

On the Context-Free Ambiguity of Emoji

Justyna Czȩstochowska,*¹ Kristina Gligorić,*¹ Maxime Peyrard,¹ Yann Mentha,¹ Michał Bień,¹
 Andrea Grütter,² Anita Auer,³ Aris Xanthos,³ Robert West¹

¹EPFL, ²University of Zurich, ³University of Lausanne
 {justyna.czȩstochowska, kristina.gligoric, maxime.peyrard, yann.mentha, michal.bien, robert.west}@epfl.ch,
 andrea.gruetter@es.uzh.ch, {anita.auer, aris.xanthos}@unil.ch

Abstract

Due to their pictographic nature, emojis come with baked-in, grounded semantics. Although this makes emojis promising candidates for new forms of more accessible communication, it is still unknown to what degree humans agree on the inherent meaning of emojis when encountering them outside of concrete textual contexts. To bridge this gap, we collected a crowdsourced dataset (made publicly available) of one-word descriptions for 1,289 emojis presented to participants with no surrounding text. The emojis and their interpretations were then examined for ambiguity. We find that, with 30 annotations per emoji, 16 emojis (1.2%) are completely unambiguous, whereas 55 emojis (4.3%) are so ambiguous that the variation in their descriptions is as high as that in randomly chosen descriptions. Most emojis lie between these two extremes. Furthermore, investigating the ambiguity of different types of emojis, we find that emojis representing symbols from established, yet not cross-culturally familiar code books (e.g., zodiac signs, Chinese characters) are most ambiguous. We conclude by discussing design implications.

1 Introduction

For over a decade, emojis have been playing an increasingly important role in online communication. As of September 2021, there are 3,633 emojis in the Unicode standard (Unicode 2021), and the number is growing, providing users with more ways to express increasingly complicated concepts. Consequently, emojis have received much attention from researchers. Various fields, including natural language processing, human-computer interaction, and Web and social media research, study the usage and function of emojis.

Beyond today’s prevalent use cases (social media and instant messaging), emojis have untapped potential for facilitating communication in other contexts as well. While letters, syllables, and words are arbitrary and highly abstract constructs that require a long time to master, emojis come already packed with richly grounded semantics. Emojis can thus be leveraged, e.g., in learning and education or to describe complex ideas to broad audiences (Gilles Doiron 2018).

However, it is unknown which emojis can be used for such goals. As a first step, it is necessary to establish how

much people agree about the context-free interpretation of individual emojis. Doing so has broad implications for social media and Web research, communication studies, education, and more. Beyond research communities, identifying which emojis are ambiguous is helpful to online communities and emoji designers to prevent introducing emojis with a high potential for miscommunication. Additionally, studying context-free emoji semantics informs us which concepts can or cannot be easily communicated with emojis.

Despite the practical importance of these questions, it is difficult to approach them with available datasets. Social media content carries inherent selection biases, and emoji studies leveraging social media have been questioned for their generalizability (Herring and Dainas 2020). Additionally, as emojis are almost only used in context, it is difficult to infer context-free interpretations. To complicate matters further, the meaning of emojis on social media evolves with time (Robertson et al. 2021), making it difficult to study their intrinsic semantics.

Here we ask: *Do individuals interpret emojis similarly? Which emojis have the potential to be used in future communication scenarios, and to what extent?* Previous work on the ambiguity of emojis focused on differences between platforms (Shurick and Daniel 2020) and their usage in context (Miller et al. 2017). Most closely related past studies focused on frequently used and anthropomorphic emojis (Shurick and Daniel 2020; Miller et al. 2017, 2016), discarding many available emojis. Whereas a lot is known about emoji sentiment and usage in context (Novak et al. 2015), less is known about emoji semantics beyond sentiment and context-free emoji interpretation. To bridge this gap, we designed and executed a crowdsourced study examining an exhaustive set of emojis, many of which are rarely used in online communication. We studied their interpretation in the absence of any textual context. Using the resulting novel dataset of emoji annotations, we address the following research questions:

RQ1: To what degree do people agree when interpreting emojis?

RQ2: What types of emojis are most and least ambiguous?

2 Related Work

Emojis: interpretation and meaning. Previous research has shown that emojis are often misunderstood (Miller et al.

*These authors contributed equally.
 Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2016, 2017). Misunderstanding is sometimes related to how the emoji’s design is interpreted in context or the way it is shown on the receiving side. In particular, in 2016, Miller et al. examined interpretations of the 25 most popular anthropomorphic emojis without context, across five popular platforms. The study compared differences in sentiment and semantics to identify the most ambiguous emojis. In 2017, Miller et al. conducted a similar study comparing sentiment variability with and without context, for 10 anthropomorphic emojis. An extensive dataset of emoji senses was created by Wijeratne et al., linking Unicode emoji representations to their meanings automatically extracted from the Web. Recent studies of emoji meaning provide a longitudinal perspective (Barbieri et al. 2018b; Robertson et al. 2021). Our work studies the intrinsic ability of emojis to convey information, independent of the textual context they are used in. In contrast to Miller et al. (2016, 2017), who focused on small subsets of anthropomorphic emojis, we consider a far more exhaustive set of emojis (see Table 1).

The aspiration of an emoji-based language. There is growing interest in the linguistic purposes of emojis (Na’aman, Provenza, and Montoya 2017) and their potential to emerge as a graphical language (Ge and Herring 2018). There have been multiple informal initiatives to create an emoji language, such as the attempt to translate *Moby Dick* into a sequence of emojis.¹ Such efforts demonstrate the potential for viewing emojis as the atomic units of graphical and intuitive language that could remove accessibility barriers inherent in standard written natural languages.

Emojis, social media, and natural language processing. Social media researchers have been examining the ways social media users use, interpret, and express emotions and information through emojis. It is well known that emojis shape online language (Feldman et al. 2021; Pavalanathan and Eisenstein 2016). Emoji usage can also be a proxy for studying human behavior; e.g., emojis are a powerful indicator in the context of crisis events (Santhanam et al. 2018) and can be used to identify group belonging (Jones, Nurse, and Li 2021). Researchers have also been analyzing the use of gender and skin-tone modifiers (Barbieri and Camacho-Collados 2018; Robertson, Magdy, and Goldwater 2020, 2021). As emojis became a standard element of online language, a need to computationally process them emerged. Creating meaningful, latent emoji representations (Eisner et al. 2016) and emoji prediction tasks (Barbieri et al. 2018a) gained importance in NLP. We thus note that our annotations can be used to compute or augment emoji representations and thus support the social media and NLP research communities.

3 Methods and Data

Emoji selection. We selected emojis for our study as follows. Starting from all available 3,633 emojis, we removed letters, numbers, flags, and gendered and skin-toned anthropomorphic emojis, considering only neutral versions of such emojis. We also removed variations of the same emoji (e.g.,

¹<https://www.kickstarter.com/projects/fred/emoji-dick>

family with three or four members). This resulted in the final set of 1,289 emojis. Furthermore, we collected emoji categories from Emojipedia.org and hand-crafted a categorization extending the seven existing categories to 20 fine-grained types, outlined in Table 1.

Annotation process. We collected emoji interpretations using the Amazon Mechanical Turk (AMT) crowdsourcing platform. Each participant was asked to “Describe emojis with a single, accurate word”. Each task consisted of 10 emojis. All participants had to be at least 18 years old, speak English, reside in the USA, have a 99% approval rate, and have completed at least 500 tasks on AMT before. Our annotator compensation was in line with ethical guidelines for AMT (Whiting, Hugh, and Bernstein 2019). Each emoji was annotated by 30 unique participants. This number was chosen via pilot studies (150 annotations for 12 emojis), showing that, as the number of annotations increases beyond 30, the word distribution remained robust. Overall, we collected 38,670 annotations, for an average of 82.5 annotations per participants. In total, there were 445 unique participants. We asked participants to provide their age, gender, and mother tongue. The majority of annotators were native English speakers (97%). Participants’ gender was well balanced (55% female, 44% male, 1% other or not stated). The average age was 38.8 (SD = 12.0).

Post-processing. To improve the quality of annotations, we performed three post-processing steps: low quality annotator detection, validation of honeypots, and spelling correction. We performed detection of annotators with low annotation quality by identifying those who used the same word for all emojis in a task. We discarded one annotator whose vocabulary size was less than 80% of the number of assigned emojis. To further ensure the quality, one unquestionably non-ambiguous emoji was placed in every task. Annotations whose answers did not match any words from the expected set (e.g., “apple” for 🍏, “pizza” for 🍕, “carrot” for 🥕) were excluded. Finally, to account for spelling mistakes, we cross-checked word validity using the PyEnchant library.

Measuring semantic variation. We use word embeddings (i.e., representations of words as numerical vectors) to quantify semantic similarity. To measure the extent to which annotators agree about emoji meaning, we calculate the dispersion of emoji annotations in a similar way to Miller et al., using GloVe vectors (Pennington, Socher, and Manning 2014) of dimensionality 200 (Řehůřek and Sojka 2011). Let V denote the set of distinct words used by respondents to annotate the considered emoji, which we will call the emoji’s *vocabulary*; f_v stands for the relative frequency of word v in the emoji’s annotations and $v^* := \arg \max_{v \in V} f_v$ is the mode annotation, i.e., the most frequent word in V . We then define the emoji’s semantic variation as the weighted average of the cosine distance between the embedding e_v of each word $v \in V$ and the embedding e_{v^*} of the mode annotation in V :

$$\text{semantic variation} = \sum_{v \in V} f_v \cdot (1 - \cos(e_v, e_{v^*})) \quad (1)$$

Category	Number	Examples
objects	202	🍷 🍷 🍷
nature	189	🌱 🌱 🌱
travel-places	129	🏰 🏰 🏰
food-drink	113	🍌 🍌 🍌
faces	111	😄 😄 😄
people	103	👨 👨 👨
activity	78	🏊 🏊 🏊
clothes & accessories	47	👉 👉 👉
symbols & signs	43	♣️ ♣️ ♣️
professions	38	👩 👩 👩
geometrical	34	🟤 🟤 🟤
hands & gestures	34	👉 👉 👉
Japanese symbols & objects	34	🍷 🍷 🍷
buttons & mobile	29	📱 📱 📱
public information	26	📢 📢 📢
letters & numbers	21	📄 📄 📄
hearts	19	❤️ ❤️ ❤️
arrows	16	➡️ ➡️ ➡️
astrological	13	♈️ ♈️ ♈️
religious	10	🕉️ 🕉️ 🕉️

Table 1: Emoji categorization. Twenty categories, number of emojis, and three examples.

4 Results

RQ1: To what degree do people agree when interpreting emojis? For each emoji, we measure the consistency among the words chosen to describe it. Consider the example of the fire emoji 🗑️, for which one annotator used the word “hot” and another, “fire”. Since the terms are different, the annotators—strictly speaking—do not agree. Yet, the words “hot” and “fire” are semantically close. We aim to capture such similarities via the notion of semantic variation (Eq. 1).

To detect semantic variations significantly different from random, we compute the semantic variation of a random baseline. We sample $n = 30$ random words from the distribution of vocabulary across all annotations and calculate the semantic variation. We repeat the process 1,000 times to compute 95% confidence intervals (CI) and obtain a baseline semantic variation of 0.69 (95% CI [0.57, 0.88]). When an emoji’s semantic variation falls in the 95% CI, its vocabulary could have equally well been a random set of words. In such cases, humans clearly do not agree in their interpretation. We repeat the same procedure to compute a random baseline with respect to vocabulary size (rather than variation), obtaining an average of 30 (i.e., all words are different), with 95% CI [27, 30]. All emojis, with their semantic variation and vocabulary size, are presented in Fig. 1.

First, we find a strong positive correlation between vocabulary size and semantic variation (Spearman’s $\rho = 0.84$, $p < 10^{-10}$). We note that emojis with a small vocabulary size can range from very ambiguous to not ambiguous at all, while the ones with a rich vocabulary have higher ambiguity.

Second, we find that 16 out of 1,289 emojis (1.2%) reach a variation of 0, i.e., they were described with a single, unique word by all 30 annotators. These include 🌈 🐶 🍷 📱 🕉️ 🕉️ 🕉️ 🕉️ 🕉️ 🕉️. For communication applications, these




Figure 1: Top: the relationship between semantic variation (x-axis), and vocabulary size (y-axis). Bottom: average semantic variation across emoji categories (cf. Table 1). Color represents the extent to which emojis within a category can be seen as belonging to an established code book of symbols. Black dashed lines represent random baselines and gray bands their bootstrapped 95% confidence intervals.

emojis are likely to be useful—they are unlikely to introduce misunderstanding.

Third, the semantic variation of 55 out of 1,289 emojis (4.3%) falls into the random baseline confidence interval. Given the intuition that emojis come with built-in semantics, it is striking that some of them exhibit such low agreement. For future communication applications, these emojis are unlikely to be useful, as they introduce high levels of ambiguity.

In summary, human agreement about the context-free meaning of emojis ranges from completely unambiguous (16 emojis) to indistinguishable from random (55 emojis), with emojis covering the whole spectrum of ambiguity. Our dataset can guide communication applications in choosing appropriate emojis to facilitate understanding.

RQ2: What types of emojis are most and least ambiguous? We further investigate discrepancies in ambiguity, expecting different categories of emojis to exhibit different average variations. We report the average semantic variation per category in Fig. 1.

Our results add nuance to the findings of Miller et al. (2016), who found that anthropomorphic emojis can be more ambiguous than emojis characterizing things. In Fig. 1, the *faces* category takes a middle place in the semantic-variation ranking. Still, it is more ambiguous than objects. Categories with the lowest average variation are *food & drink*, *clothes & accessories*, *nature*, and *hearts*.

Interestingly, the top five most ambiguous categories are the ones that emerged from further dividing the original *symbols* category. Every emoji is, of course, a symbol, but

5 Discussion

whereas some emojis derive their meaning by immediately representing the shape of commonly known objects (e.g., 🍎), those in the *symbols* category refer to entries from established code books of symbols of a less immediately pictographic nature (e.g., 🍷🍷). Such emojis are, in a sense, “symbols of symbols” or “second-order symbols”, and, without prior knowledge, may be impossible to interpret. In particular, astrological (zodiac) signs form the only category as ambiguous as the random baseline. They tended to be described with very different words or with names of other astrological signs. One could argue that astrological signs have an unambiguous mapping to their names, but without background knowledge, they yield ambiguous standalone interpretations. Similarly, emojis representing Japanese signs or having origins in Japanese culture (e.g., 🍷 🍷) occupy the third place in terms of semantic variation, likely due to annotators’ demographics and cultural background (United States residents, native English speakers). To describe emojis of Japanese, Chinese, and Korean characters (e.g., 🍷 🍷 🍷), annotators consistently used words such as: *japanese*, *chinese*, *asian*, *sign*. This was not the case for some emojis (e.g., 🍷 🍷) of whose Japanese origin annotators may not have been aware.

Based on these observations, we next investigate such “symbol-of-symbol” emojis in more detail. To quantify the degree to which an emoji belongs to an established code book of symbols (henceforth “symbolicalness”), two authors independently annotated all 1,289 emojis by indicating their level of agreement with the statement “This emoji is a symbol” on a five-point Likert scale where 1 corresponded to “absolutely disagree” and 5 to “absolutely agree”.

We pre-established an annotation framework where we assigned levels from 5 to 1, respectively, to emojis representing objects and concepts that (5) are established symbols and can be encountered in everyday or specialized activities (e.g., 🍷 🍷); (4) can have a symbolic meaning and be encountered in everyday or specialized activities (e.g., 🍷 🍷); (3) may or may not have a symbolic meaning (e.g., 🍷 🍷); (2) typically do not have a symbolic meaning (e.g., 🍷 🍷); (1) are not established symbols (faces, gestures, people) (e.g., 🍷 🍷).

Following this framework, we obtained Kendall’s $\tau = 0.8$ ($p = 1.55 \times 10^{-216}$) between the authors. We computed the average symbolicalness for each emoji and averaged the values for emojis within a category to obtain the category’s symbolicalness rating. We represent the rating with a color scale in Fig. 1. There is a weak positive correlation (Spearman’s $\rho = 0.25$, $p = 1.61 \times 10^{-19}$) between semantic variation and symbolicalness. In addition, the most ambiguous categories of emojis indeed are the ones with the highest symbolicalness rating. Even though symbols are designed to facilitate communication, our results indicate that symbolic emojis can, maybe unintuitively, be ambiguous, as their interpretation requires specific prior knowledge.

In summary, we find that human agreement about the interpretation of emojis varies across different emoji types. The symbolicalness, or the extent to which an emoji is a symbol from an established code book of symbols, is an important dimension explaining the differences.

Summary of main findings. Investigating whether people interpret emojis in the same way (RQ1), we find that emojis come with very different amounts of prepacked semantics. Some emojis are completely unambiguous, with all annotators describing them with the same word. On the opposite, others are as ambiguous as if their descriptions were drawn at random. To support the goal of using emojis to facilitate communication, the unambiguous emojis are the best candidates to bring direct benefits. Investigating what types of emojis are ambiguous (RQ2), we find that different types of emojis have very different levels of agreement in interpretation. An important dimension explaining the agreement differences is the degree to which an emoji belongs to an established code book of symbols. Emojis referring to symbols require background knowledge for interpretation and are less likely to be unambiguously recognized.

Concrete objects and things can easily be illustrated by an emoji, whereas abstract ideas and concepts are harder to represent without referring to symbolic ideas from shared cultural knowledge. Yet to support complex communication goals, it is necessary to refer to abstract ideas. This has to be via pre-established symbols, as there is no immediately pictographic way to represent abstract concepts such as peace (🕊️) or resistance (👊). Luckily, there exist ubiquitous symbols whose interpretations are widely agreed upon (e.g., ❤️ for love). These symbols can be leveraged to convey complex ideas universally.

Design implications. We highlight two mechanisms likely fueling measured variation in ambiguity and discuss their implications. First, the fact that concrete objects can more easily be universally described with emojis highlights the importance of considering the intended participants and their shared cultural background to appropriately choose the symbolic emojis to use. Second, emoji design is known to contribute to misinterpretations (Miller et al. 2017), since emojis are limited in size and need to be comprehensible even if displayed tiny; e.g., the “pine decoration” emoji 🌲 contains a fair amount of details that are difficult to display on a small scale, further jeopardizing the understanding of the concept. Our data capturing empirical emoji ambiguity can thus help make emojis more accessible and user-friendly.

Limitations and future work. Our goal is to provide initial measurements of the ambiguity of emojis. Therefore, our study is not without its limitations. All annotators provided a single word to describe an emoji. In the future, it will be interesting to extend the study to descriptions beyond one word. Also, all annotators were English speakers residing in the United States. Future work should generalize to other cultures and languages and understand how emoji ambiguity is associated with social media usage.

Code and data. Code and data are publicly available at <https://github.com/epfl-dlab/emoji-ambiguity>.

Acknowledgments. This work was funded by Collaborative Research on Science and Society (CROSS), with further support from the Swiss National Science Foundation (grant 200021_185043), Microsoft, Google, and Facebook.

References

- Barbieri, F.; Anke, L. E.; Camacho-Collados, J.; Schockaert, S.; and Saggion, H. 2018a. Interpretable Emoji Prediction via Label-Wise Attention LSTMs. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, 4766–4771. Association for Computational Linguistics.
- Barbieri, F.; and Camacho-Collados, J. 2018. How Gender and Skin Tone Modifiers Affect Emoji Semantics in Twitter. In *Proceedings of the 7th Joint Conference on Lexical and Computational Semantics*, 101–106.
- Barbieri, F.; Marujo, L.; Karuturi, P.; and Brendel, W. 2018b. Exploring Emoji Usage and Prediction Through a Temporal Variation Lens. *Proceedings of the 1st International Workshop on Emoji Understanding and Applications in Social Media*, 2130.
- Eisner, B.; Rocktäschel, T.; Bošnjak, M.; and Riedel, S. 2016. emoji2vec: Learning Emoji Representations from their Description. In *Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media*, 48–54.
- Feldman, L. B.; Barach, E.; Srinivasan, V.; and Shaikh, S. 2021. Emojis and Words Work Together in the Service of Communication. *Workshop Proceedings of the 15th International AAAI Conference on Web and Social Media*.
- Ge, J.; and Herring, S. C. 2018. Communicative functions of emoji sequences on Sina Weibo. *First Monday*, 23(11).
- Gilles Doiron, J. A. 2018. Emojis: Visual Communication in Higher Education. *International Journal of Teaching, Education and Learning J. A. Gilles Doiron*, 2(2): 1–11.
- Herring, S. C.; and Dainas, A. R. 2020. Gender and Age Influences on Interpretation of Emoji Functions. *ACM Transactions on Social Computing*, 3(2): 1–26.
- Jones, K.; Nurse, J. R.; and Li, S. 2021. The Shadowy Lives of Emojis: An Analysis of a Hacktivist Collective’s Use of Emojis on Twitter. In *Workshop Proceedings of the 15th International AAAI Conference on Web and Social Media*.
- Miller, H.; Kluver, D.; Thebault-Spieker, J.; Terveen, L.; and Hecht, B. 2017. Understanding Emoji Ambiguity in Context: The Role of Text in Emoji-Related Miscommunication. *AAAI Conference on Web and Social Media (ICWSM)*.
- Miller, H.; Thebault-Spieker, J.; Chang, S.; Johnson, I.; Terveen, L.; and Hecht, B. 2016. "Blissfully happy" or "ready to fight": Varying Interpretations of Emoji. *AAAI Conference on Web and Social Media (ICWSM)*.
- Na’aman, N.; Provenza, H.; and Montoya, O. 2017. MojiSem: Varying linguistic purposes of emoji in (Twitter) context. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics- Student Research Workshop*, 136–141.
- Novak, P. K.; Smailović, J.; Sluban, B.; and Mozetič, I. 2015. Sentiment of Emojis. *PLoS ONE*, 10(12).
- Pavalanathan, U.; and Eisenstein, J. 2016. View of More emojis, less :) The competition for paralinguistic function in microblog writing. *First Monday*, 21(11).
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global Vectors for Word Representation. In *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, 1532–1543. Association for Computational Linguistics (ACL).
- Robertson, A.; Liza, F. F.; Nguyen, D.; McGillivray, B.; and A. Hale, S. 2021. Semantic Journeys: Quantifying Change in Emoji Meaning from 2012–2018. In *Workshop Proceedings of the 15th International AAAI Conference on Web and Social Media*.
- Robertson, A.; Magdy, W.; and Goldwater, S. 2020. Emoji Skin Tone Modifiers: Analyzing Variation in Usage on Social Media. *ACM Transactions on Social Computing*, 3(2): 1–25.
- Robertson, A.; Magdy, W.; and Goldwater, S. 2021. Identity Signals in Emoji do not Influence Perception of Factual Truth on Twitter. In *Workshop Proceedings of the 15th International AAAI Conference on Web and Social Media*.
- Santhanam, S.; Srinivasan, V.; Glass, S.; and Shaikh, S. 2018. I Stand With You: Using Emojis to Study Solidarity in Crisis Events. In Wijeratne, S.; Kiciman, E.; Saggion, H.; and Sheth, A., eds., *Proceedings of the 1st International Workshop on Emoji Understanding and Applications in Social Media*.
- Shurick, A. A.; and Daniel, J. 2020. What’s behind those smiling eyes: Examining emoji sentiment across vendors. *Workshop Proceedings of the 14th International AAAI Conference on Web and Social Media*.
- Unicode, I. 2021. Emoji Counts, v14.0. <https://unicode.org/emoji/charts-14.0/emoji-counts.html>. Accessed: 2022-04-05.
- Whiting, M. E.; Hugh, G.; and Bernstein, M. S. 2019. Fair Work: Crowd Work Minimum Wage with One Line of Code. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 7: 197–206.
- Wijeratne, S.; Balasuriya, L.; Sheth, A.; and Doran, D. 2017. EmojiNet: An Open Service and API for Emoji Sense Discovery. *Proceedings of the 11th International Conference on Web and Social Media, ICWSM 2017*, 437–446.
- Řehůřek, R.; and Sojka, P. 2011. Gensim–python framework for vector space modelling. *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, 3(2).