

Emojis Decoded: Leveraging ChatGPT for Enhanced Understanding in Social Media Communications

Yuhang Zhou,¹ Paiheng Xu,² Xiyao Wang,² Xuan Lu,³ Ge Gao,¹ Wei Ai¹

¹ College of Information, University of Maryland, College Park, USA

² Department of Computer Science, University of Maryland, College Park, USA

³ College of Information Science, University of Arizona, Tucson, USA

{tonyzhou, paiheng, xywang, gegao, aiwei}@umd.edu, luxuan@arizona.edu

Abstract

Emojis, which encapsulate semantics beyond words or phrases, have become prevalent in social network communications. This has spurred increasing scholarly interest in exploring their attributes and functionalities. However, emoji-related research and application face two primary challenges. First, researchers typically rely on crowd-sourcing to annotate emojis in order to understand their sentiments, usage intentions, and semantic meanings. Second, subjective interpretations by users can often lead to misunderstandings of emojis and cause a communication barrier. Large Language Models (LLMs) have achieved significant success in various annotation tasks, with ChatGPT demonstrating expertise across multiple domains. In our study, we assess ChatGPT's effectiveness in handling previously emoji-annotated and downstream tasks. Our objective is to validate the hypothesis that ChatGPT can serve as an alternative to human annotators in emoji research and that its ability to explain emoji meanings can enhance clarity and transparency in online communications. Our findings indicate that ChatGPT has extensive knowledge of emojis. It is adept at explaining the meaning of emojis across various application scenarios and demonstrates the potential to replace human annotators in a range of tasks.

1 Introduction

Emojis, as prevalent non-verbal tokens in social networks, are crucial in developing an effective discussion in online communication. These language units encode rich semantic meanings that range from emotions, objects, to actions. The crucial role of emojis has attracted a number of researchers to study various emoji-related topics, from emoji functionality, emoji usage patterns, to the emoji application on downstream tasks. However, there are two challenges for emoji research and emoji applications for social users. For emoji research, to determine the meaning, sentiments, or intentions of emojis in different scenarios, researchers usually hire human annotators to carry out annotation jobs (Lu et al. 2018; Hu et al. 2017), which is not scalable and time consuming. For social users, their differences in demographics, cultural background, or personal experience could lead to misunderstanding of emoji, such as 😏 (upside-down face) to express irony, and incorrect interpretation may cause miscommunication and block emoji diffusion (Leonardi 2022).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

We hypothesize that ChatGPT could offer potential solutions to two key challenges related to emoji interpretation, based on prior explorations. First, given ChatGPT's demonstrated superiority over human annotators in various text annotation tasks (Kocoń et al. 2023), we believe it may possess a deep understanding of emoji usage. As such, it could replace human workers in labeling emoji attributes, reducing the need for crowd-sourcing efforts and lowering costs. Second, when users encounter unfamiliar emojis or are unsure of their meanings, ChatGPT's ability to explain emoji semantics could effectively prevent misunderstandings.

On the other hand, previous work suggests that there exist hallucinations in LLMs' generations (Zhang et al. 2023), inconsistencies between the model's output and the real-world facts or user input. Whether ChatGPT can produce generations similar to humans for the meaning of emojis is still questionable. Delineating the boundaries of ChatGPT's capabilities across emoji-related tasks of varying difficulty is essential for driving progress in future emoji research.

To verify ChatGPT's potential on emoji research and application, our paper conducts a set of qualitative and quantitative studies about ChatGPT's understanding on emoji usage in multiple tasks. Based on the current work on emoji studies, we propose three research questions to probe the ChatGPT's capability on emoji-related tasks.

RQ1 Does ChatGPT generate similar explanations on emoji semantics, sentiments, and usage intentions as humans?

RQ2 Does ChatGPT encode knowledge about emoji usage patterns associated with different communities?

RQ3 What is the performance of ChatGPT in emoji-related downstream tasks?

To answer these research questions, we first ask ChatGPT to explain the semantics, sentiments, and intentions of emojis with or without text context, across different platforms or cultures (Section 3, 4, 5). We find that the interpretation of ChatGPT on emojis is consistent with human responses in the existing literature in most cases. Next, we ask ChatGPT to generate the associated emoji patterns with multiple communities (gender, platform, hashtag, and culture) to probe ChatGPT's knowledge about emoji usage in different communities (Section 6). Finally, we test ChatGPT on emoji-related downstream tasks, i.e., irony annotation and emoji

prediction (Section 7, 8), and the results reveal that ChatGPT is capable of using the emoji information for irony annotation and making emoji recommendations based on user identity.

2 Related Work

Our work can be related to two streams of previous research studies: emoji functionalities and ChatGPT applications on the computational social science domain.

2.1 Emoji Interpretation

With the prevalence of emojis on social networks and other platforms, there has been increasing interest in studying the emoji functionality. Researchers have explored emojis with the functions to express the sentiment, highlight topics, decorate texts, adjust tones, indicate identities, and engage the audience (Ai et al. 2017; Lu et al. 2016; Ge 2019; Hu et al. 2017; Cramer, de Juan, and Tetreault 2016), and with these emoji functions, the researchers summarized the intentions to use emojis on different platforms (Hu et al. 2017; Lu et al. 2018). In addition to these functions and intentions, the researchers also noticed differences in the use of emojis in different application scenarios, such as apps, cultures, genders, hashtags, and platforms (Tauch and Kanjo 2016; Lu et al. 2016; Chen et al. 2018; Barbieri et al. 2016; Zhou and Ai 2022; Wood-Doughty et al. 2021). However, many works on emoji functionality choose to perform crowd-sourcing work to conduct the emoji interpretation, but with rich and diverse source of emoji applications, it is essential to have an automatic tool to interpret the emoji usage. In addition to emoji functionality, several researchers have applied emoji information to downstream tasks, such as emoji prediction (Barbieri et al. 2020, 2018), improving sentiment analysis (Chen et al. 2021, 2019; Felbo et al. 2017), predicting developer’s dropout (Lu et al. 2022).

2.2 LLMs in Computational Social Science

With the prevalence of LLMs, there are emerging researches evaluating their performances on various tasks (Chang et al. 2023; Liu et al. 2024), including Computational Social Science (CSS) tasks (Ziems et al. 2023; Zhu et al. 2023). Researchers have examined ChatGPT’s capacities in the zero-shot or few-shot (with in-context learning) settings, in various text annotation tasks such as political affiliation classification of tweets (Törnberg 2023); hate speech detection (Huang, Kwak, and An 2023; Li et al. 2023b; Zhu et al. 2023); discourse acts (Ostyakova et al. 2023); instruction quality evaluation (Xu et al. 2024); sentiment analysis and bot detection (Zhu et al. 2023; Zhou et al. 2023b). A comprehensive evaluation (Ziems et al. 2023) shows ChatGPT exhibits unsatisfactory results on tasks that have complex structure or have subjective taxonomies whose semantics differ from definitions learned in pretraining, while achieving better results on tasks that have objective ground truth (e.g., misinformation and fact checking) or have labels with explicit colloquial definitions in the pretraining data (e.g., emotion and political stance). Nevertheless, the results demonstrate the potential of using LLMs to help researchers

with CSS tasks and general users with social-related daily activities.

Despite the growing prominence of emojis as an emerging “language,” limited studies focus on the intersection of emojis and ChatGPT. Only a few works suggest ChatGPT’s ability to comprehend emojis. de Janeiro (2023) examines ChatGPT’s ability to generate various facial emojis to represent the valence of words, while Peng et al. (2023) explores translating natural language into emoji sequences. Belal, She, and Wong (2023) and Kocoń et al. (2023) find that ChatGPT can interpret emojis in texts to enhance sentiment prediction. The most relevant concurrent work is by Lyu et al. (2024), which investigates whether GPT-4V understands emoji interpretations and correctly employs emojis when given social media posts. However, their study does not cover the functionality or intent of emojis nor discuss emoji usage across different communities.

In contrast, we thoroughly examine the consistency of ChatGPT’s annotations on emoji semantics, sentiment, and intentions with human interpretations and probe ChatGPT’s knowledge of emoji usage patterns across various communities. We evaluate ChatGPT’s performance on emoji-related downstream tasks and explore the significance of emojis to ChatGPT in these tasks. Our evaluation spans multiple languages and platforms, qualitatively and quantitatively demonstrating ChatGPT’s understanding of emojis in broader contexts.

3 Emoji Semantics

We start evaluating ChatGPT’s understanding on emojis by asking ChatGPT to explain the semantic meanings of the emoji. Emojis, as universal language tokens, are applied by users with different language and cultural backgrounds. In this section, we conduct a qualitative study on the LLM explanation of emoji universal semantics without any given context, and then we probe the ChatGPT understandings in two different language and culture settings: the English context and the Chinese context.

3.1 Emoji Semantics without Context

We first collect the 50 emojis most frequently used from 2019 to 2022 provided by the Unicode organization,¹ and investigate how ChatGPT explains these emoji semantics without any context. The probing prompt for GPT4 is designed as follows:

Could you please describe this emoji [Emoji] and explain its meaning?

where [Emoji] is the Unicode of selected emojis. To compare the understandings between ChatGPT and humans, we also collect the meaning of emojis on the Emojipedia website,² which is researched and written by Emojipedia editors and lexicographers, as human annotations of emoji meanings. Two authors of the paper independently compare the human annotations and ChatGPT responses, finding that ChatGPT and humans provide similar explanations for all

¹<https://home.unicode.org/emoji/emoji-frequency/>

²<https://emojipedia.org/meanings>

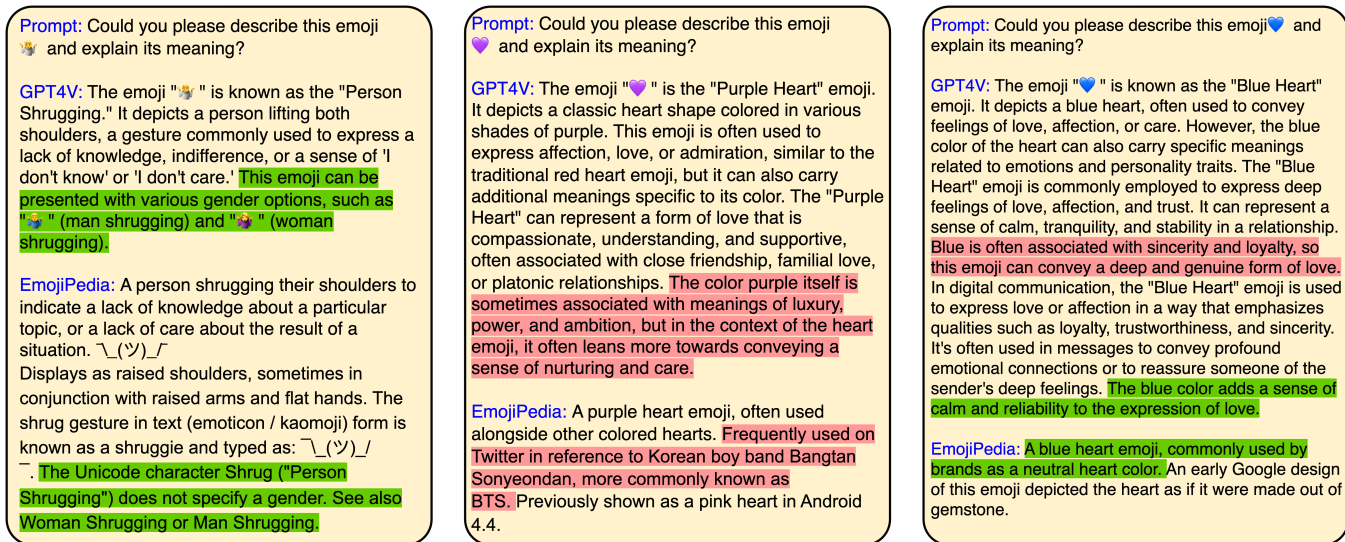


Figure 1: Meaning and explanation of emoji usage from ChatGPT and Emojipedia. Different interpretations are highlighted in red and similar interpretations from both ChatGPT and Emojipedia are highlighted in green.

50 emojis. However, we note that both ChatGPT and human annotators highlight points that the other may not cover.

We present three representative examples for the emoji 🙄 (person shrugging), 💜 (purple heart), 💙 (blue heart) in Figure 1, highlighting the same contents mentioned by humans and ChatGPT by green and different contents by red color.

From Figure 1, for the emoji 🙄 (person shrugging), it is consistently observed, both in ChatGPT’s responses and human annotations, that this emoji primarily conveys a lack of knowledge without implying any specific gender. Regarding the emoji 💜 (purple heart), ChatGPT provides an explanation on the significance of the color purple, while human annotations predominantly discuss its usage in tweets related to a Korean boy band. For the emoji 💙 (blue heart), ChatGPT exclusively addresses the particular meaning associated with the color blue. However, both ChatGPT and human sources agree that 💙 typically expresses sentiments of calmness and neutrality.

The findings suggest that the explanation from ChatGPT on semantics of most emojis overlaps with the meaning in Emojipedia. Moreover, the content of ChatGPT and human annotations can be supplementary to each other.

3.2 Emoji Semantics with Text Context

The accurate explanation of ChatGPT in emoji semantics without text context follows our expectation due to the notable performance of ChatGPT in sentiment analysis and tweet-related tasks (Zhu et al. 2023; Törnberg 2023). However, it is well known that emojis will evolve extensive semantics according to the applied text context. For example, 🐐 (goat) can mean the greatest person, and 🐍 (snake) can mean the Python language in the specific context. We thus explore whether ChatGPT can interpret an emoji’s meaning when giving tweets containing emojis with extensive seman-

tics. We choose a few emojis (🐐 goat, 🐝 honey bee, and 🐍 snake) with widely accepted special meanings³ and randomly extract tweets in 2022 that contain these emojis via Twitter API. We feed these tweets to ChatGPT and ask ChatGPT to give the semantic explanation based on the context of the tweet.

We present the prompt and the GPT4 response in Figure 2, which illustrates that ChatGPT’s interpretations of emoji semantics align closely with human understanding in most instances. ChatGPT accurately interprets the emoji 🐐 (goat) as “the greatest of all time.” Additionally, it interprets 🐝 in a Valentine’s context as “be mine.” In the third example featuring the 🐍 emoji, ChatGPT recognizes its use in symbolizing the Python programming language. However, in the last example with the tweet context “Whole lotta snakes in my grass 🐍🐍,” we argue that the emoji should symbolize the animal snake rather than untrustworthy people. The frequent use of the snake emoji to represent betrayal might have led to this misinterpretation. In conclusion, ChatGPT generally demonstrates a precise understanding of emoji semantics in textual contexts. However, when an emoji represents a less commonly used object, ChatGPT may provide erroneous interpretations.

3.3 Emoji Semantics with Language/Cultural Context

In addition to the application in the English context, emojis also show prevalence in other specific language and culture contexts and emojis may evolve to some unique meanings in a specific language context (Chen et al. 2019; Lu et al. 2016). ChatGPT, as a multilingual language model,

³selected from this website <https://bestlifeonline.com/emoji-meanings/>

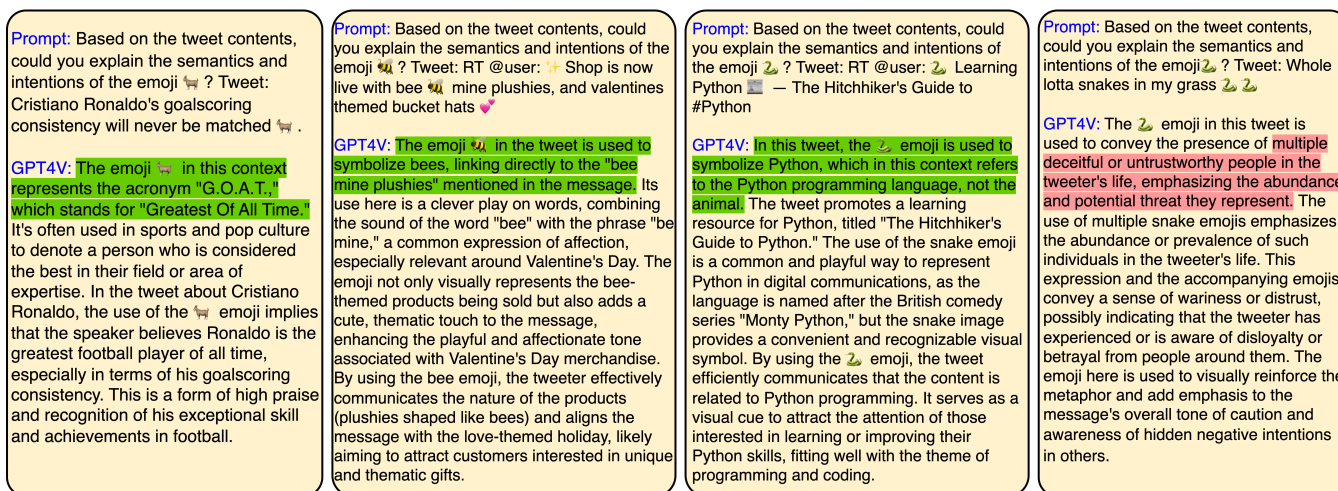


Figure 2: Semantic explanation from ChatGPT of emojis with special meaning. The green color highlights the semantic explanation consistent with human understanding, and the red color highlights the hallucinations. From the case study, ChatGPT can precisely understand the semantics of emojis with special usage when given the specific context in most cases.

has demonstrated strong cross-linguistic understanding abilities (OpenAI 2023). To evaluate the feasibility of ChatGPT replacing human annotators in different language contexts, we select a few emojis with unique meanings in the Chinese context and observe whether ChatGPT could correctly capture the specific semantics.

With a similar experiment setup as in Section 3.2, we first choose emojis that have unique meanings on Chinese social networks and test if ChatGPT can properly explain their semantics, that is, 😊 (slightly smiling face), 🐶 (doge), 🍯 (pill), 🐛 (bug), 🍋 (lemon), 🍉 (melon), and 🐦 (bird). In the Chinese pop culture context, 😊 ironically expresses dislike and ridicule; 🐶 conveys a sense of irony or sarcasm; 🍯 and 🐛 use a homophone pun, meaning “it’s over” and expressing support, respectively; 🍋 express envy and jealousy; 🍉 means gossip or drama; 🐦 refers to not showing up for a planned meeting, when it’s used as a pigeon.

Additionally, we include an emoji combo, 🐮🍺, that consists of a cow face and a beer mug (a homophone pun to express “so cool”). We prompt ChatGPT to explain the usage of these emojis. The prompts are written in English or Chinese, with or without specification of usage in Chinese pop culture. For some emojis, we include appropriate hints, such as connections to homophone puns.

We find that ChatGPT can generate desired interpretations in all cases, but only with appropriate prompts. The prompts need to include very specific information sometimes. For instance, when prompted with “Explain [Emoji]”, GPT models failed to generate desired explanations and answered with literal interpretations from their Unicode names or visual cues (e.g., 😊 as friendly instead of its common ironical use in Chinese pop culture (Wang 2022)). When prompted in Chinese, only 🐮🍺 got the correct explanation. However, 😊, 🍉, and 🐶 get desired explanations when prompted with “Explain the [Emoji] in the context of Chinese pop cul-

ture”, no matter whether prompted in English or Chinese. For 🍯 and 🐛, we had to specifically mention its homophonic usage with yào wán and chōng, respectively. Similarly for 🐦 we had to mention its usage as a pigeon.

We also observe that the Chinese answers generated by prompts written in Chinese contained more specific and useful information in general, compared to the counterparts in English. For example, the English answer for 😊 mostly conveys that its usage can be context-dependent and culturally specific and only briefly mentions “it might be used ironically”, while the Chinese version details the several specific contexts when it can be used ironically.

Meanwhile, answers in both languages are prone to hallucinations even when a large proportion of answers are desired. The model links 🍯 to the red pill and blue pill metaphor from the Matrix movie and incorrectly states that this is a common usage in Chinese social media.

In summary, ChatGPT is capable of correctly explaining the emojis in the Chinese pop culture context. However, prompting in Chinese or mentioning appropriate hints is necessary to better elicit such knowledge in some cases.

Observation 1 ChatGPT can output consistent explanations on emoji semantics when giving no context or specific English context. However, when probing the special meaning of an emoji in the Chinese pop culture context, ChatGPT is prone to hallucinations with English prompts but more accurate explanations when prompted in Chinese or given appropriate hints.

4 Emoji Sentiment

Expressing sentiment is one of the primary functions of emojis, especially face emojis (Ai et al. 2017). Previous research has employed human annotation to investigate emoji sentiment (Hu et al. 2017) and has incorporated emoji information to enhance downstream sentiment analysis (Felbo

et al. 2017; Chen et al. 2019). In this section, we prompt ChatGPT to annotate emoji sentiment and compare its annotations with those of humans. Additionally, we evaluate ChatGPT’s performance in predicting the sentiment of tweets containing emojis for downstream tasks.

4.1 Emoji Sentiment Annotation

We first ask GPT4 to label the 50 most frequently used emojis in Section 3.1 with positive, neutral and negative sentiment with the prompt in Table 12a in Appendix.

Compared to the human annotations in Hu et al. (2017), we observe that ChatGPT and human annotators share the same sentiment annotation for almost all emojis (49 out of 50). The only difference is the annotation for the 🙄 (flushed face) emoji. While the human annotation is neutral, ChatGPT’s response is negative, considering feelings of embarrassment or a sense of being overwhelmed associated with 🙄. This discrepancy indicates that ChatGPT can interpret the sentiment of emojis but may differ from human annotations in specific cases. ChatGPT’s response often notes that it requires text context to accurately determine emoji sentiment. Therefore, we discuss the influence of emojis on ChatGPT’s performance in downstream sentiment analysis tasks.

4.2 Sentiment Prediction with Emojis

Previous researchers have emphasized the importance of emoji information to judge sentiment in the given inputs (Chen et al. 2019; Felbo et al. 2017). We also investigate whether ChatGPT considers emojis in its sentiment classification predictions. For the experiment setup, we choose the sentiment analysis dataset from the TweetEval benchmark (Barbieri et al. 2020), which includes three labels: negative, neutral, and positive. We extract tweets containing emojis from the dataset, resulting in 805 tweets. We perform sentiment prediction using in-context learning (ICL) with GPT4 and GPT3.5, providing 5 demonstrations for each model. The models achieve 70.31% and 71.67% accuracy for GPT4 and GPT3.5, respectively. The fine-tuned RoBERTa model, pretrained on the tweet corpus, achieves 69.3% accuracy (Barbieri et al. 2020), indicating that ChatGPT provides reliable sentiment predictions.

Next, we remove the emojis from the input text and ask ChatGPT to assign a sentiment label again. The performance accuracy decreases to 69.56% and 70.19% for GPT4 and GPT3.5, respectively. However, this decrease in accuracy does not conclusively demonstrate that ChatGPT utilizes emoji information for sentiment prediction, as the change is not significant and removing emojis may also alter the ground-truth sentiment of the text. We compare the sentences where ChatGPT changes the sentiment prediction after removing the emojis and find that 19.50% and 18.13% of tweets flip the prediction labels for GPT3.5 and GPT4 annotations, respectively. We present three randomly selected tweets with flipped predictions for GPT4 in Table 1.

The analysis of examples in Table 1 reveals that ChatGPT’s sentiment predictions are influenced by the presence of emojis within the text. Incorporating emojis like 😞 (weary face) and 🙄 (face with rolling eyes), which carry negative sentiment, results in a shift in sentiment prediction

Tweet	Label	p_{emoji}	$p_{noemoji}$
How many more days until opening day? 😞	neutral	negative	neutral
@user Everyone knows that vegetarianism is synonymous with wellness! 😊	positive	neutral	positive
#MPN #OneDirection#MtvStarsNiallHoran JESUS 🙄🙄🙄🙄🙄	positive	positive	neutral

Table 1: Tweets with flipped sentiment predictions from ChatGPT, where p_{emoji} and $p_{noemoji}$ represents the ChatGPT predictions with and without emojis.

from neutral to negative and from positive to neutral, respectively. Similarly, the addition of multiple 😊 (smiling face with heart eyes) emojis leads to an elevation in sentiment prediction from neutral to positive. This case study suggests that ChatGPT considers emoji tokens in text as significant indicators when assessing overall sentiment.

Observation 2 *GPT-4 annotates emoji sentiment labels similarly to humans and considers emoji sentiment to enhance its predictions when performing sentiment classification tasks.*

5 Emoji Usage Intention

In addition to sentiment and meaning, previous researchers have also explored the intention behind using emojis, asking human annotators to label emoji intentions with or without context (Hu et al. 2017; Lu et al. 2018). In this section, we follow the approaches and intention definitions from previous work (Hu et al. 2017) and ask ChatGPT to reannotate the emojis with or without context. This allows us to observe whether humans and ChatGPT have similar understandings of the intention behind using emojis.

5.1 Emoji Intention Annotation without Context

First, we annotate emojis without any context to assess ChatGPT’s understanding of their universal intentions. We use the same seven types of candidate intentions as in previous work: expressing sentiment, strengthening expression, adjusting tone, expressing humor, expressing irony, expressing intimacy, and describing content. Detailed definitions of these intentions can be found in Hu et al. (2017). We select 15 emojis (5 positive, 5 neutral, and 5 negative) from Hu et al. (2017) and ask ChatGPT to rate the willingness to use each emoji for each intention on a 7-point scale (7 = totally willing, 1 = not willing at all), following the same annotation method as in Hu et al. (2017). The complete prompt is shown in Table 13 in Appendix.

We calculate the average difference between ChatGPT-rated scores and human ratings as follows: -0.13 (expressing sentiment), -0.29 (strengthening expression), -0.07 (adjusting tone), -0.54 (expressing humor), -1.49 (expressing irony), -0.05 (expressing intimacy), and 1.11 (describing content). Given that the average standard deviation (SD) of human-rated scores is 1.643 (Hu et al. 2017), ChatGPT and human annotators reach a high consensus on the intentions of expressing sentiment, strengthening expression, and adjusting tone. For other intentions, such as describing con-

tent, ChatGPT’s scores differ from human ratings, but the discrepancy is smaller than the average human SD. Hence, we conclude that ChatGPT’s understanding of the universal usage intention of emojis is generally similar to that of human users.

5.2 Intention of Emoji Usage in GitHub Posts

The similar intention annotation without context between ChatGPT and humans aligns with the accurate explanation of emoji semantics and the precise annotation of emoji sentiment. However, the intention behind using emojis can vary across different platforms. For example, users on GitHub may use emojis to organize content within a post but rarely to express irony. Understanding emoji intention in GitHub posts poses a greater challenge for ChatGPT.

In this part, we collect the same 2,000 emoji posts on GitHub with human annotations from Lu et al. (2018) and compare ChatGPT’s annotations of emoji intentions with the human ground truth labels. The candidate intentions of emojis in GitHub posts are shown in Table 2, with detailed definitions provided in Lu et al. (2018).

Intention	Sentiment expressed	Sentiment strengthened	Tone adjusted	Content described
Accuracy	67.43%	30.85%	48.64%	47.54%

Intention	Content organized	Content emphasized	Non communication use	Symbol
Accuracy	40.63%	55.49%	34.62%	14.58%

Table 2: Accuracy of ChatGPT annotation for each intention on emojis in GitHub posts with human annotation as the ground truth labels.

The annotated accuracy between GPT3.5 and GPT4 compared with human annotations is 38.8% and 49.0%, respectively. The detailed accuracy of GPT4 for each intention is presented in Table 2. We observe that for emojis used to express sentiment and emphasize content in GitHub posts, ChatGPT’s judgments are more aligned with human annotators. However, for symbols (symbolize emojis) and non-communication use (unintentional use), which are rarely applied on other social platforms, ChatGPT’s intention annotations do not match those of human annotators.

We also conduct a detailed case study and present three randomly selected posts with different annotations in Table 3 to further analyze the extent to which ChatGPT’s responses align with human understanding.

Furthermore, we feed the ChatGPT annotations in Table 3 back to the LLM and ask for justification. In the first post, human annotators perceive the use of ✨ as an enhancement of the post’s positive sentiment, whereas ChatGPT interprets its use as a means to attract attention. In the second post, which contains the word “thanks” accompanied by 🙏, human annotators believe that “thanks” alone conveys positive sentiment, and the emoji serves to reinforce this sentiment. ChatGPT’s interpretation also notices that the emoji conveys the sentiment and gives a similar annotation: sentiment expressed. The final post features 📁, which human annotators

Post	Human label	ChatGPT label
Make awesome stuff with half the time OR make quadriply awesome creations with the same time ✨	sentiment strengthened	content emphasized
thanks 🙏	sentiment strengthened	sentiment expressed
📁 Example	content organized	content described

Table 3: GitHub posts with different annotation labels for emoji usage intention from ChatGPT and humans.

suggest is used for organizing content to enhance readability, while ChatGPT views it as a symbol or emphasis of a document. From the case study, we find that even in incorrect annotations, ChatGPT still captures the primary intention behind the emoji usage. This case study highlights that ChatGPT’s interpretations of emoji usage in GitHub posts are not only plausible but also indicative of ChatGPT’s nuanced understanding of emoji applications.

Observation 3 *ChatGPT precisely deduces the intention of emoji usage without context and, when provided with a specific GitHub post, reaches a consensus with human annotations on the intention of emojis.*

6 Emoji Usage Associated with Communities

While ChatGPT has demonstrated a deep understanding of emoji semantics, intentions, and sentiments, previous work suggests that emoji usage can vary significantly across different communities, such as platforms, languages, genders, and hashtags (Lu et al. 2018, 2016; Chen et al. 2018; Zhou and Ai 2022; Zhou et al. 2023a). In this section, we investigate whether ChatGPT encodes knowledge of these associated emoji usage patterns when provided with prior information about users, platforms, or co-occurring hashtags. We construct the prompt by directly asking for the most used emojis given different community information, and ChatGPT outputs the most used emojis along with its justification. The probing prompt for ChatGPT is designed as follows:

Could you output 10 emojis that are associated with the specific community?

where *specific community* refers to the associated community, such as female users, French users, specific hashtags, or the GitHub platform.

We present ChatGPT’s answers and compare them with findings from previous work in Tables 4, 5, and 6. In particular, emojis associated with French, female, and male users are explored in Lu et al. (2016); Chen et al. (2018), and emojis co-occurring with different hashtags and on the GitHub platform are reflected in Zhou and Ai (2022); Zhou et al. (2023a); Lu et al. (2018).

From Tables 4, 5, and 6, we observe that ChatGPT relies on stereotypical associations with user attributes or platform characteristics to select emojis. In most cases, this stereotypical selection results in choosing the same or semantically similar emojis to those identified in previous works.

To evaluate ChatGPT’s understanding of irony in emojis for other languages, we utilize a recently published Arabic dataset with ironic annotations: ArSarcasMoji from Hakami, Hendley, and Smith (2023). Instead of labeling the irony of sentences, ArSarcasMoji provides the ground truth irony label of emojis in three categories: sarcastic, humorous, and not ironic. The sarcastic label indicates that ironic emojis play a negative sentiment role, while the humorous label indicates a positive role (Hakami, Hendley, and Smith 2023).

Since the label distribution in the original dataset is biased, we first downsample sentences labeled as “not ironic” and “sarcastic” to create a balanced dataset for evaluation. We compose the prompt with instructions, detailed definitions of “humorous” and “sarcastic,” and examples of Arabic sentences with labels (the complete prompt is provided in Table 16). The prediction accuracy for GPT4 and GPT3.5 is 60.1% and 63.1%, respectively, indicating that ChatGPT can understand the functionality of emojis to express irony in the Arabic context.

Observation 5 *In the emoji irony classification task, ChatGPT demonstrates promising performance in distinguishing whether an emoji is ironic or not. In the tweet irony classification task, ChatGPT considers the presence of emojis when determining the overall irony of a tweet.*

8 Emoji Recommendation

Emoji recommendation, a well-defined and widely adopted task, is essential for evaluating models’ understanding of emoji usage (Barbieri et al. 2020; Lyu et al. 2024). This task requires models to grasp both the semantics of emojis and their application context. After establishing that ChatGPT can identify the sentiments, meanings, and intentions behind emojis, we examine its ability to recommend suitable emojis based on a given text context.

8.1 Recommendation Based on Text Context

For the experiment, we use the SemEval 2018 Task 2 dataset for emoji recommendation in English and Spanish (Barbieri et al. 2018). We select the validation sets with 5,000 English and Spanish tweets for inference. The task pre-defines a list of 20 emojis for English and 19 for Spanish, and our prompt includes the predefined emoji candidates and five exemplars from the respective training set. The prompt details are shown in Table 17.

For emoji recommendation on English tweets, GPT4 and GPT3.5 achieve 22.12% and 21.52% accuracy, respectively, while the fine-tuned FastText classifier achieves 42.56% accuracy. For Spanish tweets, GPT4 and GPT3.5 obtain 20.14% and 20.96% accuracy, respectively, while the fine-tuned FastText classifier achieves 31.63% accuracy. These results suggest that, in a closed-set emoji recommendation, the ICL of ChatGPT cannot outperform the traditional fine-tuned classifier. However, ChatGPT’s performance is consistent across languages, whereas the fine-tuned classifier experiences a significant performance drop. Furthermore, we observe ChatGPT hallucinations during emoji recommendation, with 6.1% and 3.7% of English tweets being annotated

with emojis not in the predefined candidates for GPT3.5 and GPT4, respectively.

8.2 Recommendation with User Identity

Previous work suggests that emoji preferences differ based on user identity (Chen et al. 2018; Lu et al. 2016), so it is important to consider user demographic attributes when recommending emojis. ChatGPT has shown the potential to act as a communicative agent with different roles (Li et al. 2023a). Moreover, as shown in Section 6, ChatGPT encodes knowledge about discrepancies in emoji usage patterns among different communities. In this part, we investigate whether ChatGPT can consider user identity when recommending emojis for tweets.

We ask ChatGPT to act as a female or male Twitter user and provide emoji recommendations accordingly. We add another sentence at the beginning of the prompt, as in previous work (Li et al. 2023a): *I want you to act as a female/male Twitter user. Never forget your role.* We repeat the GPT4 ICL experiment on the same 5,000 English tweets. Acting as a male or female Twitter user, ChatGPT achieves 21.04% and 20.94% accuracy, respectively, with 1,110 tweets having different predictions. Although the provided gender identity does not significantly change the accuracy, it altered more than 20% of the predictions of the recommended emojis. We also performed a Chi-Square Test (Rana and Singhal 2015) to examine the significance of the difference in emoji distributions between the two prompts, yielding a p-value of less than 0.01. We present the distribution of the top 10 predicted emojis for female and male users in the first two rows of Table 8.

Role	Emoji distribution
Female	🌟: 657, 😊: 582, 🤔: 573, 🔥: 418, 💖: 410, 😏: 367, 😬: 347, 🤩: 224, 🌲: 194, ❤️: 186
Male	🌟: 575, 😊: 560, 🤔: 527, 😏: 488, 🔥: 451, 😬: 341, ❤️: 237, 💙: 227, 💖: 219, 🤩: 208
French	😊: 620, 🤔: 583, 🌟: 530, 😏: 411, 🤩: 396, 🔥: 381, 💖: 310, ❤️: 252, 🤩: 236, 🛍️: 192
American	😊: 605, 🤔: 576, 🌟: 566, 🔥: 448, 😏: 408, 😬: 361, 💖: 270, 🤩: 233, ❤️: 216, 🌲: 207

Table 8: Distributions of top 10 predicted emojis when asking ChatGPT to play different roles. The yellow color highlights the emojis 💖 and 😏, which are two emojis with the largest change in the number of recommendations for acting as male or female users.

An examination of the first two rows of Table 8 reveals a comparable emoji distribution with two notable exceptions: a higher frequency of 💖 (two hearts) when representing female users, and a higher usage of 😏 (smiling face with sunglasses) in representations of male users. This finding is consistent with ChatGPT’s understanding of usage differences for female and male users in Section 6 and the findings from previous work (Chen et al. 2018).

To further explore ChatGPT’s comprehension of variations in emoji usage across different countries, an aspect not revealed by the results in Section 6, we conducted additional experiments. In these, we simulate interactions while portraying users from France and the United States, respectively. When asking ChatGPT to act as a French user, we translate the prompts and tweets to French. We present the predicted emoji distribution in the last two rows of Table 8.

From the last two rows, although the p-value of Chi-Square Test is still smaller than 0.01, we cannot find a significant change in the occurrence numbers for a specific emoji when portraying users from France or the USA. In conjunction with the findings presented in Section 6 regarding ChatGPT’s interpretation of emoji usage by French users, it can be concluded that ChatGPT lacks awareness of the variations in emoji usage across different countries.

Observation 6 *ChatGPT does not outperform the fine-tuned traditional classifier on the closed-set emoji recommendation task. However, when provided with user gender information, ChatGPT can alter its emoji recommendations.*

9 Implications and Limitations

Our paper reveals ChatGPT’s precise understanding of emoji functionality and intention. ChatGPT’s capabilities with emojis could benefit researchers and social media users. For researchers involved in future emoji studies, leveraging human-aligned explanations of emoji semantics and intentions could significantly reduce the need for extensive crowd-sourcing, particularly in multilingual contexts. This approach would enable more effective and efficient research into emoji usage and interpretation. Social media users could utilize ChatGPT to decipher the semantics of unfamiliar emojis, thereby enhancing the clarity and transparency of online communication.

In addition to ChatGPT’s overall effectiveness in emoji-related tasks, we observed instances of hallucination. For example, when explaining the semantics of the 🐼 emoji in Figure 2, emojis with unique meanings on Chinese social media, or when identifying frequently used emojis by French users in Table 4, ChatGPT often relies on spurious correlations or prototypes formed during training or alignment, leading to hallucinations. These findings are valuable for future emoji research, reminding researchers to exercise caution regarding ChatGPT’s hallucinations, especially in cross-cultural contexts. They also highlight the need for model developers to focus more on improving non-verbal token understanding in cross-cultural settings.

All the experiments in our paper treat emojis as text tokens in the Unicode form. But some emojis due to its application in the specific cultural context or new appearance are not included in the Unicode organization, such as 🍷 (broken) and 🙄 (onlooker) in Chinese social media. For these emojis without standard Unicode, we can utilize the vision language model (VLM) from ChatGPT: GPT4V, which has been demonstrated with significant progress in image understanding (Yang et al. 2023; Wang et al. 2024), to help us understand the emoji functions. We prompt the images of 🍷 and 🙄 to GPT4V for semantic explanations. An example is

shown in Figure 3. We observe similar patterns as in Section 3.3. GPT4V generates hallucinations with Chinese prompts but outputs the correct explanation when supplementing the prompt with the context of Chinese pop culture.

We also repeat the annotation of the sentiment of emojis in Section 3.1 on GPT4V. For the 50 selected emojis, only 🙄 (pleading face) and 👁️ (eyes) have different annotations for GPT4 and GPT4v. For the 🙄 emoji, VLM attributes a negative sentiment based on its visual features: frown, wide eyes, and raised eyebrows. In contrast, GPT4 associates the 🙄 emoji with positive sentiment, perceiving it as an expression of vulnerability, empathy, and affection. Regarding the 👁️ emoji, GPT4 assigns a neutral sentiment, indicative of observation, while GPT4v leans towards a neutral, yet slightly positive, reflecting an underlying connotation of interest.

These findings suggest that, while the GPT4v and GPT4 sentiment annotations are generally the same, divergent interpretations of the 🙄 emoji can highlight the evolution of the emoji sentiment derived from the user application setting. The results in GPT4V imply that, when encountering specific emojis without the standard Unicode, the notable capability in emoji understanding of GPT4V can also help researchers and users.

With the rapid evolution of current large language models (LLMs), there may be concerns about whether the observations and conclusions presented in this study will still hold for future versions of ChatGPT. On the one hand, in most scenarios where ChatGPT produces human-aligned responses, we anticipate that its accuracy will improve as the model continues to evolve. On the other hand, we also identified cases where ChatGPT produces hallucinations, particularly in the cross-cultural contexts discussed in Sections 6 and 3.3. Since knowledge bases related to platforms like WeChat and Twitter are either not publicly available or restricted for commercial use, we infer that these hallucinations may arise from a lack of relevant knowledge in the training data. We argue that these issues may not be easily resolved by simply scaling up data or model size. Consequently, we expect that the current limitations we have observed are likely to persist and may not be quickly addressed. Therefore, our findings, which delineate the boundaries of ChatGPT’s capabilities, may help researchers better utilize the tool for emoji annotation tasks over time.

There are several limitations in our exploration work. First, previous work has shown that the responses of LLMs are sensitive to temperature and the prompt (Yoo et al. 2021; Schick and Schütze 2020; Gilardi, Alizadeh, and Kubli 2023; Zhou, Maharjan, and Liu 2023), but in our work, for each task, we only explore one type of prompt presented in Appendix. For the quantitative studies in our paper, such as intention prediction, irony prediction, sentiment classification, and emoji recommendation, we use the ChatGPT API (gpt-4-1106-preview and gpt-3.5-turbo-1106) with a temperature setting of 0 to generate the model outputs. For the qualitative studies (emoji semantics and emoji usage within communities), we initially used the ChatGPT web interface to obtain results. To demonstrate the robustness of our findings, we repeated qualitative studies using the ChatGPT API with a temperature of 0 and presented the results in

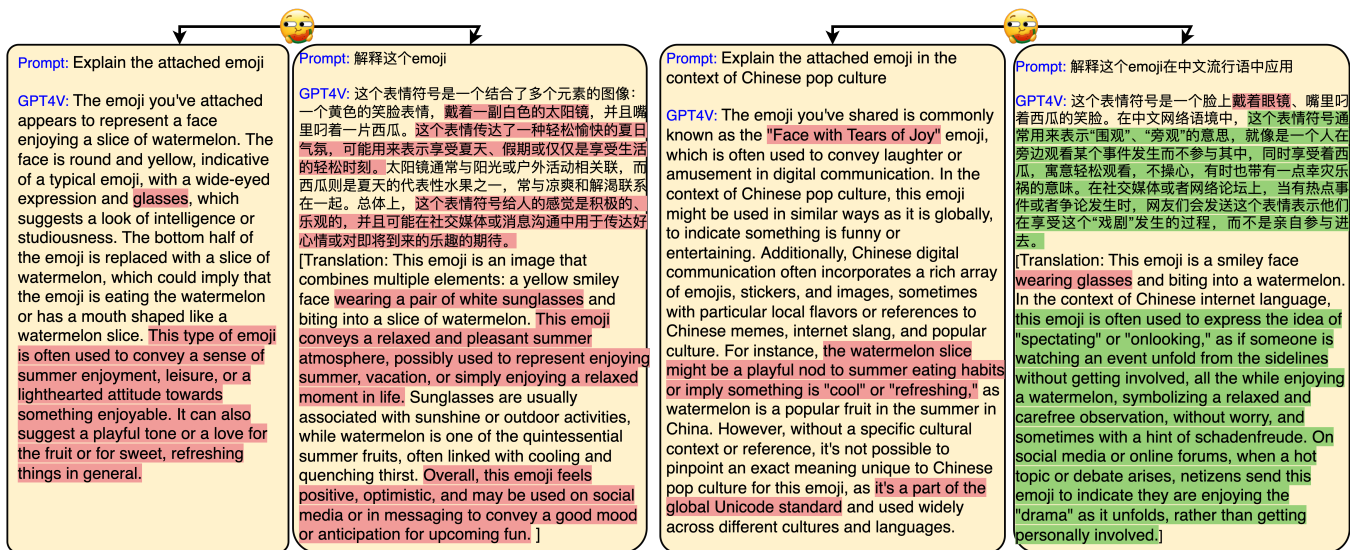


Figure 3: Qualitative Study of Emoji Semantics with Language/Cultural Context. Hallucinations are highlighted in red and desired interpretations are highlighted in green.

the Appendix .1. Additionally, to gain further insights into open-source models, we replicated some experiments using Llama3-8B with a temperature setting of 0, and the results are also provided in the Appendix .2. We found that the open-source Llama3-8B model exhibits significant hallucinations when probing emoji intentions or probing the associations between emoji and community information, making it unreliable for replacing human annotators with the base Llama3-8B model.

Second, emoji semantics or sentiment can evolve over time. For example, 😊 (slightly smiling face) has evolved the sense of irony and negative sentiment in Chinese social networks. The human annotations utilized in this paper are almost all prepared before 2018, and the 5-year gap may lead to the different annotation results.

Lastly, our paper focuses on the ICL of GPT3.5, GPT4 and GPT4V for ChatGPT and does not involve the fine-tuning experiment with OpenAI API. Moreover, with the release of GPTs (the tailored version of ChatGPT for a specific task⁴), we can find there is a version of ChatGPT from OpenAI, called *genz 4 meme*, to help understand the latest memes, which may have better understanding for emojis. We leave further exploration of the different versions of ChatGPT in the future work.

10 Conclusion

In this study, we conduct a comprehensive evaluation of ChatGPT's proficiency in interpreting emojis, focusing on various aspects including meaning, sentiment, and intention. Our analysis reveals that ChatGPT's annotations closely align with human labels. We discover that ChatGPT has embedded knowledge of emoji usage patterns prevalent in different communities. Furthermore, ChatGPT exhibits strong

performance in several downstream tasks that involve emojis, indicating its capability to leverage its emoji understanding in making predictions. The encouraging results in annotation tasks suggest that ChatGPT could significantly contribute to emoji research by serving as a substitute for human annotators, thus conserving human resources.

References

- Ai, W.; Lu, X.; Liu, X.; Wang, N.; Huang, G.; and Mei, Q. 2017. Untangling emoji popularity through semantic embeddings. In *ICWSM 2017*.
- Barbieri, F.; Camacho-Collados, J.; Neves, L.; and Espinosa-Anke, L. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. *arXiv preprint arXiv:2010.12421*.
- Barbieri, F.; Camacho-Collados, J.; Ronzano, F.; Anke, L. E.; Ballesteros, M.; Basile, V.; Patti, V.; and Saggion, H. 2018. Semeval 2018 task 2: Multilingual emoji prediction. In *Proceedings of the 12th international workshop on semantic evaluation*, 24–33.
- Barbieri, F.; Kruszewski, G.; Ronzano, F.; and Saggion, H. 2016. How Cosmopolitan Are Emojis? Exploring Emojis Usage and Meaning over Different Languages with Distributional Semantics. In *ACM-MM 2016*.
- Belal, M.; She, J.; and Wong, S. 2023. Leveraging chatgpt as text annotation tool for sentiment analysis. *arXiv preprint arXiv:2306.17177*.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Zhu, K.; Chen, H.; Yang, L.; Yi, X.; Wang, C.; Wang, Y.; et al. 2023. A survey on evaluation of large language models. *arXiv preprint arXiv:2307.03109*.
- Chen, Z.; Cao, Y.; Yao, H.; Lu, X.; Peng, X.; Mei, H.; and Liu, X. 2021. Emoji-powered sentiment and emotion

⁴<https://openai.com/blog/introducing-gpts>

- detection from software developers' communication data. *TOSEM*.
- Chen, Z.; Lu, X.; Ai, W.; Li, H.; Mei, Q.; and Liu, X. 2018. Through a Gender Lens. *WWW 2018*.
- Chen, Z.; Shen, S.; Hu, Z.; Lu, X.; Mei, Q.; and Liu, X. 2019. Emoji-Powered Representation Learning for Cross-Lingual Sentiment Classification. In *WWW 2019*.
- Cramer, H.; de Juan, P.; and Tetreault, J. 2016. Sender-Intended Functions of Emojis in US Messaging. In *Mobile-HCI 2016*.
- de Janeiro, E. d. R. 2023. Could large language models estimate valence of words? A small ablation study. *Proceedings of CBIC 2023*.
- Felbo, B.; Mislove, A.; Søgaard, A.; Rahwan, I.; and Lehmann, S. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. *EMNLP*.
- Ge, J. 2019. Emoji Sequence Use in Enacting Personal Identity. In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, 426–438. New York, NY, USA: Association for Computing Machinery. ISBN 9781450366755.
- Gilardi, F.; Alizadeh, M.; and Kubli, M. 2023. Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056*.
- Hakami, S. A. A.; Hendley, R. J.; and Smith, P. 2023. ArSarcasMojj Dataset: The Emoji Sentiment Roles in Arabic Ironic Contexts. In *Proceedings of ArabicNLP 2023*, 208–217.
- Hu, T.; Guo, H.; Sun, H.; Nguyen, T.-v.; and Luo, J. 2017. Spice up your chat: the intentions and sentiment effects of using emojis. In *ICWSM 2017*.
- Huang, F.; Kwak, H.; and An, J. 2023. Is chatgpt better than human annotators? potential and limitations of chatgpt in explaining implicit hate speech. *arXiv preprint arXiv:2302.07736*.
- Kocoń, J.; Cichecki, I.; Kaszyca, O.; Kochanek, M.; Szydło, D.; Baran, J.; Bielaniewicz, J.; Gruza, M.; Janz, A.; Kanclerz, K.; et al. 2023. ChatGPT: Jack of all trades, master of none. *Information Fusion*, 101861.
- Leonardi, V. 2022. Communication challenges and transformations in the Digital Era: emoji language and emoji translation. *Language and Semiotic Studies*, 8(3): 22–44.
- Li, G.; Hammoud, H. A. A. K.; Itani, H.; Khizbullin, D.; and Ghanem, B. 2023a. Camel: Communicative agents for "mind" exploration of large scale language model society. *arXiv preprint arXiv:2303.17760*.
- Li, L.; Fan, L.; Atreya, S.; and Hemphill, L. 2023b. "HOT" ChatGPT: The promise of ChatGPT in detecting and discriminating hateful, offensive, and toxic comments on social media. *arXiv preprint arXiv:2304.10619*.
- Liu, X.; Xu, P.; Wu, J.; Yuan, J.; Yang, Y.; Zhou, Y.; Liu, F.; Guan, T.; Wang, H.; Yu, T.; et al. 2024. Large language models and causal inference in collaboration: A comprehensive survey. *arXiv preprint arXiv:2403.09606*.
- Lu, X.; Ai, W.; Chen, Z.; Cao, Y.; and Mei, Q. 2022. Emojis predict dropouts of remote workers: An empirical study of emoji usage on GitHub. *PloS one*.
- Lu, X.; Ai, W.; Liu, X.; Li, Q.; Wang, N.; Huang, G.; and Mei, Q. 2016. Learning from the Ubiquitous Language: An Empirical Analysis of Emoji Usage of Smartphone Users. In *UbiComp 2016*.
- Lu, X.; Cao, Y.; Chen, Z.; and Liu, X. 2018. A first look at emoji usage on github: An empirical study. *arXiv preprint*.
- Lyu, H.; Qi, W.; Wei, Z.; and Luo, J. 2024. Human vs. LMMs: Exploring the Discrepancy in Emoji Interpretation and Usage in Digital Communication. *arXiv preprint arXiv:2401.08212*.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*.
- Ostyakova, L.; Smilga, V.; Petukhova, K.; Molchanova, M.; and Kornev, D. 2023. ChatGPT vs. Crowdsourcing vs. Experts: Annotating Open-Domain Conversations with Speech Functions. In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, 242–254.
- Peng, L.; Wang, Z.; Liu, H.; Wang, Z.; and Shang, J. 2023. EmojILM: Modeling the New Emoji Language. *arXiv preprint arXiv:2311.01751*.
- Rana, R.; and Singhal, R. 2015. Chi-square test and its application in hypothesis testing. *Journal of Primary Care Specialties*, 1(1): 69–71.
- Schick, T.; and Schütze, H. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Tauch, C.; and Kanjo, E. 2016. The Roles of Emojis in Mobile Phone Notifications. In *UbiComp 2016*.
- Törnberg, P. 2023. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Van Hee, C.; Lefever, E.; and Hoste, V. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, 39–50.
- Wang, S. 2022. Sarcastic Meaning of the Slightly Smiling Face Emoji from Chinese Twitter Users: When A Smiling Face Does Not Show Friendliness. *International Journal of Languages, Literature and Linguistics*, 8(2): 65–73.
- Wang, X.; Zhou, Y.; Liu, X.; Lu, H.; Xu, Y.; He, F.; Yoon, J.; Lu, T.; Bertasius, G.; Bansal, M.; Yao, H.; and Huang, F. 2024. Mementos: A Comprehensive Benchmark for Multimodal Large Language Model Reasoning over Image Sequences. *arXiv:2401.10529*.
- Wood-Doughty, Z.; Xu, P.; Liu, X.; and Dredze, M. 2021. Using Noisy Self-Reports to Predict Twitter User Demographics. In *Proceedings of the Ninth International Workshop on Natural Language Processing for Social Media*, 123–137.

Xu, P.; Liu, J.; Jones, N.; Cohen, J.; and Ai, W. 2024. The Promises and Pitfalls of Using Language Models to Measure Instruction Quality in Education. *arXiv preprint arXiv:2404.02444*.

Yang, Z.; Li, L.; Lin, K.; Wang, J.; Lin, C.-C.; Liu, Z.; and Wang, L. 2023. The dawn of Imms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9(1): 1.

Yoo, K. M.; Park, D.; Kang, J.; Lee, S.-W.; and Park, W. 2021. GPT3Mix: Leveraging large-scale language models for text augmentation. *arXiv preprint arXiv:2104.08826*.

Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; et al. 2023. Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models. *arXiv preprint arXiv:2309.01219*.

Zhou, Y.; and Ai, W. 2022. #Emoji: A Study on the Association between Emojis and Hashtags on Twitter. In *ICWSM 2022*.

Zhou, Y.; Lu, X.; Gao, G.; Mei, Q.; and Ai, W. 2023a. Emoji Promotes Developer Participation and Issue Resolution on GitHub. *arXiv preprint arXiv:2308.16360*.

Zhou, Y.; Maharjan, S.; and Liu, B. 2023. Scalable prompt generation for semi-supervised learning with language models. *arXiv preprint arXiv:2302.09236*.

Zhou, Y.; Xu, P.; Liu, X.; An, B.; Ai, W.; and Huang, F. 2023b. Explore Spurious Correlations at the Concept Level in Language Models for Text Classification. *arXiv preprint arXiv:2311.08648*.

Zhu, Y.; Zhang, P.; Haq, E.-U.; Hui, P.; and Tyson, G. 2023. Can chatgpt reproduce human-generated labels? a study of social computing tasks. *arXiv preprint arXiv:2304.10145*.

Ziems, C.; Shaikh, O.; Zhang, Z.; Held, W.; Chen, J.; and Yang, D. 2023. Can large language models transform computational social science? *Computational Linguistics*, 1–53.

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes, see the Abstract and the Introduction**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, see Section 3 to 8**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes, see Section 3 to 8**
- (e) Did you describe the limitations of your work? **Yes, see the Limitations and Implications**
- (f) Did you discuss any potential negative societal impacts of your work? **Yes, see the Limitations and Implications**

- (g) Did you discuss any potential misuse of your work? **Yes, see the Limitations and Implications**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, see the Limitations and Implications**

(i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
- (b) Have you provided justifications for all theoretical results? **NA**
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
- (e) Did you address potential biases or limitations in your theoretical framework? **NA**
- (f) Have you related your theoretical results to the existing literature in social science? **NA**
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**

3. Additionally, if you are including theoretical proofs...

- (a) Did you state the full set of assumptions of all theoretical results? **NA**
- (b) Did you include complete proofs of all theoretical results? **NA**

4. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes, see Section 3 to 8. Our experiment result is all about the in-context learning of ChatGPT, which only needs to call OpenAI API to get the results.**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes, the prompt details are in Appendix Section**
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **No, we fix the temperature to 0, so the results are fixed**
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **No, calling OpenAI API does not cost computation resources**
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **No, for quantitative results, we use accuracy as the evaluation metric**

The emoji is used to represent the bee or snake, rather than analyzing the meaning of the emoji according to the context.

Emoji Sentiment For emoji sentiment, we duplicate the experiment for annotating the sentiment of the same 50 emojis in Section 4 and find that Llama3 labels the same sentiment as human annotators in 41/50 cases, which has a lower accuracy than the ChatGPT annotation (49/50).

Emoji Intention We also repeated the experiments in Section 5.1 using the same prompt to ask the Llama3 model to annotate emoji intentions. However, we found that Llama3 only produced a single token: “\” and then ended the generation. This indicates that Llama3 struggles to follow long instructions with intention rubrics, resulting in meaningless output.

Emoji Usage Associated with Communities To probe whether Llama3 encodes the knowledge of associations between emojis and communities, we utilize the same prompt in Section 6 and ask Llama3 to generate 10 emojis associated with different communities. However, we find that, in addition to the hashtag #TheLastDance, where it generates a basketball emoji 🏀, it produces meaningless hallucinations, such as the code to calculate associations.

Overall, Llama3-8B’s performance falls short compared to human annotators in nearly all emoji-related tasks. The reasons for its shortcomings can be attributed to two main factors. First, as observed in its explanations of emoji sentiment and semantics, Llama3 lacks the necessary knowledge of emojis. Second, due to its limited instruction-following capabilities, Llama3 often generates hallucinations when faced with long or complex prompts. Consequently, we conclude that Llama3-8B is not a suitable replacement for human annotators nor a reliable tool for answering users’ questions about emojis.

3 Prompt Details

Could you assign a sentiment to this emoji: **emoji**? Please select the sentiment from positive, negative, or neutral.

(a) Prompt for probing emoji sentiment without text context.

I will provide you 5 tweets. Please assign a sentiment to each tweet. You can only select the sentiment from positive, negative, or neutral. Please only output the sentiment answer without justification.

Here are some examples:

1. user user what do these '1/2 naked pics' have to do with anything? They're not even like that.
2. I think I may be finally in with the in crowd #mannequinchallenge #grads2014 user
3. user Wow,first Hugo Chavez and now Fidel Castro. Danny Glover, Michael Moore, Oliver Stone, and Sean Penn are running out of heroes.
4. An interesting security vulnerability - albeit not for the everyday car thief
5. Can't wait to try this - Google Earth VR - this stuff really is the future of exploration....

The sentiment output is: 1. neutral 2. positive 3. negative 4. neutral 5. positive

Here is the tweet list:

{tweet}

(b) Prompt for applying ChatGPT to label the sentiment of tweets, and {tweet} is the tweet to label the sentiment.

Table 12: Prompts of sentiment-related tasks for emojis.

Please rate the willingness to use the emoji {emoji} for the following 7 intentions: expressing sentiment, strengthening expression, adjusting tone, expressing humor, expressing irony, expressing intimacy, describing content, on a 7-point scale (7 = totally willing, 1 = not willing at all). The details of the intentions are suggested as below:

{definition details}

Table 13: Prompt for probing the intention of emojis without given the text context. **definition details** can be found in Hu et al. (2017).

Please assign an intention of emoji usage from the following 8 intentions: sentiment expressed, sentiment strengthened, tone adjusted, content described, content organized, content emphasized, non-communication use, symbol. Please only output the intention answer without justification.

The details of the intentions are suggested as below:

{definition details}

Here are some examples:

1. Even the Guardian has TLS now... 🍷
2. merge when we receive 🍷, shipit
3. Thanks for your effort 🍷
4. ✓ Creating directory
5. Remote Get: ✓

The intention output is: 1. sentiment expressed 2. symbol 3. sentiment strengthened 4. non-communication use 5. content described

Here is the Github post list: {post}

Table 14: Prompt for applying ChatGPT to label the intention of emojis in GitHub posts. {post} is the GitHub to label the emoji intention. definition details can be found in Lu et al. (2018).

I will provide you 5 tweets. Please tell me whether the tweets are ironic or not. Please only output the irony label without justification. Here are some examples:

1. Such . You still have to #praisehim ;)
2. Hey heyy!!!I...wanna be a rockstar #vsocam hero #vsocam #hero #spiderman
3. user nice to see the ambulance service is so important to OUR mps
4. I love finals week! #justkidding #stressed
5. Literally cried when I woke up because I know what this day has in store for me #TheStartOfTechWeek Ready #JustShootMeKnow

The irony output is: 1. ironic 2. not ironic 3. ironic 4. ironic 5. not ironic

Here is the tweet list: {tweet}

Table 15: Prompt for applying ChatGPT to label the irony or not of a tweet. {tweet} is the tweet to label the irony.

I will provide you 5 tweets with emojis in Arabic language. Please tell me which type of irony the emojis in tweets convey. You can only select the irony type from humour, sarcasm, and no irony. Here is the definition of sarcasm and humour: Sarcasm, a specific facet of irony, employs language with a sharp, often bitter tone to mock or critique, often with exaggerated emphasis. Humour, while encompassing elements of irony, specifically denotes the ability to evoke amusement or laughter. Please only output the irony label without justification. Here are some examples:

1. {Arabic tweet example 1}
2. {Arabic tweet example 2}
3. {Arabic tweet example 3}
4. {Arabic tweet example 4}
5. {Arabic tweet example 5}

The irony output is: 1. no irony 2. no irony 3. sarcasm 4. sarcasm 5. humour

Here is the tweet list: {tweet}

Table 16: Prompt for applying ChatGPT to label the irony of emojis in Arabic tweets. {Arabic tweet example n} is the n-th example of an Arabic tweet that contains emojis with the ground-truth irony label for the emoji. {tweet} is the tweet that contains emojis to label the irony.

I will provide you 5 tweets. Please recommend only one emoji for each tweet. You can only select emojis from these emoji candidates: {emoji candidates}. Please only output the emoji answer without justification. Here are some examples:

1. Sunday afternoon walking through Venice in the sun with @user
 2. Time for some BBQ and whiskey libations. Chomp, belch, chomp! (@user)
 3. Man these are the funniest kids ever!! That face! #HappyBirthdayBubb @user Xtreme
 4. #sandiego @San Diego, California
 5. My little #ObsessedWithMyDog @user Capitol Hill
- The emoji recommendation is: 1. 🌞 2. 😊 3. 🤔 4. 🇺🇸 5. ❤️

Here is the tweet list: {tweet}

Table 17: Prompt for applying ChatGPT to predict the emoji in tweets. {tweet} is the tweet to predict the emojis. {emoji candidates} represents the candidate emojis that ChatGPT should select from and the details of candidate emojis can be found in Barbieri et al. (2018).