Semantics and Sentiment: Cross-lingual Variations in Emoji Use

Anonymous ACL submission

Abstract

Over the past decade, the use of emojis in social media has seen a rapid increase. Despite their popularity and image-grounded nature, previous studies have found that people interpret emojis inconsistently when presented in context and in isolation. In this work, we explore whether emoji semantics differ across languages and how semantics interacts with sentiment in emoji use across languages. To do so, we develop a corpus in English, Portuguese and Chinese containing definitions L1 speakers of these languages give to a set of emojis. We then use these definitions to assess whether speakers of different languages agree on whether an emoji is being used literally or figuratively, as well as whether this literal and figurative use correlates with the sentiment of the context wherein the emoji is grounded. We found that there were varying levels of disagreement on the definition for each emoji but that these stayed fairly consistent across languages. We also demonstrated a correlation between the sentiment of a tweet and the figurative use of an emoji, providing theoretical underpinnings for empirical results in NLP tasks, particularly offering insights that can benefit sentiment analysis models.

1 Introduction

Nowadays, much of our communication happens through text-based messaging on online mediums, otherwise known as computer mediated communication (CMC). Given that many natural features of language cannot be encoded in such a modality (e.g., prosody, visual context), speakers have come up with other strategies to communicate their intentions. One such strategy is to use emojis, digital icons that can be used separately or combined with text to provide extra information regarding the desired meaning of an utterance. It is hardly surprising then that the variety and popularity of emojis have increased rapidly over the past 10 years, with 3664 emojis officially encoded in the Unicode standard and used in over 22% of the tweets sent thus far (Broni, 2022).

This increase in popularity has also given rise to a growing interest in research from various domains and disciplines on emojis, their semantics, and their use in the language. To illustrate, those who work on language models have been interested in how emojis might aid such systems, e.g., in tasks such as sense disambiguation (Shardlow et al., 2022). On the other hand, psychologists and linguists have also been interested in investigating how people have integrated emojis into their language use (e.g., Gettinger and Koeszegi (2015); Braumann et al. (2010)) and the communicative functions for which they are important (e.g., Dresner and Herring (2010); Lee et al. (2016)). However, such studies are not generalisable to cultures and languages beyond English. This sole focus on English can lead to many potential harms, including technologies which are unable to be effective for a large proportion of society.

A first attempt to bridge this gap was made by Barbieri et al. (2016) who examined variation in emoji use across languages. However, their approach solely relied on the analysis of emoji vector representations, which fail to capture the complete semantic nuances of emojis. They did not incorporate the examination of human judgments in their methodology. Instead, the emoji vectors were generated based on contextual information from tweets, and subsequently, similarities were computed to assess the distinctions in emoji usage across different languages. Apart from a few other studies such as Lu et al. (2016) and Herring (2018), the cross-lingual aspect of emoji use has been relatively under-explored. This coupled with the increase in emoji use underlines the importance of further research into emoji variation and semantics as it has real-world implications in understanding online social trends, and CMC in general. Therefore, this study aims to explore the sentiment and semantics of emoji use across languages. Specifically, we will focus on the literal and figurative use of emojis used in tweets, as well as their correlation with the sentiment of the utterances in which they appear. To do so, we pose the following research questions (RQs):

RQ1: Do people disagree on an emoji's context-free interpretation within and across languages?
RQ2: Does agreement on the figurative or literal use of an emoji differ across different languages?
RQ3: Does the figurative use of emojis correlate with the sentiment of the context in which the emoji is used?

To address these questions, we carried out two online experiments in English[1], European Portuguese, and Mandarin Chinese. The first experiment aimed to collect participants' interpretation of isolated emojis (similar to the work of Częstochowska et al. (2022)) and establish the literal meaning of the emojis analysed in the second experiment. The objective of the second experiment was to gather participants' interpretations of emojis presented in textual context in regard to their sentiment and agreement with the provided literal meaning. Our overall results show that: (i) across languages, emoji meanings are fairly consistent, and (ii) there is a correlation between emoji use (literal/figurative) and sentiment (positive/negative). The data collected for our experiments will be publicly released as additional resources for the sentiment analysis and emojis' figurative use detection tasks. In the following sections, we first detail the theoretical background with which we motivate our RQs and methods, we then describe the methods used to collect data, followed by our results and a discussion. We conclude by discussing directions for future work, and the limitations of our study.

2 Background

2.1 Literal and Figurative meaning

There has been much debate in pragmatics on how to define literal and figurative language. Gibbs Jr et al. (1993) presents multiple ways that these two meanings have been defined over the years. For example, subject-matter literality, where we decide based on how often such expressions are used to talk about that subject matter, or truth-conditional literality, in which literal is defined as being anything that meets the truth conditions. Such a decision has consequences on the whole field and any future research on literal and figurative language (Giora, 1997, 2002; Gibbs Jr, 2002). To the best of our knowledge, no definitions for literal and figurative meaning exist in the context of emojis (and there are a variety of different definitions for figurative and literal meaning in general). Hence, based on one possible distinction from Gibbs Jr et al. (1993) known as *Context-Free Literality* which posits that "the literal meaning of an expression is its meaning apart from any communicative situation or its meaning in a 'null context' ", we provide the following definitions:

- **Literal meaning**: Conventional meaning given to an emoji when it is presented in isolation;
- **Figurative meaning**: any other meaning that differs from the literal meaning.

2.2 Emoji interpretation

Extracting the literal meaning of an emoji using these definitions would appear to be a trivial task. However, this is not the case. Częstochowska et al. (2022) found that, when participants are asked to give a one-word definition of an emoji, there are often quite high levels of disagreement. This varies across emojis, with some having higher levels of ambiguity than others. For example, astrological emojis (e.g., ♊ , ♈ , ♒) are the most ambiguous while heart emojis (e.g., ❤️ , 💛 , 💕) are the least. Similar trends were observed by Miller et al. (2017), who found that people often disagreed on the sentiment expressed by an emoji, both when it was presented in isolation and with its accompanying text.

Not only has there been evidence of within-language disagreement, but researchers have also demonstrated evidence of between-language variation. For example, Barbieri et al. (2016) found variation in what emojis are perceived as being similar in meaning. For example, 😭 was perceived as being highly similar to 😜 in the USA, but not in Spain. A likely reason behind such ambiguity is that emojis have multiple meanings that can be used to express one's intention (Shardlow et al., 2022). Certain emojis have more potential meanings than others, a possible explanation for why people find it harder to agree on a definition for these emojis (Częstochowska et al., 2022). In other words,

[1]We did not differentiate between American and British English

Utterance	Sentiment	Use
1. I went for a walk 😊	Positive	Literal
2. The walk was amazing 😊	Positive	Literal
3. The walk was awful 😊	Negative	Figurative
4. It's awful that she's back in the hospital 😭	Negative	Literal
5. I'm so happy. I got engaged! 😭	Positive	Figurative

Table 1: Examples of emojis' literal and figurative usage to convey sentiment.

emojis will have a literal (i.e., conventional) meaning but may also have multiple figurative meanings. This in line with research showing that emoji meanings are not static but dynamic. For example, Robertson et al. (2021) compared the word embeddings for a set of emojis over time and showed that these embeddings often changed, this demonstrates that perhaps emojis are able to shift fairly easily in terms of their meanings and that people may be aware and capable of interpreting multiple meanings for an emoji at any given moment.

2.3 Emoji Sentiment and Semantics

If emojis have multiple meanings, then it is plausible that certain meanings might become more probable in certain linguistic contexts. One such context is the sentiment of the sentence within which the emoji is placed. It has been demonstrated that a strong association between emojis and sentiment exists (e.g., Braumann et al. (2010). This is evident in the large number of emojis that have been created in order to represent different facial expressions. Furthermore, research from Hogenboom et al. (2013) has shown that emojis may have multiple uses when it comes to expressing sentiment.

Table 1 shows examples of such correlation. In sentence 1, the text itself has no clear sentiment. However, adding the emoji 😊 (which has a positive conventional meaning) provides a positive sentiment for the entire sentence. On the other hand, for sentences 2 and 3, the text itself already has either a positive or negative sentiment. In these cases, the addition of the emoji has intensified or weakened the existing sentiment respectively.

Given this relationship between emojis and sentiment, it is not unreasonable to hypothesise that certain contextual sentiments might bring out the different meanings of an emoji. In other words, the literal meaning might be used in sentences where the text has a certain sentiment, while the figurative meaning(s) might be used for other sentences with a different sentiment. For example, sentences 4 and 5 in Table 1 show texts with a negative and a positive sentiment. However, in both cases, the addition of the emoji 😭 intensifies their respective sentiment. This may be surprising given that the literal meaning of this emoji would strongly appear to be negative. Nevertheless, the emoji is able to intensify the sentiment for both sentences because it has both literal and figurative meanings. In 4, the negative literal meaning relating to sadness is the one being applied. On the other side, in sentence 5, the positive figurative meaning relating to being overcome with emotion is selected instead. Hence, the multiple uses of emojis appear to be important when it comes to sentiment.

2.4 Emojis in NLP

Despite their ubiquitous presence in CMC, the broader significance of emojis within the Natural Language Processing (NLP) domain has been relatively understudied. Given the widespread use of emojis for expressing emotions and textual nuances, previous work has showcased some of the advantages of incorporating emojis into NLP models as supportive elements for tasks such as sentiment analysis, emotion detection, and sarcasm detection, particularly emphasising their utility in multilingual contexts (Felbo et al., 2017; Subramanian et al., 2019; Duarte et al., 2019; Tomihira et al., 2020; Barbieri et al., 2022a; Manias et al., 2023).

Our investigation seeks to shed light on the foundations upon which previous work has been built, underscoring the necessity for a comprehensive evaluation of emojis in NLP. Furthermore, the data collected in our study serves as a valuable resource with potential applications in tasks such as sense disambiguation and sentiment analysis.

3 Methods

3.1 Data and Emoji selection

Ten emojis were selected from the 20 most frequently used emoji in 2021 according to the Unicode Consortium[2]. Of these, 5 face emojis and 5 non-face emojis were selected to balance between faces and non-faces. We further based our selections on ambiguity (semantic variation) scores provided by Częstochowska et al. (2022), selecting emojis with a range of scores for both the face and non-face groups. More details in Appendix A.

In order to analyse the emojis in a textual context, we collected a corpus comprised of 4000 tweets per language per emoji scraped from Twitter™with their provided API. To alleviate any strongly skewed sentiment distributions (e.g., some emojis only being shown in tweets with a positive sentiment), we queried the database using keywords that may convey the sentiment of a tweet. Following this, we used existing sentiment models to assign a sentiment to each tweet (Barbieri et al., 2020; Wang et al., 2022).[3] In addition, profanity checks were used to remove tweets with terms that were deemed explicit[4]. Finally, 1,000 tweets were randomly sampled (100 for each emoji) from the remaining tweets. For each emoji we have at least 1 positive and 1 negative examples for it, there are 10 emojis and therefore grouped into 20 conditions of 25 tweets balanced in terms of sentiment and emoji appearances.

The twitter API limits the number of tweets one can collect in total over a month so it was important to make use of the features provided by tweeter for restricting the data one collects and the main method it provides for doing so is by making use of keywords. The keyword querying is an initial step for identifying tweets with a positive and negative sentiment however we also made use of language models trained specifically for the task of sentiment analysis in these three different languages, if the language models label for the sentiment matched the sentiment intended by the filtering process then the tweet was accepted as conveying the intended sentiment, if there was a mismatch between the two, the tweet was rejected.

3.2 Experimental Design

This study conducted two experiments both involving human participants. All participants were paid on the basis of Prolific's hourly rate of £9/hour. The study received ethics approval from a departmental ethics board (details are not provided here to preserve anonymity).

Experiment 1

The objective of this experiment was to collect single-word definitions for each of the analysed emojis in English, Portuguese or Chinese, as well as their literal meaning. Similar to Częstochowska et al. (2022), participants were presented with the 10 emojis in Table 8 and asked to provide one word that they believed best conveyed the meaning of this emoji in their L1 language (example in Appendix E, Figure 3a). The task included a practice phase, with a different set of emojis, and attention checks to filter out any undesirable results (e.g., from bots and those who were not paying sufficient attention). Additionally, data regarding the participants' demographics such as age, education level, and social media usage (platform and number of hours on social media) were collected prior to the task.

Overall, 30 participants for each language were recruited through Prolific and were L1 speakers of the target language. All participants gave informed consent. The mean age of the participants was 30.5 with a range of 19 to 59. For a detailed distribution by language, see table 9.

Experiment 2

The aim of this task was to obtain results on the perception of emojis as being used figuratively or literally across sentiments.

As per Experiment 1, L1 speakers of the target language were recruited via Prolific. In each task, participants were asked to classify 25 tweets with respect to their semantics (literal or figurative) and sentiment (positive or negative). Specifically, in each trial, an emoji and its literal meaning (obtained from Experiment 1 as described in Section 3.3) was shown alongside a tweet containing the aforementioned emoji. Participants were asked whether the emoji was being used literally or figuratively, according to the literal meaning they were given (Appendix E, Figure 3b), and subsequently, the sentiment (Appendix E, Figure 3c) of the tweet. An additional option ("*I do not understand the tweet*") was given to the participants to filter out potential hard-to-understand/noisy tweets. Similar to Experiment 1, participants had a practice phase before beginning the real task, as well as attention checks. Participants completed the same demographics questionnaire as in Experiment 1.

[2] https://home.unicode.org/emoji/emoji-frequency/

[3] For Portuguese - https://github.com/Logicus03/Bert-Sentiment-Analysis

[4] https://github.com/LDNOOBW/List-of-Dirty-Naughty-Obscene-and-Otherwise-Bad-Words

Responses from 44 Chinese, 35 English and 37 Portuguese speakers were collected from Prolific. All participants were over 18 and gave informed consent. The participants were compensated based on the hourly Prolific rate of £9/hour. The study received ethics approval from the departmental ethics committee. Overall, the participants had an age range of 20 to 57 ($N = 36$, Mean = 31.8, SD = 10.0), for a full breakdown of age by language, see table 10. Overall, 2,765 data points were analysed.

3.3 Data Analysis

Literal Meaning

The literal meaning of each emoji was defined based on collected annotations. To account for variations of the same meaning, the collected one-word definitions were grouped based on their lemma (e.g., “laughing", “laugh", and “laughter" were considered the same as they share the lemma *laugh*). The word within the most frequent lemma group and with the highest relative frequency was selected as literal meaning (as per our definition of literal meaning, Section 2.1). As the concept of lemma cannot be applied to Chinese, the definitions were grouped based on shared characters *ad hoc* (e.g. 爱心 and 热爱 were grouped together as they share the character 爱).

Semantic Variation

In order to assess the agreement on the context-free emojis’ interpretations, the semantic variation metric proposed by Częstochowska et al. (2022) was used. It is defined as follows:

$$\text{sv} = 1 - \sum_{v \in V} f_v \dot{(}\cos(1 - (e_v, e_{v^*}))$$

a weighted sum of the cosine distances between the embeddings of each word v in the set V of distinct definitions for a given emoji, and the most frequent word v^* in V, where f_v and e_v are v’s frequency and embedding vector. Instead of GloVe’s English-only word representation vectors (Pennington et al., 2014) used in Częstochowska et al. (2022), we employ cross-lingual embeddings generated with XLM-T (Barbieri et al., 2022b)—an instance of XLM-R (Conneau et al., 2020)—as it was further pre-trained on Twitter data. In addition to semantic variation scores computed with XLM-T, we report results with LASER (Artetxe and Schwenk, 2019) embeddings in Appendix D.

Experiment 2

The data from experiment 2 were analysed using two logistic mixed-effects regression models in R (R Core Team, 2022, version 4.1.3 (2022-03-10), "One Push-Up"). Model 1 and Model 2 were used to address RQs 2 and 3, respectively. The models were specified using the ‘afex’ package (Singmann and Kellen, 2019) as it directly computes the p-values for the fixed effects model terms rather than the estimates for the parameters which offer an easier interpretation. Following recommendations from Barr et al. (2013), maximal models including full random effects structures were specified as justified by the design. Model 1 comprised `emoji use` as the binary response variable, and `emoji` and `language` as the main predictor variables along with an interaction term (`emoji * language`). Model 2 was specified using `sentiment` as the binary response variable, and `emoji use` and `language` as the main predictor variables along with their interaction term (`emoji use * language`). Given that not all participants reported using Twitter, both maximal models included `Twitter use` as a binary covariate. The maximal models did not converge and the model was simplified by step-by-step elimination of random effects structures until convergence was reached. This was done following Barr et al. (2013). The final models in R syntax were specified as follows:

Model 1: `emoji use ∼ emoji * language + age + (1 | participant)`
Model 2: `sentiment ∼ emoji use * language + emoji + twitter use + age + (1 | participant)`

Post-hoc tests were carried out by computing estimated marginal means and then performing pairwise analyses using the ‘emmeans’ package (Lenth, 2022). Model checks and criticisms were performed using the ‘performance’ package (Lüdecke et al., 2021). Residuals were normal and no collinearity was detected (VIF < 5) in both models.

4 Results

RQ1: Do people disagree on emoji’s contextless interpretation within and across languages?

Table 2 shows the literal meanings obtained from the one-word definitions collected in Experiment 1. Unsurprisingly, most of these meanings are consis-

Emoji	Literal Meaning		
	En	Pt	Zh
🔥	Fire	Fogo	火热
😅	Nervous	Vergonha	尴尬
😂	Laughing	Rir	笑哭
🙏	Pray	Rezar	祈祷
🎉	Party	Festa	庆祝
❤️	Love	Amor	爱心
😭	Crying	Chorar	哭泣
😊	Happy	Corado	开心
🥰	Love	Apaixonado	爱你
👍	Good	Fixe	赞

Table 2: Collected literal meanings in English (En), Portuguese (Pt) and Chinese (Zh) for the analysed emojis.

English		Portuguese		Chinese	
E	SV	E	SV	E	SV
❤️	0.0178	🎉	0.0094	🎉	0.0503
🔥	0.0370	❤️	0.0193	😊	0.0595
🥰	0.0467	🔥	0.0432	❤️	0.0624
🎉	0.0511	😭	0.0548	🙏	0.0727
😭	0.0611	👍	0.0587	😭	0.0781
👍	0.0617	🙏	0.0772	👍	0.0809
😊	0.0655	😂	0.0803	🔥	0.0895
😂	0.0667	🥰	0.0834	😂	0.0949
🙏	0.0916	😊	0.0961	😅	0.1044
😅	0.0965	😅	0.1723	🥰	0.1059

Table 3: Emojis (E) sorted by semantic variation (SV) based on definitions provided in English, Portuguese and Chinese.

	Corr.	P-value
En ↔ Pt	0.6848	0.0289
En ↔ Zh	0.1636	0.6515
Pt ↔ Zh	0.5272	0.1173

Table 4: Spearman Rank Correlation and values between emojis' semantic variation in English (En), Portuguese (Pt), Chinese (Zh). English and Portuguese are significantly positively correlated. Chinese was found not significantly correlated to English and Portuguese.

tent across all three languages, demonstrating that the literal meaning of an emoji is tied to the image-grounded nature of emojis and is somewhat impervious to cultural differences. The literal meanings of the emojis 🔥, 🙏, ❤️, and 😭 can be considered semantically equivalent for all three languages, while 😂, 🎉, and 😊 for two of the languages.

The only emojis that are semantically inconsistent across languages are 😅 (En-nervous, Pt-shame, Zh-embarrassed), 👍 (En-good, Pt-cool, Zh-like), and 🥰 (En-love, Pt-in love, Zh-love you), an inconsistency that can be attributed to the ambiguity and difficulty in defining face emojis and hand gestures (Częstochowska et al., 2022). This is confirmed by our results in Table 3, which shows the semantic variation (or ambiguity) scores for the emojis across the three languages computed on the definitions collected in Experiment 1. As one can see, 😅 was considered the most ambiguous emoji to interpret and to define for English and Portuguese participants, and second most ambiguous for Chinese participants, while 🥰 was the third and most ambiguous emoji for Portuguese and Chinese participants respectively.

Comparing the emojis' ranking based on semantic variation scores between English and Portuguese, we can see that in both languages, the emojis representing physical entities such as ❤️, 🔥, and 🎉 were deemed the least ambiguous, followed by hand gestures and face emojis. This trend is not reflected in the Chinese ranking where the emojis are equally distributed across the rank. This can be attributed to the overall higher level of Chinese semantic variations for all the emojis compared to English and Portuguese. Correlations between the rankings (Table 4) confirm that English and Portuguese participants agree to some extent on emojis' ambiguity, while no significant correlation was found between Chinese and English/Portuguese. By manually analysing the one-word definitions collected, it is notable that the high level of Chinese emoji semantic variation is caused by its less strict rules for word boundaries compared to English or Portuguese. For example, ❤️'s literal meaning 爱你 can be accepted as a single word in Chinese, while its translation "love you" would be not accepted as a single word in English.

Overall, our results show that, although disagreement on emojis' interpretation varies from emoji to emoji similar to the results obtained by Częstochowska et al. (2022), the extent to which people disagree on such interpretations seemingly depends on the linguistic features of the language in question. However, as emojis are bound to their visual icon, their literal meanings are mostly shared across languages.

Effect	df	χ^2	P-value
Language	2.00	39.08 ***	<.001
Emoji	9.00	191.49 ***	<.001
Age	1.00	0.38	.539
Language:Emoji	18.00	62.10 ***	<.001

Significance: '***' p < 0.001; '**' p < 0.01; '*' p < 0.05

Table 5: Model 3 Results for RQ2. Significant effects for Emoji, but not for language, and marginal effects for interaction between the two.

Language	Odds Ratio	SE	Z-ratio	P-value
Chinese / English	0.60	0.07	-4.518	<.0001
Chinese / Portuguese	0.82	0.09	-1.824	0.1616
English / Portuguese	1.36	0.16	2.579	0.0268

Significance: '***' p < 0.001; '**' p < 0.01; '*' p < 0.05
Tests are performed on the log odds ratio scale

Table 6: Pairwise comparisons of Estimated Marginal Means of Emoji Use by Language for RQ2.

RQ2: Does agreement on the figurative or literal use of an emoji differ across different languages?

The results of the logistic regression carried out to answer RQ2 are presented in Table 5. In terms of the main predictor variables, we found a significant effect for emoji [$\chi^2(9) = 191.49, p < 0.001$], as well as for language [$\chi^2(2) = 39.08 p < 0.001$], and a significant effect was found for the interaction between the two [$\chi^2(18) = 62.10, p < 0.001$]. Pairwise comparisons by language were performed and results in Table 6 show that only Chinese versus English emoji use is significantly different (see Table 6). These results suggest that emojis can vary in their literal and figurative use across languages, but not necessarily so. This result is perhaps unsurprising given that English and Portuguese are genetically related languages and that the majority of English and Portuguese speakers use the same social media platforms and Portuguese speakers will often view content written in English. These results also corroborate our findings in experiment 1.

Overall, the results of this model are in keeping with the results from experiment 1.

RQ3: Does the figurative use of emojis correlate with the sentiment of the context in which the emoji is used?

The results of the logistic regression carried out to answer RQ3 are presented in Table 7. We can observe a statistically significant effect with respect to emoji use [$\chi^2(1) = 136.07, p < 0.001$] and language [$\chi^2(2) = 13.66, p = 0.001$]. This suggests that the choice of employing emojis, whether in a literal or figurative manner, is closely intertwined with the sentiment conveyed. However, in contrast, the analysis did not reveal any significant effect for age [$\chi^2(1) = 0.78, p = 0.377$], nor did it reveal any interaction effect between the use of emojis and language [$\chi^2(2) = 3.20, p = 0.202$]. Furthermore, a significant difference was found for emoji [$\chi^2(9) = 114.31, p < 0.001$], reinforcing the results obtained by addressing RQ2. Finally, Twitter use was not found to be statistically significant, indicating that there was no difference in emoji interpretation between people who used Twitter and those who did not. This should help to mitigate any concerns relating to whether emojis were used differently on Twitter compared to other social media sites.

Effect	df	χ^2	P-value
Emoji Use	1.00	136.07 ***	<.001
Language	2.00	13.66 **	.001
Emoji	9.00	244.26 ***	<.001
Twitter Use	1.00	0.01	.903
Age	1.00	0.78	.377
Use:Language	2.00	3.20	.202

Significance: '***' p < 0.001; '**' p < 0.01; '*' p < 0.05

Table 7: Model 2 Results for RQ3. Significant effects were found for Emoji Use and Emojis, but not for Language.

Figure 1 shows the overall statistics of the collected data in Experiment 2. We can see that several emojis such as 😊 , 🥰 and 👍 , were much more likely to be used literally in a positive context rather than a negative one but more likely to be used figuratively in a negative context rather than a positive one, in all languages. This and the reverse pattern seems to hold for many of the other emojis (e.g., 😅 and 😭) as well, indicating that sentiment does play a role in helping speakers to identify the usage of the emoji and reduce any potential ambiguity between the multiple meanings that it may have.

5 Conclusion

This study aimed to explore the role of semantic variation and sentiment in emoji use across three languages: English, European Portuguese, and Mandarin Chinese. We conducted two separate experiments, encompassing three research questions. The first experiment involved soliciting lit-

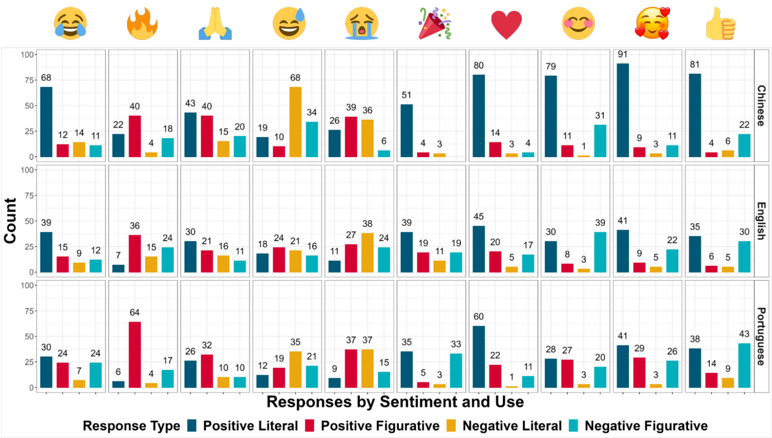

Figure 1: Counts of the annotations collected in Experiment 2, grouped by emoji in Chinese, English and Portuguese. The image shows that for most emojis, when used figuratively, their sentiment changes (e.g., 😭 from negative to positive, 😊 from positive to negative), supporting RQ3.

eral meanings of 10 carefully selected emoji stimuli in all three languages and comparing them based on a semantic variation metric. The second experiment queried participants on their understanding of the use of these emojis in tweets based on the literal meanings procured from experiment 1. Participants provided binary judgements with regard to the use (literal/figurative) of the emoji and the sentiment of the tweet (positive/negative). The results obtained from our study demonstrated that emojis exhibit variations in terms of semantic interpretation among themselves, yet their interpretations remain relatively consistent across different languages. Notably, our findings in experiment 2 corroborated the outcomes derived from experiment 1. Our results indicated that language itself does serve as a significant predictor of emoji usage or the sentiment conveyed. However, the locus of this effect seems to be driven by linguistic distance. Overall, we believe these results, while limited, pave the way for promising research directions which we discuss in the following section.

6 Future work

In this work, we gathered annotations pertaining to the sentiment and semantics of utterances that incorporate emojis, encompassing both the English and Portuguese languages. While the analysis of sentiment and the prediction of figurative use extend beyond the immediate scope of this paper, we can leverage the collected data to address the following research questions:

RQ4: To what extent can we automate the detection of whether an emoji is used in a literal or figurative sense?

RQ5: Does incorporating information about the figurative use of an emoji enhance the performance of sentiment analysis tasks?

To tackle RQ4, we posit that leveraging the capabilities of large pre-trained models, such as XLM-T, will yield reasonably effective results in discerning the figurative use of emojis. With their vast knowledge base and sophisticated language understanding, these models hold promising potential in automating the detection of nuanced emoji usage. Moreover, our study substantiated a significant correlation between figurative use and sentiment, as revealed in RQ2. Building upon this finding, we hypothesize that augmenting sentiment analysis models with explicit information regarding the usage of emojis have the potential to enhance the performance of such tasks. This could have practical applications in a variety of tasks including market research and brand interaction analysis.

Work in this domain could also be benficial to linguistic and sociolinguistc theory. Empirically investigating how speakers create alternate meanings for emojis as well as their patterns of use could provide important theoretical insights into iconicity and the interface between semantics and the pragmatics.

Limitations

Due to resource constraints, the research was limited to 10 emojis, 3 languages for Experiment 1 and 2 for Experiment 2. Given the specific nature of each emoji's relationship with figurative and literal use in different sentiments, we are only able to make conclusions about the emojis analysed in this study, making the generalisation of our findings to other emojis and languages difficult. Similarly, it is also worth noting that all the social media data used in Experiment 2 was scraped from Twitter at a specific time point (Nov 2022 - Jan 2023). Therefore, given the aforementioned flexibility of emoji use, it is important to note that only a small sample of emoji activity and use may have been represented.

Additionally, as discussed in Section 3.1, results for certain emojis might be biased due to the sentiment ratio of their occurrences. For example, the ❤️ emoji may appear much more often in tweets with a positive sentiment than those with a negative sentiment. Since the tweets were randomly sampled, the distribution of an emoji's meaning might not be balanced in the collected data. Therefore, comparisons between certain sentiments may be challenging for some emojis and languages. Although measurements have been taken to mitigate this problem, it is not possible to solve this limitation due to the sentiment analysis models' unreliability.

Potential problems can also be found when assessing the legitimacy of L1 speakers. For example, we could only control the country of residence and language spoken by the participants. Despite asking for only L1 speakers, it is plausible that some participants may not have been. Similarly, Prolific does not distinguish between European and Brazilian Portuguese. Although all the speakers of Portuguese resided in Portugal, there may have been some that were Brazilian Portuguese speakers.

Ethical Considerations

Importance of Cross-Cultural Research

The past 20 years have seen a rapid increase in the number of behavioural researchers engaging in cross-cultural research. However, recent research has shown that a lack of sample diversity in the field is still a very large problem, with 94% of Psychological Science articles having participant samples drawn from Western countries, and 71% from English-speaking countries (Rad et al., 2018).

Examining a theory cross-culturally is highly important as many older findings that were originally discovered in WEIRD[5] populations have been shown not to replicate across non-WEIRD populations (Henrich et al., 2010). For example, Fehr and Gächter (2002) found that a sample of undergraduates at the University of Zurich performed better as a group when they introduced the possibility of punishment, as the group used this to punish those who were non-cooperative. However, when the task was used with non-Western groups, this performance increase was not shown, as the group would punish both those who were non-cooperative and those who were too cooperative (Gachter et al., 2008).

As we can see from this example, findings that have been taken from only one population have very limited explanatory power. Hence, if we want to demonstrate robust findings, we need to explore our theories on much more diverse groups. Furthermore, if such findings are used in practical applications, we need to ensure that we are not causing harm to nor discriminating against a particular group. For example, the racial bias that has been seen in the AI (Fosch-Villaronga and Poulsen, 2022) and medical (El-Galaly et al., 2023; Fatumo et al., 2022) industries. While this may initially, seem to be irrelevant for emoji research, their potential use in large language models means that it is important that this data is accurate across languages.

References

Mikel Artetxe and Holger Schwenk. 2019. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7:597–610.

Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.

Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022a. XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages

[5]WEIRD: Western, Educated, Industrialised, Rich and Democratic

258–266, Marseille, France. European Language Resources Association.

Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022b. Xlm-t: Multilingual language models in twitter for sentiment analysis and beyond. In *Proceedings of the Language Resources and Evaluation Conference*, pages 258–266, Marseille, France. European Language Resources Association.

Francesco Barbieri, German Kruszewski, Francesco Ronzano, and Horacio Saggion. 2016. How cosmopolitan are emojis? exploring emojis usage and meaning over different languages with distributional semantics. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 531–535.

Dale J. Barr, Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3):255–278.

Evelyn Braumann, Olga Preveden, Sharjeel Saleem, Ying Xu, and Sabine T Koeszegi. 2010. The effect of emoticons in synchronous and asynchronous e-negotiations. In *Proceedings of the 11th Group Decision & Negotiation Conference (GDN 2010)*, pages 113–115.

Keith Broni. 2022. Global emoji use reaches new heights. *Emojipedia*.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Justyna Częstochowska, Kristina Gligorić, Maxime Peyrard, Yann Mentha, Michał Bień, Andrea Grütter, Anita Auer, Aris Xanthos, and Robert West. 2022. On the context-free ambiguity of emoji. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 1388–1392.

Eli Dresner and Susan C Herring. 2010. Functions of the nonverbal in cmc: Emoticons and illocutionary force. *Communication theory*, 20(3):249–268.

Luis Duarte, Luís Macedo, and Hugo Gonçalo Oliveira. 2019. Exploring emojis for emotion recognition in portuguese text. In *Progress in Artificial Intelligence: 19th EPIA Conference on Artificial Intelligence, EPIA 2019, Vila Real, Portugal, September 3–6, 2019, Proceedings, Part II 19*, pages 719–730. Springer.

Tarec Christoffer El-Galaly, Verena I Gaidzik, Mihnea-Alexandru Gaman, Darko Antic, Jessica Okosun, Mhairi Copland, Veronika Sexl, Adele K Fielding, Robin Doeswijk, Helen Parker, et al. 2023. A lack of diversity, equity, and inclusion in clinical research has direct impact on patient care. *HemaSphere*, 7(3).

Segun Fatumo, Tinashe Chikowore, Ananyo Choudhury, Muhammad Ayub, Alicia R Martin, and Karoline Kuchenbaecker. 2022. A roadmap to increase diversity in genomic studies. *Nature medicine*, 28(2):243–250.

Ernst Fehr and Simon Gächter. 2002. Altruistic punishment in humans. *Nature*, 415(6868):137–140.

Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1615–1625, Copenhagen, Denmark. Association for Computational Linguistics.

Eduard Fosch-Villaronga and Adam Poulsen. 2022. Diversity and inclusion in artificial intelligence. *Law and Artificial Intelligence: Regulating AI and Applying AI in Legal Practice*, pages 109–134.

Simon Gachter, Elke Renner, and Martin Sefton. 2008. The long-run benefits of punishment. *Science*, 322(5907):1510–1510.

Johannes Gettinger and Sabine T Koeszegi. 2015. More than words: the effect of emoticons in electronic negotiations. In *International conference on group decision and negotiation*, pages 289–305. Springer.

Raymond W Gibbs Jr. 2002. A new look at literal meaning in understanding what is said and implicated. *Journal of pragmatics*, 34(4):457–486.

Raymond W Gibbs Jr, Darin L Buchalter, Jessica F Moise, and William T Farrar IV. 1993. Literal meaning and figurative language. *Discourse Processes*, 16(4):387–403.

Rachel Giora. 1997. Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics*, 8(3):183–206.

Rachel Giora. 2002. Literal vs. figurative language: Different or equal? *Journal of pragmatics*, 34(4):487–506.

Joseph Henrich, Steven J Heine, and Ara Norenzayan. 2010. The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3):61–83.

Susan C Herring. 2018. Emergent forms of computer-mediated communication and their global implications. *LinguaPax Review 2017.*

Alexander Hogenboom, Daniella Bal, Flavius Frasincar, Malissa Bal, Franciska de Jong, and Uzay Kaymak. 2013. Exploiting emoticons in sentiment analysis. In *Proceedings of the 28th annual ACM symposium on applied computing*, pages 703–710.

Joon Young Lee, Nahi Hong, Soomin Kim, Jonghwan Oh, and Joonhwan Lee. 2016. Smiley face: why we use emoticon stickers in mobile messaging. In

Proceedings of the 18th international conference on human-computer interaction with mobile devices and services adjunct, pages 760–766.

Russell V. Lenth. 2022. *emmeans: Estimated Marginal Means, aka Least-Squares Means*. R package version 1.8.0.

Xuan Lu, Wei Ai, Xuanzhe Liu, Qian Li, Ning Wang, Gang Huang, and Qiaozhu Mei. 2016. Learning from the ubiquitous language: an empirical analysis of emoji usage of smartphone users. In *Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing*, pages 770–780.

Daniel Lüdecke, Mattan Ben-Shachar, Indrajeet Patil, Philip Waggoner, and Dominique Makowski. 2021. performance: An R Package for Assessment, Comparison and Testing of Statistical Models. *Journal of Open Source Software*, 6(60):3139.

George Manias, Argyro Mavrogiorgou, Athanasios Kiourtis, Chrysostomos Symvoulidis, and Dimosthenis Kyriazis. 2023. Multilingual text categorization and sentiment analysis: a comparative analysis of the utilization of multilingual approaches for classifying twitter data. *Neural Computing and Applications*, pages 1–17.

Hannah Miller, Daniel Kluver, Jacob Thebault-Spieker, Loren Terveen, and Brent Hecht. 2017. Understanding emoji ambiguity in context: The role of text in emoji-related miscommunication. In *Eleventh international AAAI conference on web and social media*.

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, Doha, Qatar. Association for Computational Linguistics.

R Core Team. 2022. A Language and Environment for Statistical Computing.

Mostafa Salari Rad, Alison Jane Martingano, and Jeremy Ginges. 2018. Toward a psychology of homo sapiens: Making psychological science more representative of the human population. *Proceedings of the National Academy of Sciences*, 115(45):11401–11405.

Alexander Robertson, Farhana Ferdousi Liza, Dong Nguyen, Barbara McGillivray, and Scott A Hale. 2021. Semantic journeys: quantifying change in emoji meaning from 2012-2018. *arXiv preprint arXiv:2105.00846*.

Matthew Shardlow, Luciano Gerber, and Raheel Nawaz. 2022. One emoji, many meanings: A corpus for the prediction and disambiguation of emoji sense. *Expert Systems with Applications*, 198:116862.

Henrik Singmann and David Kellen. 2019. An introduction to mixed models for experimental psychology. *New Methods in Cognitive Psychology*, pages 4–31.

Jayashree Subramanian, Varun Sridharan, Kai Shu, and Huan Liu. 2019. Exploiting emojis for sarcasm detection. In *Social, Cultural, and Behavioral Modeling: 12th International Conference, SBP-BRiMS 2019, Washington, DC, USA, July 9–12, 2019, Proceedings 12*, pages 70–80. Springer.

Toshiki Tomihira, Atsushi Otsuka, Akihiro Yamashita, and Tetsuji Satoh. 2020. Multilingual emoji prediction using bert for sentiment analysis. *International Journal of Web Information Systems*, 16(3):265–280.

Junjie Wang, Yuxiang Zhang, Lin Zhang, Ping Yang, Xinyu Gao, Ziwei Wu, Xiaoqun Dong, Junqing He, Jianheng Zhuo, Qi Yang, et al. 2022. Fengshenbang 1.0: Being the foundation of chinese cognitive intelligence. *CoRR*, abs/2209.02970.

Emoji	Unicode Name	Score
🔥	Fire	0.0325
🥰	Smiling face with hearts	0.1063
❤️	Red heart	0.1224
😭	Loudly crying face	0.1684
🙏	Folded hands	0.2359
😂	Face with tears of joy	0.2636
🎉	Party popper	0.2407
😅	Grinning face with sweat	0.3412
😊	Smiling face with smiling eyes	0.4583
👍	Thumbs up	0.6593

Table 8: Emojis selected for this study with their official Unicode name and semantic variation scores as reported by Częstochowska et al. (2022).

A More on emoji selection

The selection of emojis for our study was a thoughtful process driven by a combination of resource constraints, practical considerations, and a commitment to capturing a representative subset of commonly used emojis. Due to limitations in resources and the desire to manage participant annotation loads effectively, we opted for a smaller number of emojis. To ensure widespread familiarity, we rigorously chose the final set of 10 emojis based on their frequent usage. Recognising the prevalence of face emojis in the top 20 most popular emojis (😂 ❤️ 🤣 👍 😭 🙏 😘 🥰 😍 😊 🎉 😁 💕 🥺 😅 🔥 🤦 🖤 🤷 ☺️), we aimed for a balanced representation of face and non-face emojis to reflect the broader spectrum of emoji usage, as well as to counter their limited graphical variation (e.g. 😂 / 🤣 - ❤️ / 💕 / 🖤 - 😊 / 😁 / ☺️). While acknowledging the possibility of introducing some bias through this selection process, we believe it was essential to strike a fair balance and yield meaningful results in our study.

B Participant Data

Language	n	mean Age	Range	SD
Chinese	30	33.43	23-58	8.90
English	30	33.31	20-59	10.40
Portuguese	30	23.32	19-47	5.20

Table 9: Participant Age Distribution by Language for Experiment 1

Language	n	mean Age	Range	SD
Chinese	44	31.48	20-50	8.26
English	35	37.80	21-57	12.05
Portuguese	37	27.05	20-51	7.28

Table 10: Participant Age Distribution by Language for Experiment 2

C Literal meaning translation

Table 11 shows the English translations for the literal meaning of the emojis in Portuguese and Chinese.

Emoji	Literal Meaning		
	En	Pt	Zh
🔥	Fire	Fire	Fiery
😅	Nervous	Shame	Awkward
😂	Laughing	Laughing	Cry laughing
🙏	Pray	Pray	Pray
🎉	Party	Party	Celebrate
❤️	Love	Love	Love
😭	Crying	Crying	Crying
😊	Happy	Blushing	Happy
🥰	Love	Passionate	Love you
👍	Good	Cool	Thumbs up

Table 11: English translations for the emojis' literal meanings.

D Additional Experiment Results

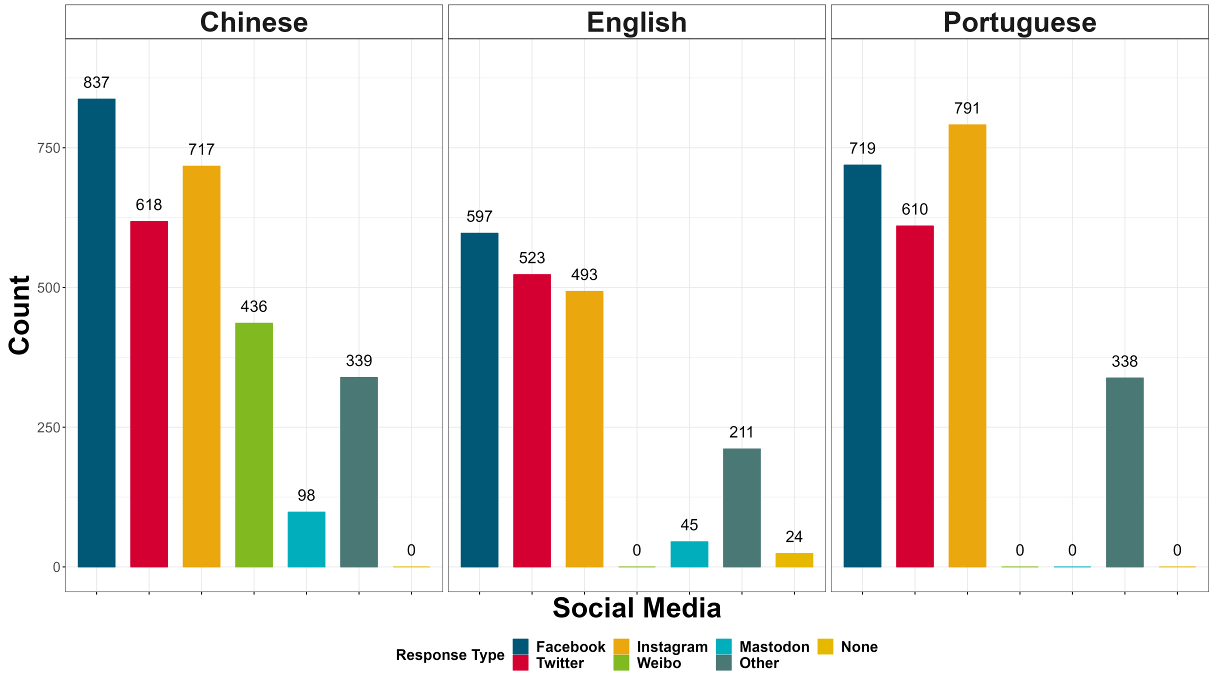

Figure 2: Experiment 2's participant responses to which social media sites they use

English		Portuguese		Chinese	
E	SV	E	SV	E	SV
❤️	0.0440	🎉	0.0171	🎉	0.0904
🎉	0.0919	❤️	0.0418	❤️	0.0931
🥰	0.1085	🔥	0.0987	🔥	0.1307
👍	0.1194	👍	0.1242	😅	0.1764
🔥	0.1253	😭	0.1248	😊	0.1764
😂	0.126	🥰	0.1321	😭	0.1868
🙏	0.1443	🙏	0.1555	🙏	0.1890
😭	0.1605	😂	0.1706	👍	0.1983
😊	0.1843	😊	0.1863	🥰	0.2502
😅	0.1929	😅	0.2474	😂	0.2700

Table 12: Emojis (E) sorted by semantic variation (SV) computed with LASER embeddings, based on definitions provided in English, Portuguese and Chinese. Compared to the ranking computed with XLM-T (Table 3), physical entities were ranked least ambiguous for all three languages.

Original		XLM-T		LASER	
E	SV	E	SV	E	SV
🔥	0.0325	🔥	0.0049	🔥	0.0209
🥰	0.1063	🥰	0.0242	🥰	0.0645
❤️	0.1224	❤️	0.0302	❤️	0.0713
😭	0.1684	🎉	0.0389	🎉	0.0892
🙏	0.2359	😭	0.0408	🙏	0.0946
🎉	0.2407	🙏	0.0582	😭	0.1033
😂	0.2636	😂	0.0689	😂	0.1624
😅	0.3412	😊	0.0764	😅	0.1651
😊	0.4583	😅	0.0796	😊	0.2129
👍	0.6593	👍	0.1094	👍	0.2434

Table 13: Emojis (E) sorted by semantic variation (SV) based on definitions provided by Częstochowska et al. (2022). Reported are the original semantic variation scores, as well as the ones computed with XLM-T and LASER embeddings. Using different encoding methods does not change significantly the emoji ranking.

	Corr.	P-value
En ↔ Pt	0.8303	0.0029
En ↔ Zh	0.3212	0.3655
Pt ↔ Zh	0.5151	0.1276

Table 14: Spearman Rank Correlation and values between emojis' semantic variation (with LASER embeddings) in English (En), Portuguese (Pt), Chinese (Zh). The correlation between English and Portuguese is stronger compared to the ones in Table 4, while the correlation remained not significant.

E Trial Samples

Here we present the screenshot of the trials' webpage shown to the participants in Experiments 1 and 2 (Figure 3).

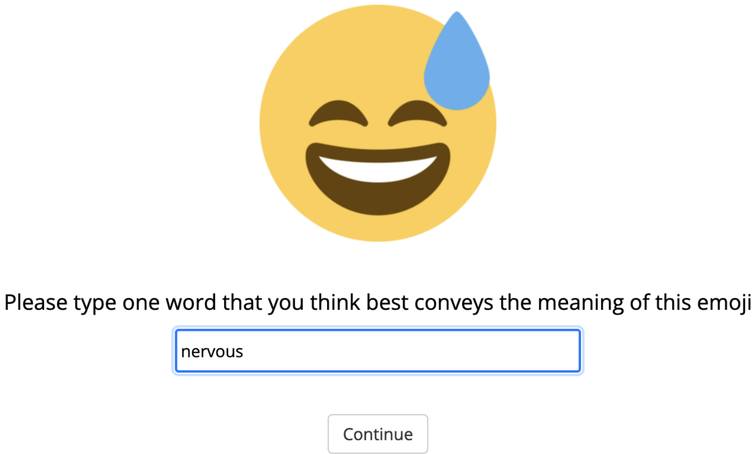

(a) Experiment 1 - One-word Definition

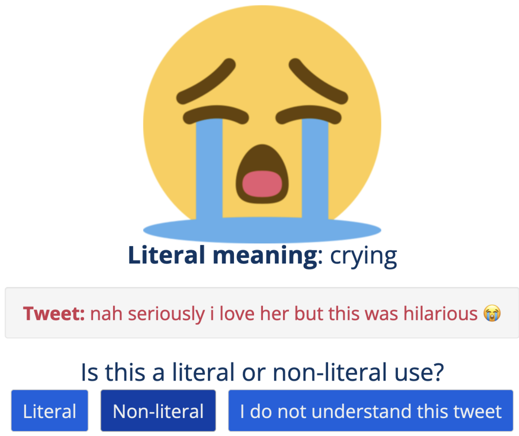

(b) Experiment 2 - Semantics

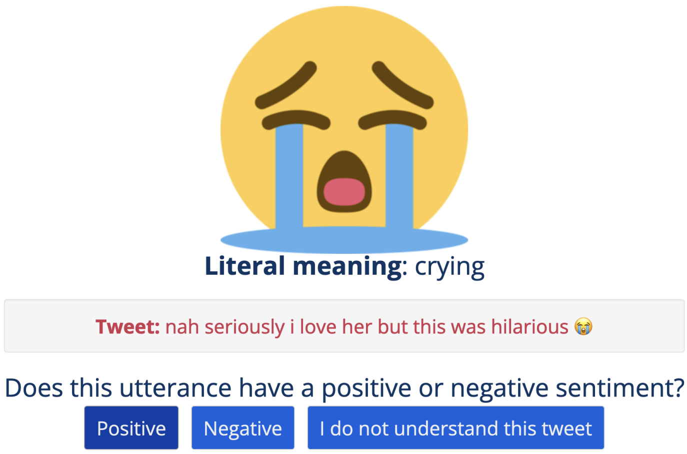

(c) Experiment 2 - Sentiment

Figure 3: Example of trials' main page for online experiments.