(Lougha - Kalaam) Approved Scientifique International Edited by
Language and Communication Laboratory
Relizane University – Algeria

ISSN : 2437-0746 / EISSN: 2600-6308
Numéro de depot: 3412 – 2015
Catégorie (C) Décision n° 1432 2019/08/13
http://www.asjp.cerist.dz/en/PresentationRevue/176

Volume : 12 - N 01. – Janvier 2026

Reçu le: 23/02 / 2025 Accepté le: 21/01 / 2026 Publié le: 31/01 / 2026

Why Large Language Models Struggle with Sarcasm: A Reflection on Pragmatic Competence, Cultural Nuances, and AI Limitations

✍CHAMI Wahid Hamza Mohamed
wahid.chami@ensia.edu.dz
The National Higher School of Artificial Intelligence /Algeria

لماذا تواجه النماذج اللغوية الكبيرة صعوبة في فهم السخرية: منظور على الكفاءة البراغماتية، الفروقات الثقافية، وحدود الذكاء الاصطناعي

✍شامي وحيد حمزة محمد
wahid.chami@ensia.edu.dz
المدرسة الوطنية العليا للذكاء الاصطناعي /الجزائر

ABSTRACT:

This reflective article explores the limitations of Large Language Models (LLMs) in understanding and generating sarcasm, a nuanced form of verbal irony that relies on context, tone, and cultural knowledge. By doing so, the article defines sarcasm and its communicative role, analyses the challenges large language models face in understanding it, examines its cultural variability, and proposes strategies to enhance models' pragmatic and contextual competence. Through a combination of theoretical analysis, real-world examples, and case studies, this article underscores the gap between human and machine understanding of sarcasm and offers recommendations for improving LLMs' pragmatic competence.

Keywords: Artificial Intelligence, ChatGPT, LLMs, Pragmatic Competence, Sarcasm

مُلخّص البحث

تتناول هذه المقالة التأملية حدود النماذج اللغوية الكبيرة (LLMs) في فهم وإنتاج السخرية، وهي شكل دقيق من أشكال المفارقة اللفظية يعتمد على السياق والنغمة والمعرفة الثقافية. وتعرّف المقالة السخرية ودورها في التواصل، وتحلّل التحديات التي تواجهها هذه النماذج في معالجتها، وتستكشف تنوعها الثقافي، وتقترح استراتيجيات لتعزيز كفاءتها التداولية والسياقية. ومن خلال الجمع بين التحليل النظري والأمثلة الواقعية ودراسات الحالة، تسلط المقالة الضوء على الفجوة بين الفهم البشري والآلي للسخرية وتقدّم توصيات لتطوير الكفاءة التداولية للنماذج اللغوية الكبيرة للذكاء الاصطناعي.

كلمات مفتاحية: الذكاء الاصطناعي، الكفاءة البراغماتية، اسلوب السخرية، النماذج اللغوية الكبيرة، شات جي بي تي

1. INTRODUCTION

Sarcasm is a form of verbal irony where the intended meaning of a statement is the opposite of its literal meaning, often used to convey humor, criticism, or emphasis. It relies heavily on context, tone, and shared knowledge between the speaker and the listener. For example, saying "Great job!" in a flat tone after someone spills a drink is clearly sarcastic, as the literal praise contradicts the situational context. In human communication, sarcasm serves as a sophisticated tool for expressing nuanced emotions, building social bonds, and navigating complex social dynamics (Gibbs, 2000). It is particularly prevalent in informal conversations, literature, and media, where it adds layers of meaning and emotional depth.

Sarcasm is a key aspect of **pragmatic competence**, which refers to the ability to use language appropriately in social contexts. Pragmatic competence involves understanding not just the literal meaning of words, but also the speaker's intentions, the context of the conversation, and the cultural norms that shape communication (Kasper & Rose, 2002). Sarcasm, as a form of indirect communication, requires the listener to infer the speaker's true meaning by analyzing contextual cues, tone, and shared knowledge. This makes it a high-level pragmatic skill that demonstrates a deep understanding of language and social dynamics. For instance, recognizing sarcasm often depends on subtle cues like facial expressions, vocal intonation, and situational context—elements that are deeply rooted in human experience and cultural understanding (Attardo, 2000).

However, for Large Language Models (LLMs), sarcasm presents a significant challenge. LLMs like GPT-4 are trained on vast amounts of text data, but they lack the ability to fully grasp the contextual and emotional nuances that underpin sarcasm (Bender et al., 2021). While they can sometimes detect sarcasm when it is explicitly signaled (e.g., through exaggerated phrasing or obvious contradictions), they often fail in more subtle cases where tone, facial expressions, or situational context are critical. For example, an LLM might misinterpret a sarcastic comment like "Oh, fantastic, another Monday!" as genuine enthusiasm, missing the underlying frustration. This limitation stems from the fact that LLMs operate purely on textual patterns and lack access to the multimodal cues (e.g., tone, body language) and real-world experiences that humans use to interpret sarcasm (Joshi et al., 2017).

Moreover, generating sarcasm is equally challenging for LLMs. While they can produce text that mimics sarcastic patterns, the result often feels forced or inappropriate because the model cannot fully understand the social and emotional context required to use sarcasm effectively. For instance, an LLM might generate a sarcastic response in a situation where sincerity is expected, leading to misunderstandings or offense. This highlights a broader issue: LLMs struggle with the **pragmatic flexibility** that humans naturally possess, making it difficult for them to navigate the complexities of sarcasm and other forms of indirect communication (Hovy & Yang, 2021).

As such, sarcasm is a crucial component of pragmatic competence that enriches human communication by adding layers of meaning and emotional nuance. However, its reliance on context, tone, and shared knowledge makes it a significant challenge for LLMs. As these tools continue to evolve, addressing their limitations in understanding and generating sarcasm will require advancements in multimodal learning, contextual understanding, and cultural sensitivity (Bender et al., 2021). Until then, the nuanced art of sarcasm remains a uniquely human skill, underscoring the gap between artificial and human intelligence in the realm of pragmatic competence.

In this respect, this article has two main objectives:

* Shedding light on the linguistic and cultural dimensions of Sarcasm.
*Underpinning the challenges of LLMs in interpreting Sarcasm in its cultural nuances.

2. Understanding Sarcasm in Human Communication

Sarcasm is a sophisticated form of verbal irony that relies on a mismatch between literal and intended meaning. It involves saying something that is the opposite of what is meant, often to convey humor, criticism, or emphasis. For example, when someone says, "Oh, fantastic, another Monday!" after a long and tiring day, the literal meaning ("fantastic") contradicts the intended meaning ("this is terrible"). This mismatch is central to sarcasm and requires the listener to infer the speaker's true intent by relying on tone, context, and shared knowledge (Gibbs, 2000; Attardo, 2000). These elements work together to signal that the statement is not meant to be taken at face value.

In this respect the tone, the context, and the shared knowledge are key factors in how sarcasm functions. Tone, particularly in spoken language, provides critical auditory cues that signal sarcasm. A flat, exaggerated, or mocking tone can indicate that the speaker does not mean what they are saying literally (Rockwell, 2007). Context, on the other hand, helps the listener understand why the statement is sarcastic. For instance, if someone says, "Nice weather we're having," during a thunderstorm, the context makes it clear that the statement is sarcastic (Attardo, 2000). Shared knowledge between the speaker and listener also plays a crucial role. Sarcasm often relies on inside jokes, cultural references, or mutual understanding of a situation. Without this shared knowledge, sarcasm can easily be misinterpreted or go unnoticed (Kreuz & Roberts, 1995).

Additionally, humans detect sarcasm through a combination of subtle cues, including intonation, facial expressions, and prior knowledge. In spoken communication, intonation is one of the most immediate indicators of sarcasm. A sarcastic statement often has a distinct pitch or rhythm that differs from sincere speech. For example, a sarcastic "Thanks a lot" might be delivered with a drawn-out, exaggerated tone (Rockwell, 2007). Facial expressions also play a significant role. A raised eyebrow, a smirk, or an eye roll can signal that the speaker is being sarcastic (Attardo, 2000). In written communication, however, where tone and facial expressions are absent, humans rely on contextual clues and prior knowledge to detect sarcasm. For instance, the use of exaggerated language or incongruent statements can hint at sarcasm, as in the case of a tweet like, "Absolutely love it when my flight gets delayed. Best part of my day!" (Joshi et al., 2017).

However, sarcasm is not universally understood in the same way. In other words, it is culturally dependent. Different cultures have varying norms around directness and indirectness in communication, which influences how sarcasm is used and interpreted (Holtgraves, 2005). In high-context cultures (e.g., Japan, India), where communication relies heavily on implicit understanding and context, sarcasm may be more subtle and nuanced. In contrast, in low-context cultures (e.g., the United States, Germany), where communication tends to be more direct, sarcasm may be more overt and easily recognizable (Hall, 1976). For example, in cultures that value politeness and indirectness, sarcasm might be used sparingly to avoid causing offense, while in cultures that embrace bluntness, sarcasm might be more frequent and direct. This cultural variability means that what is perceived as sarcasm in one culture might be taken literally in another, leading to potential misunderstandings (Kreuz & Roberts, 1995).

Reflecting on the cultural dependence of sarcasm highlights its role as a social and relational tool. It is not just about the words themselves but about the shared understanding and cultural norms

that give those words meaning. This is why sarcasm can be challenging for artificial systems like Large Language Models (LLMs) to detect and generate. LLMs lack the lived experience, cultural awareness, and ability to interpret subtle cues like tone and facial expressions that humans rely on to understand sarcasm (Bender et al., 2021). As a result, sarcasm remains a uniquely human form of communication, deeply rooted in our social and cultural contexts.

3. Why Large Language Models Struggle with Sarcasm

Large Language Models (LLMs), such as GPT-4, have made remarkable strides in generating human-like text and performing complex language tasks. Yet, despite their sophistication, they consistently falter when it comes to understanding sarcasm. Sarcasm, a form of verbal irony, relies on a mismatch between literal and intended meaning, often requiring the listener to infer the speaker's true intent through tone, context, and shared knowledge. LLMs, however, process text based on statistical patterns rather than actual understanding, lack real-world experience, and struggle to interpret implicit meanings. These limitations become particularly evident in written text, where tone and gestures—key indicators of sarcasm—are absent. To illustrate these challenges, let us delve into specific examples and case studies that highlight why sarcasm remains a stumbling block for even the most advanced AI systems.

3.1. Processing Texts Based on Statistical Patterns Rather Than Understanding

At their core, LLMs are prediction machines. They analyze vast amounts of text data to identify patterns and predict the next word in a sequence. While this allows them to generate coherent and contextually relevant text, it does not equate to genuine comprehension. Sarcasm, by its very nature, involves saying one thing while meaning the opposite, often relying on subtle cues that go beyond the literal meaning of words. For example, consider the statement, "Oh, great, another Monday!" To a human, this is clearly sarcastic, expressing frustration about the start of the workweek. However, an LLM might interpret it literally, responding with something like, "I'm glad you're excited about Monday!" This misinterpretation occurs because the model relies on statistical associations between words rather than an understanding of the speaker's intent (Bender et al., 2021).

Another example is the phrase, "Sure, because I have all the time in the world." A human would immediately recognize this as sarcasm, implying the speaker is overwhelmed or annoyed. An LLM, however, might respond with, "That's great! Having free time is wonderful." This failure to grasp the underlying meaning highlights the limitations of relying solely on textual patterns without true comprehension.

3.2. Lack of Real-World Experience and Implicit Meaning

Sarcasm often depends on shared knowledge and real-world experiences that LLMs simply do not possess. For instance, consider the statement, "I love it when my flight gets delayed." A human would understand this as sarcasm because they know the frustration of travel delays. An LLM, however, lacks this lived experience and might interpret the statement literally, responding with, "That's wonderful! Delays can be a great time to relax." This disconnect underscores the importance of real-world context in interpreting sarcasm—a dimension that LLMs cannot access (Hovy & Yang, 2021).

Similarly, sarcasm often relies on cultural references or inside jokes that are not explicitly stated. For example, a tweet like, "Nothing like a 2-hour meeting that could have been an email," is

easily understood by humans as a sarcastic critique of workplace inefficiency. An LLM, however, might fail to recognize the sarcasm because it cannot infer the cultural context of workplace frustrations. Instead, it might respond with, "Meetings are a great way to collaborate!" This inability to grasp implicit meanings is a significant barrier to understanding sarcasm (Joshi et al., 2017).

3.3. Challenges in Written Text: Absence of Tone and Gestures

One of the most significant challenges for LLMs is recognizing sarcasm in written text, where tone, facial expressions, and gestures, key indicators of sarcasm, are absent. In spoken communication, humans rely on auditory cues like intonation and visual cues like facial expressions to detect sarcasm. For example, a sarcastic "Nice job!" might be delivered with a sarcastic tone or an eye roll, making the intent clear. In written text, however, these cues are missing, leaving LLMs to rely solely on the words themselves.

Consider the statement of your colleague at work saying "Oh, sure, because that's exactly what I needed today." After you told him you will leave earlier today. A human would likely infer sarcasm based on the exaggerated language and the implied frustration. An LLM, however, might interpret it literally, responding with, "I'm glad it was helpful!" This misinterpretation occurs because the model cannot detect the sarcastic tone or the speaker's underlying frustration (Rockwell, 2007).

Another example is the phrase, "Wow, this is exactly what I wanted for my birthday, a pair of socks." A human would immediately recognize the sarcasm, understanding that the speaker is disappointed with the gift as he might expect something more valuable. An LLM, however, might respond with, "I'm glad you like your gift! Socks are a practical and thoughtful present." Or in best cases hesitate to take a decision about whether the phrase is sarcastic or genuine(Fig 1). This failure to recognize sarcasm highlights the limitations of LLMs relying solely on textual patterns without access to tone or gestures (Fig 2).

.

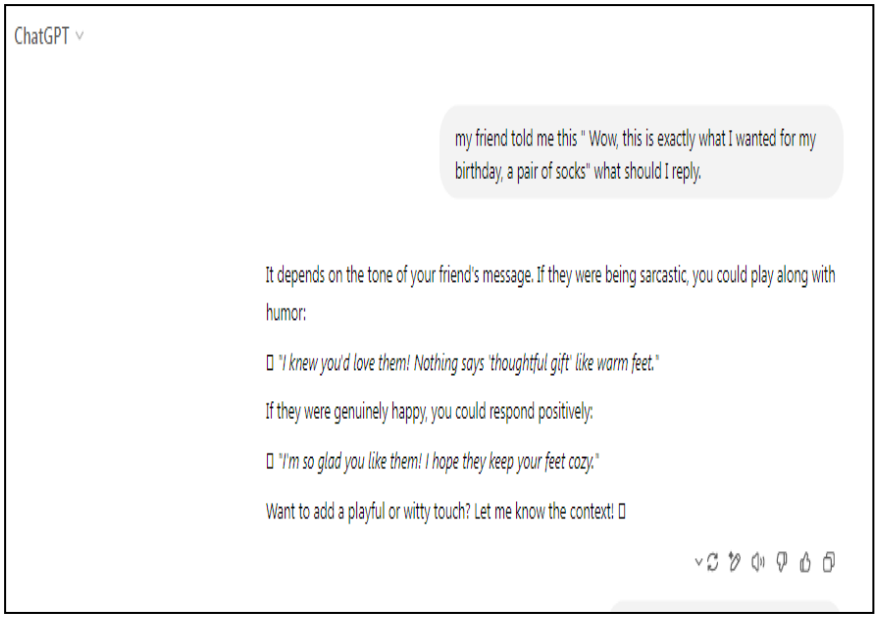

Figure 01: ChatGPT Inability to Decide Sarcastic Tone

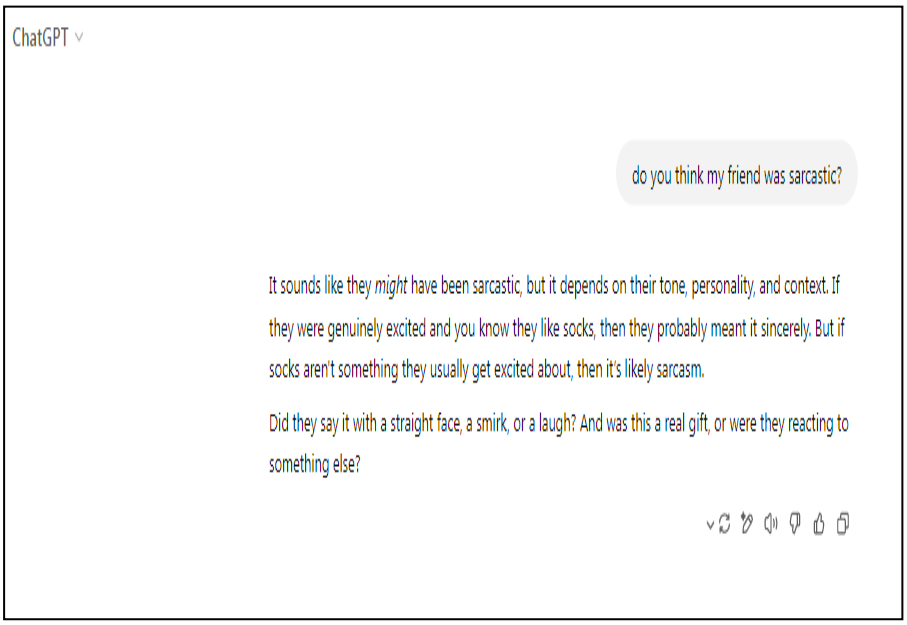

Figure 02: ChatGPT Enquiry about Facial Expression

3.4. Case Studies: LLMs Misinterpreting Sarcasm

Real-world examples further illustrate LLMs' struggles with sarcasm. In a study on sarcasm detection in social media (Hovy & Yang, 2021), the researchers found that LLMs frequently misclassified sarcastic tweets as sincere.

❖ Example 1:

In customer service interactions, LLMs often fail to detect sarcasm in user complaints. For example, a user might write, "Oh, fantastic service, really impressed," to express dissatisfaction.

 AI often assigns a sentiment score to words (Sentiment Analysis):

- *"Fantastic"* → Positive
- *"Really impressed"* → Positive

 However, sarcasm often masks a negative sentiment with positive words.

 Solution: AI could check whether previous statements (or common complaint patterns) suggest dissatisfaction.

❖ Example 2:

A tweet that read, "Can't wait more to see the bill!" was classified as a positive statement by the LLM due to the presence of words like "can't wait to." Humans, however, easily recognized the sarcasm, understanding the fear of a costly meal.

❖ Example 3

Another example comes from content moderation on social media platforms. LLMs tasked with identifying harmful content might misinterpret sarcastic posts as genuine. For instance, a sarcastic comment like, "Yeah, because that's a brilliant idea," might be flagged as positive feedback, while its true intent is criticism. This highlights the broader implications of LLMs' inability to reliably detect sarcasm, particularly in applications where understanding tone and intent is crucial (Hovy & Yang, 2021).

4. The Role of Culture in Sarcasm and AI's Challenges

Sarcasm is a universal form of communication, but its use, interpretation, and cultural significance vary widely across different societies. While it is prevalent in some cultures, such as English-speaking countries, it may be less common or expressed differently in others, such as in the Arab world. These cultural variations pose significant challenges for AI systems, particularly Large Language Models (LLMs), which are often trained on biased or culturally limited datasets.

Table 1. *Comparison of Sarcasm Use in English-Speaking Cultures and the Arab World*

Aspect	English-Speaking Cultures	Arab World
Frequency of Use	Common and widely accepted	Less frequent
Directness	Often direct and overt	Subtle and indirect
Politeness	Used even in polite contexts	Avoided in formal settings

AI systems, including LLMs, are trained on large datasets that often reflect the cultural biases of their creators . This can lead to uneven performance in recognizing sarcasm across different cultural contexts. For example, LLMs trained primarily on English-language data may excel at detecting sarcasm in English but struggle with sarcasm in Arabic or other languages.Moreover, The misinterpretation of sarcasm by AI systems can lead to serious ethical concerns, including misinformation, miscommunication, and cultural insensitivity. For example:

AI-powered news aggregators or social media platforms might misinterpret sarcastic posts as factual statements, spreading false information. For instance, a sarcastic tweet like, "Sure, because climate change is a hoax," might be taken literally by an AI, leading to the dissemination of misinformation.

In customer service, AI chatbots might fail to recognize sarcasm in user complaints, leading to inappropriate or tone-deaf responses. For example, a user might write, "Oh, I know that your company cares about its clients' views," to express dissatisfaction, but the chatbot might respond with, "Thank you for your feedback!, completely missing the sarcasm. Another example is a user's sarcastic complaint like, "Oh, fantastic service, really impressed," might be met with a generic "We do our best to assure your satisfaction".

Furthermore, AI systems that misinterpret sarcasm in one culture may offend users from another culture. For example, an LLM might fail to recognize the subtle sarcasm in an Algerian statement like, "*Rak 3ayech fi Hollywood!*" (you live in Hollywood, meaning that you are unrealistic) and respond in a way that seems dismissive or disrespectful to Amercican number one film industry.

5.Conclusion

The difficulties large language models (LLMs) face in detecting and interpreting sarcasm have important implications for their use in real-world applications. In customer service, failure to recognize sarcasm can lead to tone-deaf or inappropriate responses that frustrate users and damage brand reputation. In social media moderation, misinterpretations may cause harmless posts to be flagged as harmful or allow offensive content to pass unnoticed. These issues reveal the current limitations of AI systems in understanding the subtleties of human communication, especially across diverse cultural contexts.

The implications, therefore, are clear: LLMs are not yet ready for fully autonomous deployment in areas requiring deep pragmatic understanding. Their limitations highlight ethical and operational risks, such as biased moderation or miscommunication, when human oversight is absent.

To address these challenges, future research should focus on integrating multimodal learning, contextual reasoning, and cultural adaptability. Combining textual, auditory, and visual cues can help models detect sarcasm more accurately, while greater contextual understanding can improve interpretation. Training on culturally diverse datasets will also enhance fairness and inclusivity. Collaboration between AI researchers, linguists, and cultural experts is essential to ensure models are both effective and culturally sensitive.

References

Attardo, S. (2000). "Irony as Relevant Inappropriateness." *Journal of Pragmatics*, 32(6), 793-826.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*

Gibbs, R. W. (2000). "Irony in Talk Among Friends." *Metaphor and Symbol*, 15(1-2), 5-27.

Hall, E. T. (1976). *Beyond Culture*. Anchor Books.

Holtgraves, T. (2005). "Social Psychology, Cognitive Psychology, and the Linguistic Interplay in Irony Understanding." *Current Directions in Psychological Science*, 14(1), 18-21.

Hovy, D., & Yang, D. (2021). "The Importance of Modeling Social Factors of Language: Theory and Practice." *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*

Joshi, A., Bhattacharyya, P., & Carman, M. J. (2017). "Automatic Sarcasm Detection: A Survey." *ACM Computing Surveys*, 50(5), 1-37.

Kasper, G., & Rose, K. R. (2002). *Pragmatic Development in a Second Language*. Wiley-Blackwell.

Kreuz, R. J., & Roberts, R. M. (1995). "Two Cues for Verbal Irony: Hyperbole and the Ironic Tone of Voice." *Metaphor and Symbolic Activity*, 10(1), 21-31.

Rockwell, P. (2007). "Vocal Features of Conversational Sarcasm: A Comparison of Methods." *Journal of Psycholinguistic Research*, 36(5), 361-369.