Instagram Reactions to a Virtual Dining Companion: Qualitative Coding vs. Automated Sentiment Analysis

Hunter Fong
University of Genoa
Genoa, Italy
hunter.fong@edu.unige.it

Radoslaw Niewiadomski
University of Genoa
Genoa, Italy
radoslaw.niewiadomski@unige.it

Magdalena Igras-Cybulska
Interactive Crafts sp. z o.o.
Lublin, Poland
magda.igras@gmail.com

Abstract

Technologies designed to support solo eaters by providing them company are an emerging design space at the intersection of AI, affective computing, and everyday dining rituals. This study explores public reactions to such novel systems by analyzing Instagram comments on a viral post depicting a man dining with a virtual reality partner. Using a mixed-methods approach, we first applied automated sentiment tools—including TextBlob, VADER, NRCLex, and BERT—to quantify surface-level sentiment. However, these analyses failed to capture the layered, performative, and culturally nuanced nature of user responses. To address this, we conducted manual thematic coding of 719 comments, revealing a wide spectrum of emotional tones—from sarcasm and judgment to empathy, curiosity, and self-reflection. Our findings demonstrate that many seemingly "positive" comments were in fact mocking or ambivalent, and that public discourse about emerging technologies like virtual companionship is deeply shaped by generational expression, meme culture, and shifting social norms. Thus, social media can be a valuable source of opinions and preferences regarding novel interactive technologies—even speculative prototypes—but the collected data needs to be analyzed on multiple levels.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Computing methodologies** → *Natural language processing*.

Keywords

virtual companions, sentiment analysis, qualitative coding, HCI, social media, VR, emoji, interpretive analysis, social perception, virtual agents

ACM Reference Format:
Hunter Fong, Radoslaw Niewiadomski, and Magdalena Igras-Cybulska. 2025. Instagram Reactions to a Virtual Dining Companion: Qualitative Coding vs. Automated Sentiment Analysis. In *CHItaly 2025: 16th Biannual Conference of the Italian SIGCHI Chapter (CHItaly 2025), October 06–10, 2025, Salerno, Italy*. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3750069.3750352

1 Introduction

The rapid advancement of artificial intelligence (AI) and immersive technologies—including Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR)—has enabled the development of virtual characters designed to simulate human interactions in increasingly sophisticated ways. These systems are now being explored for their potential to support emotional well-being, reduce loneliness, and enhance everyday experiences, social skills, productivity, and entertainment [11, 29, 38, 40]. Designers and researchers alike often frame these technologies with optimistic expectations—hoping they will offer comfort, emotional support, or social scaffolding in contexts where human interaction is limited.

In particular, the PRIN 2022 Project, COmputational Models of COmmensality for artificial Agents (COCOA)[1], responds to this challenge by investigating different interactive technologies that support the commensality experience. The project activities range from defining theoretical models to real-time interactive systems that simulate or scaffold commensal experience. One specific application area is that of Artificial Commensal Companions (ACCs) [37], which are designed to provide social presence during meals. These may take the form of virtual characters [31], or physically embodied robots [17]. Prior studies have investigated public expectations around such systems, often revealing ambivalent attitudes: while users may express skepticism in theoretical scenarios [21], their in-situ responses can be more open, or even enthusiastic [16]. Understanding these responses is crucial—not only to refine the design of companion systems, but also to anticipate the sociocultural dynamics that shape their adoption.

In this study, we explore public reactions to such a novel system by analyzing Instagram comments on a post depicting a man dining with a virtual reality partner. This methodological choice has at least two advantages: first, exploring social media allows us to collect a relatively large number of spontaneous reactions towards technology that might not even be realized (e.g., so-called speculative prototypes), providing important insights for the subsequent development phase. It is in line with previous studies that exploit social media—often on a larger scale—to analyze attitudes towards emergent technologies such as delivery robots [27], VR systems in healthcare [25] or generative AI [42]. Second, while it is hard to argue that social media users are the perfect representation of the entire population, they represent a more diverse group than the student population typically used in many studies. Our expectation is that it allows us to learn the perspectives of a broader group, and is appropriate for an exploratory study such as this. Furthermore, social media platforms such as Instagram are expected to contain comments that are more spontaneous, authentic and less constrained by social regulations [26].

The specific Instagram post used in this study presented a particularly rich opportunity for public perception analysis because

[1] https://cocoa-research.github.io/

ACM ISBN 979-8-4007-2102-1/25/10
https://doi.org/10.1145/3750069.3750352

of its context: solo dining with a virtual partner. Eating alone is often associated with stigma, isolation, or perceived lack of social support [5], and the act of publicly dining with a virtual companion directly challenges existing social norms around commensality—the shared practice of eating together [20, 45]. Prior work on ACCs has explored their potential to reduce loneliness and support emotional well-being during meals [32], making this post a compelling and ecologically valid case study.

We began our analyses by applying a suite of automated tools for sentiment and emotion detection, including VADER, TextBlob, NRCLex, and BERT-based transformer models, along with part-of-speech tagging (SpaCy), emoji/GIF frequency analysis, and lexical emotion detection. Initial outputs suggested a broadly neutral-to-positive reception. However, based on prior research [1, 4, 12], we anticipated that these tools might misclassify culturally coded or emoji-heavy commentary—particularly on platforms like Instagram, where communication is often performative, ironic, and stylized. To address this, we applied both automated and manual methods in parallel, using qualitative coding to interpret the kinds of rhetorical nuance that surface-level tools frequently miss. This mixed-methods approach reflects a broader challenge in HCI and social computing: understanding how people react to emerging technologies often requires culturally fluent, interpretive analysis, especially in domains involving identity, intimacy, and ritual—such as eating.

In more detail, we manually coded all comments using a 25-category scheme developed inductively through thematic analysis. This process revealed a spectrum of reactions—ranging from confusion, disgust, and judgment to admiration, curiosity, and relational self-reflection. Many users used humor, performance, and irony not to obscure their feelings, but to *frame them socially*. The most common themes were not celebration or support, but disapproval in its many forms—especially mocking and judgmental commentary. As will be discussed in more detail later, the large majority of this derision is aimed specifically at the user, and not at the new technology itself.

The rest of the paper is structured as follows. We begin by reviewing prior work on sentiment analysis, qualitative methods in social media research, and cultural/platform-specific interpretation. We then describe the dataset, outline our mixed-methods procedure, and present results from both automated and manual analyses. Finally, we discuss the social implications of our findings, limitations of the current approach, and design recommendations for those developing or studying ACCs.

2 Related Work

2.1 Sentiment Analysis and Emotion Detection

Automated sentiment analysis has become a foundational method in social media research, human-computer interaction, and marketing. Models like VADER [22], TextBlob [30], and NRCLex [35] are frequently used to evaluate public attitudes, emotional responses, and social trends at scale. Studies have applied these tools to analyze sentiment in product reviews [36], political discourse [46], chatbot interactions [28], and platform-specific phenomena like TikTok reactions [23].

In the context of Human-Computer Interaction (HCI) and emerging technologies, sentiment models are often used as proxies for user acceptance or cultural legitimacy. These tools offer speed and scalability, but their interpretive power is limited—especially in informal, emoji-rich, or sarcastic discourse [1, 19]. Studies have shown that models tend to overclassify laughter or emoji-heavy comments as positive, even when they convey critique, mockery, or irony [10, 39], even when leveraging feature-engineering techniques tailored for sarcasm detection on platforms like Twitter [9]. While efforts have been made to improve sarcasm detection [14, 24], most systems still struggle with short, highly contextual, or culturally-coded inputs.

Social media sentiment analysis must grapple with linguistic nuances such as sarcasm, memes, and emoji usage. Content that is sarcastic or ironic can invert the expressed sentiment, posing challenges for traditional sentiment classifiers. Recent research emphasizes that sarcasm – a "significant aspect of human sentiment" – is often overlooked in sentiment analysis pipelines [13]. Users often express sentiment via culturally coded humor rather than literal statements. These findings suggest that sentiment analysis of public opinion must go beyond literal text, accounting for contextual and performative aspects of social media expression. Advanced multimodal sentiment analysis techniques are beginning to do so – for instance by combining vision and NLP to interpret meme content [8, 43] – but this remains an open research frontier.

2.2 Qualitative and Mixed-Methods Approaches to Attitudinal Research

Studying public perceptions of emerging technologies benefits from mixed-methods approaches that combine large-scale data analysis with in-depth qualitative insights. One common strategy is to analyze social media discourse as a window into public opinion, then complement those findings with interviews or surveys. For example, Lee and Toombs mined Twitter and Instagram posts from a university campus to understand students' reactions to delivery robots operating on campus pathways [27]. By qualitatively coding the content of these posts (e.g. identifying themes of excitement, fear, or humor) and quantifying their prevalence, they could gauge overall sentiment and specific concerns people voiced about the robots. This kind of social media content analysis, when paired with follow-up user interviews, can reveal not just what people say about a technology in public, but *why* they feel that way. Indeed, mixed-methods research has been advocated in HCI and social computing as a way to capture both the breadth of public opinion (via big data analytics) and the depth of individual experiences (via qualitative inquiry). Recent studies illustrate this approach clearly. Miyazaki et al. analyzed 3 million tweets about generative AI (e.g., ChatGPT), finding generally positive sentiment with notable exceptions, such as negative reactions from artists concerned about misuse of their work [34].

In HCI research, qualitative coding is widely used to capture subtleties of social judgment, humor, irony, and emotional ambivalence that quantitative tools often miss [33, 44]. These methods are especially useful when responses are stylized or performative, as is common on platforms like Instagram. By combining automated sentiment pipelines with manual thematic analysis, our study aims to reveal not just what people say about virtual companions—but how, why, and with what social performance techniques they express it.

2.3 Cultural and Platform-Specific Expressions

Users' reactions to technology are deeply influenced by the cultural codes and platform norms of the social media environments in which those reactions occur. What people choose to express – and how they express it – can vary dramatically between, say, TikTok and LinkedIn, or between one community and another. Prior work has shown that each social platform cultivates its own "vernacular" of communication. TikTok, for example, encourages a norm of candid and playful expression. Barta and Andalibi's interview study of TikTok users found that the platform's ethos of "fun" and authenticity (encouraging users to "just be you") actually normalizes sharing of both positive and negative experiences in a supportive atmosphere [3]. This suggests that on TikTok, users might respond to a new AI gadget or companion by creating humorous, personal videos – perhaps meme-filled or emotionally frank – because the culture there rewards relatable authenticity.

In contrast, platforms like Instagram historically emphasize polished aesthetics and positivity [6]. Although recent trends have moved toward more "authentic" content on Instagram as well, users may still be less inclined to post overt criticisms or raw emotions there, compared to TikTok. Such platform norms influence how technologies are performatively received: a quirky home robot might become the subject of affectionate parody on TikTok, yet the same robot might be presented in a carefully filtered photo on Instagram with a witty caption, aligning with each platform's style.

Cultural context also plays a key role in technology perception. Norms, values, and humor are not universal – an ironic joke about an AI helper that resonates in one culture or language might fall flat (or be misinterpreted) in another. Researchers have observed cross-cultural differences in attitudes toward social robots [7] and chatbots [15], often mediated by familiarity and societal narratives about AI. For instance, Japanese popular culture, which includes friendly robot characters, may prime Japanese users to welcome companion robots, whereas Western media's frequent portrayal of rogue AI could engender skepticism elsewhere [47]. These cultural frames can manifest in social media reactions: Eastern audiences might use cute or respectful language for an AI companion, while Western users might respond with sarcasm or dystopian memes. Understanding these subtleties is crucial. As Sharma et al. note [43], online expressions (like memes) are often "derived from prior social and cultural experiences", and thus decoding a community's reaction to technology requires fluency in its cultural references.

In practice, this means that studies of user sentiment must interpret content in light of both the platform's performative conventions and the community's cultural background. Recognizing an emoji-based joke or a viral meme format specific to a platform can be the difference between misreading user sentiment and truly grasping the public's nuanced reaction to a technology. The related work in this area encourages designers and analysts to be sensitive to these context factors – for example, tailoring an AI companion's persona to align with local social norms, or choosing the right social media channels to engage with users. By accounting for cultural and platform-specific modes of expression, we can better interpret and predict the social reception of new interactive technologies.

3 Materials and Methods

3.1 The Case Study Post and Dataset

Our dataset was derived from an Instagram Reel posted on July 25, 2024 by the account @thebestoftiktok, which features viral content aggregated from TikTok and elsewhere. The video — available at https://www.instagram.com/reel/C900KFNOH5l — shows a man dining alone at a sushi restaurant while wearing a mixed reality (MR) headset, appearing to interact with a someone else. The caption reads: "The Guy brought his virtual girlfriend out for sushi…" The man actively interacts with a virtual companion unseen on the roll, gesturing, smiling, and enthusiastically conversing, as if fully engaged in a shared dining experience. At one point, he even pretends to feed the invisible figure with chopsticks, vividly reinforcing the presence of the virtual companion. The table in front of him contains sushi dishes and drinks. Other customers seated nearby can clearly observe his actions.

As of April 14, 2025, the post had received over 65,000 likes and 922 comments, making it a high-engagement example of public response to social VR. For analysis, we manually reviewed and cleaned the comment dataset, removing duplicates, non-English comments with no translation, and replies that contained only tagging or minimal interpersonal exchange. The final dataset included 719 unique comments, of which 235 were replies.

This dataset provides a naturalistic opportunity to examine social reactions to virtual companionship in the context of dining—an activity traditionally marked by co-presence and social bonding. It also reflects wider cultural conversations about AI, intimacy, and performative identity online.

3.2 Automated Quantitative Analysis

To establish a quantitative baseline, we applied a suite of natural language processing tools:

- **TextBlob:** provided sentence-level polarity (from -1.0 to +1.0) and subjectivity (from 0.0 to 1.0), capturing both valence and personal tone.
- **VADER (Valence Aware Dictionary and sEntiment Reasoner):** tailored for social media, VADER returned proportions of negative, neutral, and positive sentiment, along with a compound score for overall emotional intensity.
- **NRCLex:** counted words associated with basic emotions (e.g., joy, anger, sadness), based on the NRC Emotion Lexicon.
- **Transformer-based Emotion Models (BERT):** used context-aware models to assign probabilistic emotion scores (e.g., joy: 0.56, sadness: 0.12).
- **Part-of-Speech Tagging (SpaCy):** identified the most frequent nouns, verbs, and adjectives.
- **Emoji and GIF Analysis:** recorded the frequency and type of emojis and GIFs, including their presumed emotional valence.

These tools were chosen for their relevance to short-form, social-media text and their capacity for large-scale processing. Lexicon-based systems such as VADER are widely used in HCI, marketing, and social media research to provide quick sentiment estimates at scale [19, 36, 46]. However, they are known to overestimate positivity in the presence of emojis or informal expressions—particularly in sarcastic or culturally coded contexts [10, 14, 24, 39]. Because these tools remain the default in many computational pipelines, we applied them as a baseline. We also conducted a supplemental analysis using a RoBERTa-based sentiment model trained on social media text, which is known to better handle contextual cues and irony [18, 41]. This comparison is discussed in the Results section.

3.3 Qualitative Coding Procedure

To uncover emotional and social nuance missed by automated methods, we conducted a full manual coding of the dataset using a 25-category scheme developed inductively through thematic analysis. These categories were later clustered into four broader groups for interpretability: **Approval**, **Disapproval**, **Interpretive/Ambiguous**, and **Meta/Self-Referential Commentary**. Key subcategories included "Disapproval – Mocking," "Romantic Self-Projection," and "Gendered Resentment," which emerged as especially frequent and socially meaningful. The original scheme was created through an iterative open coding process guided by grounded theory principles. One author conducted several close readings of the entire dataset, initially tagging comments with provisional, descriptive labels such as "lonely," "judgy," "making a joke," or "feels bad for him." As the coding progressed, these initial codes were refined through constant comparison—grouping related concepts, merging overlapping labels, and splitting ambiguous ones based on rhetorical function or affective tone. For example, an early category labeled "emotional reaction" was later divided into more specific categories such as "Pity," "Sympathy," and "Bravery," depending on whether the comment expressed sadness, support, or admiration. Conversely, the category "Mocking" was initially too broad, and later split into distinct forms such as "Judgment" (moral evaluation) and "Romantic Self-Projection" (using humor to project one's own relationship frustrations).

Some categories, such as "Focus on Social Constructs," were created to capture comments that did not directly judge the person or the technology, but instead reflected on societal implications or norms (e.g., "This is what the world is coming to" or "We've gone too far with tech"). This interpretive distinction was important for identifying layered commentary that blends personal opinion with broader cultural critique.

The final scheme was developed and refined over multiple coding passes. Each comment was then reviewed again using the finalized 25-category set. Coding was binary and non-exclusive, meaning each comment could be tagged with zero or more categories depending on its rhetorical structure and affective tone. The full category list and frequencies are presented in Table 1.

While this coding scheme was tailored to this dataset, many categories (e.g., 'Mocking,' 'Judgment,' 'Pity,' 'Confusion') reflect recurring patterns in public responses to other emergent or controversial interactive technologies and may be generalizable to other social media platforms or case studies.

Table 1: Frequency of manually coded comment themes (N=719).

Theme	Frequency
Approval – Bravery	24
Approval – Defensive	61
Approval – Jealousy	20
Approval – Live & Let Live	105
Approval – Sympathy	36
Disapproval – Bad for Society	113
Disapproval – Concern for Mental Health	89
Disapproval – Disgust	70
Disapproval – Judgment	411
Disapproval – Mocking	389
Disapproval – Pity	150
Focus on Social Constructs	74
Focus on bystander	19
Gendered Resentment	76
Lack of Understanding or Confusion	31
Objective Judgment	5
Pop Culture Reference	11
Privacy Concern	1
Relate to Self	43
Romantic Self-Projection	27
Suspected Spam	1
Thought it was VR	2
Thought it was a Real Person	5
Thought Post was Spam	5
Vague Curiosity or Interest	19

4 Results

4.1 Automated Sentiment Analysis

We first applied a suite of natural language processing tools (VADER, TextBlob, NRCLex, and transformer-based emotion models) to our dataset. The results of our automated quantitative analysis, including sentiment distribution and emotional classification by tool, are summarized in Table 2. Below, we discuss some specific examples in more detail:

- **VADER** classified comments as predominantly neutral. Specifically, 71.9% of comments were labeled as *neutral*, 14.6% as *negative*, and 13.6% as *positive*. Despite the apparently balanced ratio of positive to negative posts, a closer look at the classified posts reveals a tendency to overestimate the positivity of certain posts. For example, comments consisting only of "💀💀💀" or "😂" were interpreted by VADER as positive, despite clearly signaling mockery in context.
- **TextBlob** calculated polarity and subjectivity scores for each comment, yielding an average polarity of 0.025 and an average subjectivity of 0.273. These scores suggest a neutral-to-slightly-positive tone with relatively low personal expression.
- **RoBERTa-based sentiment classification** was also conducted using a transformer model fine-tuned for social media sentiment [2]. While more context-sensitive than VADER or TextBlob, RoBERTa also fails to recognize mocking tone or

sarcastic content. For instance, comments like “he’s braver than the troops” or “finally a girl who won’t leave you” were labeled as positive, despite clearly conveying social critique or mockery. Overall, RoBERTa sentiment scores aligned more closely with VADER than with our manual codes, highlighting the limitations of even advanced models in interpreting stylized, context-rich online discourse.

- **NRCLex** analysis showed that the most frequently detected emotions were *positive* (21.2%), *trust* (14.0%), and *joy* (10.2%). Negative emotions such as *anger* (5.4%) and *disgust* (5.2%) were significantly underrepresented given the tone of the dataset.
- **Part-of-speech (POS) tagging**, using SpaCy, revealed that most comments were noun- or emoji-heavy, offering little syntactic structure to support deeper inference.
- **Emoji and GIF sentiment analysis** interpreted visual shorthand (e.g., “cry laughing”) as joyful, despite frequent use in mocking or dismissive contexts.

Table 2: Automated sentiment analysis summary (N=719).

Tool/Model	Results
VADER	Positive: 13.6%, Neutral: 71.9%, Negative: 14.6%
TextBlob	Polarity: 0.025 (scale -1 to 1); Subjectivity: 0.273 (0 to 1)
NRCLex (top emotions)	Positive: 21.2%, Trust: 14.0%, Joy: 10.2%, Anger: 5.4%, Disgust: 5.2%
RoBERTa Sentiment	Mostly Neutral; Sarcasm frequently misclassified as positive or neutral
Emoji Analysis	Most frequent emojis: 😂, 💀; Misclassified as joyful (actual meaning: sarcasm, mockery)
POS Tagging (SpaCy)	Frequent nouns: girl, life, bill; adjectives: virtual, real, brave

4.2 Limitations of Automated Tools

To assess the performance of automated sentiment tools, we manually coded ten emoji-heavy comments that exemplify the dataset’s most common expressive forms. As shown in Table 3, VADER frequently interpreted sarcastic or mocking comments as *positive*, particularly when composed primarily of laughter emojis (e.g., "😂😂😂😂😂😂😂"). RoBERTa, though more context-sensitive, defaulted to *neutral* sentiment in 7 of the 10 cases. In contrast, manual coding identified frequent use of mocking, judgment, gendered resentment, and relational self-projection—none of which were detectable through lexicon-based sentiment classification alone.

These examples were not randomly selected; instead, they were chosen to typify high-frequency comment patterns. Our goal was to show concrete mismatches between sentiment models and manual coding, not to quantify the accuracy of the model, but to demonstrate interpretive failure in key rhetorical forms.

This comparison illustrates a broader pattern: while automated tools may handle full sentences with clear affective markers, they struggle with short, image-heavy, and socially coded expressions.

Table 3: Sentiment classification for 10 illustrative emoji-rich comments. These examples were purposefully selected to reflect the most common expressive forms in the dataset and highlight where automated sentiment models (VADER, RoBERTa) diverged from manual interpretation.

Comment (shortened)	Manual	VADER	RoBERTa
"😄"	Mocking	Positive	Neutral
“@user maybe he doesn’t want to deal with real girls and lose his peaceful life 😂"	Gendered Resentment	Neutral	Negative
😂😂😂😂😂😂😂😂😂😂😂	Mocking	Positive	Neutral
"@user😅 💀 "	Mocking	Neutral	Neutral
“He is going to ask her to split the bill 😂"	Judgment + Mocking	Neutral	Neutral
“She paid with her virtual card 😂"	Mocking	Neutral	Neutral
“This is so cringe 😭"	Judgment	Negative	Negative
😭 🤦‍♀️	Judgment	Negative	Negative
"@user 😂😂😂"	Mocking	Positive	Neutral
"@user 😭😭"	Judgment	Negative	Neutral

4.3 Manual Coding Results

Manual coding was conducted on all 719 Instagram comments using a 25-theme scheme developed inductively through close thematic reading. The most frequent themes were overwhelmingly disapproving in tone.

Disapproving Themes. The most common categories conveyed social judgment, mockery, and relational discomfort. **Disapproval – Judgment** was the most frequent (411 comments), followed by **Disapproval – Mocking** (389), and **Disapproval – Pity** (150). Other prominent themes included **Disapproval – Bad for Society** (113), **Disapproval – Concern for Mental Health** (89), and **Disapproval – Disgust** (70). These categories often overlapped within the same comment, expressing layered disapproval—for instance, mocking someone while also expressing pity or concern for their mental state.

Several relational and gendered sub-themes also emerged. **Gendered Resentment** (76 comments) involved references to dating dynamics or generalized frustration toward women (e.g., “at least she won’t nag”). **Romantic Self-Projection** (27 comments) captured projection of personal relationship histories, often masked as humor (e.g., “at least she won’t leave him like my ex did”).

Ambiguous and Mixed-Valence Themes. A smaller set of categories reflected interpretive ambiguity or socially performative stance-taking. **Lack of Understanding or Confusion** (31 comments) and **Vague Curiosity or Interest** (19) captured moments where users were unsure how to interpret the post—or tentatively intrigued. **Focus on Social Constructs** (74) reflected responses framed not at the individual, but at larger themes like loneliness, gender roles, or societal trends.

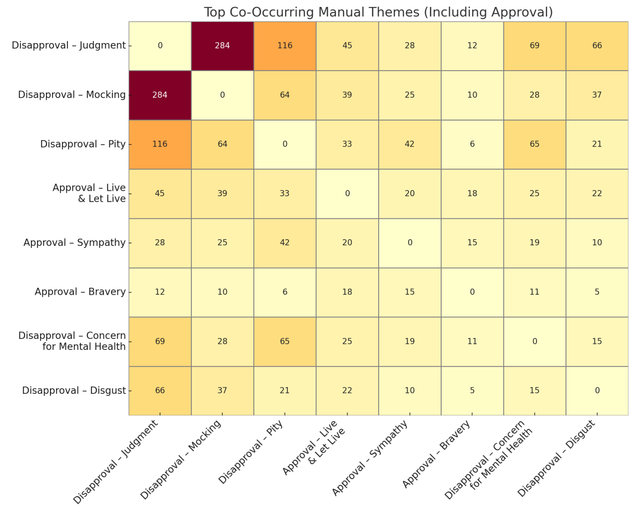

Figure 1: Top co-occurring manual coding themes, including both disapproval and approval categories (N=719). Cell values represent the number of comments coded with both themes. Diagonal values are zeroed for readability. This visualization highlights how certain emotional tones, such as "Mocking" and "Judgment" or "Pity" and "Sympathy," frequently co-occur within individual comments, reflecting the layered nature of public reactions.

Approving Themes. Clear positive sentiment was rare. **Approval – Live & Let Live** appeared in 105 comments and often signaled tolerance rather than enthusiasm (e.g., "he's not hurting anyone"). Other approving themes were even less frequent: **Approval – Defensive** (61), **Approval – Sympathy** (36), **Approval – Bravery** (24), and **Approval – Jealousy** (20). These tended to be performative, ironic, or layered—for instance, expressing support "in theory" while signaling discomfort with actual use.

Each comment in our dataset could be assigned to multiple themes, reflecting its tone, rhetorical structure, and layered emotional stance. The co-occurrence matrix in Figure 1 captures how often pairs of themes were coded within the same comment. This reveals not only expected pairings—such as "Mocking" and "Judgment" (284 co-occurrences) or "Pity" and "Concern for Mental Health" (65)—but also more complex intersections across valence. Some comments, for example, combined disapproval with reluctant empathy, or framed curiosity in the language of critique. These patterns suggest that public responses are rarely flat in tone; instead, they compress contradiction, irony, and ambivalence into the dense expressive forms of social media.

In sum, the manual coding revealed that while automated tools interpreted most comments as neutral or mildly positive, the dominant affective tone of the dataset was disapproval, judgment, and emotionally charged projection. These findings support the need for interpretive frameworks that capture intention, ambiguity, and cultural meaning beyond the reach of sentiment lexicons.

Quantifying Interpretive Divergence. To better understand the gap between automated and manual analysis, we compared sentiment labels across methods. Based on the final manual coding scheme, 549 of the 719 comments (76.4%) included one or more disapproval-related themes, 149 (20.7%) included one or more approval-related themes, and only 29 (4.0%) were coded as interpretive, ambiguous, or neutral. In contrast, VADER labeled 71.9% of all comments as *neutral*, 14.6% as *negative*, and 13.6% as *positive*. Of the 549 manually disapproval-coded comments, 387 (70.5%) were labeled *neutral* by VADER. Only 101 were correctly labeled as negative. This divergence reveals a major interpretive gap: while three-quarters of comments expressed some form of disapproval in manual coding, sentiment tools flattened most reactions into neutrality. It also underscores the challenge of interpreting stylized, sarcastic, or multi-layered online commentary using off-the-shelf sentiment tools.

Moreover, only 113 of the 549 disapproving comments (20.6%) focused on the technology itself; the remaining 80% directed judgment, sarcasm, or ridicule toward the user.

These findings underscore the limitations of lexicon-based sentiment classifiers in detecting social critique, sarcasm, or stylized negativity—especially in the performative and image-rich medium of Instagram. While neither method offers "ground truth," the contrast between them highlights how platform fluency, irony, and affective complexity challenge computational approaches.

4.4 Illustrative Examples of Divergent Interpretation

To further demonstrate where automated sentiment models and human interpretation diverge, we present examples from two representative comment clusters. These clusters were selected to reflect common misalignments in tone interpretation, rather than isolated edge cases.

Cluster 1: Sarcasm Labeled as Positive or Neutral

- "Finally a girl who won't leave you 😭😭" — Manual code: Mocking + Romantic Self-Projection; VADER: Positive.
- "He's braver than the troops 💀💀💀" — Manual code: Mocking; VADER: Positive.

Cluster 2: Sympathy Labeled as Neutral

- "Maybe this helps with his loneliness. Let him be." — Manual code: Sympathy + Live and Let Live; VADER: Neutral.
- "As long as he's happy, who are we to judge?" — Manual code: Approval – Defensive; VADER: Neutral.

These examples illustrate how even subtle shifts in tone or emoji use can lead to divergent interpretations—especially when irony, generational expression, or social framing are involved. In such cases, human coders with platform fluency are better equipped to capture meaning than automated models.

Taken together, these quantitative and qualitative comparisons reveal the limitations of standard sentiment analysis in this domain—not because the tools fail entirely, but because the nature of expression on social media complicates the very notion of emotional valence. This challenges us to rethink how we measure "public perception," especially in relation to socially embedded technologies like ACCs.

5 Discussion

Our manual coding revealed that public responses to the post presenting a person eating with a virtual companion were complex and layered—combining disapproval, relational frustration, humor, and subtle forms of curiosity. Emojis such as "😂" and "💀" were frequently used to signal mockery or judgment but were misread as joyful or neutral by sentiment models. This aligns with prior work showing that lexicon-based tools like VADER and transformer models often misclassify sarcastic, emoji-rich content as positive or neutral, missing the performative and culturally coded nature of commentary on Instagram, where emoticons often signal irony or mockery rather than literal positivity [1, 4, 12].

These findings reinforce a broader insight: online commentary is often performative. On platforms like Instagram, users may adopt ironic, sarcastic, or hyperbolic tones not only to express opinion but to manage identity and audience perception. Consequently, the apparent tone of public comments—especially those relying on meme logic or emoji shorthand—may obscure genuine interest or acceptance.

Popular media may contribute to this disconnect. Dystopian portrayals of AI—as emotionless, threatening, or manipulative—have shaped public imagination long before users encounter such systems in everyday life. These cultural scripts frame emerging technologies like ACCs as inherently suspicious, making public mockery more socially acceptable than tentative curiosity or admiration. And yet, despite this resistance, subtle signs of acceptance are present. Themes such as "Approval – Live & Let Live," "Sympathy," and "Bravery," though infrequent, suggest that some users are already reframing their responses. This aligns with broader adoption patterns seen across technologies, where early adopters and trend-setters—especially those fluent in platform culture—help reshape public perception over time. We argue that public mockery may obscure private curiosity. Understanding the future of commensal technologies requires interpreting these subtle signals of engagement—not just the loudest or most ironic reactions.

At the same time, we admit, that differently from our expectations, the commenters tended to focus more on the users of the technology rather than the technology itself, often through a lens of social stigmatization and cultural stereotypes. Thus, having users free to post their own comments instead of using more focused research tools such as questionnaires, we gain only limited knowledge about the potential benefits of this technology, as well as only a few suggestions for their developers. This is a particularly interesting observation: although the post highlights a novel and emerging use of technology, observers tend to focus primarily on the people using it. This might be specific to the nature of the social media platform used in this study. We can only hypothesize that the same post, if published on a more specialized or technology-oriented social network, would attract a different kind of comments. This observation complements our previous studies on the potential risks and concerns associated with this technology [21], in which respondents indicated social stigmatization and threat of negative perception by others as main concerns.

Nevertheless, certain comments address the technology itself, and the potential it represents. For example: "can't really blame him, im sure his virtual partner offers something humans can't...", "He may be onto something......", and " Crazy no! ? It's Time!"

6 Conclusion

In this study, we explore public reactions to a system that allows the person to eat with a virtual partner in public space by analyzing Instagram comments. While automated tools suggested a broadly neutral-to-positive sentiment profile, our qualitative analysis revealed a far more nuanced reception. Public responses combined disapproval, confusion, self-reflection, and ambivalence—often layered in ironic or socially performative ways. These findings caution against interpreting social media sentiment at face value, particularly in spaces where humor and judgment co-exist. A minority of users expressed sympathy, admiration, or even latent curiosity—though often framed through defensiveness or meme logic. While this study centers on a single instance of virtual commensality, its insights speak to the specific social and emotional terrain that ACCs must navigate.

Our work contributes to research in HCI, AI-mediated interaction, and computational social science by combining standard sentiment analysis pipelines with detailed, qualitative coding. In doing so, we demonstrate the value of layered interpretation in design research and propose practical design guidelines for those building socially embedded AI technologies. To sum up, the social media are a valuable source of information on the social perception of novel interactive technologies and their potential users, but analyses need to be performed on multiple levels. Our framework thus represents a methodological contribution: an interpretable, socially grounded alternative to conventional sentiment labels, tailored for the complexities of online commentary.

6.1 Limitations and Future Work

This study has several limitations. First, we did not have access to demographic information about the commenters, which limits our ability to analyze reactions by age, gender, cultural background, or social identity. Instagram comments are anonymous, stylized, and ephemeral, and interpretation depends heavily on platform-specific conventions and audience dynamics.

Second, the dataset centers on a single viral post. While its popularity and visual clarity made it a compelling case study, its uniqueness may limit generalizability to other types of virtual companionship. It also introduces potential framing effects: the post's caption explicitly identifies the virtual partner as a "girlfriend," a relational framing that may have elicited stronger gendered or moral judgments than other labels might have (e.g., "a long-distance partner," "a companion," or "a deceased relative"). Future work could explore how altering textual framing—while holding video content constant—affects public response.

Third, all manual coding was conducted by a single author with American, millennial cultural fluency. While this enabled informed interpretation of emoji and meme-driven content, it introduces subjectivity. A multi-coder or cross-cultural team could yield additional nuance and inter-rater reliability.

Fourth, our analysis does not account for variation in the form or modality of the companion. Prior studies have shown that reactions to artificial agents often depend on factors such as embodiment, gender presentation, voice, and relational framing. Because the virtual companion in this case was not directly visible or audible, commenters likely projected their own assumptions onto the situation—some imagining an anime girlfriend, others a technological surrogate for a partner. Future work should explore how perceptions differ across modalities (e.g., chatbots, avatars, robots) and how impersonation or relational framing affects emotional response.

Last but not least, the analysis is limited to a single type of social network, which has a strong identity that shapes its users, their behavior, and the roles they assume. Therefore, these results cannot be generalized to the entire population or to other contexts. We would expect different outcomes if a different medium were used (such as Linkedin).

Despite these limitations, the emotional and rhetorical patterns that emerged were consistent across hundreds of comments, and they provide important insight into how commensal technologies are judged, mocked, or cautiously explored in public discourse.

In future work, we plan to generate a series of controlled posts varying key contextual elements—such as the eater's gender, social setting, or relational framing of the virtual character (e.g., friend, partner, relative) as well as the type of social network. This would allow us to investigate how social norms, empathy, and judgment are shaped not just by the behavior shown, but by the narrative context provided. We also aim to examine the use of targeted sentiment analysis methods that distinguish between opinion about a user versus their technology. Finally, we hope to include richer metadata where available to explore platform-level trends and demographic variation in reception.

Acknowledgments

This work is supported by the PRIN 2022 project COCOA, PRIN 2022T8ZNNM, funded by the European Union - Next Generation EU (NGEU) Programme, Mission 4, Component 1, CUP D53D23008850001.

Data Availability

The full dataset of Instagram comments used in this study is not shared publicly. The coded dataset (with comment IDs but not full text) is available upon reasonable request for research verification.

References

[1] Malak Abdullah Alsabban, Mark Weal, and Wendy Hall. 2024. Embracing emojis in sarcasm detection to enhance sentiment analysis. In *2024 International Conference on Computer and Applications (ICCA)*. IEEE, 1–7.

[2] Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 1644–1650. doi:10.18653/v1/2020.findings-emnlp.148

[3] Kristen Barta and Nazanin Andalibi. 2021. Constructing authenticity on TikTok: Social norms and social support on the" Fun" Platform. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29.

[4] Naman Bhargava, Mohammed I Radaideh, O Hwang Kwon, Aditi Verma, and Majdi I Radaideh. 2025. On the Impact of Language Nuances on Sentiment Analysis with Large Language Models: Paraphrasing, Sarcasm, and Emojis. *arXiv preprint arXiv:2504.05603* (2025).

[5] Amanda Björnwall, Ylva Mattsson Sydner, Afsaneh Koochek, and Nicklas Neuman. 2021. Eating Alone or Together among Community-Living Older People: A Scoping Review. *International Journal of Environmental Research and Public Health* 18, 7 (2021), 3495. doi:10.3390/ijerph18073495

[6] Octavia Calder-Dawe, Cherie Lacey, Margaret Wetherell, and Michael S Daubs. 2025. Being a positive influence (r): Exploring affective pedagogies of wellbeing and positivity on Instagram. *Feminist Media Studies* 25, 3 (2025), 576–592.

[7] Noah Castelo and Miklos Sarvary. 2022. Cross-cultural differences in comfort with humanlike robots. *International Journal of Social Robotics* 14, 8 (2022), 1865–1873.

[8] Dushyant Singh Chauhan, Gopendra Vikram Singh, Aseem Arora, Asif Ekbal, and Pushpak Bhattacharyya. 2022. An emoji-aware multitask framework for multimodal sarcasm detection. *Knowledge-Based Systems* 257 (2022), 109924.

[9] Zheng Lin Chia, Michal Ptaszynski, Fumito Masui, Gniewosz Leliwa, and Michal Wroczynski. 2021. Machine Learning and feature engineering-based study into sarcasm and irony classification with application to cyberbullying detection. *Information Processing & Management* 58, 4 (2021), 102600.

[10] Dmitry Davidov, Oren Tsur, and Ari Rappoport. 2010. Enhanced Sentiment Learning Using Twitter Hashtags and Smileys. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*. Coling 2010 Organizing Committee, Beijing, China, 241–249.

[11] David DeVault, Ron Artstein, Grace Benn, Teresa Dey, Ed Fast, Alesia Gainer, Kallirroi Georgila, Jon Gratch, Arno Hartholt, Margaux Lhommet, Gale Lucas, Stacy Marsella, Fabrizio Morbini, Angela Nazarian, Stefan Scherer, Giota Stratou, Apar Suri, David Traum, Rachel Wood, Yuyu Xu, Albert Rizzo, and Louis-Philippe Morency. 2014. SimSensei kiosk: a virtual human interviewer for healthcare decision support *(AAMAS '14)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1061–1068.

[12] Raghu Dhumpati, Archana Sasi, Shaik Johny Basha, Suryanarayana Vadhri, P Saroja, Swathi Alluri, and Ramesh Vatambeti. 2025. Enhancing sarcasm detection in sentiment analysis for cyberspace safety using advanced deep learning techniques. *Scientific Reports* 15, 1 (2025), 22681.

[13] Junlin Fang, Wenya Wang, Guosheng Lin, and Fengmao Lv. 2024. Sentiment-oriented Sarcasm Integration for Video Sentiment Analysis Enhancement with Sarcasm Assistance. In *Proceedings of the 32nd ACM International Conference on Multimedia*. 5810–5819.

[14] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using Millions of Emoji Occurrences to Learn Any-Domain Representations for Detecting Sentiment, Emotion and Sarcasm. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, 1616–1626.

[15] Dunigan P Folk, Chenxi Wu, and Steven J Heine. 2025. Cultural variation in attitudes toward social chatbots. *Journal of Cross-Cultural Psychology* 56, 3 (2025), 219–239.

[16] Hunter Fong, Selim Soufargi, Maurizio Mancini, and Radoslaw Niewiadomski. 2025. *Designing Interactions for Artificial Commensal Companions*. Association for Computing Machinery, New York, NY, USA, 306–311. https://doi.org/10.1145/3715668.3736351

[17] Ayaka Fujii, Kei Okada, and Masayuki Inaba. 2021. A Basic Study for Acceptance of Robots as Meal Partners: Number of Robots During Mealtime, Frequency of Solitary Eating, and Past Experience with Robots. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. 73–80. doi:10.1109/RO-MAN50785.2021.9515451

[18] Debanjan Ghosh, Alexander R. Fabbri, and Smaranda Muresan. 2015. Sarcastic or Not: Word Embeddings to Predict the Literal or Sarcastic Meaning of Words. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, 1003–1012.

[19] Anastasia Giachanou and Fabio Crestani. 2016. Like It or Not: A Survey of Twitter Sentiment Analysis Methods. *Comput. Surveys* 49, 2 (2016), 28:1–28:41. doi:10.1145/2938640

[20] Claude Grignon. 2001. Commensality and Social Morphology: An Essay of Typology. In *Food, Drink and Identity: Cooking, Eating and Drinking in Europe Since the Middle Ages*, Peter Scholliers (Ed.). Berg, Oxford, UK, 23–33.

[21] Albana Hoxha, Hunter Fong, and Radoslaw Niewiadomski. 2024. Do We Need Artificial Dining Companions? Exploring Human Attitudes Toward Robots in Commensality Settings. In *Companion Proceedings of the 26th ACM International Conference on Multimodal Interaction (ICMI Companion '24)*. Association for Computing Machinery, New York, NY, USA, 223–229. doi:10.1145/3686215.3686220

[22] Clayton J. Hutto and Eric Gilbert. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media (ICWSM '14)*. Association for the Advancement of Artificial Intelligence, Palo Alto, CA, USA, 216–225. https://ojs.aaai.org/index.php/ICWSM/article/view/14550

[23] Mahmud Isnan, Gregorius Natanael Elwirehardja, and Bens Pardamean. 2023. Sentiment analysis for TikTok review using VADER sentiment and SVM model. *Procedia Computer Science* 227 (2023), 168–175.

[24] Aditya Joshi, Pushpak Bhattacharyya, and Mark James Carman. 2017. Automatic Sarcasm Detection: A Survey. *Comput. Surveys* 50, 5 (2017), 1–22. doi:10.1145/3124420

[25] Michelle Sophie Keller, Hannah J Park, Maria Elena Cunningham, Joshua Eleazar Fouladian, Michelle Chen, and Brennan Mason Ross Spiegel. 2017. Public Perceptions Regarding Use of Virtual Reality in Health Care: A Social Media Content Analysis Using Facebook. *J Med Internet Res* 19, 12 (19 Dec 2017), e419. doi:10.2196/jmir.7467

[26] Rebekka Kreling, Adrian Meier, and Leonard Reinecke. 2022. Feeling authentic on social media: Subjective authenticity across Instagram stories and posts. *Social Media+ Society* 8, 1 (2022), 20563051221086235.

[27] Ahreum Lee and Austin L Toombs. 2020. Robots on campus: understanding public perception of robots using social media. In *Companion Publication of the 2020 Conference on Computer Supported Cooperative Work and Social Computing*. 305–309.

[28] Eleonora Lima and Tiffany Morisseau. 2025. Can chatbots teach us how to behave? Examining assumptions about user interactions with AI assistants and their social implications. *Frontiers in Artificial Intelligence* Volume 8 - 2025 (2025). doi:10.3389/frai.2025.1545607

[29] Rui Liu and Tomoo Inoue. 2014. Application of an Anthropomorphic Dining Agent to Idea Generation. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication* (Seattle, Washington) *(UbiComp '14 Adjunct)*. ACM, New York, NY, USA, 607–612. doi:10.1145/2638728.2641342

[30] Steven Loria. 2018. TextBlob: Simplified Text Processing. https://textblob.readthedocs.io/en/dev/. Accessed: 2025-05-13.

[31] Maurizio Mancini, Radoslaw Niewiadomski, Gabriele De Lucia, and Francesco Maria Longobardi. 2024. A Virtual Agent as a Commensal Companion. In *Proceedings of the 24th ACM International Conference on Intelligent Virtual Agents* (GLASGOW, United Kingdom) *(IVA '24)*. Association for Computing Machinery, New York, NY, USA, Article 26, 4 pages. doi:10.1145/3652988.3673963

[32] Maurizio Mancini, Radoslaw Niewiadomski, Gijs Huisman, Merijn Bruijnes, and Conor Patrick Gallagher. 2020. Room for One More? Introducing Artificial Commensal Companions. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–8. doi:10.1145/3334480.3383027

[33] Sarah McRoberts, Haiwei Ma, Andrew Hall, and Svetlana Yarosh. 2017. Share First, Save Later: Performance of Self through Snapchat Stories. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 6902–6911. doi:10.1145/3025453.3025771

[34] Kunihiro Miyazaki, Taichi Murayama, Takayuki Uchiba, Jisun An, and Haewoon Kwak. 2024. Public perception of generative AI on Twitter: an empirical study based on occupation and usage. *EPJ Data Science* 13, 1 (2024), 2.

[35] Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets. In *Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*. Association for Computational Linguistics, Atlanta, Georgia, USA, 321–327.

[36] Mohamed M. Mostafa. 2013. More than words: Social networks' text mining for consumer brand sentiments. *Expert Systems with Applications* 40, 10 (2013), 4241–4251.

[37] Radoslaw Niewiadomski, Gabriele De Lucia, Gabriele Grazzi, and Maurizio Mancini. 2022. Towards Commensal Activities Recognition. In *Proceedings of the 2022 International Conference on Multimodal Interaction (ICMI '22)* (Bengaluru, India). Association for Computing Machinery, New York, NY, USA, 549–557. doi:10.1145/3536221.3556566

[38] Radoslaw Niewiadomski, Jennifer Hofmann, Jérôme Urbain, Tracey Platt, Johannes Wagner, Bilal Piot, Huseyin Cakmak, Sathish Pammi, Tobias Baur, Stephane Dupont, Matthieu Geist, Florian Lingenfelser, Gary McKeown, Olivier Pietquin, and Willibald Ruch. 2013. Laugh-aware Virtual Agent and Its Impact on User Amusement. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems* (St. Paul, MN, USA) *(AAMAS '13)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 619–626.

[39] Petra Kralj Novak, Jasmina Smailović, Borut Sluban, and Igor Mozetič. 2015. Sentiment of Emojis. *PLOS ONE* 10, 1–22. doi:10.1371/journal.pone.0144296

[40] Magalie Ochs and Philippe Blache. 2016. Virtual Reality for Training Doctors to Break Bad News. In *Adaptive and Adaptable Learning*, Katrien Verbert, Mike Sharples, and Tomaž Klobučar (Eds.). Springer International Publishing, Cham, 466–471.

[41] Shereen Oraby, Vrindavan Harrison, Lena Reed, Ernesto Hernandez, Ellen Riloff, and Marilyn Walker. 2016. Creating and Characterizing a Diverse Corpus of Sarcasm in Dialogue. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Association for Computational Linguistics, Los Angeles, USA, 31–41.

[42] Weihong Qi, Jinsheng Pan, Hanjia Lyu, and Jiebo Luo. 2024. Excitements and concerns in the post-ChatGPT era: Deciphering public perception of AI through social media analysis. *Telematics and Informatics* 92 (2024), 102158. doi:10.1016/j.tele.2024.102158

[43] Chhavi Sharma, Deepesh Bhageria, William Scott, Srinivas Pykl, Amitava Das, Tanmoy Chakraborty, Viswanath Pulabaigari, and Bjorn Gamback. 2020. SemEval-2020 Task 8: Memotion Analysis–The Visuo-Lingual Metaphor! *arXiv preprint arXiv:2008.03781* (2020).

[44] Limor Shifman. 2014. *Memes in Digital Culture*. MIT Press, Cambridge, MA.

[45] Wendy Spettigue and Katherine A. Henderson. 2004. Eating Disorders and the Role of the Media. *Journal of the Canadian Academy of Child and Adolescent Psychiatry* 13, 1 (2004), 16–19. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2533817/

[46] Mike Thelwall. 2017. Gender Bias in Sentiment Analysis. *Online Information Review* 42 (11 2017), 00–00. doi:10.1108/OIR-05-2017-0139

[47] Sunny Yoon. 2024. Imaginaries of humanoids and evolutions of technological visions of AI in Eastern and Western media. *AI & SOCIETY* (2024), 1–10.