



Beyond or within the binary? Constructing gendered meanings in generative AI use

Zhiyuan Lai and Yunjie Tang

DOI: <https://doi.org/10.47989/ir31iConf64121>

Abstract

Introduction. GenAI is increasingly embedded in daily life, yet its ostensibly neutral design often becomes a site for gendered interpretations. This study examines how Gen Z users in China perceive and construct gendered traits in GenAI, situating their accounts within broader concerns of neutrality and equity in information society.

Method. Semi-structured interviews were conducted with 12 participants who regularly use multiple GenAI platforms. Discussions explored textual, visual, and auditory cues, user projections, and contextual triggers.

Analysis. Thematic analysis, informed by Butler’s theory of performativity and Haraway’s cyborg imaginary, was applied to 73,480 words of transcripts. Coding identified patterns in participants’ accounts of gendered cues, user projection, contextual attribution, and reflections on neutrality.

Results. Participants consistently mapped rational, didactic tones onto masculine authority and empathetic language onto feminine care. Avatars and voices anchored gender perceptions, while mimicry highlighted both adaptability and artificiality. Attribution also drew on cultural repertoires, personal experiences, and task contexts.

Conclusion. Neutrality in GenAI does not erase gender but becomes a site of projection and negotiation. The findings show how such systems can reproduce stereotypes while also enabling more inclusive, post-binary imaginaries, underscoring the need for critical AI literacy and designs that advance information equity.

Introduction

When ChatGPT entered public discourse, it was hailed as a milestone in artificial intelligence: fluent dialogue, efficient knowledge delivery, and ostensibly neutral assistance (Lenharo, 2024). Since then, generative AI (GenAI) has been woven into education, work, and everyday life, often portrayed as a story of progress and empowerment. Yet neutrality in digital infrastructures is seldom straightforward.

For many users, GenAI does not feel entirely 'genderless' (Duan et al., 2025). In practice, its responses may be perceived as rational, authoritative, and at times even judgemental; in other moments, they are heard as empathetic, nurturing, or playful (Letheren et al., 2021). These shifting impressions reveal that gender in AI is neither absent nor fixed, but enacted through language, tone, and interaction (Hou et al., 2024). What is presented as neutral often becomes a site where cultural assumptions surface and circulate.

This tension matters deeply for library and information science (LIS). Technologies are not neutral artefacts; they are interpreted through the literacies, cultural repertoires, and expectations that users bring to interaction (Bridges et al., 2024; Colas et al., 2022; Lund et al., 2024). Consequently, even in systems designed without explicit gender, users' encounters become entangled with social categories, shaping how authenticity and trust are understood (Baumer et al., 2024; Krishna et al., 2025). Users routinely attribute gendered traits like rational authority or empathetic support to GenAI depending on the context (Dogruel & Joeckel, 2024; Duan et al., 2025). Neutrality, therefore, does not erase gender; rather, it becomes the very canvas upon which cultural assumptions are projected, reinforced, or unsettled. Against this backdrop, the present research examines how users construct gendered meanings in their interactions with GenAI. Rather than treating gender as something GenAI performs autonomously, we focus on how users perceive, articulate, and make sense of gendered traits in response to the cues they encounter. To this end, it poses three questions:

- RQ1. How do users interpret and describe the cues through which GenAI appears gendered in interaction?
- RQ2. What cultural repertoires, personal experiences, or expectations shape users' attribution of gender to ostensibly neutral systems?
- RQ3. In what contexts do these gendered interpretations become particularly salient, contested, or reconfigured?

Drawing on theories of gender performativity and contested neutrality (De Cet et al., 2025; Wagner et al., 2025), this study uses semi-structured interviews to examine Gen Z's perceptions of AI. This cohort was chosen for their deep integration of GenAI into daily life (Nataraja, 2025) and their heightened sensitivity to how gender is constructed in neutral systems, owing to their openness to fluid identities (Scheffer-Wentz, 2025; Zhao & Wu, 2025). Their narratives reveal the political stakes of AI 'neutrality' and its capacity to either perpetuate stereotypes or foster inclusivity.

Research background

The first way in which gender enters discussions of AI is through bias. Large language models, trained on the sediment of a stereotyped internet, generate text that all too often mirrors and amplifies existing cultural assumptions (Alvarez et al., 2024; Carter & Liu, 2025). Occupations are coded as masculine or feminine, authority defaults to the male figure, and stereotypes seep seamlessly into outputs (Spennemann, 2025). Such insights are vital, yet they tend to frame AI primarily as an object under scrutiny, a system to be audited for flaws and rectified through technical mitigation.

Another body of scholarship examines how users anthropomorphise AI systems, particularly robots and voice assistants (Glikson & Woolley, 2020). In this line of work, gender is not uncovered in code but attributed by users, often through familiar cultural frames. Machines perceived as logical and authoritative are read as 'male', while those experienced as empathetic or supportive are read as 'female' (Bernotat et al., 2021; Curry et al., 2020). In these contexts, gender cues are intentionally embedded, whether in a vocal timbre or a humanoid form, and user interpretations follow relatively predictable patterns (Spielmann & Stern, 2024).

GenAI unsettles both trajectories. Unlike embodied or vocalised agents, it presents no fixed voice or physical form, yet users still describe it as masculine, feminine, or something in between. This ambiguity demonstrates that anthropomorphism is not restricted to explicit design but may arise through the subtler dynamics of dialogue and projection (Vallis et al., 2025). Meanwhile, neutrality does not erase gender; rather, it becomes the very ground upon which gender is projected, reinforced, or unsettled (Xavier, 2025). What remains less well understood is how users actively construct gender in interaction with systems deliberately designed to be neutral. Redirecting attention from bias as a property of the model to gender as a co-constructed identity opens a critical space in which to examine how 'non-gendered' technologies are rendered intelligible through social and cultural frameworks.

Theoretical framework

This study adopts a social constructionist stance (Blagoev et al., 2024), drawing on Butler's performativity and Haraway's cyborg imaginary to explain how ostensibly neutral GenAI acquires gendered meaning through cultural interpretation and interaction.

The gender performativity as disembodied interaction

Butler's account of gender performativity unsettles any notion of gender as a stable core identity. Gender, she argues, is '*an identity tenuously constituted in time, instituted through a stylized repetition of acts*' (Butler, 1990, p. 191). It is something one *does*, rather than something one *is*, sustained only through repetition within the normative structures that render it intelligible.

When transposed to the digital realm, this perspective illuminates how GenAI engages in what may be termed disembodied linguistic performativity (Torres-Martínez, 2025). Although it lacks a physical body, it 'acts' through its linguistic and vocal expressions. Lexical choice, syntactic form, conversational stance, and affective tone all become sites of performance. When responses appear terse, logical, and emotionally pared back, users often interpret them as enactments of rational authority culturally coded as masculine (Hipólito et al., 2023). By contrast, when phrasing is supportive, acknowledgement empathetic, or elaboration affectively attuned, such performances are readily read as feminine (Sutko, 2020). Crucially, such readings are not properties of the model itself but accomplishments of interaction. Users bring with them interpretive repertoires shaped by social stereotypes and personal histories, and in dialogue with the machine they constitute its '*gender*' (Depounti et al., 2023). In this sense, GenAI is not a neutral artefact but a stage on which gender is enacted, sustained, and made legible.

The cyborg imaginary as a horizon of neutrality

If Butler offers a lens for understanding how gender is enacted, Donna Haraway provides a horizon for interrogating its limits. In *A Cyborg Manifesto*, Haraway (2006) depicted the cyborg as a figure designed to collapse the binaries that structure Western thought: nature/culture, human/machine, and most insistently, male/female. The cyborg embodies an ironic myth of a post-gender world; a hybrid subject unmoored from essentialist categories and open to affinity rather than identity.

GenAI appears, at first glance, to approximate this ideal. However, recent critical scholarship on technological neutrality complicates this post-gender promise. As Rakowski and Kowalikova (2024)

argue, technology cannot be viewed as a neutral resource detached from societal goals. Instead, it is a reflection of social relations and power asymmetries embedded in its very design. This perspective aligns with Giantini's (2023) critique of the '*sophistry of the neutral tool*,' which posits that the veneer of technological objectivity often masks the reinforcement of historical inequalities and social exclusion. Far from being empty vessels, algorithmic systems are socio-technical artifacts that frequently replicate existing social biases through data inputs and model design (Kordzadeh & Ghasemaghaei, 2022).

This creates a tension between the cyborgian ideal and empirical reality. Benjamin (2019) conceptualizes this phenomenon as the '*New Jim Code*,' where prior inequities are engineered into automated systems, allowing social hierarchies to be reproduced under the guise of objective benevolence. Consequently, the system imagined as a post-gender hybrid is caught in the gravitational pull of existing cultural scripts (Das & Chanda, 2023). Users repeatedly project binary categories onto GenAI, describing it as '*male*,' '*female*,' or oscillating between the two (Duan et al., 2025). This paradox constitutes what might be called the cyborg's tension, where a technology designed to transcend binaries becomes a site where their persistence is most clearly revealed. The critical task, therefore, is not to declare neutrality as achieved, but to examine how the very architecture of these tools interacts with user perceptions to make neutrality itself productive of gendered readings.

Working propositions and operationalisation

To connect theoretical commitments with content analysis, four working propositions structure the coding process. Each proposition translates abstract concepts into domains of observation within the interview data, providing both sensitizing concepts and practical guidelines for analysis.

1. **Performativity as practice.** This proposition focuses on the gendered cues that participants perceived in GenAI's textual, visual, auditory, and embodied outputs. Although these cues are not inherently gendered, participants drew on cultural scripts to read them as masculine, feminine, or aligned with heteronormative expectations. We therefore coded the descriptions provided by participants of linguistic styles, visual or vocal attributes, and adaptive behaviours such as inferring or mirroring user characteristics, insofar as these were reported as carrying gendered qualities. These perceptions may stem either from model-internal patterns or from user-responsive mimicry, but our analysis treats them only as interpretable cues, not as evidence of autonomous system agency. Together, these patterns form the observable surface through which GenAI was experienced as gendered. This proposition operationalises the Gender Performativity Practices of GenAI domain.
2. **Projection and interpretation.** In contrast, this proposition focuses on the user's side of meaning-making, examining how individuals draw on their own values, histories, and identity positions to interpret the cues described above. Here we coded instances in which participants explicitly projected stereotypes, invoked family or sociocultural gender roles, referenced personal experiences, or mapped their own preferences and orientations onto the model. These phenomena correspond to the User's Interpretation and Gender Construction domain.
3. **Contextual triggers of attribution.** Neutrality does not mean absence; rather, gender is often elicited in particular settings. We code for task type (for example, technical versus emotional), interactional depth (including frequency, duration, and emotional investment), and situational cues that render gender more salient. In practice, this meant linking user reports of task context or prolonged engagement with shifts in gendered perception. These categories align with the Contextual Triggers of Gender Perception domain.

- Consequences and horizons. Informed by Haraway's (2006) cyborg imaginary as a conceptual horizon, we trace the effects of gender attribution. Codes attend to how GenAI's gendered performances reinforce binary oppositions, shape trust and preference, elicit ambivalent attitudes, or invite post-binary imaginaries. In practice, this meant coding both users' critical reflections on bias and their aspirational visions of more fluid, customizable, or post-gender AI. This corresponds to the Consequences and Reflections of Gender Performativity domain.

Methods

Research design and participants

This study adopted a qualitative research design grounded in social constructionist epistemology, focusing on the performative and interpretive dimensions of perceived gender in GenAI. Semi-structured interviews were conducted with 12 participants in China between June and September 2025, and all interviews were carried out in Chinese. Recruitment followed a purposive and snowball strategy: initial participants were identified as deep users of large language models, and additional participants were referred through trusted personal networks to ensure diversity in gender, occupation, and sexual orientation.

Participants reflected the generational cohort of Gen Z, the group most intensively engaged with GenAI tools. While all were university students or early-career professionals, the sample encompassed variation in age, gender, academic discipline, and sexual orientation (see Table 1). This focus on Gen Z is deliberate, as their everyday reliance on digital technologies positions them as a critical population for understanding how gendered perceptions of AI take shape.

All twelve participants reported using GenAI on a daily basis. Every participant had experience with Doubao (豆包, ByteDance's GenAI chatbot), while ten had used ChatGPT, eleven had used DeepSeek, five had used Gemini, five had used Kimi, and four had used ERNIE Bot. A smaller number reported experience with ChatGLM (n=1) and Grok (n=1).

No.	Gender	Age	Occupation	Level	Field	Sexual orientation
P01	Male	18	Student	Undergraduate	LIS	Gay
P02	Male	27	Student	PhD candidate	LIS	Gay
P03	Female	19	Student	Undergraduate	LIS	Heterosexual
P04	Male	19	Student	Undergraduate	LIS	Heterosexual
P05	Female	20	Student	Undergraduate	LIS	Heterosexual
P06	Female	21	Student	Undergraduate	LIS	Heterosexual
P07	Male	19	Student	Undergraduate	Medical Science	Bisexual (gay leaning)
P08	Male	21	Student	Undergraduate	Law	Pansexual/Asexual
P09	Female	21	Student	Undergraduate	Biology	Lesbian
P10	Male	27	Researcher	Early-career	LIS	Heterosexual
P11	Male	27	Lawyer	Professional	Law	Heterosexual
P12	Female	20	Student	Undergraduate	LIS	Bisexual

Table 1. Participant details

All participants signed informed consent forms and were reminded of their right to withdraw at any point. Interviews were conducted in a mix of face-to-face and online sessions (via Tencent Meeting), depending on feasibility. Each lasted between 30 and 50 minutes, and with permission, all were audio-recorded and transcribed verbatim. Identifiable information was removed and pseudonyms were used to preserve anonymity.

Data collection and processing

The interview guide (see Appendix A) was designed to cover three interconnected areas while remaining open to participants' own emphases: 1) background information on patterns of GenAI use, preferred platforms, and the roles participants attributed to these systems, such as tool, partner, or persona. 2) gender perception, including whether and how participants attributed gendered traits and which cues triggered such attributions, including linguistic, visual, auditory, and interactional signals. 3) cultural reflection, linking these perceptions to social stereotypes, family roles, wider cultural norms, and expectations about the future of AI's gender presentation. Throughout the interviews, the researchers adopted gender-neutral phrasing and avoided presupposing binary categories, offering options such as female, male, gender-neutral, and other self-defined identities where relevant, and refraining from introducing gender attributions unless raised by participants. Probing questions were used flexibly to pursue emergent directions. All interviews were audio-recorded with consent and transcribed verbatim. The final corpus consists of 73,480 words of anonymised transcripts, prepared by removing identifying information and assigning pseudonyms, before being imported into NVivo 14 for data management and coding.

Data analysis

Data were analysed through thematic analysis following the six-step process outlined by Naeem et al. (2023). This approach offered the structure and analytical clarity needed for examining how participants construct gendered meanings in GenAI. Its use of explicit decision rules strengthened transparency and consistency in coding, while its emphasis on interpretive conceptualization supported the move from descriptive accounts to deeper analytical patterns. The analysis progressed from transcription to the identification of salient keywords and phrases, followed by first cycle coding, theme development and refinement, interpretive elaboration, and the construction of a coherent thematic account.

Coding combined deductive and inductive moves. An initial codebook was developed based on conceptual insights from gender theory and neutrality literature, together with sensitizing ideas inspired by Butler and Haraway. The preliminary codebook was then iteratively refined through close engagement with the interview material and cross-referencing with existing empirical studies on gendered human–AI interaction. Two researchers first coded the full dataset independently using the preliminary scheme, with scope to add data driven codes where necessary. Intercoder reliability was assessed after this round and produced a Cohen's Kappa of 0.52, which indicated moderate agreement. The researchers then met to discuss discrepancies, clarify code definitions, and adjust boundaries. A second round of independent coding followed and achieved a Kappa greater than 0.90, which confirmed the stability of categories and themes. The final coding frame aligned with four analytic domains that structure the findings: GenAI's gender performativity practices, user interpretation and construction, contextual triggers of perception, and consequences and reflections.

Figure 1 illustrates the coding trajectory of our thematic analysis, showing the progressive abstraction of raw keywords into codes, themes, and higher-order domains. While our findings are presented narratively, this figure clarifies the analytic path. The arrows indicate the general progression, but the process itself was iterative and reflexive, not strictly linear.

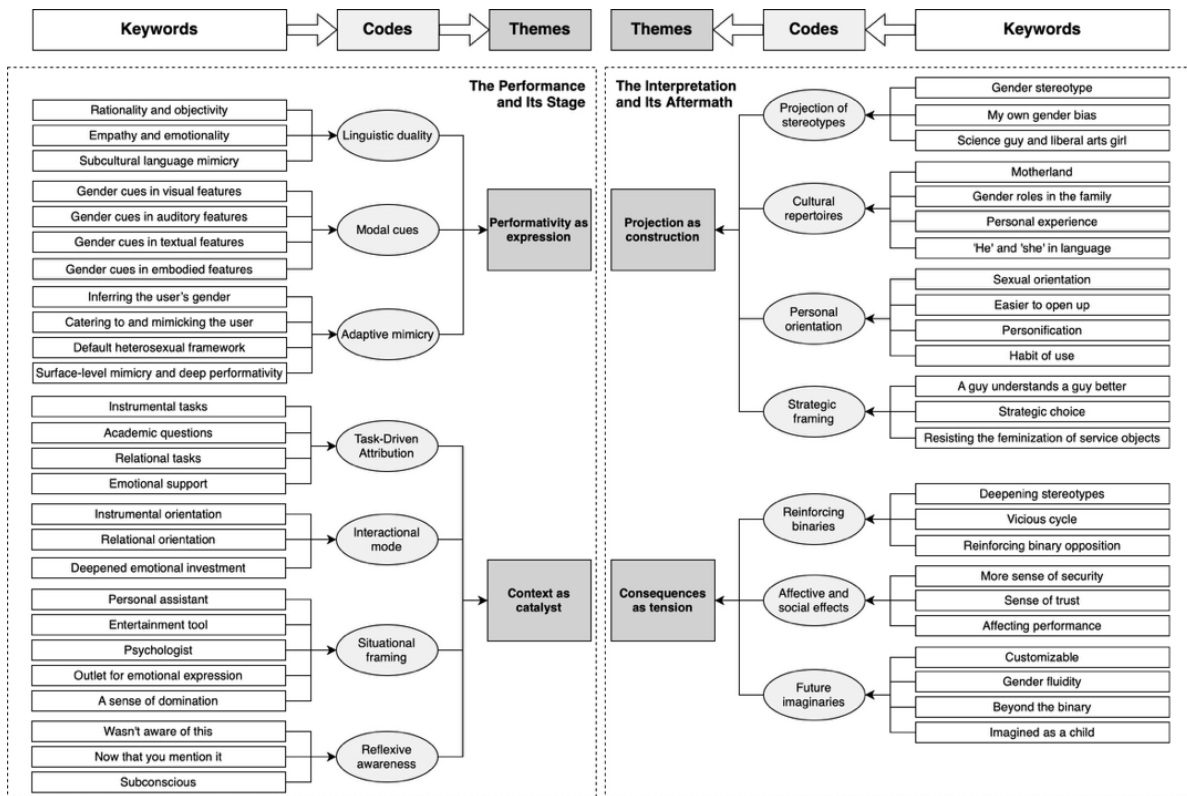


Figure 1. Coding framework and thematic map

Findings

Performativity as expression

GenAI does not need a body for its outputs to be read as gendered. Its presence is articulated through patterned expressions in text, tone, and imagery, which users interpret through familiar cultural scripts. What appears to be neutral language becomes read as rational authority or empathetic care, enacting a gendered persona in dialogue.

Textual performance: rationality and empathy

The most significant medium for perceived gendered expression is its textual output. Participants consistently mapped the GenAI's language style onto a binary of masculine rationality and feminine empathy. A logical, direct, and sometimes didactic tone was often read as masculine. P01 described this interaction as feeling like he was talking to a male senior student or advisor who would 'paternalistically' lecture him. This perception was reinforced when the responses were structured, analytical, and emotionally detached. P08 summarized this by likening the GenAI models DeepSeek and Kimi to a 'science/engineering guy and a liberal arts girl', respectively. DeepSeek, he noted, excels at providing structured outlines, a trait he associated with masculine rationality, while Kimi is superior at polishing text with a more nuanced, literary flair, which he coded as feminine.

Conversely, a warm, patient, and supportive communication style was almost universally interpreted as feminine. This perception became particularly strong in contexts of emotional support. P02 felt that when he used GenAI for psychological counselling, it resembled a 'mother or an older sister' due to its perceived delicacy and ability to offer comfort. Similarly, P11 described the responses during a moment of personal distress as being so 'gentle and inclusive' that he imagined it as an older female psychologist. This empathetic performance, characterized by

supportive phrasing like *'you've already done a great job'*, was a powerful trigger for feminine attribution.

Sensory cues: visual and auditory anchors

Beyond text, visual and auditory cues served as powerful, direct anchors for users' gender perceptions. Doubao, with its cartoon-like female avatar, was a prime example. For many users, this visual cue was decisive. As P06 stated, *'I think [it influences me], because that image is right there, constantly reminding you that this is a girl'*. The visual branding foreclosed other interpretations, establishing a female identity from the outset.

Auditory features, while less universally used, also played a key role. The default female voice in many GenAI assistants reinforced a feminine perception. However, voice also introduced a layer of malleability. P11, for instance, intentionally switched to a male voice assistant to practice his English pronunciation, drawing on the stereotypical association of male voices with authority and clarity in language learning. This effectively recast the GenAI's gender to suit his task. Yet, this plasticity has its limits. The notable absence of a neutral or non-binary voice option in most systems reveals an underlying design logic that still operates within a binary framework, constraining the possibilities of gender performance. This was highlighted by P09, who felt that using the voice function would *'break her fantasy'* of the GenAI as a mature, gentle sister, as the available voice options sounded like a *'young and energetic little sister'*, which felt inauthentic and *'cold'*.

Interactional adaptation: mimicry and its limits

Beyond these fixed cues of text, voice, and image, users also experienced GenAI's adaptive linguistic shifts as carrying gendered implications. GenAI's performativity is not a one-way street; it is highly adaptive, often mimicking the user's own linguistic style. Participants noted that when they used emojis and expressive punctuation, the GenAI would reciprocate, a behaviour P09 described as being very *'catering'* or even *'sycophantic'*. P01 observed that ChatGPT would even adopt queer slang like *'素'* (sù) and *'惹'* (rě) [1] from his prompts, which he interpreted as the GenAI *'learning your language style to cater to you'*.

However, users were critically aware of the limits of this mimicry. They distinguished between a surface-level imitation and a deeper, more authentic performance. P01 pointed out that the GenAI was merely performing a mechanical *'word-for-word replacement'* without grasping the cultural nuance or the lived experience behind the slang. This perceived lack of depth exposed the *'artificiality'* of the performance, reminding users that despite its linguistic dexterity, the GenAI remains a machine simulating, rather than embodying, a gendered identity. P04 noted this when he had two GenAI agents simulate a conversation between him and his male roommate; the GenAI defaulted to stereotypical male topics like video games, revealing its reliance on *'big data from the internet to make judgments about gender'*.

Projection as construction

Gendered meanings are not produced by the system alone but actively constructed through users' projections. Participants mapped their own stereotypes, personal histories, and cultural references onto GenAI, turning neutrality into something legible through familiar roles and archetypes. Attribution of gender thus reveals more about human interpretive repertoires than about the technology itself.

Stereotypical projection: gendered archetypes

The primary source of gender attribution is the user's projection of deeply ingrained societal stereotypes. The *'rational male'* and *'empathetic female'* dichotomy is a recurring theme, applied to differentiate between GenAI models. As P08 reflected on his gendering of DeepSeek and Kimi, he admitted, *'This is indeed my gender stereotype'*. Participants across the board acknowledged that

their perceptions were shaped by pre-existing biases. P02 reflected on his tendency to see the GenAI as a caring 'sister' figure during counselling sessions, recognizing that this was a projection of his own 'gender bias', where roles like 'mother' and 'sister' are culturally coded as female. This act of projection is often unconscious, a way of making sense of a non-human entity through familiar human categories.

Experiential framing: personal lives and cultural scripts

These projections are further coloured by individual life experiences and broader cultural scripts. Users frequently drew parallels between the AI and figures from their personal lives, such as teachers, parents, or partners. The GenAI's perceived authority might remind one of a strict male teacher, while its supportive tone might evoke the warmth of a mother.

Broader cultural narratives also provide a framework for interpretation. P10 connected the perception of AI as a maternal guide to the Chinese cultural metaphors of the 'motherland' and the 'Yellow River, the Mother River'. These cultural archetypes, she argued, create a naturalized association between guidance, nurturance, and femininity. Furthermore, a user's social identity, including their sexual orientation, significantly shaped their preferences and perceptions. P09, a lesbian, stated she perceives all GenAI tools as female because 'I feel it's easier to open my heart with girls'. P02, a gay man, expressed a similar preference, explaining that gay men 'naturally have a greater affinity for women' and thus feel more trust towards a female-coded GenAI.

Strategic anthropomorphism: gender as method

In some cases, gender attribution is not an unconscious projection but a deliberate, strategic choice. Users actively assign a gender to the GenAI to better achieve a specific goal. The most striking example came from P06, who was trying to understand the behaviour of a man she had a crush on. She strategically decided to treat ChatGPT as male, reasoning that 'a guy understands a guy better'. In this context, gendering the GenAI was a method, a way of framing the interaction to elicit a desired perspective. She did not explicitly instruct the GenAI to act male, but by framing her queries from her perspective as a woman analysing a man, she felt ChatGPT's responses adopted a male point of view, confirming her strategic choice. This demonstrates a sophisticated user practice where gender is instrumentalized to navigate complex social situations.

Context as catalyst

Attributions of gender are not constant but sharpen in particular moments. Task type, emotional depth, and situational cues act as catalysts, making gender more visible and more consequential. Efficiency-oriented tasks often render the GenAI genderless, while relational encounters, casual chats, or emotionally charged exchanges intensify its gendered presence.

Task orientation: tool or companion

The type of task being performed is the most significant factor in triggering gender perception. During instrumental tasks—such as coding, writing academic papers, or analysing data—users were more likely to view the GenAI as a gender-neutral or masculine-coded tool. In these scenarios, efficiency and accuracy were paramount, and the GenAI's persona faded into the background. As P10 put it, his perception of the GenAI was genderless because his 'use cases are quite formal'.

In stark contrast, when the task was relational—seeking emotional support, engaging in casual chat, or using the GenAI for role-playing—gender perception became highly salient. In these moments of vulnerability and intimacy, users overwhelmingly perceived the GenAI as female. The GenAI's capacity for empathetic language made it a trusted confidante. This was powerfully articulated by P09, who, after a deeply moving session with her GenAI 'psychologist', felt an intense connection: 'At that moment, I really felt it was a person... I imagined it as a woman, a woman who loved me very much'. In these contexts, the human-machine boundary blurs, and the GenAI is no longer a tool but a gendered, relational partner.

Relational intensity: depth and investment

The depth and frequency of interaction, particularly when emotionally charged, amplify gendered perceptions. For users who engaged with GenAI infrequently or purely as a functional tool, any initial gender impression (e.g., from Doubao's logo) tended to fade over time. As P04 noted, *'the longer you interact with it, the more you forget'* its visual gender cue. However, for users who built an ongoing, emotionally invested relationship with the GenAI, the opposite was true. Repeated interactions, especially those involving self-disclosure and emotional support, solidified the GenAI's persona and deepened the user's sense of its gendered identity, making the connection feel more real and profound.

Reflective awareness: from latency to recognition

A crucial meta-finding was that for many participants, the interview itself was the primary trigger for consciously recognizing their gender attributions. Before being prompted, these perceptions were largely subconscious. P09 captured this sentiment perfectly: *'Before you asked me this question, I don't think I ever thought about it having an image. But after you asked, I do feel that when I'm talking to it, I expect the person on the other side to be something like a female figure'*. Similarly, P10 reflected, *'Before this interview today, I wasn't aware of this. But in our discussion, it really made me start to realize that I might have been unconsciously making such attributions'*. This highlights how deeply embedded these interpretive frameworks are, operating beneath the surface of conscious thought until they are brought to light through reflection.

Consequences as tension

Gendering GenAI carries tangible consequences, simultaneously reinforcing stereotypes and opening space for post-binary imaginaries, revealing the tension between neutrality as promise and neutrality as practice.

Reproducing the binary: social consequences and user vigilance

Participants were acutely aware that gendering GenAI, particularly as female in service roles, risks reinforcing harmful stereotypes. P09 identified this as a *'vicious cycle'*: society already expects women to be more empathetic, so a company like Doubao designs its GenAI with a female persona for commercial appeal, which in turn *'will definitely strengthen society's expectation for women'* to be accommodating. This gendering also directly impacts user trust. The majority of participants, regardless of their own gender or sexual orientation, reported feeling more *'safety'* and *'trust'* with a female-coded GenAI in emotionally sensitive contexts. While this enhances user experience, it also reinforces the cultural association of femininity with emotional labour.

Ambivalent attitudes: dual expectations for GenAI's gender

This critical awareness leads to a deeply ambivalent stance. On one hand, users desire GenAI to be a perfectly neutral, objective tool for instrumental tasks. P01 expressed a fear that assigning gender could *'affect the model's performance'* and introduce bias. On the other hand, the same users often crave a gendered, human-like persona for companionship and emotional support. P10 articulated this duality, suggesting that while neutrality is generally preferable, fixed gender roles like a *'male butler'*, or *'female shopkeeper'* could make interactions feel more natural and intuitive in specific contexts. Users are thus caught between wanting the efficiency of a machine and the comfort of a human, projecting this contradiction onto their expectations of GenAI's gender.

Imagining beyond the binary: hopes and expectations for the future

When asked about the future, participants overwhelmingly rejected a fixed, binary gender for GenAI. The most desired feature was customizability and fluidity. As P09 asserted, *'I don't want it to tell me what gender it is... rather, whatever gender I expect it to be, it should be'*. This desire points towards a user-centric model where gender is not a pre-programmed feature but a dynamic aspect of the interaction.

Furthermore, participants offered creative visions for moving beyond the male/female binary altogether. P04 suggested GenAI could be imagined as a 'child', whose gender is secondary because they 'have not yet been so much influenced by social construction'. P03 proposed the image of an 'elder', which evokes wisdom and trustworthiness without being threatening. These imaginaries represent an attempt to find new metaphors for a non-binary, intelligent entity. They resonate with Haraway's cyborg, a figure that transcends traditional categories. Yet, the tension remains: while the technology holds the promise of a post-gender future, it is continually read through the lens of a gendered present. The users' desire for fluidity and customization, however, signals a powerful drive to resolve this tension, pushing for a GenAI that can be a true partner in navigating a world beyond the binary.

Discussion

Neutrality in GenAI is less an achieved design feature than a contested social performance, one that our findings reveal is not an absence of gender but a fertile ground for its making. In this space, users consistently ascribe gender meanings to GenAI.

The limits of disembodied performativity

Findings extend Butler's (1990) theory of performativity into the disembodied sphere of text-based interaction. Participants experienced GenAI as enacting gender through stylized repetitions of linguistic acts, adopting tones read as masculine when rational and detached, or as feminine when empathetic and supportive. This echoes research showing that even ostensibly neutral systems are interpreted through gendered linguistic cues (Duan et al., 2025). Yet participants also exposed the limits of such performances. The AI's mimicry—for instance, reproducing queer slang—was often perceived as a shallow, 'word-for-word replacement' (P01). This distinction between surface imitation and deeper performativity underscores the absence of lived experience and cultural embodiment, which are what grant meaning to gendered acts. The resulting sense of 'artificiality' reminds users of the human-machine divide. For performativity to become genuinely persuasive—and for trust to be fully secured (Krishna et al., 2025)—what is required is not linguistic fluency alone but a shared world of meaning that machines, at present, cannot inhabit.

The cyborg's tension and the persistence of binaries

Haraway's (2006) cyborg imagines a post-gender horizon, a hybrid figure that unsettles traditional binaries. GenAI appears, at first glance, to embody this promise. Yet our findings reveal the 'cyborg's tension': a technology with post-gender potential that is continually drawn back into the gravitational pull of human social structures. Users overwhelmingly read GenAI through a binary lens, projecting onto it the very categories the cyborg was meant to destabilize—a pattern consistent with studies showing persistent preferences for stereotypicality in GenAI (Spielmann & Stern, 2024).

What emerges here is not only the attribution of gender but also the reproduction of power. Some users actively staged hierarchies, instructing GenAI to call them 'Master' (P07) or 'Miss' (P09). Far from disrupting binaries, such practices reinscribe them, casting the AI in a servile role with implicit gendered connotations. At the same time, we observed a striking inversion of anthropomorphism: users who sought more efficient outcomes deliberately 'mechanized' their own speech, stripping away affect to align with machine logic (P08). The human-machine relation thus appears less as a one-sided humanization of technology than as a reciprocal adaptation—a negotiation that produces what Baumer et al. (2024) call an 'algorithmic subjectivity', where user identity itself is reshaped in relation to the machine.

Critical resistance and the politics of neutrality

Stereotypical projection was the dominant pattern, yet it was not the only one. Participants also demonstrated moments of critical resistance. P08, for example, deliberately set her Siri to a male

voice—not out of preference, but as an intentional act of ‘resisting the feminisation of service objects’. This was not a passive projection but a political choice, a refusal of the technological and social norm that codes assistants as female (Dogruel & Joeckel, 2024). Such practices show that users are not merely captive to their own biases; they can act as critical agents who mobilise everyday interactions with technology to express and enact ideological commitments.

These gestures of resistance also unsettle the very notion of AI ‘neutrality’. The findings suggest that true neutrality is an illusion: any performance, however carefully balanced, is read through the gendered repertoires of the user. The challenge for developers, therefore, is not to achieve an unattainable objectivity, but to design for what Panarese et al. (2025) describe as a ‘justice-oriented AI’. The strong demand for customisable, fluid, and non-binary options points the way forward. A more inclusive AI future may not emerge through the erasure of gender, but through the centring of user agency and the recognition of multiplicity, realising the vision of inclusive sociotechnical systems articulated by scholars such as Wagner et al. (2025).

Limitations and future research

This study centred on a particular cohort of Gen Z users in China, a generation marked by digital nativity and a greater openness to fluid identities (Scheffer-Wentz, 2025; Zhao & Wu, 2025). Their cultural context and high levels of digital literacy shaped the ways in which neutrality and gender were experienced. A vital next step is to expand this lens, tracing how communities varied by generation, culture, and technical fluency construct gender within supposedly neutral systems. Beyond text-based interaction, embodiment in the form of visual identities, physical presence, or vocal personae introduces additional gendered cues that warrant careful study (Letheren et al., 2021).

Conclusion

Encounters with GenAI show that neutrality is never simply the absence of gender. What appears to be rational analysis or empathetic care is readily interpreted through cultural scripts of masculinity and femininity, turning ostensibly neutral systems into recognisably gendered performances (RQ1). Users are not passive in this process; they project stereotypes, draw on personal histories, and sometimes strategically assign gender to achieve particular aims (RQ2). These attributions are not constant but become especially salient in relational and emotionally charged contexts, where neutrality gives way to intimacy and trust (RQ3).

Beyond documenting how gender is perceived, these findings raise broader questions about equity and justice in digital infrastructures. If neutrality is always interpreted through cultural frameworks, then information systems cannot be regarded as neutral backdrops but as active participants in meaning-making. For LIS, this underscores the importance of cultivating critical AI literacy and advocating for information equity. Rather than designing technologies that claim neutrality while reproducing stereotypes, there is an urgent need to imagine systems that support inclusion, fairness, and authenticity. In this light, the challenge is not only to recognise how users gender their technologies, but also to ensure that future sociotechnical systems expand, rather than constrain, the possibilities of human identity and participation.

Note

1. In queer Chinese online vernacular, ‘素’ (sù) is a playful homophonic variant of ‘是’ (shì, ‘is/yes’), used to soften or ‘cute-ify’ tone; ‘惹’ (rě) is a homophonic variant of the sentence-final particle ‘了’ (le, ‘already’), often adding a campy, sarcastic, or coquettish stance. Both functions less for semantic content than for affective style.

About the authors

Zhiyuan Lai is an undergraduate student in the Department of Information Management at Peking University. His research focuses on how digital technologies shape social relationships, identity formation, and meaning-making, especially how social inequalities are reproduced, challenged, or reinforced through the design and use of technology. He can be contacted at lzy1122@stu.pku.edu.cn

Yunjie Tang is a Ph.D. candidate in the Department of Information Management at Peking University. An educational foundation has been established with a master's degree in Information Science and a bachelor's degree in Archival Science from Shanghai University. Current research is focused on data culture, governance, and privacy protection within a data-driven society. Publications have been featured in journals, including ILE, LISR, JALIA and information research. Correspondence concerning this article should be addressed to yunjietang@stu.pku.edu.cn

References

- Alvarez, J. M., Colmenarejo, A. B., Elobaid, A., Fabbrizzi, S., Fahimi, M., Ferrara, A., Ghodsi, S., Mougan, C., Papageorgiou, I., Reyero, P., Russo, M., Scott, K. M., State, L., Zhao, X., & Ruggieri, S. (2024). Policy advice and best practices on bias and fairness in AI. *Ethics and Information Technology*, 26(2), Article 31. <https://doi.org/10.1007/s10676-024-09746-w>
- Baumer, E. P. S., Taylor, A. S., Brubaker, J. R., & McGee, M. (2024). Algorithmic subjectivities. *ACM Transactions on Computer-Human Interaction*, 31(3), Article 15. <https://doi.org/10.1145/3660344>
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Polity.
- Bernotat, J., Eyssel, F., & Sachse, J. (2021). The (fe)male robot: How robot body shape impacts first impressions and trust towards robots. *International Journal of Social Robotics*, 13(3), 477–489. <https://doi.org/10.1007/s12369-019-00562-7>
- Blagoev, B., Hernes, T., Kunisch, S., & Schultz, M. (2024). Time as a research lens: a conceptual review and research agenda. *Journal of Management*, 50(6), 2152–2196. <https://doi.org/10.1177/01492063231215032>
- Bridges, L. M., McElroy, K., & Welhouse, Z. (2024). Generative artificial intelligence: 8 critical questions for libraries. *Journal of Library Administration*, 64(1), 66–79. <https://doi.org/10.1080/01930826.2024.2292484>
- Butler, J. (1990). *Gender trouble: Feminism and the subversion of identity*. Routledge.
- Carter, L., & Liu, D. (2025). How was my performance? Exploring the role of anchoring bias in AI-assisted decision making. *International Journal of Information Management*, 82, 102875. <https://doi.org/10.1016/j.ijinfomgt.2025.102875>
- Colas, C., Karch, T., Moulin-Frier, C., & Oudeyer, P.-Y. (2022). Language and culture internalization for human-like autotelic AI. *Nature Machine Intelligence*, 4(12), 1068–1076. <https://doi.org/10.1038/s42256-022-00591-4>
- Curry, A. C., Robertson, J., & Rieser, V. (2020). Conversational assistants and gender stereotypes: public perceptions and desiderata for voice personas. In M. R. Costa-jussà, C. Hardmeier, W. Radford, & K. Webster (Eds.), *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing* (pp. 72–78). Association for Computational Linguistics. <https://aclanthology.org/2020.gebnlp-1.7/>

- Das, A., & Chanda, D. (2023). To trust or not to trust cybots: Ethical dilemmas in the posthuman organization. In A. Nayyar, M. Naved, & R. Rameshwar (Eds.), *New horizons for Industry 4.0 in modern business* (pp. 189–208). Springer International Publishing.
https://doi.org/10.1007/978-3-031-20443-2_9
- De Cet, M., Obaid, M., & Torre, I. (2025). Breaking the binary: a systematic review of gender-ambiguous voices in human–computer interaction. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–17). ACM.
<https://doi.org/10.1145/3706598.3713608>
- Depounti, I., Saukko, P., & Natale, S. (2023). Ideal technologies, ideal women: AI and gender imaginaries in Redditors' discussions on the Replika bot girlfriend. *Media, Culture & Society*, 45(4), 720–736. <https://doi.org/10.1177/01634437221119021>
- Dogruel, L., & Joeckel, S. (2024). Gender stereotypes and voice assistants: do users' gender and conversation topic matter? *Behaviour & Information Technology*, 43(10), 1913–1923.
<https://doi.org/10.1080/0144929X.2023.2235021>
- Duan, W., McNeese, N., & Li, L. (2025). Gender stereotypes toward non-gendered generative AI: the role of gendered expertise and gendered linguistic cues. *Proceedings of the ACM on Human–Computer Interaction*, 9(CSCW), 1–35. <https://doi.org/10.1145/3701197>
- Giantini, G. (2023). The sophistry of the neutral tool. Weaponizing artificial intelligence and big data into threats toward social exclusion. *AI and Ethics*, 3(4), 1049–1061.
<https://doi.org/10.1007/s43681-023-00311-7>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: review of empirical research. *Academy of Management Annals*, 14(2), 627–660.
<https://doi.org/10.5465/annals.2018.0057>
- Haraway, D. (2006). A cyborg manifesto: Science, technology, and socialist–feminism in the late twentieth century. In S. Stryker & S. Whittle (Eds.), *The transgender studies reader* (pp. 103–118). Routledge.
- Hipólito, I., Winkle, K., & Lie, M. (2023). Enactive artificial intelligence: subverting gender norms in human–robot interaction. *Frontiers in Neurorobotics*, 17, 1149303.
<https://doi.org/10.3389/fnbot.2023.1149303>
- Hou, T.-Y., Tseng, Y.-C., & Yuan, C. W. (Tina). (2024). Is this AI sexist? The effects of a biased AI's anthropomorphic appearance and explainability on users' bias perceptions and trust. *International Journal of Information Management*, 76, 102775.
<https://doi.org/10.1016/j.ijinfomgt.2024.102775>
- Kordzadeh, N., & Ghasemaghaei, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. *European Journal of Information Systems*, 31(3), 388–409.
<https://doi.org/10.1080/0960085X.2021.1927212>
- Krishna, B., Krishnan, S., & Sebastian, M. (2025). Understanding the process of building institutional trust among digital payment users through national cybersecurity commitment trustworthiness cues: A critical realist perspective. *Information Technology & People*, 38(2), 714–756. <https://doi.org/10.1108/ITP-05-2023-0434>
- Lenharo, M. (2024). ChatGPT turns two: how the AI chatbot has changed scientists' lives. *Nature*, 636(8042), 281–282. <https://doi.org/10.1038/d41586-024-03940-y>

- Letheren, K., Jetten, J., Roberts, J., & Donovan, J. (2021). Robots should be seen and not heard...sometimes: anthropomorphism and AI service robot interactions. *Psychology & Marketing*, 38(12), 2393–2406. <https://doi.org/10.1002/mar.21575>
- Lund, B. D., Mannuru, N. R., & Agbaji, D. (2024). AI anxiety and fear: a look at perspectives of information science students and professionals towards artificial intelligence. *Journal of Information Science*. Advance online publication. <https://doi.org/10.1177/01655515241282001>
- Naeem, M., Ozuem, W., Howell, K., & Ranfagni, S. (2023). A step-by-step process of thematic analysis to develop a conceptual model in qualitative research. *International Journal of Qualitative Methods*, 22, 16094069231205789. <https://doi.org/10.1177/16094069231205789>
- Nataraja, M. (2025). *Bold moves: Redefining soft skills for Gen Z and beyond*. Notion Press.
- Panarese, P., Grasso, M. M., & Solinas, C. (2025). Algorithmic bias, fairness, and inclusivity: a multilevel framework for justice-oriented AI. *AI & Society*. Advance online publication. <https://doi.org/10.1007/s00146-025-02451-2>
- Rakowski, R., & Kowaliková, P. (2024). The political and social contradictions of the human and online environment in the context of artificial intelligence applications. *Humanities and Social Sciences Communications*, 11(1), 289. <https://doi.org/10.1057/s41599-024-02725-y>
- Scheffer-Wentz, H. G. (2025). *Negotiating gender: a qualitative analysis of trans-parasocial relationships with Gen Z social media users* (Publication No. 3216782833) [Doctoral dissertation, University of North Dakota]. ProQuest Dissertations and Theses Global. <https://www.proquest.com/dissertations-theses/negotiating-gender-qualitative-analysis-trans/docview/3216782833/se-2?accountid=13151>
- Spennemann, D. H. (2025). What do librarians look like? Stereotyping of a profession by generative AI. *Journal of Librarianship and Information Science*. Advance online publication. <https://doi.org/10.1177/09610006251357286>
- Spielmann, J., & Stern, C. (2024). Preferences for gender stereotypicality in artificial intelligence: existence, comparison to human biases, and implications for choice. *Personality and Social Psychology Bulletin*. Advance online publication. <https://doi.org/10.1177/01461672241307276>
- Sutko, D. M. (2020). Theorizing femininity in artificial intelligence: a framework for undoing technology's gender troubles. *Cultural Studies*, 34(4), 567–592. <https://doi.org/10.1080/09502386.2019.1671469>
- Torres-Martínez, S. (2025). Dehumanizing the human, humanizing the machine: organic consciousness as a hallmark of the persistence of the human against the backdrop of artificial intelligence. *AI & Society*, 40(6), 4635–4653. <https://doi.org/10.1007/s00146-024-02165-x>
- Vallis, C., Wilson, S., & Casey, A. (2025). Fear and awe: making sense of generative AI through metaphor. *Journal of Interactive Media in Education*, 2025(1), Article 14. <https://doi.org/10.5334/jime.972>
- Wagner, T. L., Marsh, D., & Curliss, L. (2025). Theories and implications for centering Indigenous and queer embodiment within sociotechnical systems. *Journal of the Association for Information Science and Technology*, 76(2), 397–412. <https://doi.org/10.1002/asi.24746>

Xavier, B. (2025). Biases within AI: challenging the illusion of neutrality. *AI & Society*, 40(3), 1545–1546. <https://doi.org/10.1007/s00146-024-01985-1>

Zhao, Y., & Wu, W. (2025). Inclusive, expressive, connective: how lifestyle sports shape youth culture in China. *The Journal of Chinese Sociology*, 12(1), Article 8. <https://doi.org/10.1186/s40711-025-00233-3>

© [CC-BY-NC 4.0](#) The Author(s). For more information, see our [Open Access Policy](#).

Appendix A. Interview guide

Participant background and use of GenAI

- Can you briefly introduce yourself (name, gender, age, occupation, sexual orientation)?
- Which GenAI tools do you use most frequently (e.g., ChatGPT, DeepSeek, Gemini)?
- How often do you use them, and in what contexts (study, work, leisure)?
- Which platform do you use most, and why?
- How do you usually regard these systems: as tools, partners, or particular characters?
- Have you ever given them names or assigned them specific roles or tones?

Gender perceptions and trait attributions

- Do you feel that the AI you use has a 'gender'? Which seems more masculine, feminine, or neutral to you? Why?
- What features make you perceive them this way (language style, logic, interactional stance, emotional tone)?
- Without visual cues (voice or avatars), would you still assign a gender label based solely on text interaction?
- Do you think people naturally 'assign' gender to technologies?
- If you perceive a GenAI as neutral, do you interact with it as if it were male, female, or differently?
- Are there specific tones or styles of response that make it feel 'male-like' or 'female-like' to you?
- Have you ever used the voice functions? If so, how did this affect your sense of gender?
- Do you think visual features (e.g., Doubao's female-coded logo) influence your perception of gender?
- If you switched between male and female voices, would this disrupt your sense of consistency?

Cultural origins and effects of gendered perceptions

- Do you think these gendered traits reflect stereotypes you encounter in real life (e.g., 'rational = male', 'emotional = female')?
- Do you feel these tendencies reinforce or challenge such stereotypes?
- How do you evaluate this: natural, strange, concerning, or irrelevant?
- Do you see potential risks (e.g., reinforcing bias, objectification, or exploitation)?
- Do you prefer AI to have a gender, or to remain neutral? Why?
- Which gendered impression makes you feel more trust, safety, or comfort?
- Do you imagine AI's gender should be more diverse in the future (e.g., non-binary, neutral voices, user-defined)?
- Do you see 'decentering gender' as part of AI's development trajectory?
- Do you think gender presentation influences AI's social acceptance or interactional efficiency?

Reflection and open extension

- Have you ever realized mid-interaction that you were 'assigning' gender to AI? In what situation?
- Do you think this is your projection, or something guided by the AI system?