



## Original-Forschungsarbeit

# Iranischer digitaler Diskurs, affektive Ausrichtungen und die Geopolitik der KI

Mahsa Havsson<sup>1</sup>, Mandana Sajjadi<sup>2\*</sup>

<sup>1</sup> MSc, Computer und Systemwissenschaften; Spezialisierung auf Informationssicherheit, Universität Stockholm, Stockholm, Schweden

<sup>2</sup> Assistenzprofessorin, Abteilung für Global Studies, Fakultät für Weltstudien, Universität Teheran, Teheran, Iran

**Empfangen:** 28. Februar 2025 **Akzeptiert:** 4. Juni 2025

### Zusammenfassung:

Diese Studie untersucht, wie persischsprachige Nutzer auf X das von China entwickelte große Sprachmodell DeepSeek interpretieren und emotional darauf reagieren. Auf der Grundlage eines kuratierten Korpus von 1.112 Beiträgen iranischer Nutzer verwendet die Forschung einen Mixed-Methods-Ansatz, der Sentiment-Analyse, Themenmodellierung und Co-Occurrence-Netzwerkanalyse kombiniert. Die Ergebnisse zeigen eine geschichtete diskursive Landschaft, in der DeepSeek nicht nur als technisches Produkt fungiert, sondern als symbolischer Ort zur Aushandlung geopolitischer Ausrichtungen, epistemischen Vertrauens und technologischer Ambitionen dient. Sechs zentrale affektive Ausrichtungen – Neutralität, Skepsis, Hoffnung, Stolz, Angst und Ablehnung – strukturieren das Engagement der Nutzer mit dem Modell und spiegeln ambivalente, jedoch politisch informierte Reaktionen wider. Die thematische Analyse identifizierte acht wiederkehrende Themen, darunter Leistungsbewertungen, chinesische Souveränität, KI-Ethik und kulturelle Identität, die oft in komplexen rhetorischen Konfigurationen gemeinsam auftreten. Diese Ergebnisse deuten darauf hin, dass iranische Nutzer DeepSeek als Stellvertreter einsetzen, um über inländische technologische Einschränkungen, Plattformpolitik und die sich verschiebenden Konturen der globalen KI-Hegemonie nachzudenken.

**Schlüsselwörter:** iranische soziale Medien, große Sprachmodelle, digitale Diskursanalyse, KI-Geopolitik

---

\* Korrespondierender Autor [✉ sajjadi.m@ut.ac.ir](mailto:sajjadi.m@ut.ac.ir) [🌐 https://orcid.org/0000-0003-3297-3950](https://orcid.org/0000-0003-3297-3950)

#### Wie dieser Artikel zu zitieren ist:

Havsson, M., & Sajjadi, M. (2025). Iranian digital discourse, affective alignments, and the geopolitics of AI. *Spektrum Iran*, 38(2), 187-212.

[📄 https://doi.org/10.22034/spektrum.2026.551119.1040](https://doi.org/10.22034/spektrum.2026.551119.1040)

© 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](http://www.spektrumiran.com)

## گفتار دیجیتالی ایرانی، گرایش‌های عاطفی و ژئوپلیتیک هوش مصنوعی

مهسا هاوسون<sup>۱</sup>، ماندانا سجادی<sup>۲\*</sup>

<sup>۱</sup> کارشناس ارشد، علوم کامپیوتر و سیستم‌ها؛ دانشگاه استکهلم، پردیس کیستا، استکهلم، سوئد

<sup>۲</sup> استادیار، گروه مطالعات جهانی، دانشکده مطالعات جهان، دانشگاه تهران، تهران، ایران

دریافت: ۱۴۰۳/۱۲/۱۰؛ پذیرش: ۱۴۰۴/۳/۱۴

### چکیده:

این مطالعه بررسی می‌کند که کاربران فارسی‌زبان در شبکه اجتماعی X چگونه مدل زبانی بزرگ DeepSeek، که توسط چین توسعه یافته، را تفسیر کرده و نسبت به آن واکنش عاطفی نشان می‌دهند. با بهره‌گیری از یک مجموعه داده منتخب شامل ۱,۱۱۲ پست جمع‌آوری‌شده از کاربران ایرانی، پژوهش از رویکردی ترکیبی شامل تحلیل احساسات، مدل‌سازی موضوعی و تحلیل شبکه هم‌رخدادها استفاده می‌کند. یافته‌ها نشان‌دهنده یک چشم‌انداز گفتمانی لایه‌ای هستند که در آن DeepSeek صرفاً به عنوان یک محصول فناوری عمل نمی‌کند، بلکه به عنوان یک فضای نمادین برای مذاکره در مسائل هم‌رستایی ژئوپلیتیکی، اعتماد معرفتی و آرمان‌های فناوری ظاهر می‌شود. شش گرایش عاطفی اصلی—خنثی‌گرایی، تردید، امید، غرور، اضطراب و بی‌تفاوتی—تعامل کاربران با مدل را شکل می‌دهند و پاسخ‌هایی دوپهلوی اما آگاهانه سیاسی را منعکس می‌کنند. تحلیل موضوعی هشت موضوع تکرارشونده شامل مقایسه عملکرد، حاکمیت چین، اخلاق هوش مصنوعی و هویت فرهنگی را شناسایی کرد که اغلب در ساختارهای بلاغی پیچیده با هم رخ می‌دانند. این نتایج نشان می‌دهند که کاربران ایرانی از DeepSeek به عنوان یک نماینده برای تأمل درباره محدودیت‌های فناوری داخلی، سیاست‌های پلتفرم و تغییر شکل‌های هژمونی جهانی هوش مصنوعی استفاده می‌کنند.

**واژگان کلیدی:** شبکه‌های اجتماعی ایرانی، مدل‌های زبانی بزرگ، تحلیل گفتمان دیجیتال، ژئوپلیتیک هوش مصنوعی



Original Research Paper

# Iranian digital discourse, affective alignments, and the geopolitics of AI

Mahsa Havsson<sup>1</sup>, Mandana Sajjadi<sup>2\*</sup>

<sup>1</sup> MSc, Computer and Systems Sciences; Stockholm University, Stockholm, Sweden

<sup>2</sup> Assistant Professor, Department of Global Studies, Faculty of World Studies, University of Tehran, Tehran, Iran

Received: Feb. 28, 2025 Accepted: Jun. 4, 2025

## Abstract

This study investigates how Persian-speaking users on X interpret and emotionally respond to DeepSeek, a Chinese-developed large language model. Drawing on a curated corpus of 1,112 posts collected from Iranian users, the research employs a mixed-method approach incorporating sentiment analysis, topic modeling, and co-occurrence network analysis. The findings reveal a layered discursive landscape in which DeepSeek serves not merely as a technological product but as a symbolic site for negotiating issues of geopolitical alignment, epistemic trust, and technological aspiration. Six major affective orientations – neutrality, skepticism, hope, pride, anxiety, and dismissiveness – structure user engagement with the model, reflecting ambivalent yet politically informed responses. Thematic analysis identified eight recurring topics, including performance comparisons, Chinese sovereignty, AI ethics, and cultural identity, which often co-occurred in complex rhetorical configurations. These results suggest that Iranian users deploy DeepSeek as a proxy to reflect on domestic technological constraints, platform politics, and the shifting contours of global AI hegemony.

**Keywords:** DeepSeek, Iranian social media, large language models, digital discourse analysis, AI geopolitics

---

\* Corresponding Author

✉ [sajjadi.m@ut.ac.ir](mailto:sajjadi.m@ut.ac.ir)

🌐 <https://orcid.org/0000-0003-3297-3950>

### How to Cite this Article:

Havsson, M., & Sajjadi, M. (2025). Iranian digital discourse, affective alignments, and the geopolitics of AI. *Spektrum Iran*, 38(2), 187-212.

🔗 <https://doi.org/10.22034/spektrum.2026.551119.1040>

© 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](http://www.spektrumiran.com)

## **1. Introduction**

Artificial Intelligence has introduced a radically new ontological and epistemological condition for human societies. It transforms not only how tasks are automated but also how knowledge, power, and subjectivity are constructed and distributed (Zamani et al., 2025). The accelerated development of large-scale machine learning systems—particularly large language models (LLMs)—has disrupted nearly every sphere of human activity, from knowledge production to interpersonal communication. In fields such as healthcare, AI now supports diagnostic systems (Toosi et al., 2025), predictive modeling for disease outbreaks (Tomraee et al., 2022), and individualized treatment planning (Esteva et al., 2019). In architecture and urban planning, generative design systems powered by AI optimize material use and spatial efficiency (Sakhaei et al., 2024). Education is being transformed through AI tutors, plagiarism detection systems (Rahmatian & Sharajsharifi, 2022), and adaptive learning platforms that respond to student behavior in real time (Luckin, 2017). Child-rearing practices, too, have been reshaped by the integration of smart toys, AI-driven educational apps, and surveillance technologies embedded in household devices (Holloway & Green, 2016). In the legal domain, AI tools assist in predictive policing, case outcome forecasting, and even sentencing suggestions (Salehi et al., 2026)—raising critical concerns about fairness, accountability, and legal subjectivity (Surden, 2019). Similarly, in corporate management and public administration, decision-making is increasingly mediated by algorithmic systems that model risk, allocate resources, and track performance (Kellogg et al., 2020). Far from remaining neutral instruments, these technologies reorganize the relationships between human agency, institutional authority, and computational governance (Shahghasemi, 2025).

Yet, beyond transforming specific domains, AI systems are now deeply entangled in global political rivalries and ideological formations. AI has emerged not merely as a technical frontier but as a geopolitical battleground, where questions of technological sovereignty, epistemic control, and digital colonization are increasingly foregrounded. The race for AI dominance between China and the United States has catalyzed a new form of Cold War, in which control over data, algorithms, and compute power is seen as tantamount to economic and military supremacy (Allen, 2019; Ding, 2018). Nations are investing in sovereign AI infrastructures, censoring foreign

platforms, and designing regulatory frameworks to assert control over digital epistemologies. As Bender et al. (2021) argue, the development of foundational models—those trained on enormous corpora and generalized across tasks—is not only a technical innovation but a political act, embedding particular linguistic, cultural, and normative assumptions within algorithmic infrastructures. Meanwhile, global AI ethics debates reveal deep asymmetries in whose values are encoded and whose vulnerabilities are ignored (Crawford, 2021). In this global landscape, AI is not just a tool for computation but a terrain for ideological contestation and symbolic projection.

Within such a context, public perceptions become politically consequential (Shahghasemi, 2017). The social legitimacy of AI systems, their patterns of adoption, and the degree of public trust they command all hinge on how these systems are represented, narrated, and contested in everyday discourse (Sharifi Poor Bgheshmi & Sharajsharifi, 2025). Public perception shapes the reception of AI technologies in real-world contexts, mediating user behavior, regulatory responses, and institutional decisions (Fast & Horvitz, 2017). Importantly, public perceptions of big-tech corporations—such as OpenAI, Google, Baidu, or DeepSeek—are not formed in isolation but through historically embedded, culturally specific, and politically inflected communicative practices. These corporations are not perceived as neutral service providers but as actors situated within global power hierarchies. Consequently, discourse around AI platforms often serves as a proxy for broader debates about surveillance, censorship, cultural hegemony, and epistemic legitimacy (Zuboff, 2019). Understanding these discursive formations is essential to grasp how AI is domesticated, contested, or resisted across different sociopolitical contexts.

Social media platforms have become the primary sites for the articulation and circulation of such perceptions. Unlike traditional media, social media enables users not only to consume but also to produce and contest narratives about technological systems (Bahmani & Javaheri Tehrani, 2025). Public discourse on platforms like X, Telegram, or Instagram is characterized by its immediacy, dialogical structure, and algorithmic amplification, making it a rich site for observing how technological imaginaries are constructed in real time (Milan, 2015). Hashtag ecologies, meme cultures, and comment threads function as discursive infrastructures through which users frame, question,

or endorse emerging technologies. These platforms facilitate affective and argumentative expression, allowing for a wide spectrum of emotional, ideological, and epistemic positions to emerge (Shahghasemi, 2016). Moreover, digital publics in politically constrained environments often use social media not merely for communication but for strategic articulation—crafting narratives that reflect complex negotiations between local experience and global phenomena.

In the Iranian context, social media discourse occupies an especially charged and revealing role. Despite heavy state censorship, filtering, and surveillance, Persian-language digital publics remain active and discursively potent, particularly on platforms like X that are technically blocked but still widely accessed through circumvention tools. Iranian users are among the most politically literate and rhetorically sophisticated participants in online discourse within the region, engaging with global events, technological trends, and political crises through an idiom shaped by satire, subtext, and coded dissent (Rahimi, 2011). The digital sphere in Iran functions as a semi-autonomous zone of commentary and critique, often serving as a space where political frustration, technological aspiration, and cultural identity converge. Discourse on foreign technologies—especially those developed in geopolitical rivals such as the United States or China—provides a particularly rich site for observing how Iranians articulate their place within global technological hierarchies and reimagine their national trajectories.

It is within this politically layered and rhetorically rich context that the present study is situated. The focus of this research is on the Persian-language digital discourse surrounding DeepSeek, a Chinese-developed large language model that emerged as a rival to Western counterparts such as OpenAI's ChatGPT. Drawing on a curated dataset of 1,112 Persian-language posts from X, the study undertakes a multi-method analysis combining sentiment classification, topic modeling, and co-occurrence network analysis. The aim is not merely to map opinions about DeepSeek's technical capabilities but to understand how Iranian users interpret, appropriate, and contest the model as a symbolic object embedded in geopolitical, cultural, and epistemic narratives. Specifically, the study seeks to answer: What affective registers structure Iranian users' engagement with DeepSeek? What thematic patterns recur in their discourse, and how are these topics

interconnected? how do these patterns reflect broader concerns around technological sovereignty, surveillance, and platform politics?

## **2. Methodology**

This study employed a multi-stage, mixed-method approach to investigate how Iranian users discursively engage with DeepSeek on Persian-language social media. The platform X (formerly Twitter) was selected as the primary data source due to its relative openness, text-centric format, and its discursively active Iranian user base, which—despite state-level filtering—remains a key site for technologically literate, politically engaged publics. A targeted query strategy was developed to capture a broad spectrum of relevant discourse. Search terms included transliterations of “DeepSeek” (e.g., “دیپسیک”) alongside categorical and comparative expressions such as “مدل زبان چینی” and “رقیب چتجی‌پی‌تی,” thereby ensuring thematic inclusivity beyond direct brand mentions. Data collection occurred during two months following DeepSeek’s public release, yielding an initial corpus of 4,872 posts.

To enhance the semantic coherence and analytical reliability of the corpus, a rigorous preprocessing pipeline was implemented. Duplicate posts—primarily retweets and copy-pasted messages—were identified and removed using exact and fuzzy string-matching algorithms, reducing the dataset to 2,844 unique entries. A second filtration phase excluded non-commentary content such as image-only posts, spam, and bot-like behavior, resulting in a refined corpus of 1,432 posts. Orthographic normalization was then applied to address the variability and inconsistency of Persian digital writing. This involved standardizing script forms, removing zero-width non-joiners, and preserving non-standard punctuation and spelling when they conveyed meaningful rhetorical or emotional nuance. Posts were further cleaned of user mentions, URLs, and hashtags unless they contributed directly to the semantic content. After this final phase, a high-integrity dataset of 1,112 posts was retained for analysis.

This curated corpus served as the basis for three computational analyses. First, sentiment and emotion profiling was conducted using a hybrid model combining a Persian-language lexical sentiment classifier with manual coding on a stratified sample to validate and refine affective categories. This

yielded six affective orientations—neutral, skeptical, hopeful, proud, anxious, and dismissive—tailored to the contextual nuances of AI discourse in Iran. Second, Latent Dirichlet Allocation (LDA) was used to extract eight latent thematic clusters from the corpus, determined via coherence optimization. These topics were interpreted as discursive repertoires encompassing performance comparison, technological sovereignty, ethics, censorship, and cultural identity. Third, a co-occurrence network analysis was employed to map how topics clustered and intersected within individual posts. Posts were coded for up to three high-probability topics. Edges between topics were calculated based on their frequency of co-occurrence, producing a weighted semantic network.

### **3. Findings**

The study began with the construction of a large-scale dataset to examine Persian-language social media discourse surrounding DeepSeek. This process was guided by the objective of identifying meaningful, user-generated commentary rather than indiscriminately collecting digital chatter. Data were sourced primarily from X, which remains one of the most discursively rich platforms for Iranian users, particularly those engaging in discussions about global technology and AI. Despite the platform's inaccessibility within Iran without circumvention tools, it continues to host a demographically skewed yet analytically valuable user population—digitally literate, politically aware, and linguistically performative.

To capture a broad sample of relevant discourse, a targeted query strategy was employed. A list of search terms was compiled to maximize lexical variation and thematic inclusivity, covering both direct mentions of DeepSeek and broader technological references. The search included transliterations such as “دیپ‌سیک” and categorical phrases such as “مدل زبان” or “چت‌جی‌پی‌تی” or “رقیب چت‌جی‌پی‌تی”. These terms were selected to index not only mentions of the model itself but also how users positioned it within comparative, geopolitical, or national frameworks. The query was run over a two-month window following DeepSeek’s major public release, a period characterized by a notable uptick in digital commentary on artificial intelligence across Persian-speaking online publics. This initial query yielded a raw dataset of 4,872 posts.

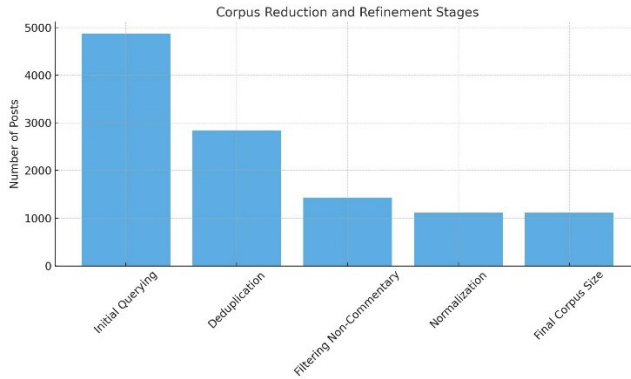


Figure 1. Corpus reduction and refinement stages

The raw corpus contained a significant amount of duplicated content due to user behaviors such as retweeting and copy-pasting popular opinions. To ensure analytical validity, a deduplication protocol was implemented using text-matching techniques that identified exact and near-exact repetitions. This step reduced the corpus to 2,844 unique entries. Posts that had been retweeted verbatim or recycled with minimal alteration were excluded, while quote tweets and derivative commentary were retained. This filtering helped preserve discursive heterogeneity while eliminating distorting repetition.

Following deduplication, a second round of filtration targeted posts lacking substantive commentary. Items that consisted solely of images, emojis, links, or bot-like promotional forwards were excluded on the grounds that they did not contain meaningful textual content. The objective here was not to impose an evaluative hierarchy on content types but to ensure that the retained corpus allowed for linguistic and affective analysis. After this filtering step, the dataset was reduced to 1,432 entries.

The remaining corpus underwent an orthographic normalization phase. Persian social media text is marked by a high degree of non-standardized spelling, informal grammar, and platform-specific encoding. These features pose challenges to computational processing and topic modeling. Normalization involved standardizing the Persian script by aligning Arabic and Persian variants of characters (e.g., different forms of “ye” and “kaf”), removing diacritics and zero-width non-joiners, and applying consistent punctuation. Non-linguistic elements such as hashtags, user mentions, and

URLs were removed unless semantically informative. Emojis, though affectively salient, were excluded to prioritize purely textual affect markers. Crucially, this process avoided excessive sanitization and preserved informal spelling, intensification, and sarcasm indicators, which are common in affect-laden posts.

After normalization, a small number of structurally ambiguous or residual low-information entries were removed, resulting in a final dataset of 1,112 unique posts. This final corpus was deemed sufficient in scale for computational modeling while maintaining high semantic integrity. It includes a broad range of discursive formats, from short aphoristic statements to longer reflections, as well as a wide spectrum of affective tones. The distribution of posts across preprocessing steps is presented in Table 1; the visual equivalent can be seen in Figure 1.

*Table 1. The distribution of posts across preprocessing steps*

<b>Step</b>	<b>Post Count</b>
Initial Querying	4,872
Deduplication	2,844
Filtering Non-Commentary	1,432
Normalization	1,112
Final Corpus Size	1,112

This methodical reduction reflects a commitment to both analytical rigor and discursive representativeness. The goal was not simply to retain posts but to curate a corpus in which the patterns, tensions, and emotional textures of Iranian talk about DeepSeek could be meaningfully identified and interpreted. Posts selected for inclusion reflect spontaneous user engagement and address themes ranging from performance evaluation to geopolitical suspicion, from digital nationalism to ethical apprehension. The textual data, while modest in quantity relative to the initial scrape, provides sufficient granularity for both computational modeling and interpretive reading.

An important contextual factor shaping the corpus is the nature of the user base itself. Iranian users on X typically represent a technologically literate, often urban demographic with experience accessing foreign platforms through VPNs. This demographic skew introduces certain interpretive constraints: the discourse captured is not representative of the general population but is instead indicative of a politicized sub-public engaged with international technological developments. Nevertheless, this segment is precisely the one most actively forming and circulating narratives around global AI trends. Their posts, therefore, function not just as opinion but as discourse-in-action: negotiating Iran's place in global technological hierarchies, articulating comparisons between Western and Eastern models, and projecting hopes or anxieties onto digital artifacts like DeepSeek.

In focusing exclusively on X for this round of analysis, the study deliberately sidestepped the broader Persian digital ecosystem, which includes highly active AI-related discourse on Telegram, Instagram, and YouTube. These platforms, particularly Telegram, are known for hosting informal discussions, viral content, and channel-driven opinion shaping. While incorporating such content would broaden the representational scope, it also introduces complexities related to multimodality, moderation variability, and text-image interdependencies. Future studies may extend the dataset to include multimodal sources, but for the present purpose, the focus on text-rich, commentary-heavy posts from X offers a manageable and analytically tractable starting point.

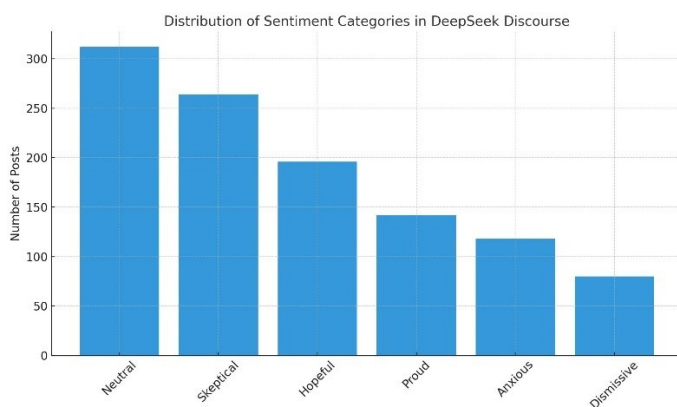
The resulting corpus allows for downstream sentiment modeling, topic detection, and co-occurrence analysis, all grounded in carefully cleaned and interpretable text. The preprocessing pipeline was not merely a technical step, but a substantive intervention in defining what kinds of speech count as relevant, expressive, and analytically productive in the study of AI discourse. By eliminating repetition and noise while preserving the rhetorical and emotional granularity of user contributions, the study establishes a firm empirical foundation for analyzing how DeepSeek is being made legible, contestable, and meaningful within the Iranian digital public sphere.

### **3.1. Sentiment and Emotion Analysis**

Understanding how Iranian users emotionally engage with DeepSeek requires more than simply identifying positive or negative opinions. In a

context marked by political ambivalence, technological aspiration, and geopolitical rivalry, affective responses to foreign AI systems like DeepSeek are often layered, ambivalent, and rhetorically complex. This section analyzes the emotional patterns embedded in the 1,112 posts collected in the final corpus, drawing on a hybrid model of sentiment profiling that captures not only polarity but also affective function. The goal is to move beyond conventional sentiment binaries and instead chart how emotions like skepticism, anxiety, pride, or curiosity circulate within the digital discourse surrounding DeepSeek.

The sentiment analysis proceeded in two stages. First, an unsupervised lexical sentiment classifier was applied using a custom Persian-language dictionary fine-tuned for AI discourse. This allowed for the categorization of posts into broad emotional registers. Second, manual coding was conducted on a stratified sample to refine category boundaries and validate ambiguous cases. Through this combined method, six primary sentiment categories were identified as thematically and affectively meaningful within the context of the study: Neutral, Skeptical, Hopeful, Proud, Anxious, and Dismissive. These categories do not represent universal emotional states but rather context-specific affective modes through which Iranian users relate to DeepSeek as a technological, political, and symbolic object. Figure 2 presents the distribution of sentiment categories in DeepSeek discourse.



*Figure 2. Distribution of sentiment categories in DeepSeek discourse*

Neutral posts were the most frequent, accounting for 312 out of 1,112 entries. These messages typically employed a descriptive or observational tone, relaying performance comparisons, announcement summaries, or technical details without overt affective charge. However, neutrality in this context often functioned as a discursive strategy rather than emotional absence. Many neutral posts featured a restrained tone—commenting on DeepSeek’s capabilities, speed, or data sources in ways that indexed curiosity or latent skepticism. The high frequency of neutrality suggests that many users are in a state of evaluative suspension, awaiting further information or collective consensus before committing to a strong position. This pattern reflects broader features of Persian digital discourse, where indirectness and diagnostic commentary often mask deeper emotional investments.

Skeptical posts comprised the second-largest category, with 264 instances. These messages conveyed doubt about DeepSeek’s reliability, motivations, or claims of performance superiority. Common themes included suspicion of Chinese data practices, comparisons to ChatGPT framed to discredit DeepSeek, and concerns about the authenticity of reported benchmarks. Skepticism here is not merely cognitive but affective—it signals a defensive posture rooted in historical experiences of foreign technology, media manipulation, and propaganda exposure. These posts often deployed irony or sarcasm to question the legitimacy of DeepSeek’s presentation as a “rival” to Western models. In this affective frame, users are not necessarily hostile to the technology but remain resistant to its framing by state actors or media platforms they perceive as politicized.

Hopeful sentiments appeared in 196 posts. These messages expressed cautious optimism about the rise of Eastern alternatives to Western AI, the possibility of improved access to language models in the region, or the technological momentum symbolized by DeepSeek. Posts in this category often emphasized the future-oriented potential of the model, portraying it as part of a global diversification of AI power. For some users, hope was linked to a desire for escape—from sanctions, from dependence on U.S.-based platforms, or from linguistic exclusion. In this sense, hopeful sentiment was not merely about DeepSeek itself but about the broader imaginaries it invokes: of a multipolar AI ecosystem, of local agency, or of non-Western technological sovereignty. These posts frequently contained conditional language, hedging their optimism with caveats about performance reliability or ethical concerns, which distinguishes them from unqualified enthusiasm.

Proud sentiments were appeared in 142 posts. This category reflects expressions of admiration or nationalist identification with DeepSeek's perceived achievements. While the model is Chinese, many Iranian users expressed pride in what they saw as a challenge to Western digital dominance. Some posts framed DeepSeek's emergence as part of a larger narrative of Eastern innovation, aligning it with broader geopolitical shifts toward Asia. In these discourses, pride becomes a displaced expression of national aspiration: even if the model is not Iranian, its success is read as symbolic of a shared resistance to technological dependency. These affective framings serve as a mechanism of identification through analogy – users locate their own hopes for national development in the success of a culturally proximate, politically sympathetic alternative to U.S.-led AI hegemony.

Anxious sentiment was visible in 118 posts. This category captured unease about the implications of DeepSeek and generative AI more broadly. Users expressed concern about surveillance, job displacement, epistemic manipulation, and the future of language, education, and social trust. Anxieties often took the form of rhetorical questions or dystopian allusions, invoking fears of an AI-driven future where control is centralized and human agency undermined. Importantly, anxiety was not limited to DeepSeek as a specific model – it was frequently bundled with broader fears about algorithmic power and digital governance. Some posts framed DeepSeek as merely the latest instance of a wider trend toward opaque and potentially coercive technologies. These expressions reflect the lingering presence of techno-political trauma in Iranian digital life, where innovation is frequently interpreted through lenses of control, risk, and asymmetry.

Dismissive posts were the least frequent, totaling 80 entries. These messages rejected DeepSeek either as overhyped, irrelevant, or fundamentally inferior to existing models. Dismissiveness in this context served as a form of boundary policing – articulating who counts as a credible actor in the global AI landscape and who does not. Some dismissive posts adopted a performative tone, mocking DeepSeek's name, interface, or media presentation. Others expressed frustration with what they perceived as propaganda surrounding the model's capabilities. This sentiment category, though numerically small, plays a significant discursive role: it marks the limits of empathy and recognition, delineating the affective boundary between "serious" and "symbolic" technology. The distribution of these sentiment categories is presented in Table 2.

*Table 2. Distribution of sentiment categories in DeepSeek discourse*

Sentiment	Post Count
Neutral	312
Skeptical	264
Hopeful	196
Proud	142
Anxious	118
Dismissive	80

This sentiment landscape reveals several important features of Iranian digital discourse on DeepSeek. First, affective polarization is not dominant. While there are pockets of enthusiasm and rejection, most users adopt mid-spectrum emotional positions such as skepticism, cautious hope, or ambivalent observation. This suggests a communicative culture that favors provisional alignment, affective layering, and interpretive flexibility rather than sharp ideological camps. Second, emotion is not merely reactive but constitutive: users deploy emotional registers to organize their perception of DeepSeek, not simply to express personal feelings. For instance, pride is not just felt—it is performed to establish solidarity with broader geopolitical imaginaries. Similarly, skepticism functions rhetorically to critique not just the model, but the structures that promote it.

Third, the coexistence of hopeful and anxious posts points to the ambivalent promise of AI in the Iranian context. On the one hand, DeepSeek is welcomed as a possible vector of technological empowerment and geopolitical rebalancing. On the other, it is feared as a vehicle for surveillance, control, or epistemic substitution. This tension is particularly acute in a country where both foreign technologies and domestic infrastructures are often viewed through the lens of politicized risk. In this environment, emotion serves as a way to manage uncertainty—allowing users to claim a position, to express stakes, and to signal membership in particular interpretive communities.

Finally, sentiment categories often co-occur or transition across time and post sequences. Some users who initially expressed skepticism later adopted hopeful tones as technical demonstrations improved. Others moved from

pride to anxiety as discussions of data usage and user control gained visibility. These shifts underscore the importance of treating sentiment not as static but as temporally dynamic and discursively situated. Future longitudinal analyses may track these transitions more formally, but even in this snapshot, the evidence suggests that public opinion on DeepSeek is not hardened but in flux—waiting to be shaped by further interaction, technological performance, and discursive framing.

### **3.2. Topic Modeling and Thematic Structuring**

The thematic structure of Persian-language discourse on DeepSeek reflects a layered and multifaceted engagement with artificial intelligence, one that exceeds narrow technical evaluation and enters the domains of geopolitics, cultural identity, epistemic trust, and technological sovereignty. To analyze this structure, topic modeling was applied to the cleaned corpus of 1,112 posts using Latent Dirichlet Allocation (LDA), a probabilistic model designed to uncover latent thematic clusters within a corpus based on word co-occurrence patterns. Preprocessing ensured that orthographic consistency and linguistic variability in Persian digital writing were accounted for, and the number of topics was determined using coherence scores across models ranging from six to fifteen topics. A model with eight topics yielded the best balance between semantic resolution and interpretive clarity, as larger models fragmented coherent discursive strands while smaller ones collapsed distinct affective and argumentative motifs into overly generic clusters.

The resulting eight-topic model captures the major thematic constellations around which Iranian users articulate their perceptions of DeepSeek. These topics are not simply clusters of lexical similarity but represent discursive repertoires—bundles of framing devices, analogical reasoning, moral positions, and emotional orientations. Each topic consists of a heterogeneous array of posts, but the dominant thematic focus is stable enough to allow for clear interpretive labels. The frequency distribution of these topics (Table 3), along with a horizontal bar chart (Figure 3).

Table 3. Distribution of topics in DeepSeek discourse

Topic	Post Count
Performance Comparisons with ChatGPT	198
Chinese Tech and Sovereignty	176
AI Ethics and Surveillance	162
Domestic Tech Inferiority Complex	140
Access, Censorship, and Control	134
Platform Trust and Propaganda	122
Educational and Economic Impact	102
Language Models and Cultural Identity	78




Figure 3 Distribution of topics in DeepSeek discourse

The largest topic centers on performance comparisons with ChatGPT. These posts evaluate DeepSeek in relation to OpenAI’s flagship model, drawing on user experiences, benchmark reports, and anecdotal demonstrations. Comparisons are not limited to technical functionality but extend to user interface, response coherence, and linguistic versatility. In these posts, DeepSeek is rarely treated in isolation; rather, it becomes legible through the lens of global competition in AI, with ChatGPT serving as a default standard against which new entrants are measured. Posts in this cluster express mixed sentiment, with some users praising DeepSeek’s speed or coding features while others question its reliability, contextual understanding, or Persian-language proficiency. This topic reflects the normalization of AI benchmarking as a form of public discourse and highlights how technical comparison becomes a socially distributed method of epistemic judgment.

The second-largest topic, labeled Chinese Tech and Sovereignty, captures posts that frame DeepSeek as a symbolic expression of China's technological ascendance and geopolitical strategy. Many users interpret the model not as a neutral tool but as a national artifact – an expression of state ambition and civilizational rivalry. In these framings, DeepSeek becomes a proxy through which Iranian users negotiate their own position in relation to both China and the West. Some posts express admiration for China's ability to produce competitive models despite restrictions and censorship, while others voice concern about dependency on Chinese digital infrastructure. This topic is saturated with affective ambiguity: pride in Eastern innovation coexists with wariness about techno-authoritarianism. Posts in this category also often include analogies to Iran's own digital ambitions, questioning whether a similar model could be developed domestically or whether political and infrastructural constraints preclude such efforts.

The authority and legitimacy of AI in more sensitive areas like law have already been studied (Salehi et al., 2025). The third topic, AI Ethics and Surveillance, comprises posts concerned with the potential misuse of DeepSeek and similar models for surveillance, misinformation, and behavioral control. These users focus on the socio-political affordances of AI rather than its capabilities per se. Issues such as data privacy, consent, algorithmic bias, and the ideological framing of responses dominate this cluster. Some posts speculate about the use of large language models for population monitoring in authoritarian contexts, invoking Chinese practices as a cautionary template. Others reflect on the broader implications of relying on opaque systems to mediate knowledge, interaction, and decision-making. This topic is marked by heightened emotional intensity, with expressions of anxiety and critique emerging alongside calls for transparency and regulation. The discourse suggests a relatively sophisticated understanding of the ethical stakes of generative AI, even among non-specialist users.

A fourth topic emerges around what can be described as a Domestic Tech Inferiority Complex. Here, DeepSeek becomes a reference point for Iranian users to articulate dissatisfaction with the domestic technological landscape. Posts in this cluster highlight the lack of investment, transparency, and coordination in Iranian AI development, often contrasting the perceived sophistication of DeepSeek with what users describe as the fragmented or performative nature of

local initiatives. Some posts adopt a self-critical tone, lamenting brain drain and institutional inefficiency. Others use sarcastic or ironic registers to critique the state's rhetorical commitment to technological advancement without corresponding material outcomes. This topic does not merely express envy or resignation but functions as a discursive space where national aspirations are measured against external achievements, often in implicitly moral terms.

Another important thematic formation revolves around access, censorship, and control. Posts in this topic interrogate the infrastructural and political conditions that govern who gets to use AI, under what conditions, and with what limitations. Many users express frustration at the difficulty of accessing DeepSeek due to regional restrictions or language barriers. Others speculate about potential censorship embedded in the model, given its Chinese origin. Crucially, these concerns are not limited to DeepSeek itself but are often articulated through comparisons to Iran's own history of internet filtering, platform blocking, and informational gatekeeping. The model thus becomes a lens through which larger questions of digital rights, informational sovereignty, and epistemic autonomy are raised. These posts often frame AI not just as a technical service but as a contested political infrastructure, one that mediates and potentially restricts the flow of knowledge.

The sixth topic, Platform Trust and Propaganda, focuses on how users interpret DeepSeek in relation to media narratives, official endorsements, and perceived propaganda. Posts in this category exhibit high levels of skepticism toward claims made about the model's performance or independence. Some users challenge the neutrality of media reporting, pointing to what they see as state-coordinated efforts to frame DeepSeek as superior to Western alternatives. Others criticize what they perceive as uncritical adoption of Chinese narratives in Iranian media. This topic reveals a critical awareness of the role that platforms and governments play in shaping the discursive environment around AI. It also demonstrates how trust in technology is filtered through trust in institutions—a dynamic especially pronounced in societies marked by information asymmetries and media politicization.

The educational implications of AI technologies have been documented in previous research (Khodabin et al., 2023). The seventh topic, educational and economic impact, includes posts that discuss the implications of DeepSeek for academic integrity, skill development, and labor markets. Some users

highlight the potential of AI to democratize education and reduce knowledge gaps, especially in under-resourced contexts. Others raise concerns about automation, job displacement, or the erosion of human expertise. Posts in this topic often adopt a practical tone, reflecting the everyday stakes of AI integration in sectors such as education, freelance work, or digital services. While less emotionally charged than other topics, this cluster provides insight into how users domesticate AI—situating it not only in global competition or abstract ethics but also in the material conditions of Iranian life.

The final topic, language models and cultural identity, is the smallest but conceptually rich. Posts in this cluster explore the role of language models in shaping cultural representation, translation accuracy, and linguistic politics. Some users express concern that DeepSeek, like many LLMs, marginalizes non-standard dialects or misrepresents Persian cultural references. Others reflect on the symbolic significance of AI fluency in Persian, whether as a sign of inclusion or a tokenistic gesture. In a context where linguistic identity is politically charged and often suppressed in both domestic and foreign media systems, these posts articulate a subtle form of cultural anxiety. They raise the question of what it means for a model to “know” a language, and whether such knowledge entails recognition or mere data capture.

Taken together, these eight topics form a thematic architecture that reveals how DeepSeek is discursively constructed in Iran not merely as a product, but as a prism for articulating broader social, political, and epistemic concerns. Topics do not operate in isolation; many posts traverse multiple clusters, combining performance comparison with ethical critique or geopolitical analysis with emotional skepticism. This multidimensionality reflects the richness of public discourse in hybrid media environments, where affect, ideology, and platform affordances co-constitute the structure of attention.

The topic model thus serves not just as a technical partition of the corpus but as a map of the discursive field through which DeepSeek becomes meaningful. It illuminates the patterned ways in which Iranian users frame foreign AI, highlighting the affective, analogical, and argumentative strategies they use to position themselves in a global technological imaginary. These patterns will be further interrogated in the following section, which maps the co-occurrence of topics and sentiment categories to reveal the connective tissue of Iranian digital discourse on artificial intelligence.

### 3.3. Co-occurrence Network Analysis

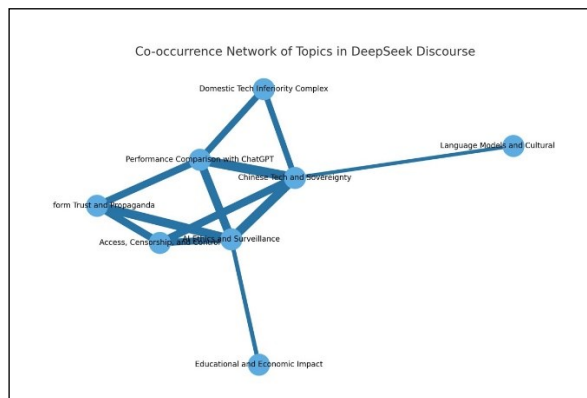
While topic modeling reveals the major thematic axes around which discourse about DeepSeek is organized, it does not, on its own, capture how these themes interact within user discourse. In reality, social media posts are not thematically discrete; they frequently blend multiple lines of argument, emotion, and analogy. A co-occurrence network analysis helps visualize this multidimensionality by mapping how often different topics appear together within the same post. Rather than treating topics as isolated partitions, this approach treats them as nodes in a dynamic semantic system, where links represent shared discursive terrain. The resulting network structure offers a more granular view of how Iranian users relate topics like technological performance, censorship, ethics, and cultural identity to each other in their evaluations of DeepSeek.

The co-occurrence matrix was constructed by coding each of the 1,112 posts for up to three topics based on their highest topic probabilities, then aggregating pairwise counts of how frequently topics appeared together across the corpus. The matrix was filtered to retain only the strongest connections—those with a minimum co-occurrence frequency of 24—allowing for a clearer and more interpretable network structure. In the final network, each node represents a topic, and each edge represents a non-trivial frequency of co-occurrence. The thickness of each edge is proportional to the number of times the two topics appeared in the same post. A visual representation of this network is provided in Figure 4.

Table 4. The co-occurrence matrix

Topic Pair	Co-occurrence Count
Performance Comparison with ChatGPT – Chinese Tech and Sovereignty	64
Performance Comparison with ChatGPT – AI Ethics and Surveillance	55
Performance Comparison with ChatGPT – Platform Trust and Propaganda	47
Chinese Tech and Sovereignty – AI Ethics and Surveillance	60
Chinese Tech and Sovereignty – Access, Censorship, and Control	51
AI Ethics and Surveillance – Platform Trust and Propaganda	58
AI Ethics and Surveillance – Access, Censorship, and Control	46
Domestic Tech Inferiority Complex – Performance Comparison with ChatGPT	43
Domestic Tech Inferiority Complex – Chinese Tech and Sovereignty	39
Platform Trust and Propaganda – Access, Censorship, and Control	49
Educational and Economic Impact – AI Ethics and Surveillance	28
Language Models and Cultural Identity – Chinese Tech and Sovereignty	24

The resulting network is both interpretable and analytically rich. At its center lie three highly interconnected topics: Performance Comparison with ChatGPT, Chinese Tech and Sovereignty, and AI Ethics and Surveillance. These three nodes form the core of the discourse, representing the primary frames through which DeepSeek is interpreted. Their high co-occurrence suggests that Iranian users often evaluate DeepSeek not simply by its functionality, but in terms of its national origins and its ethical implications. When users talk about performance, they often do so in the same breath as questioning the motives behind the model's development, its political alignment, or its potential misuse. This central triad captures the affective and geopolitical density of the discourse: DeepSeek is not just an application; it is a political artifact whose meaning is stabilized only in relation to multiple overlapping concerns.



*Figure 4. The co-occurrence network*

The edge between performance comparison and Chinese Technological Sovereignty, with 64 co-occurrences, is the strongest in the network. This bond reflects a recurring pattern in which users juxtapose DeepSeek's technical capabilities with its symbolic position in a broader geopolitical contest. For some, outperforming Western models is seen as proof of Eastern ascendancy; for others, comparisons are used to question the legitimacy of performance claims and to highlight state-driven narratives. In either case, the coupling of performance and sovereignty functions discursively to politicize the benchmark itself. What is ostensibly a technical evaluation becomes a proxy for geopolitical alignment and technological pride or skepticism.

Closely related is the connection between Chinese Technological Sovereignty and AI Ethics and Surveillance, with 60 instances of co-occurrence. This pairing highlights a deeper tension in how users process the political origins of AI. Posts in this cluster often praise China's rapid technological progress while simultaneously expressing concern about the authoritarian applications of that progress. The discourse does not appear polarized into pro- or anti-China camps but is instead marked by ambivalence: admiration is frequently tempered by suspicion, and excitement by critique. The overlapping of these two themes suggests that ethical concerns are not tangential but central to how foreign technological development is interpreted in Iran. Particularly in a society with its own experiences of surveillance and state control, users remain sensitive to the moral and political affordances of AI, especially when those affordances are mediated through another state's infrastructure.

A similarly strong bond is found between AI Ethics and Platform Trust or Propaganda. With 58 posts linking these two themes, this edge underscores how epistemic suspicion and ethical concern often travel together. Many users question whether DeepSeek's outputs can be trusted, not only due to technical limitations but because of perceived bias, omission, or ideological filtering. These critiques are not unique to DeepSeek – similar concerns have been raised about Western models – but the linkage between ethics and media narratives suggests that Iranian users are highly attuned to the ways in which AI is framed by institutions. Whether the model is framed by Chinese, Western, or Iranian media, users often assume that such framings are not neutral but motivated. This epistemological caution reflects a broader digital literacy shaped by years of navigating politicized information environments.

The next set of dense co-occurrences involves access, censorship, and control. This theme connects frequently with both Chinese Technological Sovereignty and AI ethics, forming a subnetwork focused on infrastructural and informational asymmetries. Posts here interrogate who gets to access DeepSeek, under what constraints, and with what tradeoffs. Co-occurrences between this topic and platform trust (49) or ethical concerns (46) indicate a communicative space where users reflect not only on what AI can do, but who it serves and what it excludes. These are not mere technicalities; they are expressions of social justice, user autonomy, and geopolitical agency. Within

these posts, AI is framed as a gatekeeping mechanism as much as a generative tool – one that may reinforce hierarchies even as it promises disruption.

Peripheral yet meaningful connections also emerge. Posts linking the Domestic Tech Inferiority Complex to performance comparison (43) and Chinese sovereignty (39) illustrate how national self-evaluation is structured by external reference points. Iranian users often measure their own capabilities not in absolute terms, but relative to others – most notably China and the West. These co-occurrences demonstrate that DeepSeek functions not just as a symbol of foreign achievement but as a mirror in which users reflect on domestic technological aspirations, failures, and contradictions. It becomes a comparative object through which Iran's developmental trajectory is debated and reimagined.

Further out on the network periphery are the themes of educational and economic impact and language models as tools of cultural identity. Their lower co-occurrence frequencies do not imply irrelevance, but rather thematic specificity. Educational and economic impact appears most frequently alongside AI ethics, indicating that concerns about AI's societal influence are anchored in normative questions about fairness, displacement, and skill obsolescence. Meanwhile, language models and cultural identity co-occur with Chinese sovereignty discourse, reflecting concerns about linguistic representation, cultural bias, and symbolic inclusion. These less frequent yet targeted connections serve to enrich the discursive field, offering perspectives that extend beyond the central triad of performance, politics, and ethics.

Taken together, the co-occurrence network reveals that Iranian discourse on DeepSeek is structured not by topic silos but by entangled thematic constellations. Users do not isolate technical appraisal from moral concern, nor do they separate geopolitical evaluation from domestic self-reflection. Instead, posts weave together performance, ethics, ideology, and access in complex ways. This discursive entanglement illustrates the depth of public reasoning in Persian-language digital spaces and challenges simplistic assumptions that user engagement with AI is either naive or purely instrumental.

What the co-occurrence network revealed was a communicative structure built not on binary positions but on negotiated tensions. Users hold multiple,

sometimes conflicting, affective and epistemic orientations simultaneously. They admire and distrust, hope and fear, critique and identify. These layered stances are not a sign of incoherence but a form of cognitive and affective sophistication, shaped by the hybrid media environment of Iran and the lived history of technological mediation under constraint. By mapping how topics are linked within the same utterances, the co-occurrence network gives us a window into how Iranian users not only talk about DeepSeek, but how they think with and through it – using it to reflect on the politics of innovation, the ethics of automation, and the very nature of communicative agency in an increasingly AI-mediated world.

#### **4. Conclusion**

This study set out to examine how Iranian digital publics articulate their affective, epistemic, and geopolitical orientations toward artificial intelligence through discourse on DeepSeek – a Chinese-developed large language model that entered the global AI scene as a symbolic counterpoint to Western platforms. What emerged from this inquiry was not merely a picture of user opinion about a particular technology, but a broader discursive formation in which DeepSeek functioned as a proxy, a mirror, and at times a screen onto which a wide range of political, cultural, and existential anxieties were projected. Rather than existing in isolation, user sentiments and thematic references clustered into patterned, relational formations. These patterns reflect not only attitudes toward DeepSeek, but also more deeply rooted narratives about technological sovereignty, platform politics, epistemic justice, and national developmental aspiration.

At the emotional level, the corpus was defined less by polarized endorsement or rejection than by a complex distribution of ambivalent and layered affect. While sentiments such as hope, pride, and skepticism were visibly present, they rarely appeared in isolation. Instead, posts frequently folded together affective registers that reflected deep ambivalence – hope tempered by doubt, pride laced with anxiety, or neutrality masking latent ideological critique. This affective heterogeneity is particularly significant when situated within the Iranian context, where digital communication is both politically constrained and rhetorically inventive. Users deploy emotion not merely as expression, but as strategy: a way to navigate censorship,

articulate veiled dissent, and manage geopolitical uncertainties. The range of emotional orientations toward DeepSeek thus indexes not only reactions to a foreign AI model, but the lived experience of negotiating technological modernity from a position of infrastructural and epistemic marginality.

Thematic analysis further revealed that DeepSeek discourse is structured around eight recurring topics that traverse domains of performance, politics, ethics, education, and identity. These topics are not isolated silos, but mutually referential discursive repertoires, which Iranian users mobilize to interrogate both the model's capabilities and its symbolic freight. Particularly striking was the centrality of geopolitical framing—especially the entanglement of DeepSeek with narratives of Chinese technological sovereignty and U.S.-China rivalry. For many users, DeepSeek was neither a neutral tool nor simply a foreign model; it was perceived as a political artifact, embedded in a web of global power asymmetries. This framing enabled users to situate themselves in relation to broader ideological currents, often through analogical reasoning: China's trajectory was read as a mirror, a warning, or an aspirational template for Iran. In these discourses, technology is not merely functional—it is moralized, politicized, and historicized.

Co-occurrence network analysis added a third dimension to this interpretive structure by revealing the connective logic of Iranian discourse on DeepSeek. Topics such as performance, ethics, and sovereignty were not merely co-present in the dataset—they were tightly entangled in individual posts, suggesting that user evaluations were embedded in a dense semantic network of overlapping concerns. What this reveals is a communicative field organized not by categorical binaries, but by negotiated tensions. Users simultaneously evaluated DeepSeek's utility and interrogated its origins; they assessed its fluency in Persian while questioning the cultural politics of representation; they noted its accessibility while flagging its potential for surveillance. These entanglements speak to a discursive sophistication that resists simplistic accounts of either techno-enthusiasm or techno-skepticism. Instead, they suggest that Iranian digital publics are engaging in what might be called critical domestication—the process by which foreign technologies are interpreted, appropriated, and re-coded through local ideological grammars.

## References

- Allen, G. C. (2019). *Understanding China's AI strategy*. Center for a New American Security. <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy>
- Bahmani, H. and Javaheri Tehrani, F. (2025). Cancer, Suffering, and the Role of Social Media: Insights from Iranian Nurses. *Journal of Cyberspace Studies*, 9(2), 487-503. <https://doi.org/10.22059/jcss.2025.394220.1142>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*(pp. 610–623). ACM. <https://doi.org/10.1145/3442188.3445922>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Ding, J. (2018). *Deciphering China's AI dream: The context, components, capabilities, and consequences of China's strategy to lead the world in AI*. Future of Humanity Institute, University of Oxford.
- Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29. <https://doi.org/10.1038/s41591-018-0316-z>
- Fast, E., & Horvitz, E. (2017). Long-term trends in the public perception of artificial intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 31, No. 1). <https://doi.org/10.1609/aaai.v31i1.10635>
- Holloway, D., & Green, L. (2016). The Internet of toys: A posthuman and multimodal analysis. *Media International Australia*, 170(1), 27–37. <http://dx.doi.org/10.1080/22041451.2016.1266124>
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
- Khodabin, M., Zibaei, F. and Piriyaee, F. (2023). AI Literacy and Digital Readiness in Iranian Media. *Journal of Cyberspace Studies*, 7(2), 299-320. <https://doi.org/10.22059/jcss.2025.396155.1166>
- Luckin, R. (2017). *Towards artificial intelligence-based assessment systems*. *Nature Human Behaviour*, 1(3), 1–4. <http://dx.doi.org/10.1038/s41562-016-0028>
- Milan, S. (2015). From social movements to cloud protesting: The evolution of collective identity. *Information, Communication & Society*, 18(8), 887–900. <https://doi.org/10.1080/1369118X.2015.1043135>
- Rahimi, B. (2011). The agonistic social media: Cyberspace in the formation of dissent and consolidation of state power in postelection Iran. *The Communication Review*, 14:3,158-178. <http://dx.doi.org/10.1080/10714421.2011.597240>

- Rahmatian, F., & Sharajsharifi, M. (2022). Reimagining MBA education in the age of artificial intelligence: A meta-synthesis. *Socio-Spatial Studies*, 6(1). <https://doi.org/10.22034/soc.2022.223610>
- Sakhaei, S., Soroori Sarabi, A., Tomraee, S., Khodabin, M., & Sharajsharifi, M. (2024). Disaster risk management and AI: A grounded theory approach to epidemic response. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(3), 1699-1708. <https://doi.org/10.62225/2583049X.2024.4.3.4420>
- Salehi, K., Habib Zadeh Khiyaban, S. and Sabbar, S. (2025). Artificial Intelligence and the Future of International Law and Power. *Journal of World Sociopolitical Studies*, 9(4), 923-958. doi: 10.22059/wsps.2025.401951.1552
- Salehi, K., Habib Zadeh Khiyaban, S. and Sabbar, S. (2026). Artificial Intelligence and Crime Detection: A Critical Review. *Journal of Cyberspace Studies*, 10(1), 181-197. <https://doi.org/10.22059/jcss.2025.402206.1179>
- Shahghasemi, E. (2017). Cultural schema theory. In, Kim, Y. Y. *The international encyclopedia of intercultural communication*. New York: Wiley. <https://doi.org/10.1002/9781118783665.ieicc0019>
- 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. (October 26, 2016). Human Rights against Human Rights: Sexism in Human Rights Discourse for Sakineh Mohammadi. *Society*, 53(6), 614-618. <https://doi.org/10.1007/s12115-016-0073-x>
- Sharifi Poor Bgheshmi, M. S., & Sharajsharifi, M. (2025). Managing the crisis: AI and the demise of national sovereignty? *Journal of World Sociopolitical Studies*, 9(4).
- Surden, H. (2019). Artificial intelligence and law: An overview. *Georgia State University Law Review*, 35(4), 19-22. <https://ssrn.com/abstract=3411869>
- Tomraee, S., Hosseini, S. H. and Toosi, R. (2022). Doctors for AI? A systematic review. *Socio-Spatial Studies*, 6(1), 13-26. <https://doi.org/10.22034/soc.2022.219431>
- Toosi, R., Tomraee, S., Khodabin, M. and Sakhaei, S. (2025). Telemedicine: An AI solution, at last?. *Code, Cognition & Society*, 1(1), 59-87. <https://doi.org/10.22034/ccsr.2025.526987.1002>
- Zamani, M., Nosraty, N. and Soroori Sarabi, A. (2025). Towards a business healthy lifestyle: Reducing risks while increasing efficiency?. *Code, Cognition & Society*, 1(1), 29-58. <https://doi.org/10.22034/ccsr.2025.526977.1001>
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.