories (anger, fear, joy, and sadness), and the ground truth intensity values range between 0 and 1.

The Affective Tweets (AT) package was provided to all participants as a baseline for the competition (Mohammad and Bravo-Marquez, 2017b), providing a rich set of features constructed based on several emotion and sentiment lexicons such as NRC-EmoLex, NRC10E, etc. (Mohammad and Bravo-Marquez, 2017a).

Model. We rely on a deep neural network to predict the emotion intensity for each tweet, adopting a similar CNN-LSTM architecture as that of IMS (Köper et al., 2017), the 2nd-ranked system among all participants in the competition, with the CNN architecture based on that proposed by Kim (2014). In training, each tweet is represented by a matrix of size $m \times d$, where d is the dimensionality of the pre-trained word vectors and $m = 50$ is the maximal token sequence length considered for

Method	A	F	J	S	Avg	d
<i>Interpretable</i>						
AT	0.65	0.66	0.60	0.69	0.65	n/a
EmoTag	0.70	0.73	0.69	0.75	0.72	620
<i>Non-Interpretable</i>						
Random Init.	0.68	0.72	0.66	0.73	0.70	300
Google News	0.70	0.72	0.67	0.75	0.71	300
GloVe	0.70	0.73	0.68	0.76	0.72	300
GloVe-Twitter	0.72	0.74	0.68	0.76	0.73	200
IMS	0.71	0.74	0.71	0.71	0.72	300

Table 3: Comparing with other methods, with regard to anger (A), fear (F), joy (J), sadness (S), average (Avg), dimensionality (d).

a tweet. We can thus feed in either regular word vectors or our interpretable emoji-based EmoTag vectors for the series of words in the tweet. We applied a dropout rate of 0.25. The obtained matrix then serves as input to a convolutional layer with a window size of 3, followed by a max-pooling layer (size 2) and an LSTM (Hochreiter and Schmidhuber, 1997) to predict a numerical output for each tweet. This numerical value was then added as a feature along with other auxiliary features, and passed to a Random Forest regressor to obtain the final intensity score for a particular emotion. The IMS team used a total of 142 features, including the 45 baselines features from Affective Tweets. Since we are comparing our results with both the baseline features and the features used by the IMS team, our classifier is also fed with the 142 features. All features were passed to a random forest regressor with 800 trees for identifying the intensity of a given emotion. A separate model is trained for each of the four target emotions.

Results. Table 3 summarizes the results for the EmoInt task, providing Pearson correlations for each emotion as well the average Pearson correlation for all four emotions, along with the dimensionality of the respective word vectors used in the experiments. The results show that the interpretable word vectors in EmoTag are able to yield results that are comparable with those of other dense word representations that are not interpretable. It should be kept in mind that EmoTag was built based on a very small corpus, i.e., only 20M tweets, comparing to the massive size of the corpora used for pretrained word vectors such as the two GloVe models. For further comparison, we also report results on just the AffectiveTweets (AT) features, as well as the original IMS system. In some cases, for example for the *sadness* emo-

Unicode	Emoji	Novak Score	EmoTag Score	Description
U+1F61C	🙄	0.455	0.482	Face with stuck-out tongue and winking eye
U+1F617	😘	0.611	0.591	Kissing face
U+1F49A	💚	0.656	0.654	Green heart
U+1F48B	💋	0.691	0.744	Kiss mark

Table 4: Comparison of emoji sentiment scores from EmoTag and Novak et al. (2015).

tion, EmoTag actually outperforms the IMS team’s baseline.

5.2 Evaluating the Sentiment of Emojis

Next, we evaluate to what extent our interpretable word-emoji vectors can aid in revealing the sentiment of emojis.

Method. For obtaining sentiment scores, we rely on the NRC Emotion Lexicon EmoLex (Mohammad and Turney, 2013), a list of English words and their associations with eight basic emotions (anger, fear, anticipation, trust, surprise, sadness, joy, and disgust) and two sentiments (negative and positive). The associations are merely given as Boolean labels (0 or 1). To obtain a sentiment score for an individual emoji, we first consider all words with a sentiment score of 1 in EmoLex. Then, we rank all words associated with the given emoji based on their similarity score according to our interpretable word vectors, where a higher similarity score results in a higher rank. According to the ranking, the top $K = 3$ words are picked and their similarity scores are aggregated using a simple addition, which becomes the ultimate sentiment score for the given target emoji.

Results. To evaluate the sentiment score of emojis, we measure the Pearson correlations for several groups of emojis treating the scores by Novak et al. (2015) as gold scores. Table 5 summarizes the Pearson correlations for several groups of emojis. The first row of the table represents Novak’s top 100 positive sentiment emojis. We also consider additional groups based on the Unicode standard emoji descriptions, particularly those with a face and those with monkey faces.

Note that we observed a high positive sentiment score for all emojis with kiss symbol or kissing face in our data, compared to Novak’s scores. For some emojis, our model obtains a high sentiment score such as 0.991 for 🐱 “Kissing Cat Face

Emoji Group	Correlations
Top 100 positive emojis	0.71
Emojis with face	0.45
Monkey face emojis	0.53
Emojis with kissing	0.14

Table 5: Pearson Correlations for Sentiment Score

with Closed Eyes” U+1F63D, whereas the score by Novak et al. (2015) is 0.571. This can happen for several reasons. In some cases, the sentiment scores they propose may be misleading for certain emojis, especially if they are less frequent in their dataset. An example is 🙄 U+1F63D, which has an occurrence frequency of 88 only, compared to emojis such as 😄 “Face with Tears of Joy” U+1F602, which occurred 14,622 times. Thus, in some cases, their results may not be reliable.

Still, the results often show a strong agreement, although our method produces sentiment scores for emojis only indirectly via their associations with words. Table 4 provides examples of such sentiment scores generated by EmoTag and Novak et al. (2015).

5.3 Evaluating Emotion Profiles of Emojis

Finally, we use our data to evaluate to what extent emojis are associated with certain emotions. For this, we again rely on our emoji-based word vectors in conjunction with EmoLex, the NRC Emotion Lexicon (Mohammad and Turney, 2013). EmoLex provides a set of words along with a set of binary labels, where 1 signifies that the word carries a particular association, while 0 represents the negative case.

Method. First, for each emoji, we identify the top K in EmoLex according to their cosine similarity with the emoji, as obtained in our interpretable word vectors, where a higher similarity score entails a higher rank. For the top K words, we compute a weighted average of emotion labels. The emotion labels are taken from EmoLex, while the similarity scores are used as weights. This weighted average then serves as the final emotion score of the emoji. The same process is followed for all emojis. **Results.** We evaluate our induced emoji emotion scores indirectly by using them to reproduce emotion intensity scores for words, for which we have ground truth intensity scores in the Affect Intensity lexicon by (Mohammad, 2018). This lexicon comes with 6K tokens, where tokens

are grouped by the four emotions *anger*, *fear*, *joy*, and *sadness*. It provides crowdsourced emotion intensity scores, which range between 0 and 1, with 1 meaning that the word exhibits the highest degree of association with a particular emotion and 0 referring to the lowest degree. Note that this ground truth resource is distinct from the NRC Emotion Lexicon used in inducing our scores. The latter merely provides Boolean labels for word–emotion pairs, and thus it is non-trivial to derive affect intensity scores from it, particularly via emojis.

To reproduce word emotion intensities based on our emoji emotion scores, we proceed as follows. For a given word w , we rank the top K emojis based on their similarity score in the EmoTag word vectors, where higher scores entail a higher rank. Once the top K emojis have been identified, we then compute the arithmetic mean of the emotion scores of those related emojis, which yields the final emotion score for the target word w . We chose $K = 10$, which led to better results than alternative values.

Table 8 depicts the Pearson correlations for different subsets of the Affect Intensity lexicon. These correlations reveal how close we are in predicting the emotion score for a given word based on our emoji emotion scores. The first row shows the scores for words that are common to all four emotion groups, whereas the last row includes all words. Table 7 provides examples of emotion scores for a few select emojis.

Analysis. For further analysis, we compare our scores with the classification obtained by Rakhmetullina et al. (2018). Table 6 compares the emotional label that their classification provides against our emotion scores for *anger*, *joy*, *sadness*. Note that this is the complete set of emoji results provided in their paper, apart from one additional emoji for the emotion *surprise*, which our method currently does not support, due to its omission in EmoLex. Their labeling did not include the emotion *fear*, so we omit it in our comparison. The bold scores in the last three columns indicate what emotion labeling we would obtain if we had to select a single label for an emoji based on our obtained emotion intensity scores. For example, in our case, emoji 🙏 “Folded Hands” U+1F64F obtains the highest score 0.485 for the emotion *joy*, which is labeled as being in the *joy* category in their study as well. There are three cases (high-

Unicode	Name	Emoji	E2E Label	Anger	Joy	Sadness
U+1F612	Unamused face	😏	anger	0.418	0.119	0.333
U+1F602	Face with tears of joy	😂	joy	0.381	0.099	0.326
U+1F60D	Smiling face with heart-eyes	😍	joy	0.307	0.308	0.137
U+1F60A	Smiling face with smiling eyes	😊	joy	0.067	0.248	0.247
U+1F495	Two hearts	❤️	joy	0.172	0.383	0.142
U+1F601	Beaming face with smiling eyes	😄	joy	0.091	0.123	0.079
U+263A	Smiling face	😊	joy	0.095	0.245	0.176
U+1F604	Grinning face with smiling eyes	😁	joy	0.184	0.188	0.149
U+1F618	Face blowing a kiss	😘	joy	0.233	0.215	0.144
U+1F64F	Folded hands	🙏	joy	0.187	0.485	0.351
U+1F62D	Loudly crying face	😭	sadness	0.246	0.198	0.272
U+1F629	Weary face	😓	sadness	0.236	0.186	0.234
U+1F622	Crying face	😢	sadness	0.284	0.210	0.333

Table 6: A comparison between Emoji2Emotion (E2E) and EmoTag

Emoji	Name	A	F	J	S
U+1F620 😡	Angry face	0.49	0.36	0.07	0.44
U+1F46E 👮	Police officer	0.34	0.49	0.16	0.27
U+1F492 💒	Wedding	0.09	0.14	0.63	0.25
U+1F4A9 💩	Pile of poo	0.35	0.34	0.15	0.47

Table 7: Emotion scores of emojis for anger (A), fear (F), joy (J), sadness (S).

Tokens	A	F	J	S	Avg
Common Words	0.51	0.41	0.19	0.50	0.40
All Words	0.45	0.40	0.20	0.45	0.37

Table 8: Pearson Correlations of gold scores and our predicted scores for Affect Intensity lexicon

lighted in red) at which our scoring would fail. According to their labeling system and results, emoji 😓 “Weary Face” U+1F629 should have obtained its highest score for *sadness* (0.234) instead of *anger* (0.236), though both scores are very close in our case.

The emoji 😂 “Face With Tears of Joy” U+1F602 scored 0.381 on *anger*, which is the highest among all scores for it, although the authors of Emoji2Emotion marked it as belonging to the *joy* category. This may stem from the phenomenon of people frequently confusing this emoji with the 😭 “Loudly Crying Face” U+1F62D emoji. In Table 1, we observe that both appear together very often, which results in a strong association with a negative emotion (anger) for an emoji that intrinsically ought to be more associated with joy.

6 Conclusion

The characteristics of a medium profoundly affect the way that people express themselves using said medium. While written communication lacks the non-verbal cues that make face-to-face communication particularly effective for problem-solving (Bordia, 1997), modern social media, and messaging platforms have unique properties that are interesting in their own right. Among these, the use of emojis stands out as meriting very special consideration, not least due to their ability to compensate for some of the shortcomings of written language as a medium in conveying emotion and affect.

While research in social science and social media analytics has extensively studied the use of emojis in everyday communication, previous work has not fully explored the connection between emojis and emotion. This paper presents a detailed analysis of how emojis and words co-occur in social media, including their connection to emotions. It also shows how an interpretable word embedding can be formed with the help of emojis, which shows promise as an additional ingredient in emotion detection-related tasks.

Another key contribution of this work is the creation of a large resource, consisting of several different sub-lexicons that describe connections among emoji, words, and other items, as well as emotion scores for emojis, which are released to the public². We hence believe that this work will substantially benefit other researchers in several different fields.

²<http://emoji.nlproc.org>

References

- Francesco Barbieri, Francesco Ronzano, and Horacio Saggion. 2016. What does this emoji mean? a vector space skip-gram model for twitter emojis. In *Language Resources and Evaluation conference, LREC*. Portoroz, Slovenia.
- Prashant Bordia. 1997. Face-to-face versus computer-mediated communication: A synthesis of the experimental literature. *The Journal of Business Communication (1973)* 34(1):99–118. <https://doi.org/10.1177/002194369703400106>.
- Andres Campero, Bjarke Felbo, Joshua B. Tenenbaum, and Rebecca Saxe. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. *arXiv 1708.00524* <https://arxiv.org/abs/1708.00524>.
- Suzanne Chevalier-Skolnikoff. 1973. Facial expression of emotion in nonhuman primates. In *Darwin and Facial Expression: A Century of Research in Review*, Academic Press New York, NY, pages 11–89.
- Charles Darwin. 1872. *The Expression of the Emotions in Man and Animals*. Appleton. The original was published 1898 by Appleton, New York. Reprinted 1965 by the University of Chicago Press, Chicago and London,.
- Dmitry Davidov, Oren Tsur, and Ari Rappoport. 2010. Enhanced sentiment learning using twitter hashtags and smileys. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*. Association for Computational Linguistics, Stroudsburg, PA, USA, COLING '10, pages 241–249. <http://dl.acm.org/citation.cfm?id=1944566.1944594>.
- Gerard de Melo. 2017. Multilingual vector representations of words, sentences, and documents. In *Proceedings of IJCNLP 2017*. <http://www.aclweb.org/anthology/I17-5002>.
- Xin Dong and Gerard de Melo. 2018. Cross-lingual propagation for deep sentiment analysis. In *Proceedings of AAAI 2018*. AAAI Press.
- Xin Dong and Gerard de Melo. 2018. A helping hand: Transfer learning for deep sentiment analysis. In *Proceedings of ACL 2018*.
- Ben Eisner, Tim Rocktäschel, Isabelle Augenstein, Matko Bosnjak, and Sebastian Riedel. 2016. emoji2vec: Learning emoji representations from their description. In *Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media*. Association for Computational Linguistics, Austin, TX, USA, pages 48–54. <http://aclweb.org/anthology/W16-6208>.
- Paolo Ferragina, Francesco Piccinno, and Roberto Santoro. 2015. On analyzing hashtags in twitter. In *International AAAI Conference on Web and Social Media*.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *Processing* pages 1–6.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Samuel Juni and Julie S Gross. 2008. Emotional and persuasive perception of fonts. *Perceptual and motor skills* 106(1):35–42.
- Linda K. Kaye, Stephanie A. Malone, and Helen J. Wall. 2017. Emojis: Insights, affordances, and possibilities for psychological science. *Trends in Cognitive Sciences* 21(2):66 – 68. <https://doi.org/https://doi.org/10.1016/j.tics.2016.10.007>.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, pages 1746–1751. <https://doi.org/10.3115/v1/D14-1181>.
- 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*. Workshop at Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Copenhagen, Denmark.
- Tugba Kulahcioglu and Gerard de Melo. 2018. Fontlex: A typographical lexicon based on affective associations. In Nicoletta Calzolari (Conference chair), Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga, editors, *Proceedings of the 11th Language Resources and Evaluation Conference (LREC 2018)*. European Language Resources Association (ELRA), Paris, France.
- Tugba Kulahcioglu and Gerard de Melo. 2019. Paralinguistic recommendations for affective word clouds. In *Proceedings of ACM IUI 2019*. ACM, New York, NY, USA, pages 132–143. <https://doi.org/10.1145/3301275.3302327>.
- 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*, Curran Associates, Inc., pages 3111–3119. <http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>.

- Saif Mohammad. 2018. [Word affect intensities](#). In *Proceedings of the 11th Language Resources and Evaluation Conference*. European Language Resource Association, Miyazaki, Japan. <https://www.aclweb.org/anthology/L18-1027>.
- Saif Mohammad and Felipe Bravo-Marquez. 2017a. [WASSA-2017 shared task on emotion intensity](#). In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Association for Computational Linguistics, Copenhagen, Denmark, pages 34–49. <https://doi.org/10.18653/v1/W17-5205>.
- Saif M. Mohammad and Felipe Bravo-Marquez. 2017b. Emotion intensities in tweets. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM)*. Vancouver, Canada.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence* 29(3):436–465.
- Petra Kralj Novak, Jasmina Smailović, Borut Sluban, and Igor Mozetič. 2015. [Sentiment of emojis](#). *Plos One* 10(12). <https://doi.org/10.1371/journal.pone.0144296>.
- Aisulu Rakhmetullina, Dietrich Trautmann, and Georg Groh. 2018. [Distant supervision for emotion classification task using emoji 2 emotion](#). In *Proceedings of the 1st International Workshop on Emoji Understanding and Applications in Social Media (Emoji2018)*. Stanford, CA, USA. <http://ceur-ws.org/Vol-2130/>.
- Yikun Xian, Zuohui Fu, S. Muthukrishnan, Gerard de Melo, and Yongfeng Zhang. 2019. [Reinforcement knowledge graph reasoning for explainable recommendation](#). In *Proceedings of SIGIR 2019*. ACM, New York, NY, USA, pages 285–294. <https://doi.org/10.1145/3331184.3331203>.
- Xianda Zhou and William Yang Wang. 2017. [Mojitalk: Generating emotional responses at scale](#). *arXiv 1711.04090* <https://arxiv.org/abs/1711.04090v1>.