

Children, consent, and control in Persian social media

Shaho Sabbar*

Department of Iranian Studies, Faculty of World Studies, University of Tehran,
Tehran, Iran. (*✉ shaho.sabbar@ut.ac.ir, <https://orcid.org/0000-0001-5801-7137>)

Article Info	Abstract
<p>Original article</p> <p>Main Object: Humanities & Social Sciences</p> <p>Received: 23 November 2025</p> <p>Revised: 25 November 2025</p> <p>Accepted: 26 November 2025</p> <p>Published online: 20 December 2025</p> <p>Keywords: digital rights, platformized childhood, Persian social media, sharenting, vernacular governance.</p>	<p>Background: Social media platforms such as Facebook, Instagram, TikTok, and X have redefined the norms of sociality, identity performance, and participation in the public sphere.</p> <p>Aims: This study examines how Persian-language users on X negotiate the visibility of children in online spaces through affective discourse and vernacular governance.</p> <p>Methodology: Analyzing 2,392 posts, we identify seven thematic formations—ranging from family blogging and sharenting to screenshot-based mockery, celebrity child cultures, and rights-based critiques. Using a hybrid methodological approach combining high-recall data retrieval, supervised multi-label topic modeling, and sentiment–intensity analysis, we map how practices like quote-tweeting and screenshotting structure public debates around parental branding, childhood agency, privacy, and consent. Central to this ecosystem is the culturally specific figure of “Arat’s father”, a discursive shorthand for the commodification of childhood under platform economies.</p> <p>Findings: The findings reveal a layered affective landscape where humor, outrage, and pedagogical neutrality coexist, enabling users to police age norms and negotiate ethical boundaries in real time.</p> <p>Conclusion: This study reveals how ordinary users in Iran and Persian-speaking contexts regulate childhood visibility through platform affordances, emotional repertoires, and normative claims. It also proposes a reproducible pipeline for analyzing culturally specific digital publics with methodological transparency and ethical sensitivity.</p>

Cite this article: Sabbar S. (2026). “Children, consent, and control in Persian social media”. *Cyberspace Studies*. 10(1): 239-256. doi: <https://doi.org/10.22059/jcss.2025.406842.1207>.



Creative Commons Attribution-NonCommercial 4.0 International License
 Website: <https://jcss.ut.ac.ir/> | Email: jcss@ut.ac.ir |
 EISSN: 2588-5502
 Publisher: University of Tehran

1. Introduction

The advent of new communication technologies has continually reshaped the structural, cultural, and emotional contours of human societies. From the invention of the printing press to the emergence of television, each wave of media innovation has produced new forms of connectivity, reconfigured social hierarchies, and altered the temporal and spatial dimensions of public life (McLuhan, 1964). In particular, the digitization of communication has accelerated the pace and expanded the scale at which information circulates, collapsing the boundaries between private and public spheres (Salehi et al., 2025). The Internet, and later mobile digital platforms, have not only enabled users to consume information but also to produce and disseminate it, thereby democratizing content creation while introducing new asymmetries of visibility, surveillance, and control (van Dijck, 2013). Networked societies do not simply mediate existing social processes but transform them, creating a new social morphology defined by decentralized interactions and algorithmic governance (Sabbar & Khiyaban, 2023). Within this context, platforms are not neutral tools; they are infrastructures that mediate affect, shape discourse, and reconstitute relations of power (Castells, 2009).

The proliferation of social media in the early 21st century represents a paradigmatic shift within this broader technological evolution. Social media platforms such as Facebook, Instagram, TikTok, and X have redefined the norms of sociality, identity performance, and participation in the public sphere. Unlike earlier broadcast media, social media invite constant interaction, co-presence, and self-surveillance, rendering everyday life into a stream of data and images (Boyd, 2014). The participatory affordances of these platforms have encouraged the rise of “networked publics” in which individuals are not only consumers but also producers (“prosumers”) of cultural content (Jenkins et al., 2015). Simultaneously, these spaces are shaped by platform logics such as algorithmic curation, metrics of engagement, and monetization structures, which incentivize visibility and virality over deliberation and care (Gillespie, 2018). The result is a digitally mediated environment in which personal disclosures, affective performances, and identity markers become key currencies for attention and influence. These transformations have deeply affected familial, professional, and political practices, inviting users to navigate complex ethical terrains around representation, labor, and authenticity in real time.

Among the most notable cultural consequences of social media is the intensification and democratization of celebrity culture (Harvey, 2018; Ellcessor, 2016; Sobande et al., 2022; Elliott, 2011; Mortensen & Jerslev, 2013; Lalancette & Raynauld, 2017; Stewart & Giles, 2019; Shahghasemi, 2025; Koh, 2021; Couldry, 2006). Traditional celebrity, once defined by media gatekeeping and institutional endorsement, has given way to what Marwick and Boyd (2011) call “micro-celebrity”—

a mode of self-presentation and self-branding that leverages social media to cultivate followers, monetize personal narratives, and sustain parasocial relations. This shift has broadened the field of visibility, enabling ordinary users to attain varying degrees of fame through strategic content production and affective labor. The rise of influencers, lifestyle bloggers, and family vloggers illustrates how private lives, including those of children, can be packaged for public consumption under a regime of platform capitalism. Notably, children have become prominent figures within this new ecology—not only as passive subjects but as active symbols in the construction of familial and aspirational narratives. From sponsored toy reviews to curated family feeds, the child figure has emerged as both content and commodity, participating in what Duffy (2017) describes as a gendered and affective “aspirational labor” economy.

The visibility of children on social media has raised significant ethical, legal, and emotional concerns. On the one hand, digital platforms have enabled parents to document, archive, and share their children's lives in unprecedented ways, often under the guise of care, pride, or community participation. On the other hand, these practices—commonly referred to as “sharenting”—risk exposing minors to surveillance, exploitation, and identity complications before they are capable of informed consent (Blum-Ross & Livingstone, 2017). Scholars have noted that children featured in social media content often become nodes of symbolic value, positioned to accrue likes, follows, and sponsorships, sometimes without meaningful safeguards or consideration of long-term consequences. Furthermore, the boundaries between documentation, performance, and commodification are increasingly blurred, particularly when platform metrics incentivize regular posting and emotional resonance. In certain cases, children are not merely included in content but positioned as the primary source of a family's digital capital—a dynamic that raises questions about agency, labor, and the intergenerational ethics of visibility. The implications extend beyond the domestic sphere; public responses to children's content—from fan engagement to ridicule and shaming—constitute a form of participatory governance that is often culturally coded and emotionally charged.

Within this matrix of visibility, performance, and governance, children's online presence becomes a site of vernacular negotiation—where users, platforms, and cultural scripts collide. In Persian-language digital publics, especially on platforms like X, these negotiations are mediated not only by global platform dynamics but also by localized cultural, religious, and legal frameworks that shape norms around family, modesty, and public decorum. Iranian and Persian-speaking users, navigating a semi-censored and diasporic media ecology, often articulate their concerns through a combination of humor, irony, and normative critique. The figure of “Arat's father”, for instance, has

become a symbolic shorthand in Persian social media for the parental overreach and commercialization of children's digital identities. This vernacular governance—expressed through quote-tweets, screenshots, and affective reactions—constitutes a decentralized, yet potent, mode of social regulation in the absence of formal legal enforcement or platform intervention.

This study examines how Persian-language users on X negotiate the visibility of children in online spaces through affective discourse and platform-native practices. Using a dataset of 2,392 public posts, the research deploys a hybrid methodological approach combining high-recall data retrieval, supervised topic modeling, and sentiment-intensity analysis to identify recurring themes and affective patterns. The study explores seven thematic formations, including family blogging, sharenting, screenshot-based mockery, celebrity child cultures, and rights-based critiques. Central to this inquiry is an investigation of how ordinary users articulate ethical boundaries, express ambivalence, and enact bottom-up forms of control over childhood visibility in the digital sphere. By situating Persian social media practices within broader debates on platformized childhood, digital rights, and affective publics, the study contributes to a more nuanced understanding of how care, consent, and commodification are negotiated in culturally specific contexts. It also proposes a reproducible analytic pipeline that respects linguistic nuance and ethical complexity, offering a methodological template for future research on vernacular digital governance.

2. Methodology

The corpus for this study was assembled from Persian-language activity on X with the specific aim of tracing how children's presence on social media is framed, evaluated, and morally negotiated. We began with a broad retrieval protocol that combined hand-crafted Persian seed queries, transliterated variants, and colloquial hashtags associated with family influencing, "sharenting", screenshot-based ridicule of minors, safety and age verifications, celebrity and aspirational cultures around children, early sports and performance pressure, child-rights advocacy, and the culturally salient "Arat-father" phenomenon. This strategy prioritized recall over precision to minimize topical blind spots. The initial harvest comprised 10,886 posts.

Cleaning proceeded in stages to preserve affective and discursive signal while removing mechanical noise. Exact duplicates and near-duplicates were identified using character-normalized n-gram similarity and eliminated; link-only, image-only, emoji-only, and auto-generated reposts without added text were filtered out; and a lightweight language gate retained posts with sufficient Persian content despite common code-switching. Text normalization harmonized Persian orthography (including Arabic/Persian yeh and kaf, zero-width non-joiners, and

diacritics), collapsed elongations, and removed URLs and @mentions. Emojis were not discarded wholesale; instead, they were abstracted to polarity-aware placeholders so that valence cues would remain available to the sentiment model without tethering analysis to glyph idiosyncrasies. After these steps, the analytic dataset comprised 2392 unique posts. All data were drawn from public posts and are reported in aggregate. Identifiers and screenshots of minors are not reproduced, and the figure of “Arat” is treated analytically as a discursive construct rather than an identifiable child. The resulting pipeline high-recall retrieval, careful denoising, orthographic normalization, multi-label topic tagging, sentiment/intensity modeling, and impact scoring yields a corpus that is both faithful to Persian vernacular expression and amenable to reproducible computational analysis. This foundation enables the study to link platform practices to thematic formations and affective repertoires with methodological transparency and interpretive rigor.

3. Findings

The distribution of text lengths, as shown in the histogram, reveals a multimodal pattern in how Persian-language users engage with the platform X when posting about children’s online presence. The majority of posts cluster between 20 and 60 words, with notable peaks around 25, 35, and especially near 55 words—suggesting a tendency toward mid-length, expressive content that balances emotional nuance with platform constraints. Posts shorter than 20 words are relatively rare, indicating that minimal expressions or reactions are not the dominant mode of engagement in this discourse. Instead, the length distribution points to a discursive style where users often elaborate on their perspectives, particularly in contexts such as ridicule, moral critique, or affective storytelling—each of which demands sufficient textual space to convey tone, irony, or argument. The long tail beyond 60 words, while less populated, reflects instances of extended commentary or narration, often tied to intense affect or detailed observations. Overall, this distribution supports the study’s claim that Persian social media users do not merely react passively to child-centered content but actively negotiate meaning through emotionally and morally charged narratives that require space for articulation.

Using the five-way affect schema, the corpus of 2392 Persian-language posts yields the following distribution: Happy accounts for 912 posts (38.13%), Neutral for 537 (22.45%), Furious for 462 (19.31%), Angry for 341 (14.26%), and Delighted for 140 (5.85%). Two features stand out. First, positive tenor is visibly present but not unambiguous: the dominance of “Happy” suggests that much of the conversation is framed through everyday affirmation, humor, or culturally valued scripts of pride and care; yet its distance from “Delighted” indicates that outright celebration is comparatively rare. In

Persian vernaculars on X, “happy” registers frequently perform conviviality and irony at once users can signal approval or social ease while preserving a posture of knowing skepticism. Second, negative affect is substantial: combining “Furious” and “Angry” yields 33.57% of all posts, a magnitude consistent with grievance-driven dynamics in affective publics. This polarity often surfaces around perceived overexposure of minors, anxieties about consent and privacy, and the disciplining gaze of screenshot culture; anger here is not merely impulsive but a rhetorical resource that polices norms, calls out alleged exploitation, and negotiates the limits of parental authority and peer mockery.




Figure 1. Distribution of text lengths (in words)

The Neutral band roughly one-fifth of the dataset does not imply emotional absence. Rather, it frequently carries observational critique, dry reportage, and template-like commentary that platform cues encourage (linking, quoting, noting a profile), allowing users to index stance through minimal propositional content while delegating much of the affective work to the embedded artifact (a bio, a story, a statistic).

In aggregate, then, the emotional field is asymmetrically positive but thick with friction: convivial “Happy” posts sit alongside a sizable stratum of indignation and outrage, while a neutral middle keeps the discourse legible and transmissible. This structure is typical of platformed deliberation where humor softens edges without dissolving conflict, and where indignation travels quickly through low-cost acts of quoting and screenshotting. For the analyses that follow, this distribution justifies a dual focus on (a) how positivity is performed without tipping into celebratory unanimity and (b) how negative affect underwrites vernacular governance bottom-up attempts to regulate what children should post, what adults may amplify, and who gets to judge in the first place.




Figure 2. Dataset sentiment distribution

4. Topic modeling and thematic analysis

To recover the latent thematic structure of Persian-language discourse about children on social media, we implemented a Latent Dirichlet Allocation (LDA) pipeline on the cleaned corpus described above. Preprocessing was optimized for colloquial Persian on X: orthographic normalization (Persian–Arabic yeh/kaf, zero-width non-joiners), de-longation, tokenization preserving multiword expressions and hashtags, and selective retention of affective markers while stripping URLs and mentions. This yielded a vocabulary in which vernacular cues memetic phrasings, ironic particles, and platformed formulae were preserved rather than flattened.

The resulting document–term matrix was pruned with conservative frequency thresholds to reduce sparsity without erasing minority registers central to rights and safety talk. Model selection proceeded by sweeping the number of topics from 5 to 20 and computing standard semantic coherence for each specification. Coherence increased sharply from five to seven topics and attained its global maximum at 7 topics ($C=0.5181$). Additional local peaks appeared at 14 topics ($C=0.5041$) and 19 topics ($C=0.4954$), with intervening solutions showing either modest declines or instability (e.g., 8 topics at 0.4464). Substantively, the seven-topic solution cleanly separated the major formations visible in exploratory reading: parental exploitation centered on the “Arat-father” figure, screenshot-based mockery of minors’ profiles and Stories, ethics/safety concerns around underage presence and privacy, family blogging and sharenting, celebrity/aspirational cultures around children, early sports/body conditioning and performance pressure, and a rights/abuse discourse while minimizing cross-topic leakage. In contrast, higher-K models tended to split these cores along stylistic seams (e.g., ironic vs. declarative mocking; Instagram-centric vs. cross-platform sharenting), marginally raising coherence but eroding interpretability and inflating redundancy. Given this profile, we adopt

the 7-topic model as the primary analytic frame and treat the 14-topic solution as a robustness check to probe sub-facets when needed (for instance, distinguishing safety talk focused on data privacy from that focused on reputational harm). Qualitative inspection of top-weighted documents per topic confirmed face validity: salient keywords cohered with vernacular labels, and the most representative posts exhibited clear intra-topic homogeneity alongside the expected inter-topic contrasts. At the same time, LDA surfaced hybrid utterances that straddled clusters especially posts that combine playful ridicule with normative boundary-setting underscoring that Persian discourse around children online is a braided formation rather than a set of sealed compartments.

Methodologically, the coherence landscape justifies a parsimonious taxonomy that maximizes semantic clarity while preserving space for overlap. We therefore map subsequent analyses sentiment profiles, co-occurrence structures, and engagement dynamics onto the seven topics as the communicative backbone of the corpus, and draw on the finer 14-topic segmentation only where it advances interpretation without sacrificing readability. This dual resolution allows us to describe the thematic field with conceptual economy while acknowledging the granular textures through which affect, morality, and platform practice interlock in Persian-language publics on X.




Figure 3. Coherence score vs number of topics (with n-grams)

4.1. Arat-father / parental exploitation

Posts referencing the “Arat-father” figure cohere around tensions between paternal branding and a child’s autonomy. Within this topic ($n=65$), positive registers (“Happy” 26; “Delighted” 5; $\approx 48\%$) sit alongside a sizable tranche of neutrality (20; $\approx 31\%$) and concentrated indignation (“Furious” 11; “Angry” 3; $\approx 22\%$). The positive and neutral bands often narrate the everyday labor of child-centred content backstage vignettes of prompting, staging, and playful cajoling framed

as benign family entrepreneurship. Yet even casual admiration tends to be hedged by self-awareness about platform incentives. Critical posts mobilize the language of abuse, consent, and digital labor, describing the child as a performer under adult direction and the father as a micro-manager of attention. One user, reacting to a clip of a nine-year-old being directed for a room-tour, glosses: “This is textbook child exploitation; why is his father’s income tied to the kid’s face?” Another, more sardonic, translates parental hustle into moral shorthand: “Cute, yes but it’s still the dad’s business model.” The topic thus functions as a diagnostic lens on platformized parenthood: affectively ambivalent, narratively intimate, and normatively contested. It surfaces a vernacular ethics in which “care” and “conversion” (of intimacy into visibility) are negotiated in real time, with users oscillating between indulgence and censure.

Table 1. Topic titles and distribution of posts

Topic Title	Total Posts	Furious	Angry	Neutral	Happy	Delighted
Family blogging with children	1737	363	264	336	691	83
Mocking minors profiles/ stories	168	12	23	67	60	6
Child rights & abuse discourse	100	25	10	15	42	8
Fame/ celebrity cultures around kids	87	17	11	22	34	3
Arat-father/ parental exploitation	65	11	3	20	26	5
Early sports & performance pressure	45	12	6	6	18	3
Ethics & safety (underage/ privacy/ risk)	40	8	8	10	14	0

4.2. Mocking minors’ profiles/stories (screenshot-based shaming)

This is the corpus’s most populous and stylistically distinctive formation (n=168), anchored in screenshot culture and age policing. The modal tone is observational or jocular (“Neutral” 67; “Happy” 60; together ≈76%), while explicit condemnation remains a minority (“Angry” 23; “Furious” 12; ≈21%); outright celebration is rare (“Delighted” 6). Posts typically stage a minor’s bio or Story as an object of public taste judgment an ironic theater where cuteness, precocity, and cringe are sorted. A typical voice reads: “My eleven-year-old cousin posted his report card ‘Mom must come sign’ and he added a swagger caption.” Another quips about a teen aunt’s profile anniversary: “Third year together with Bob Sponge I’m replying, I’m thrilled =)))”. Humor lubricates surveillance: ridicule enforces boundaries on what is “age-appropriate,” while allowing posters to disclaim cruelty through playfulness. Counter-voices surface intermittently, calling the practice bullying or public shaming, but the dominant script is performative:

users accrue likes through wry exposure while reaffirming communal norms of maturity and self-presentation. The result is a soft disciplinary apparatus light on moralizing, heavy on irony that nonetheless shapes minors' vernacular self-fashioning.

4.3. Ethics & safety (underage, privacy, risk)

Although smaller in volume ($n=40$), this topic exhibits a sharper negative tilt (Angry=8; Furious=8; 40%) and a conspicuous absence of unalloyed delight (Delighted=0). "Happy" posts (14; 35%) often reflect relief at protective action accounts locked, faces blurred, or adults intervening whereas "Neutral" commentary (10; 25%) circulates procedural advice on reporting, age checks, and privacy settings. The negative band is lexically dense with risk terms: grooming, doxxing, blackmail, reputational harm, and "future self". One instance observes: "A 14-year-old dancing on Instagram faces public, comments wide open; how is this anyone's idea of safety?" Another laments an underage engagement: "We normalized adult scripts for kids; now we panic about the fallout." Here, platform affordances are reframed as threat surfaces: visibility becomes exposure, virality becomes loss of control, and archives become liabilities that outlive juvenile intent. Compared with screenshot-mockery, the tone is less theatrical and more diagnostic; users appeal to child-rights language, invoke platform policies, and imagine concrete harms traveling from online scenes into offline futures. This cluster therefore anchors the corpus's normative horizon, translating scattered anxieties into a minimally shared ethic of consent, proportionality, and reversibility in children's digital footprints.

4.4. Family blogging with children

Family-centric content is the largest formation in the corpus ($n=1,737$), and its affective profile is ambivalent in a patterned way. Positive registers are prominent Happy $\approx 40\%$ (691) and Delighted $\approx 5\%$ (83) often celebrating cuteness, kinship pride, and the routinized craft of staging posts. Yet positivity rarely equals naïveté. A sizeable Neutral band $\approx 19\%$ (336) reads like platform-literate commentary on aesthetics and metrics, while the negative tail Angry $\approx 15\%$ (264) and Furious $\approx 21\%$ (363) foregrounds concerns about consent, overexposure, and the conversion of intimacy into content. Typical appreciative remarks praise a child as "a bundle of her parents' charms," but even admiring voices index the production apparatus behind the feed. Critical posts, by contrast, frame the same scenes as reputational risk or unpaid digital labor for the minor. One celebratory user writes, "Saw a little girl's page she's a doll; the whole account radiates warmth." A counter-voice laments, "A pretty teen cousin turns into everyone's algorithm great for likes, bad for the girls watching." Read together, these strands reveal a vernacular negotiation between care and conversion: followers reward

family branding with attention while simultaneously policing its ethical perimeter.

4.5. Fame/celebrity cultures around kids

This smaller cluster (n= 87) tracks how children are drawn into celebrity logics mini-influencers, look-alikes, and affiliative fame. The tone is mixed but skewed light: Happy \approx 39% (34) and Neutral \approx 25% (22) outnumber the negatives (Angry \approx 13% (11); Furious \approx 20% (17); Delighted \approx 3% (3)). Much of the positivity is playful fandom: users applaud precocious performance, clever edits, or kin who “manage” a child’s online persona. Yet the critique is sharp when adult scripts are mapped onto minors romance tropes, parasocial teasing, and “crush” narratives. A typical post chuckles, “My 13-year-old cousin set up a gaming channel; my aunt claps like a pro manager honestly, kudos.” A more sardonic one notes, “New followers need a PSA the ‘bride’ joke wasn’t a confession; when kids copy celebrity storylines, adults should not.” The cluster thus stages aspiration versus appropriateness: celebration of visibility is never far from anxiety about scripting childhood through fame templates.

4.6. Early sports & performance pressure

Although modest in size (n= 45), this topic concentrates debate about training regimes, branded academies, and projecting professional futures onto children. The affective split is telling: Happy \approx 40% (18) and Delighted \approx 7% (3) praise discipline, teamwork, and hope; Neutral \approx 13% (6) recounts routines and facilities; Angry \approx 13% (6) and Furious \approx 27% (12) call out premature specialization and parental overreach. Admiring posts frame sport as character-building “Sign him early; he loves the drills”. Critical ones highlight instrumentalization “Dad lays out training mats and calls the boy a football prodigy; whose dream is this?” A diaspora voice recalls shutting down a page after realizing “the scene was too heavy for kids.” What emerges is a platform vernacular of risk/ benefit accounting: communities endorse structured play but resist pipelines that compress leisure into brandable performance, especially when the child’s agency is opaque and the future self bears the costs of today’s publicity.

4.7. Child rights & abuse discourse

Rights-oriented talk is smaller in volume yet normatively dense. In our corpus (n= 100), roughly half the posts convey supportive or advocacy tones (Happy= 42; Delighted= 8), a fifth adopt informational or procedural neutrality (Neutral= 15), and just over a third express grievance or alarm (Angry= 10; Furious= 25). The positive subset rarely signals cheerfulness in the colloquial sense; rather, it frames protection, consent, and redress as collectively desirable ends (“lock the account; blur the face; report, don’t repost”). Neutral entries circulate

guidance links to reporting flows, reminders about minimum-age policies, and checklists for privacy settings performing a pedagogical role that keeps the conversation transmissible across networks. The negative band is lexically saturated with harm vocabularies assault, grooming, doxxing, stigma, future self and often narrates specific risk scenarios to underline irreversibility. One widely echoed sentiment paraphrases an Instagram thread: “If a ten-year-old is assaulted and forced into public visibility, whose rights are we protecting when we circulate the clip?” A counter-example from the advocacy side stresses proportionality: “Document, report, and preserve evidence don’t make a child’s face your content.” Taken together, this cluster re-centers the debate on the child as rights-bearing subject rather than a vehicle for parental branding or peer entertainment. It reframes platform affordances as governable infrastructures feeds, archives, virality as objects of ethical choice, insisting on consent, minimization, and reversibility as baseline norms for Persian-language publics engaging with minors online.

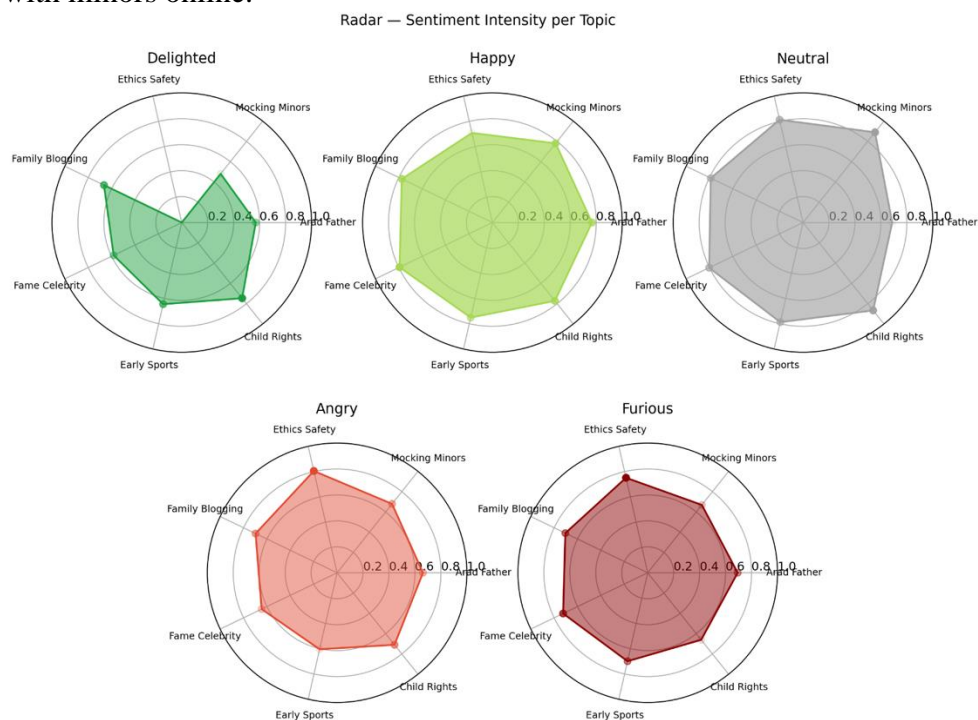


Figure 4. Radar— sentiment intensity per topic

The seven-topic network derived from the CSV resolves into a hub-and-spoke formation anchored by Family Blogging with Children ($n=1,737$). In raw co-tags, the strongest edge links Family Blogging to Child Rights & Abuse Discourse ($n=89$), followed by ties to Fame/Celebrity Cultures ($n=29$) and Ethics & Safety ($n=21$). This pattern indicates that the routine circulation of child-centred content is

where normative contestation most visibly accretes: rights-based arguments, safety heuristics (age, privacy, reversibility), and celebrity logics are not parallel conversations but converge on the family-influencer feed. Importantly, the Family Blogging ↔ Child Rights edge remains strong even under a relative metric (Jaccard ≈ 0.051), which counters the trivial explanation that the hub merely inflates co-occurrence by size.




Figure 5. Sankey diagram of topic-sentiment relationship

A second layer appears when we privilege relative overlap rather than raw counts. The highest Jaccard is Arat-Father ↔ Early Sports (0.10; 10 co-tags), followed by Family Blogging ↔ Child Rights (0.051), Ethics & Safety ↔ Child Rights (0.045), and Arat-Father ↔ Child Rights (0.038). These ratios show where topics are mutually constitutive: posts about the “Arat-father” figure frequently travel with performance pipelines and training regimens; when users discuss underage visibility in safety terms, they also articulate a rights grammar; and rights talk, in turn, is disproportionately triggered by family-produced content. By contrast, Mocking Minors’ Profiles/Stories despite co-tagging with Family Blogging ($n=11$) registers no measurable overlap with Child Rights in this corpus ($n=0$). That absence is analytically meaningful: screenshot-based ridicule functions as a vernacular disciplinary script (taste, maturity, cringe) more than a rights-claiming discourse.

The Fame/Celebrity node sits between admiration and alarm. It co-occurs with Family Blogging ($n=29$; Jaccard ≈ 0.016) and registers smaller but non-trivial bridges to Child Rights ($n=3$; 0.016) and Ethics & Safety ($n=2$; 0.016), suggesting that the celebrity template is a magnifier rather than an independent sphere: when children are framed as look-alikes, mini-influencers, or affiliates of adult fame, users are

more likely to invoke rights and risk. Finally, the Arat-Father node links into both Child Rights ($n=6$; 0.038) and Family Blogging ($n=3$; ≈ 0.0017), but its most distinctive bridge is to Early Sports a small yet tight dyad that foregrounds instrumentalization and performance pressure. Taken together, the network depicts an economy where everyday family content is the gravitational center, rights/safety act as normative counterweights, and mockery polices boundaries largely outside a rights vocabulary. This structure helps explain why affective flare-ups cluster around familiar, high-visibility practices (family feeds), while smaller, high-Jaccard couplings (e.g., Arat-father with training regimes) mark sites where users renegotiate the line between care, ambition, and exploitation.

7-Topic Co-occurrence Network — Kids on Social Media




Figure 6. Kids on social media; 7-topic co-occurrence network

7. Conclusion

This study has examined how Persian-language users on X negotiate the visibility of children in digital publics. The paper reveals an affectively ambivalent and structurally complex discourse ecosystem. Across 2,392 posts, we identified a layered field of vernacular governance where users both participate in and critique the cultural production of childhood online. Through a hybrid methodology that combined high-recall corpus building, topic modeling, sentiment analysis, and co-occurrence mapping, this research has made visible the implicit norms, emotional repertoires, and ethical tensions that shape user engagements with child-centered content in Persian social media. The findings invite a broader reflection on how childhood is being

reimagined, regulated, and instrumentalized in platform societies—particularly in contexts where formal regulation is limited, and informal, culturally encoded modes of oversight proliferate.

The first major finding of this study lies in its reconceptualization of affect as a mode of governance. Rather than treating emotions as merely expressive or reactive, the analysis demonstrates how affect operates as an instrument of norm enforcement and social ordering. “Happy” posts, for instance, often perform conviviality and irony in tandem—signaling social approval while hedging against full endorsement. Conversely, the significant presence of “Angry” and “Furious” posts points to an emergent moral economy in which grievance is a communicative resource used to police perceived violations of consent, propriety, and child autonomy. These affective distributions are not random; they correlate systematically with topic-specific concerns, as seen in the heightened negativity surrounding the Arat-father phenomenon and early sports training, and the ambiguous positivity surrounding family blogging and mockery. Together, they form what could be termed an “affective infrastructure”—a discursive scaffolding through which Persian users make ethical judgments, perform solidarity or censure, and regulate the boundaries of acceptable child visibility.

Second, the study advances the concept of vernacular governance in digital childhoods. In the absence of clear institutional protocols or enforceable child rights protections online, ordinary users develop their own forms of ethical reasoning and boundary-making. These vernacular scripts often rely on platform-native practices—quote-tweeting, screenshotting, ironic captioning—and culturally specific reference points like the figure of “Arat’s father.” These mechanisms perform more than critique; they instantiate soft architectures of accountability, signaling to others what is tolerable, shameful, exploitable, or empowering. For instance, screenshot-based ridicule does not explicitly invoke legal or rights frameworks, but it functions as a potent disciplining device, steering children and their caregivers toward normative models of online self-presentation. Similarly, mockery, praise, and outrage become tools through which the Persian-speaking public navigates the tension between parental authority and child agency, between visibility and vulnerability.

The study’s thematic mapping through topic modeling confirms that Persian social media discourse on children is neither homogenous nor dichotomous. Instead, it is braided: topics such as family blogging, celebrity aspiration, and child rights do not exist in isolation but intersect in dynamic and sometimes contradictory ways. The hub-and-spoke structure of the thematic network, with family blogging as the gravitational center, reflects the centrality of parental content in public debates about children. Around this hub, we observe a range of ethical counterweights—from the rights discourse to the Arat-father critique, from safety concerns to disciplinary humor—each representing a

different axis of contestation. Importantly, these formations are not merely descriptive categories but sites of affective investment and moral labor. Users do not passively consume child content; they curate, reframe, judge, and re-circulate it, often with a self-aware irony that marks the discourse as both performative and reflexive.

This study also surfaces a critical insight about the platformization of childhood. As children increasingly appear as content producers, branded figures, or symbolic assets in their parents' digital performances, their identities become entangled in the attention economy and its extractive logics. The posts analyzed in this study capture the cultural ambivalence surrounding this shift. On one hand, children are sources of pride, aesthetic pleasure, and familial intimacy; on the other, they are seen as exploited laborers, targets of ridicule, or unwitting symbols of parental ambition. The ambivalence is not merely attitudinal—it reflects structural contradictions in how digital platforms incentivize visibility while eroding consent. In this context, the figure of the “mini-influencer” or “football prodigy” becomes emblematic of a broader social condition: the commodification of potential and the outsourcing of aspiration to the bodies and faces of children. The public's mixed reactions—ranging from admiration to satire to indignation—underscore a collective struggle to name and navigate this new moral terrain.

Moreover, the study's co-occurrence analysis highlights which combinations of topics activate stronger ethical reflexes. The most notable pairings—such as Arat-father with Early Sports, or Family Blogging with Child Rights—suggest that it is not visibility per se that provokes concern, but the confluence of visibility with instrumentalization, labor, and reputational risk. By contrast, the absence of co-tagging between Child Rights and Mockery underscores the limits of ethical recognition in certain affective registers. Humor, it appears, can displace or defer moral scrutiny even when the same content—such as a child's image or bio—is circulated in more explicitly critical contexts. This asymmetry is analytically significant: it shows that ethical salience is not uniformly distributed across discursive genres, but is shaped by tone, context, and platform norms. Thus, interventions aimed at protecting children's rights online must attend not only to content but to the modalities through which that content travels and is interpreted.

From a methodological standpoint, this study contributes a replicable and culturally sensitive pipeline for analyzing affective publics in non-Western digital environments. By calibrating its data collection strategies to Persian orthographic idiosyncrasies and affective idioms, and by integrating topic modeling with sentiment-intensity profiling and co-occurrence logic, the research demonstrates how large-scale analysis can preserve nuance without sacrificing rigor. In doing so, it avoids the pitfalls of both surface-level quantification and

purely anecdotal critique. Future research could extend this pipeline to other linguistic contexts or platforms, as well as incorporate longitudinal or participatory components to trace how norms around child visibility evolve over time or are negotiated by children themselves.

Ethically, the study emphasizes that public visibility is not equivalent to ethical transparency. Even when content is technically “public,” its reuse, circulation, and interpretation are morally fraught—especially when it involves minors. Analysts, commentators, and platforms must remain vigilant about the temporal and relational asymmetries embedded in such content: the fact that today’s post may haunt tomorrow’s adolescent, or that a father’s branding decision may foreclose a child’s future digital autonomy. The notion of a “future self,” repeatedly invoked in the corpus, points to an emergent ethics of reversibility—an awareness that visibility cannot be undone, and that digital traces can outlive the affective and developmental contexts that produced them. As such, this study calls for a child-centered framework of digital rights that prioritizes consent, proportionality, and temporality—not merely in legal terms, but as vernacular principles of everyday governance.

Conflict of interest

The author declared no conflicts of interest.

Ethical considerations

The author has completely considered ethical issues, including informed consent, plagiarism, data fabrication, misconduct, and/or falsification, double publication and/or redundancy, submission, etc. This article was not authored by artificial intelligence.

Data availability

The dataset generated and analyzed during the current study is available from the author on reasonable request.

Funding

This research did not receive any grant from funding agencies in the public, commercial, or non-profit sectors.

References

- Blum-Ross, A. & Livingstone, S. (2017). “Sharenting, parent blogging, and the boundaries of the digital self”. *Popular Communication*. 15(2): 110-125. <https://doi.org/10.1080/15405702.2016.1223300>.
- Boyd, D. (2014). *It’s Complicated: The Social Lives of Networked Teens*. Yale University Press.
- Castells, M. (2009). *Communication Power*. Oxford University Press.
- Couldry, N. (2006). “Culture and citizenship: The missing link? The missing link?”. *European Journal of Cultural Studies*. 9(3): 321-339.

- <https://doi.org/10.1177/1367549406066076>.
- Duffy, B.E. (2017). *(Not) Getting Paid to Do What You Love: Gender, Social Media, and Aspirational Work*. Yale University Press.
- Elkessor, E. (2016). “‘One tweet to make so much noise’: Connected celebrity activism in the case of Marlee Matlin”. *New Media & Society*. 20(1): 255-271. <https://doi.org/10.1177/1461444816661551>.
- Elliott, A. (2011). “‘I Want to Look Like That!’: Cosmetic surgery and celebrity culture”. *Cultural Sociology*. 5(4): 463-477. <https://doi.org/10.1177/1749975510391583>.
- Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and The Hidden Decisions that Shape Social Media*. Yale University Press.
- Harvey, A. (2018). “The fame game: Working your way up the celebrity ladder in Kim Kardashian: Hollywood”. *Games and Culture*. 13(7): 652-670. <https://doi.org/10.1177/1555412018757872>.
- Jenkins, H.; Ito, M. & Boyd, D. (2015). *Participatory Culture in a Networked Era: A Conversation on Youth, Learning, Commerce, and Politics*. Polity Press.
- Koh, W. (2021). “A new day for Hulk Hogan: Celebrity selves and racial diversity in contemporary professional wrestling”. *European Journal of Cultural Studies*. 25(2): 778-796. <https://doi.org/10.1177/13675494211028865>.
- Lalancette, M. & Raynald, V. (2017). “The power of political image: Justin Trudeau, Instagram, and celebrity politics”. *American Behavioral Scientist*. 63(7): 888-924. <https://doi.org/10.1177/0002764217744838>.
- Marwick, A. & Boyd, D. (2011). “To see and be seen: Celebrity practice on Twitter”. *Convergence*. 17(2): 139-158. <https://doi.org/10.1177/1354856510394539>.
- McLuhan, H.M. (1964). *Understanding Media: The Extensions of Man*. McGraw-Hill.
- Mortensen, M. & Jerslev, A. (2013). “Taking the extra out of the extraordinary: Paparazzi photography as an online celebrity news genre”. *International Journal of Cultural Studies*. 17(6): 619-636. <https://doi.org/10.1177/1367877913503425>.
- Sabbar, S. & Khiyaban, S.H.Z. (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. (2025). “Artificial Intelligence and the future of international law and power”. *Journal of World Sociopolitical Studies*. 9(4): 923-958. <https://doi.org/10.22059/wsps.2025.401951.1552>.
- Shahghasemi, E. (2025). “‘Woke’ in translation: Persian perspectives on Platform X”. *Discov Glob Soc*. 3. 104. <https://doi.org/10.1007/s44282-025-00253-x>.
- Sobande, F.; Kanai, A. & Zeng, N. (2022). “The hypervisibility and discourses of ‘wokeness’ in digital culture”. *Media, Culture & Society*. 44(8): 1576-1587. <https://doi.org/10.1177/01634437221117490>.
- Stewart, S. & Giles, D. (2019). “Celebrity status and the attribution of value”. *European Journal of Cultural Studies*. 23(1): 3-17. <https://doi.org/10.1177/1367549419861618>.
- van Dijck, J. (2013). *The Culture of Connectivity: A Critical History of Social Media*. Oxford University Press.