



Original-Forschungsarbeit

KI als Grenzobjekt: der persische X-Diskurs

Shaho Sabbar¹

1 Assistenzprofessor, Abteilung für Iranistik, Fakultät für Weltstudien, Universität Teheran, Teheran, Iran

Empfangen: 1. März 2025 **Akzeptiert:** 7. Juni 2025

Zusammenfassung:

Diese Studie untersucht, wie persischsprachige Nutzer auf der Social-Media-Plattform X generative künstliche Intelligenz als sozio-technisches und diskursives Phänomen verhandeln. Auf der Grundlage eines Datensatzes von 24.215 persischsprachigen Beiträgen verwenden wir ein Multi-Label-Topic-Modeling-Verfahren sowie affektives Profiling, um den öffentlichen Diskurs über KI-Werkzeuge, ihre wahrgenommenen Implikationen und normative Bewertungen ihrer Nutzung zu analysieren. Anstatt Stimmung als statischen Indikator von Meinungen zu betrachten, interpretieren wir affektiven Ausdruck als kommunikativen Akt, der durch plattformspezifische Anreize und kulturelle Kontexte geprägt ist. Unsere Ergebnisse zeigen, dass KI nicht nur als technisches Artefakt positioniert wird, sondern auch als Grenzobjekt, das mit Debatten über Expertise, Ethik und institutionelle Legitimität verflochten ist. Der Diskurs ist in praktischen Anliegen verankert – insbesondere in Bezug auf Arbeit, Bildung und Vergleiche zwischen KI-Werkzeugen –, erweitert sich jedoch häufig zu kulturspezifischen Narrativen über Risiko, Fairness und epistemische Autorität. Emotional ist die Diskussion durch pragmatischen Optimismus, kritische Intensität und ein beträchtliches neutrales Spektrum gekennzeichnet, das eher Orientierung als Bewertung widerspiegelt. Diese Studie trägt zu aktuellen Debatten in der Kommunikationswissenschaft, der KI-Ethik und den Plattformstudien bei, indem sie eine nicht anglophone, kulturell verankerte Analyse dafür liefert, wie Öffentlichkeiten eine alltagsprachliche Governance über aufkommende Technologien praktizieren.

Schlüsselwörter: künstliche Intelligenz, Sentimentanalyse, digitale Öffentlichkeiten, alltagsprachliche Governance, persischer Diskurs

* Korrespondierender Autor

✉ shaho.sabbar@ut.ac.ir

🌐 <https://orcid.org/0000-0001-5801-7137>

Wie dieser Artikel zu zitieren ist:

Sabbar, S. (2025). AI as a boundary object: The Persian X discourse. *Spektrum Iran*, 38(2), 61-82.

🔗 <https://doi.org/10.22034/spektrum.2026.569202.1059>

© 0 0 Copyright © Der/die Autor(en); Dieses Werk ist lizenziert unter einer Creative Commons Namensnennung - Nicht kommerziell - Keine Bearbeitungen 4.0 International (CC-BY-NC) Lizenz. Homepage: www.spektrumiran.com

هوش مصنوعی به مثابه ابژه مرزی: گفتمان فارسی در ایکس

شاهو صبار^۱

۱ استادیار، گروه مطالعات ایران، دانشکده مطالعات جهان، دانشگاه تهران، تهران، ایران

دریافت: ۱۴۰۳/۱۲/۱۱ پذیرش: ۱۴۰۴/۰۳/۱۷

چکیده:

این مطالعه بررسی می‌کند که کاربران فارسی‌زبان در پلتفرم رسانه اجتماعی ایکس چگونه با هوش مصنوعی مولد به عنوان پدیده‌ای اجتماعی-فنی و گفتمانی تعامل می‌کنند. ما با بهره گرفتن از مجموعه داده‌ای شامل ۲۴۲۱۵ پست فارسی، از چارچوب مدل‌سازی موضوعی چندبرچسبی و نیز پروفایل‌سازی عاطفی برای تحلیل گفتمان عمومی درباره ابزارهای هوش مصنوعی، پیامدهای ادراک شده آن‌ها و داوری‌های هنجاری درباره کاربردشان استفاده کرده‌ایم. به جای آنکه احساسات را به عنوان شاخصی ایستا از نگرش‌ها در نظر بگیریم، بیان عاطفی را به مثابه کنشی ارتباطی تفسیر می‌کنیم که تحت تأثیر مشوق‌های پلتفرمی و بافت فرهنگی شکل می‌گیرد. یافته‌های ما نشان می‌دهد که هوش مصنوعی نه تنها به عنوان یک مصنوع فنی، بلکه به عنوان ابژه‌ای مرزی بازنمایی می‌شود که با مباحث مربوط به تخصص، اخلاق و مشروعیت نهادی درهم تنیده است. این گفتمان در نگرانی‌های عملی — به‌ویژه در حوزه کار، آموزش و مقایسه میان ابزارهای هوش مصنوعی — ریشه دارد، اما اغلب به روایت‌های فرهنگی خاص درباره ریسک، انصاف و مرجعیت معرفتی بسط می‌یابد. از نظر عاطفی، این گفت‌وگو با نوعی مثبت‌اندیشی عمل‌گرایانه، شدت انتقادی و طیف قابل توجهی از موضع خنثی مشخص می‌شود که بیشتر نشان‌دهنده جهت‌گیری است تا ارزیابی. این مطالعه با ارائه تحلیلی فرهنگی و غیرانگلیسی از نحوه اعمال حکمرانی عامیانه از سوی عموم بر فناوری‌های نوپدید، به مباحث جاری در حوزه ارتباطات، اخلاق هوش مصنوعی و مطالعات پلتفرم‌ها کمک می‌کند.

واژگان کلیدی: هوش مصنوعی، تحلیل احساسات، عمومیت‌های دیجیتال، حکمرانی عامیانه، گفتمان فارسی



Original Research Paper

AI as a boundary object: The Persian X discourse

Shaho Sabbar^{1*}

¹ Assistant Professor, Department of Iranian Studies, Faculty of World Studies, University of Tehran, Tehran, Iran

Received: Mar., 1, 2025 Accepted: Jun. 7, 2025

Abstract

This study investigates how Persian-speaking users on the social media platform X engage with generative artificial intelligence as a sociotechnical and discursive phenomenon. Drawing on a dataset of 24,215 Persian-language posts, we employ a multi-label topic modeling framework and affective profiling to analyze public discourse surrounding AI tools, their perceived implications, and normative judgments about their use. Rather than treating sentiment as a static indicator of opinion, we interpret affective expression as a communicative act shaped by platform incentives and cultural context. Our findings show that AI is positioned not only as a technical artifact but as a boundary object entangled with debates over expertise, ethics, and institutional legitimacy. The discourse is anchored in practical concerns – especially labor, education, and comparisons among AI tools – but frequently extends into culturally specific narratives about risk, fairness, and epistemic authority. Emotionally, the conversation is marked by pragmatic positivity, critical intensity, and a sizable neutral band reflecting orientation rather than evaluation. This study contributes to ongoing debates in communication, AI ethics, and platform studies by offering a non-Anglophone, culturally grounded analysis of how publics perform vernacular governance over emerging technologies.

Keywords: artificial intelligence, sentiment analysis, digital publics, vernacular governance, Persian discourse

* Corresponding Author

✉ shaho.sabbar@ut.ac.ir

🌐 <https://orcid.org/0000-0001-5801-7137>

How to Cite this Article:

Sabbar, S. (2025). AI as a boundary object: The Persian X discourse. *Spektrum Iran*, 38(2), 61-82.

📄 <https://doi.org/10.22034/spektrum.2026.569202.1059>

© 2025 Copyright © The Author(s); This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC-BY-NC) License. Homepage: www.spektrumiran.com

1. Introduction

Throughout history, the emergence of new communication technologies has reconfigured the conditions of public life. From the printing press to the smartphone, such technologies have not only extended communicative capacity but also shaped how individuals perceive themselves, their social roles, and their relationship to knowledge and power. These perceptions are neither uniform nor apolitical; rather, they are deeply entangled with affect, authority, and cultural imaginaries. As Harold Innis (1951) and Marshall McLuhan (1964) argued, media technologies do not merely transmit content—they exert structural influence, favoring particular modes of expression and societal organization. Today, the rise of generative artificial intelligence represents a comparable epistemic shift, reframing intimacy, cognition, and expertise within algorithmically mediated systems of meaning-making (Elliott, 2023; Seaver, 2022).

Public reactions to AI are shaped not only by functional evaluations but also by emotional and moral concerns. This affective response is especially visible in the context of "algorithmic intimacy" — the perceived closeness and social presence that some users experience when interacting with AI-driven agents (Bozdağ, 2025; Reineke, 2022). While earlier technologies inspired utopian or dystopian imaginaries about mass democratization or social alienation, AI technologies invite a subtler blend of fascination, ambivalence, and critique than earlier waves of media innovation (Salehi et al., 2026). Studies of AI-mediated relationships show how users anthropomorphize chatbots, frame them as companions, or even as emotional partners (Balazadeh & Kajonius, 2025; George et al., 2023; Sjoraida, 2025). These engagements reveal that users' understanding of AI is shaped not merely by its accuracy or productivity, but by its capacity to signal empathy, authority, and social intention (Chu et al., 2025; Turkle, 2011).

How people imagine and emotionally respond to AI will, in turn, shape the trajectory of the technology itself (Sabbar & Habib Zadeh Khiyaban, 2023). As Papacharissi (2015) has argued, the circulation of sentiment in digital spaces helps constitute "affective publics" — collectives organized not only around shared beliefs but around shared intensities. In this light, public perception is not simply a reflection of technological progress; it is a force that can accelerate, redirect, or constrain it. Expectations about AI's ethical

boundaries, usefulness, and risks inform not only personal adoption but institutional policies and commercial design choices. Intimate forms of interaction—such as venting to a chatbot or querying it about political beliefs—can crystallize

into collective norms about what AI should and should not do (Illouz, 2007; Lupton, 2022; Shahghasemi, 2025). These expectations travel through discourse, anchoring future imaginaries in contemporary discourse.

Social media platforms play a central role in mediating this discourse. They operate not only as amplifiers of information but as infrastructures of affective and epistemic coordination. The platform X, in particular, has become a key site for the articulation of public sentiment about AI. With its rapid tempo, algorithmic virality, and entanglement of personal and professional identities, X facilitates the emergence of vernacular expertise—informal yet consequential claims about how AI works, what it means, and what it should become. On X, users post not only evaluations of AI systems but narratives, warnings, jokes, and moral commentaries that collectively shape public understanding. This is especially true in non-Anglophone contexts, where localized meanings of AI emerge through culturally embedded discourse practices and political framings.

In the Persian-language digital sphere, the convergence of algorithmic tools and platform discourse has produced a particularly rich and understudied communicative field. As recent studies suggest, Persian-speaking publics increasingly use platforms like X to articulate vernacular forms of technological governance—public deliberations, critiques, and norm-setting that occur outside formal regulatory channels. These discourses are shaped by hybrid media ecologies, where access to global tools is uneven, state censorship remains active, and users often navigate both technical and moral complexities in evaluating digital systems. Generative AI, in this context, is not merely an innovation but a cultural problem: a figure of fascination, suspicion, and speculative projection.

Despite the richness of this public conversation, existing scholarship has largely neglected the ways in which non-Western publics discuss and contest the meanings of AI. Much of the literature on AI discourse remains Anglocentric, focusing primarily on elite media narratives or institutional policy debates. In contrast, everyday public talk—especially in languages

other than English—remains under-theorized, even though such talk plays a formative role in shaping user norms, perceptions of legitimacy, and trust in AI systems (Papacharissi, 2015; Shahghasemi et al., 2025). There is a pressing need to understand how publics outside the Global North interpret AI, not as an abstract technological domain, but as a lived, affectively charged, and socially negotiated object.

This study addresses that gap by analyzing a large-scale corpus of Persian-language posts from X, focusing on how users talk about generative AI across technical, emotional, and normative registers. Building on theories of affective publics, algorithmic intimacy, and vernacular governance, we examine AI discourse not only for its content but for its communicative form: how users perform expertise, express sentiment, and negotiate legitimacy. We treat affective expressions not as raw indicators of approval or disapproval but as socially situated acts—moves within a field of public reasoning, identity performance, and epistemic claim-making (van Dijk, 1993; Bourdieu, 1991).

Our approach combines multi-label topic modeling with affective profiling, allowing us to map both the thematic structure and emotional tone of AI discourse among Persian-speaking users. Rather than imposing normative assumptions about technological progress or ethical risk, we seek to surface how publics themselves articulate what matters—whether in concerns about labor displacement, academic integrity, misinformation, or emotional authenticity. These findings contribute to ongoing conversations in AI ethics, communication studies, and cultural sociology by grounding analysis in the actual speech practices of users navigating new technological realities.

Specifically, this study is guided by the following research questions:

1. What thematic concerns dominate Persian-language discourse about generative AI on the X platform?
2. How do users express affect in their evaluations of AI tools, and what emotional patterns characterize the discourse?
3. In what ways do users perform vernacular governance—norm-setting, critique, and pedagogy—through their interactions with AI discourse?

By addressing these questions, we aim to illuminate how publics not only react to AI but participate in shaping its social meaning. In doing so, we position generative AI as a boundary object—flexible enough to accommodate multiple interpretations but structured enough to anchor contestation (Star & Griesemer, 1989). The discourse around AI is not merely a reflection of public opinion; it is a site of cultural labor where values, expectations, and technological futures are being actively negotiated.

2. Methodology

To examine how Persian-speaking users discuss and evaluate artificial intelligence across every day, institutional, and speculative registers, we compiled a large corpus of public posts from X (formerly Twitter). Data collection followed a high recall retrieval logic: the goal was to minimize thematic blind spots by casting a wide net around AI-related talk rather than narrowly sampling only explicit technical discussions. We therefore combined multiple families of Persian and English keywords and their common transliterations, including general AI terms (e.g., هوش مصنوعی, ai), model and platform names (e.g., ChatGPT, Grok, Gemini, Copilot), interaction and use terms (e.g., پرامپت/Prompt, prompt engineering), and risk or governance cues (e.g., جعل, حریم خصوصی, کپی رایت, مقررات, دیپ فیک/deepfake). Because AI discourse on X frequently involves code switching and brand shorthand, retrieval terms were expanded iteratively during pilot checks to capture vernacular variants, spacing differences, and orthographic alternations. This process produced an initial pool of 26,111 posts. Preprocessing was designed to preserve discursive and affective signals while removing mechanical noise and reducing distortions produced by platform repetition. First, duplicate and near duplicate posts were removed, along with items that were not analytically usable for text-based modeling (for example, posts with no substantive text). Second, we filtered out posts that were not meaningfully about AI despite containing ambiguous trigger terms, using a combination of rule-based exclusions and manual spot checks to prevent systematic drift into irrelevant domains. Third, we normalized Persian orthography to improve the stability of downstream matching and classification. This included harmonizing Arabic/Persian character variants (e.g., ی/ی and ک/ك), handling zero-width non-joiners, and stripping diacritics.

Links and user mentions were removed to reduce sparsity and to avoid overweighting platform-specific identifiers. Throughout cleaning, we prioritized retaining the expressive content of posts, including colloquial phrasing and common mixed language forms, because these features carry much of the interactional meaning on X.

After deduplication, relevance filtering, and normalization, the final analytic dataset contained 24,215 posts. All records were drawn from public content and are reported only in aggregate. In reporting examples, we avoid including user identifiers or content that could facilitate re-identification. This curated corpus provides the empirical foundation for subsequent steps in the analysis, including multi-label topic tagging across ten thematic formations and the mapping of co-occurrence structures and affective profiles.




Figure 1. Distribution of text lengths (in words)

3. Findings

Applying the five-category affect schema to the full analytic corpus ($n = 24,215$) produces a distribution in which evaluative warmth is present but far from uniform. The largest share of posts falls into Happy (8,271; 34.16%), followed closely by Neutral (7,838; 32.37%). More emphatic positivity, coded as Delighted, accounts for 2,400 posts (9.91%). On the negative side, Angry comprises 2,848 posts (11.76%), and Furious 2,858 (11.80%). Taken together, Angry + Furious represent 5,706 posts (23.56%), while Happy + Delighted

represent 10,671 posts (44.07%). This pattern suggests a public conversation that is neither celebratory nor uniformly alarmist. The prominence of Happy relative to Delighted points to a largely pragmatic register: users often signal approval, usefulness, or everyday satisfaction without escalating into unequivocal enthusiasm. At the same time, the sizeable Neutral band indicates that much AI talk on X functions as circulation and orientation rather than overt evaluation: sharing updates, noting capabilities or failures, relaying comparisons, and positioning the self as an observer in a fast-moving informational environment. Negative affect is nonetheless substantial. The combined quarter of posts coded as Angry or Furious signals that AI is also a recurring site of friction and moral contestation, where users mobilize indignation and outrage as communicative resources, not simply as spontaneous reactions. In the analyses that follow, this distribution motivates treating AI discourse as an effectively mixed field: routine approval and curiosity coexist with sustained critique, and platform dynamics help both travel and settle into recognizable, repeatable stances.




Figure 2. Dataset sentiment distribution

3.1. Topic Modeling and Thematic Analysis

To map how AI is talked about in Persian-language publics on X, we used an unsupervised topic modeling step as a diagnostic lens rather than as the final measurement instrument. After standard text preparation for Persian and code-switched writing (orthographic harmonization, removal of URLs and user mentions, and normalization of common spacing and character variants), we estimated a series of LDA models while varying the number of topics from 5 to 20. We then compared solutions using a semantic coherence criterion to identify a topic resolution that balances empirical fit with interpretability for communication analysis. Coherence values rose unevenly across the sweep and reached their maximum at $K = 10$ (coherence = 0.5241). Smaller local improvements appeared at higher resolutions (e.g., $K = 14$: 0.4903; $K = 19$: 0.4793), but these did not surpass the ten-topic solution. Substantively, the ten-topic model was preferred because it produced categories that correspond to recognizable communicative formations in the corpus while avoiding over-fragmentation. Lower K solutions tended to fuse distinct arenas of AI talk (for example, merging tool comparison chatter with work-related anxiety, or folding governance debates into misinformation talk), which made it harder to track how different kinds of claims and emotions travel. Higher K specifications, by contrast, often split coherent clusters into narrow subthemes that were difficult to name cleanly and were less useful for doctoral-level theorization about platformed publics, affect, and epistemic authority. We therefore adopt the ten-topic solution as the primary thematic frame for the analyses that follow.

Because posts on X frequently span multiple domains in the same utterance (for example, a model comparison embedded in a story about workplace use, or prompting advice paired with ethical boundary testing), we operationalized the thematic frame as multi-label topic tagging rather than forcing each post into a single bucket. The final corpus contains 24,215 posts, and 74.78% of posts are linked to at least one of the ten themes (mean topics per post = 1.21). Overlap is common: 41.39% of posts activate exactly one topic label, 23.27% activate two, and 8.16% activate three; smaller shares activate four or more (1.61% for four, 0.31% for five, 0.04% for six). A minority (25.22%) remains untagged under this topical frame, reflecting either highly generic AI references or content that does not meaningfully align with the ten

major formations. The ten themes and their prevalence in the corpus (non-exclusive) are as follows: Tools and model comparisons (6,324; 26.12%), Work, economy, and business impacts (5,954; 24.59%), Education and university/school contexts (5,648; 23.32%), Future narratives and predictions (3,131; 12.93%), Ethics, privacy, regulation, and governance (2,284; 9.43%), Content creation and creativity (2,135; 8.82%), Deepfakes, misinformation, and verification (1,504; 6.21%), Programming and software development (1,434; 5.92%), Sensitive social and mental health discussions (480; 1.98%), and Prompting practices and strategies (316; 1.30%). Framed this way, the thematic backbone of Persian-language AI discourse is anchored in practical, comparative talk about tools and their everyday implications, while still reserving distinct space for governance concerns, authenticity crises, and the occasional movement into sensitive domains.




Figure 3. Coherence score vs number of topics

3.2. AI Tools Ecosystem and Model Comparisons

Within the corpus, tool-centered discussion constitutes a large and highly performative stream (n = 6,324 posts). Affective tone is split between Neutral (2,190; 34.63%) and Happy (2,140; 33.84%), with a substantial critical layer (Angry 924; 14.61%, Furious 592; 9.36%) and comparatively fewer fully celebratory posts (Delighted 478; 7.56%). Reading the Texts suggests that “comparison” is rarely a calm, laboratory-style evaluation; instead, users

narrate tools through everyday trials, viral anecdotes, and rapid verdicts that travel well on X. Posts juxtapose models by name, circulate claims about new versions, and treat outputs as evidence in microdebates about competence, bias, and reliability. The same thread can move from an impromptu benchmark to a joke about social etiquette with machines, as in the recurring anthropomorphic register (e.g., “My mom apologizes to ChatGPT for taking its time”). Negative affect often clusters around perceived regressions, hallucinations, access and cost frustrations, or the sense that hype is outpacing trustworthy verification, captured in short evaluative utterances such as “Grok is a mess today.” Communication-wise, these posts do reputational work: they elevate some tools as “serious,” demote others as “broken,” and establish vernacular criteria for credibility (speed, accuracy, language handling, refusal behavior) that circulate through quote posting and quick comparisons rather than formal reviews.

3.3. Prompting Practices and Strategies

Although smaller in volume (n = 316), prompting talk is disproportionately instruction-oriented and tends to carry an upbeat tone (Happy 117; 37.03%, Delighted 58; 18.35%) alongside a sizable Neutral band (103; 32.59%) and relatively limited negativity (Angry 20; 6.33%, Furious 18; 5.70%). Qualitatively, the Texts read like a folk pedagogy of interaction: users share templates, warn against vague requests, and present prompting as a communicative skill that turns “a chatbot” into a cooperative worker. Many posts emphasize specifying role, constraints, and format, often in the voice of practical coaching (e.g., “Give it a role, your goal, and the output format, then iterate”). At the same time, prompting discourse becomes a site where norms of legitimacy and accountability are negotiated, especially when prompts are used for coding or workplace outputs; users not only exchange techniques but they also propose boundaries (review, disclosure, and responsibility for errors) that reassert human agency over delegated writing. In communication terms, this topic is metacommunication: it is about how to talk to AI, and it produces a lay vocabulary of clarity, control, and alignment that helps participants position themselves as competent users rather than passive consumers.

3.4. Future Narratives, Scenarios, and Predictions

Future-facing posts (n = 3,131) display a mixed emotional profile, dominated by Neutral (1,098; 35.07%) and Happy (1,062; 33.92%) but with a notable share of high intensity reactions (Furious 334; 10.67%) and a meaningful layer of excitement (Delighted 490; 15.65%), while Angry is comparatively lower (147; 4.69%). In the texts, “the future” functions as a rhetorical device for making present stakes legible: users forecast labor market disruption, reimagine professional identities (especially for developers and knowledge workers), and debate whether the key resource will be skill, access, or infrastructure. Some posts frame AI as a coming competence gap and counsel adaptation through fundamentals and steering capacity (e.g., “Those who master the basics and can direct models will be ahead”), whereas others articulate a more conflictual imaginary in which geopolitical competition, compute, and platform power set the terms of possibility. The future topic also carries an affective oscillation typical of platformed publics: optimism travels via narratives of productivity and opportunity, while outrage concentrates around fears of deception, deskilling, or institutional unpreparedness. Analytically, these posts are less about forecasting accuracy than about social coordination: they align audiences around what to worry about, what to learn, and who should be held responsible as AI is narrated into collective horizons.

3.5. Education, Universities, and Schools

Education-oriented AI talk forms one of the densest thematic streams in the corpus (n = 5,648). Its affective profile is led by Neutral (1,968; 34.84%) and Happy (1,719; 30.44%), with Delighted present but secondary (653; 11.56%) and a substantial negative layer (Angry 579; 10.25%, Furious 729; 12.91%). Reading the texts, education is discussed less as abstract policy and more as a lived site of friction where institutional routines collide with generative tools: posts revolve around classroom practices, teaching, and assignment design, and disputes over what counts as legitimate learning (teaching/class references are frequent in this topic, and cheating or plagiarism cues appear repeatedly). Users describe AI as both tutor and shortcut, often in the same breath, narrating how students draft essays, solve problems, or translate readings while instructors respond by changing task formats, tightening assessment, or demanding process evidence. A recurring communicative

move is moral boundary work: some posts normalize AI as a study aid, while others frame it as an integrity threat that undermines evaluation and devalues credentials. In English paraphrase, a typical stance reads: "If the exam only tests memorization, of course, students will outsource it; redesign the assessment." Another counter-stance warns: "AI makes copying effortless; universities need new verification norms." In short, education discourse is simultaneously pedagogical, regulatory, and affectively ambivalent, with neutrality carrying practical reportage and negative affect marking disputes over fairness, discipline, and institutional lag.

3.6. Labor Market, Economy, and Business Impacts

Work and economic implications constitute a similarly large formation (n = 5,954) and display a mixed but clearly consequential emotional structure: Neutral (1,915; 32.16%) and Happy (1,889; 31.73%) dominate, Delighted is meaningful (782; 13.13%), and high intensity negativity is nontrivial, especially Furious (891; 14.96%) relative to Angry (477; 8.01%). Qualitatively, the Texts frame AI as a pressure point for livelihoods and organizational power. Some posts treat AI as a competitive advantage and a pathway to new income streams, emphasizing speed, output, and market positioning, while others interpret it as a mechanism for restructuring labor, concentrating value, or hollowing out entry-level pathways. Users frequently talk about hiring, job search, and pay in concrete terms, and business-oriented posts often evaluate whether firms should replace contractors, compress teams, or demand higher output with the same wages. The tone shifts depending on who is imagined as the beneficiary: entrepreneurs and managers are more likely to speak in an optimizing register, whereas workers and recent graduates more often voice anxiety or indignation. A common English paraphrase is: "AI will not replace you, but someone using AI will," contrasted with: "Companies will use AI to squeeze workers and call it innovation." From a communication perspective, this topic is where economic imaginaries become moral narratives about responsibility, fairness, and who absorbs risk when productivity is reframed as a personal obligation rather than an institutional decision.

3.7. AI for Programming and Software Development

Developer-facing AI discussion (n = 1,434) centers on the practicalities of coding with models and the evolving identity of the programmer. Its

affective profile is anchored in Neutral (506; 35.29%) and Happy (440; 30.68%), with Delighted (173; 12.06%) and a comparable share of negative sentiment (Angry 161; 11.23%, Furious 154; 10.74%), suggesting a pragmatic field where appreciation coexists with ongoing frustration. Reading the Texts, users treat AI as a pair programmer that excels at boilerplate, debugging assistance, and explanation, but fails in ways that matter for production work: hallucinated APIs, brittle logic, missing edge cases, and confident wrong answers. Debug and bug-related language is common, and many posts emphasize the need for verification, tests, and human judgment rather than blind copying. The discourse repeatedly draws a line between code generation and software engineering: users note that models can propose functions quickly, yet struggle with requirements gathering, architecture, and maintaining coherence across a larger system. A typical English paraphrase is: "It wrote the code in seconds, then I spent an hour fixing subtle mistakes," or "Great for scaffolding, risky for final logic unless you know what you are doing." Importantly, this topic also hosts future-oriented claims about developer status, but expressed through everyday practice: debates over whether fundamentals become more or less important, and whether AI shifts the valued skill from writing code to specifying, reviewing, and integrating it.

3.8. Content Creation and Creativity

Posts tagged in this theme (n = 2,135) frame generative AI primarily as a practical cultural instrument for writing, translating, designing, and remixing, with affect clustering around Neutral (739; 34.61%) and Happy (671; 31.43%), and a smaller but visible Delighted band (262; 12.27%). Negative affect is present yet not dominant (Angry 215; 10.07%, Furious 248; 11.62%), often surfacing when users discuss quality collapse, sameness, or perceived devaluation of human craft. Reading the Texts suggests three recurring communicative genres: first, AI as a fast stylistic assistant (captions, resumes, formal letters, academic paraphrase), where users foreground efficiency and controllability; second, AI as a vernacular studio for images and short videos, frequently described through trial-and-error narratives about what the model can and cannot render; third, playful creativity (poetry, song-like lines, comedic remixes) that positions AI outputs as shareable artifacts in the attention economy. A typical English paraphrase of these posts is instruction-like and performative: "Rewrite this in a formal tone, then give

me three punchy versions,” or “I asked it to generate a poster concept, and it nailed the layout, but the details were wrong.” Across these patterns, creativity is not treated as an inner essence but as a platformed workflow: users negotiate originality, attribution, and taste in public, using AI outputs as both tools and talking points.

3.9. Deepfakes, Misinformation, and Verification

This topic (n = 1,504) carries one of the most conflictual affective profiles in the study, with negativity unusually concentrated (Angry 275; 18.28% and Furious 284; 18.88%, totaling 37.16%), alongside Neutral (418; 27.79%) and Happy (393; 26.13%). Qualitatively, the Texts portray synthetic media as an epistemic stress test for platform-publics: users circulate warnings about manipulated videos and voice cloning, dispute authenticity in real time, and debate what forms of evidence remain trustworthy when “seeing” and “hearing” become cheap to counterfeit. Many posts read as rapid vernacular fact-checking, either advising caution before resharing or offering informal heuristics (reverse search, source triangulation, “wait for confirmation”). An English example that captures the tone is: “Assume it is AI until proven otherwise, do not amplify it,” or “If a voice note can be generated, what counts as proof anymore?” At the same time, a smaller share treats deepfakes as spectacle or novelty, which helps explain the continued presence of Happy and Neutral posts. From a communication perspective, this theme shows how verification labor becomes distributed: responsibility is pushed outward to users and communities through low-cost alerts, while outrage functions as a mechanism for disciplining careless amplification and reasserting norms of evidentiary responsibility on X.

3.10. Ethics, Data Governance, Privacy, and Regulation

Governance-oriented discussion (n = 2,284) is comparatively balanced but persistently normative, led by Neutral (753; 32.97%) and Happy (740; 32.40%), with Delighted (244; 10.68%) and a meaningful Furious layer (319; 13.97%) that exceeds Angry (228; 9.98%). Reading the Texts indicates that “ethics” is rarely abstract philosophy here; it is articulated through concrete disputes over accountability, safety, and institutional lag. Users argue about whether platforms and states should regulate AI more aggressively, whether training data practices violate consent or ownership, and how to treat harms ranging from privacy leakage to discriminatory or dangerous outputs. Some

posts adopt a policy talk register (“We need clear rules, audits, and enforcement”), while others personalize risk through everyday cautionary advice, as in: “Do not upload private documents to a model you do not control,” or “If a system can be biased, governance is not optional.” There is also a recurring tension between openness and restriction: a pragmatic pro-access stance framed around learning and productivity coexists with calls for limits framed around safety, surveillance, and rights. Analytically, this topic is where Persian-language users most explicitly perform public reasoning about sociotechnical order, distributing blame and responsibility across developers, platforms, regulators, and end users while negotiating what “responsible use” should mean in practice.

3.11. Sensitive Social Topics and Mental Health/Psychology

This is the smallest thematic cluster in the dataset ($n = 480$), yet it is among the most affectively polarized, with Happy (135; 28.13%) and Angry (135; 28.13%) tied as the modal categories, followed by Neutral (115; 23.96%), Furious (68; 14.17%), and a small Delighted share (27; 5.63%). Reading across the posts, AI is positioned as a moral and intimate interlocutor: users test models on culturally charged issues (especially around sexuality, religion, and social norms), circulate model answers as proof of bias or “ideological programming,” and debate whether a chatbot can be trusted as an epistemic authority in domains where legitimacy is contested. Alongside this, a distinct strand treats AI as an informal mental health support tool, with users describing chatbots as a low-barrier “listener” for stress, loneliness, or everyday coping, while others react sharply against what they perceive as an irresponsible or unsafe substitution for professional care. In English paraphrase, posts often take the form of public experiments (“I asked the model about X; here is what it said”) or warnings (“Do not treat a chatbot like a therapist”). Communication-wise, this topic shows how AI talk becomes a proxy for broader struggles over values and vulnerability: affect intensifies when AI is imagined not as a tool but as a social actor whose answers may legitimate, stigmatize, or emotionally steer users.

AI as a boundary object: The Persian X discourse

Table1. Distribution of topics by sentiments

Topic Title	Total Posts	Furious	Angry	Neutral	Happy	Delighted
AI Tools Ecosystem and Model Comparisons	6324	592	924	2190	2140	478
Labor Market, Economy, and Business Impacts	5954	891	477	1915	1889	782
Education, Universities, and Schools	5648	729	579	1968	1719	653
Future Narratives, Scenarios, and Predictions	3131	334	147	1098	1062	490
Ethics, Data Governance, Privacy, and Regulation	2284	319	228	753	740	244
Content Creation and Creativity	2135	248	215	739	671	262
Deepfakes, Misinformation, and Verification	1504	284	275	418	393	134
AI for Programming and Software Development	1434	154	161	506	440	173
Sensitive Social Topics and Mental Health/Psychology	480	68	135	115	135	27
Prompting Practices and Strategies	316	18	20	103	117	58




Figure 4. Sankey diagram of topic-sentiment relationships

Radar — Sentiment Intensity per Topic

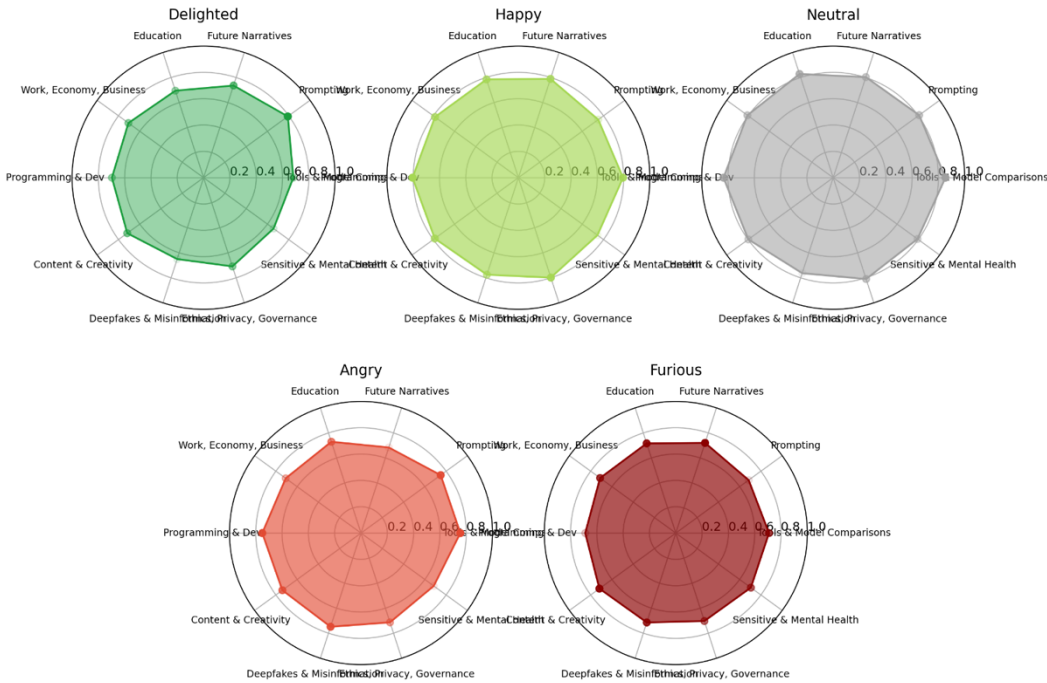


Figure 5. Radar-sentiment intensity per topic

3.12. Topic Co-occurrence Network

The topic co-tag network, built from multi-label assignments across the 24,215-post corpus, shows that Persian-language AI discourse on X is organized less as isolated “issue publics” and more as an interconnected field in which practical talk repeatedly folds into institutional and political-economic stakes. In raw co-tags, the strongest tie links Education with Work/Economy (n = 1,841), followed by Future Narratives with Work/Economy (n = 1,249) and Tools/Model Comparisons with Work/Economy (n = 1,139); Work/Economy also connects strongly to Ethics/Privacy/Governance (n = 853) and, at a slightly lower level, Tools with Education (n = 774). These counts already suggest that “work” functions as the discourse’s structural backbone: even when posts begin as tool chatter or classroom talk, they are frequently reframed through labor, productivity, credential value, and organizational advantage. To ensure this pattern is not merely an artifact of topic size, we also inspected relative overlap using the

Jaccard coefficient, which confirms the same ordering: Education ↔ Work/Economy shows the highest overlap (Jaccard ≈ 0.189), followed by Future ↔ Work/Economy (≈ 0.159) and Work/Economy ↔ Ethics/Privacy/Governance (≈ 0.116); Tools ↔ Work/Economy remains substantial on this scale as well (≈ 0.102). Interpreted through communication theory, this clustering indicates that users treat AI as a skills-and-institutions problem: debates about classroom use are rhetorically tethered to employability and inequality, and governance talk is pulled toward workplace consequences rather than remaining a purely abstract “ethics” register. A second, more content-specific bridge appears around authenticity: Content Creation co-occurs with Deepfakes/Misinformation at a comparatively high relative rate ($n = 305$; Jaccard ≈ 0.091), suggesting that creative adoption and verification anxiety travel together, as when synthetic media is discussed simultaneously as a tool for production and as a threat to evidentiary trust. Finally, some ties are near-absent, marking the network’s boundaries: Prompting rarely overlaps with Sensitive Social/Mental Health content ($n = 1$; Jaccard ≈ 0.0013) and only weakly with Deepfakes ($n = 7$; ≈ 0.0039), implying that “how-to” prompting craft tends to circulate as a specialized competence discourse rather than as a gateway into the most contentious moral or psychological domains.




Figure 6. Co-occurrence network

4. Conclusion

The discourse surrounding artificial intelligence is not merely a collection of opinions—it is a living archive of how a society thinks, feels, and argues about its future. This study has shown that Persian-language users on the X platform engage with AI not as passive observers of technological change, but as active participants in the social negotiation of meaning, legitimacy, and moral consequence. These publics do not merely adopt or reject AI; they narrate it into being, fold it into institutional rhythms, and contest its significance within the constraints and affordances of platformed expression.

In this sense, AI operates less as an object and more as a mirror, a medium through which people project fears, hopes, grievances, and competencies. The pragmatic focus on tools and skills, the anxious debates about education and employment, the fascination with synthetic creativity, and the contentious disputes over governance and verification—all reflect the broader work of cultural sense-making under conditions of epistemic volatility. AI becomes a boundary object not because of its technical ambiguity, but because it is emotionally and socially overdetermined. It sits at the intersection of what is known and what is imagined, what is experienced and what is speculated.

The affective structure of the discourse reveals this ambivalence. The preponderance of pragmatic approval and the persistence of indignation are not contradictory; they are co-constitutive. They mark the simultaneous domestication and disruption of cognitive authority: AI is useful, but it is also unsettling. It accelerates productivity while destabilizing the moral architecture of effort, credit, and expertise. Publics do not simply ask what AI can do—they ask who is authorized to speak about it, to benefit from it, to be displaced by it. They ask what counts as knowledge in a world where synthetic fluency begins to mimic human intention.

But perhaps most importantly, this study underscores that meaning is not inherent in the machine; it is forged in discourse. Through hashtags and quote tweets, jokes and jeremiads, users craft a vernacular epistemology—one that does not merely receive expert framings of AI but reconstitutes them in situated, emotionally charged, and politically aware terms. These acts of everyday reasoning, though often informal, are consequential. They are how

publics rehearse governance in the absence of regulation, how they scaffold legitimacy in the absence of trust.

In analyzing a non-Anglophone, under-theorized context, this study also intervenes in a dominant tendency within AI scholarship: the unspoken universalism of English-language data, elite policy framings, and normative ethics. What Persian-language discourse reveals is not only difference, but *difference as method*—a reminder that cultural specificity is not noise but signal. It is through the granular, the colloquial, and the affectively charged that the global contours of AI adoption and resistance take shape.

The questions animating this discourse—what should count as truth, labor, risk, and authority in the age of intelligent machines—are not merely technological. They are ontological. They ask what it means to be human when cognition can be simulated, when language can be generated, and when attention can be automated. These are not questions that platforms can answer, but they are questions that platforms force us to ask—repeatedly, publicly, and often urgently.

References

- Balazadeh, K., & Kajonius, P. (2025). Exploring Intimacy with Artificial Intelligence: Validation of Robot Intimacy Receptivity Scale (RIRS). *International Journal of Social Robotics*, 1–13.
- Bourdieu, P. (1991). *Language and symbolic power* (J. B. Thompson, Ed.; G. Raymond & M. Adamson, Trans.). Harvard University Press.
- Bozdağ, A. A. (2025). The AI-mediated intimacy economy: A paradigm shift in digital interactions. *AI & Society*, 40(4), 2285–2306.
- Chu, W., Skirpan, M., Mir, K., & Gray, M. L. (2025). Illusions of intimacy: Analyzing emotional discourse in large-scale AI-user interactions. In *Proceedings of the 2025 ACM Conference on Human Factors in Computing Systems (CHI)*. <https://arxiv.org/abs/2505.11649>
- Elliott, A. (2023). Algorithmic intimacy: Digital companionship and the new emotional order. *Information, Communication & Society*. <https://doi.org/10.1080/1369118X.2023.2282554>
- George, A. S., George, A. H., Baskar, T., & Pandey, D. (2023). The allure of artificial intimacy: Examining the appeal and ethics of using generative AI for simulated relationships. *Partners Universal International Innovation Journal*, 1(6), 132–147.
- Illouz, E. (2007). *Cold intimacies: The making of emotional capitalism*. Polity Press.
- Innis, H. A. (1951). *The bias of communication*. University of Toronto Press.
- Lupton, D. (2022). “Sharing Is Caring:” Australian self-trackers' concepts and practices of personal data sharing and privacy. *Frontiers in Digital Health*, 3, 649275.
- McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill.
- Papacharissi, Z. (2015). *Affective publics: Sentiment, technology, and politics*. Oxford University Press.
- Reineke, M. J. (2022). The touching test: AI and the future of human intimacy. *Contagion: Journal of Violence, Mimesis, and Culture*, 29, 123–146.
- Sabbar, S., & Habib Zadeh Khiyaban, S. (2023). Algorithms of Displacement: Emotional and Rhetorical Responses to AI-Driven Job Loss in Digital Public Discourse. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(4), 1324-1331. <https://doi.org/10.62225/2583049X.2023.3.4.5012>
- Salehi K, Habib Zadeh Khiyaban S, Sabbar S. (2026). Artificial Intelligence and Crime Detection: A Critical Review. *Cyberspace Studies*. 10(1): 181-197. <https://doi.org/10.22059/jcss.2025.402206.1179>
- Seaver, N. (2022). *Computing taste: Algorithms and the makers of music recommendation*. University of Chicago Press.
- Shahghasemi, E. (2025). AI; A Human Future. *Journal of Cyberspace Studies*, 9(1), 145-173. <https://doi.org/10.22059/jcss.2025.389027.1123>

- Shahghasemi, E., Gholami, F., & Alikhani, Z. (2025). Global patterns of social media use and political sentiment. *Discover Global Society*, 3, 36. <https://doi.org/10.1007/s44282-025-00171-y>
- Sjoraida, D. F. (2025). AI and Emotional Intimacy: Exploring Romantic Bonds with Artificial Companions. *Humanexus: Journal of Humanistic and Social Connection Studies*, 1(8), 333-342.
- Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, "translations" and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, 19(3), 387-420.
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each other*. Basic Books.
- van Dijk, T. A. (1993). Principles of critical discourse analysis. *Discourse & Society*, 4(2), 249-283. <https://doi.org/10.1177/0957926593004002006>