Algorithmic curation, information security, and public trust on social platforms: Case studies of TikTok and YouTube

Dinara Akbergen[1,*,†], Aidiye Aidarbekov[1,2,†] and Ha Jin Hwang[1,†]

[1] *Astana IT University, Mangilik El Avenue 55/11, 010000 Astana, Kazakhstan*
[2] *Maqsut Narikbayev University, Korgalzhyn Highway 8, 010000 Astana, Kazakhstan*

Abstract

This paper examines how algorithmic curation on social platforms affects information security and public trust. We synthesize recent findings on exposure drift, homogeneity, amplification, coordinated inauthentic behavior, and limits of user control, focusing on YouTube and TikTok. We outline an audit and forensics toolkit that combines black-box and counterfactual experiments with provenance and integrity checks, and we propose an operational workflow for oversight: detect, assess, mitigate, and report. Case studies highlight platform-specific dynamics: on YouTube, risks concentrate in narrow topical corridors and extended recommender-only sessions, with faster adaptation in the sidebar than on the homepage; on TikTok, short video affordances enable rapid niche lock-in, stronger coordination signals, and persistence of unwanted content for some users. We discuss governance options, including exposure diversity constraints, external auditability, and privacy-preserving transparency, and we conclude with priorities for reproducible evaluation.

Keywords

algorithmic curation; cybersecurity; digital forensics; recommendation systems; amplification; echo chambers; coordinated inauthentic behavior; TikTok; YouTube

1. Introduction

Algorithmic curation is now the default gateway to information on major social platforms. TikTok's For You Page (FYP) and YouTube's recommendations optimize for engagement and predicted relevance, reshaping how users encounter facts and viewpoints. This has direct implications for cyber risk management: recommender systems may interact with adversarial tactics (e.g., coordinated inauthentic behavior) and organic dynamics (e.g., homophily), leading to amplification, selective exposure, and erosion of public trust.

On YouTube, evidence points to a nuanced risk profile. User-facing and audit studies report mild ideological echo chambers and a modest right-leaning drift over longer sessions that follow recommendations, while finding limited systematic "rabbit holes" into extreme content for average users [1]. Causal experiments with counterfactual bots (post-2019) further suggest that recommendations can, on average, moderate consumption relative to user-driven trajectories; notably, sidebar suggestions "forget" prior far-right preferences after about 30 videos when users switch to moderate content [2]. These findings imply that risks depend on topic, session behavior, and platform design.

On TikTok, the risk surface differs due to short-video formats, rapid trend cycles, and the central role of the FYP. Recent computational work documents coordinated inauthentic behavior adapted to

[1] *CISN 2025: Workshop on Cybersecurity, Infocommunication Systems and Networks, November 19-20, 2025, Almaty, Kazakhstan*

* Corresponding author.
† These authors contributed equally.

242885@astanait.edu.kz (D. Akbergen); aidiye.aidarbekov@astanait.edu.kz (A. Aidarbekov); hjhwang@astanait.edu.kz (H. Hwang)
0009-0005-3359-3900 (D. Akbergen); 0000-0001-7197-1638 (A. Aidarbekov); 0000-0001-8752-1906 (H. Hwang)

video-first affordances, synchronized posting, content reuse, and hashtag-sequence overlaps, which creates distinct detection challenges [3]. Qualitative accounts complement this picture with reports of algorithmic persistence, unwanted content recurring despite negative feedback, raising questions about user control, trust, and exposure diversity in interest-driven feeds [4].

These observations motivate combining audit and measurement with digital forensics and governance. The audit literature systematizes harm classes (discrimination, distortion, exploitation, misjudgment) and distills effective methods such as sock puppets, scrapes, and crowd studies, while noting under-audited domains such as TikTok [5]. Proposals for platform-supported auditing argue that vetted researcher access to relevance estimators can reconcile transparency with privacy and intellectual-property protection, enabling routine oversight [6]. In parallel, media forensics provides provenance and integrity tools that are essential for attribution and incident response across platforms [7].

This paper makes three contributions to the study of algorithmic risk and platform governance. First, we consolidate fragmented findings from auditing and media forensics into a unified operational taxonomy, identifying how amplification, drift, and coordination signals manifest across TikTok and YouTube. Second, we propose a reproducible oversight workflow that links detection, assessment, mitigation, and reporting, bridging auditing methods with digital forensics rather than treating them as separate research tracks. Third, we extend existing governance models by outlining implementation-ready levers, such as exposure-diversity constraints and privacy-preserving transparency, and by specifying the complementary roles of platforms, regulators, and researchers. Together, these contributions move beyond narrative review and provide a transferable framework for risk management across recommendation systems.

2. Related work

2.1. Audits of recommendation systems

Prior work on recommendation audits has focused on three recurring questions: (i) what harms are produced or amplified by ranking and personalization, (ii) how to measure them externally with black-box or user-task methods, and (iii) how to connect audit findings to platform or regulatory responses. Early studies on YouTube showed that controlled “recommender-only” sessions can lead to ideological or topical narrowing compared to mixed navigation, and that different surfaces (sidebar vs homepage) adapt at different speeds [1], [2], [9].

2.2. Short-video and TikTok-style feeds

Later work extended audit techniques to short-video platforms, especially TikTok, where content is shorter, signals are denser, and coordination is easier to hide. These studies add CIB-relevant indicators such as media reuse, synchronized posting in short windows, and repeated hashtag sequences, and show that unwanted content can reappear despite explicit user feedback [3], [4].

2.3. Media forensics and platform-supported auditing

In parallel, media-forensics research developed provenance, integrity, and cross-platform lineage tools for incident investigation on social media; platform-supported auditing proposals added privacy-preserving access for vetted researchers, but this line is often treated separately from recommendation audits [6], [7]. Our paper joins the two by using forensics as the evidentiary layer inside an audit-driven risk workflow.

2.4. Compliance-oriented social media risk models

Compliance-oriented models link incident detection with identity abuse, mis/disinformation, and reporting obligations, but they rarely model recommendation-specific dynamics such as surface-level adaptation speed or feedback-suppression efficacy. We make this link explicit so that audit findings on YouTube and TikTok can be integrated into organizational or platform governance [14].

3. Methodology

We use a scoping review with two concise case studies (YouTube, TikTok). The objective is to synthesize recent empirical findings and practical auditing/forensics methods relevant to cyber-risk contexts.

Corpus and selection: the review is based on the set of papers you provided plus references cited within those papers. We applied backward and forward snowballing inside this corpus. No additional database queries were run outside the provided materials.

Synthesis approach: we use narrative synthesis. For each study we extract platform, study design, key measures, main findings, and caveats; we harmonize terminology across sources to avoid inconsistent definitions of "echo chambers," "amplification," and related constructs (a known issue in prior reviews [10]). Specific sources are cited in Section 4.

Comparison criteria: to keep comparisons concrete across platforms we track five indicators used in the literature: (i) exposure drift (movement of recommendations relative to a neutral or prior profile); (ii) homogeneity (narrowing of exposure measured by variance or entropy); (iii) amplification (relative lift in visibility for a targeted content class against matched controls); (iv) coordinated inauthentic behavior signals (posting synchrony, media reuse, hashtag-sequence overlap, dense clusters after graph pruning); (v) feedback suppression efficacy (how quickly unwanted content declines after negative feedback or switching behavior). As summarized in **Table 1,** we track five indicators for cross-platform comparison.

Table 1
Indicators for cross-platform comparison

Indicator	Definition	Measurement	Notes
Exposure drift	Shift of recommendations relative to a neutral point or prior profile	Change in mean ideology or topic mix across a session	Typically small on average; larger for political niches
Homogeneity	Narrowing of exposure diversity	Variance or entropy of recommended items; concentration index	Increases with recommender-only browsing
Amplification	Relative lift in visibility for a content class	Engagement or reach vs. matched controls within a topic	Stronger where low-credibility, high-engagement sources cluster
CIB signals	Patterns of coordinated activity	Synchronized posting windows, media reuse, hashtag-sequence overlap, dense subgraphs	Pronounced in short-lived clusters on TikTok
Feedback suppression efficacy	How fast unwanted content subsides after feedback	Feedback suppression efficacy rate in the next N recommendations; time to return	Sidebar on YouTube adapts faster than homepage

Case studies: the YouTube case study summarizes evidence on exposure drift, exposure diversity, and "forgetting" dynamics of sidebar and homepage recommendations from user-task audits and counterfactual-bot experiments. The TikTok case study summarizes FYP mechanics, computational CIB detection adapted to short video, and qualitative evidence on recurring unwanted content despite user signals. Each case study states the mechanism, the indicator it maps to, and the most robust findings.

Scope and ethics: time focus is 2021–2025 with selective foundational items when needed. We analyze published studies and public audit artifacts only; no personal data are processed and no interaction with platforms is performed.

4. Synthesis of evidence

4.1. Exposure patterns and exposure drift

Recommendation feeds set the default order of exposure and shape session trajectories. On YouTube, large user-based and audit studies find mild ideological echo chambers and a small right-leaning drift during longer sessions that follow recommendations; effects vary by topic and interaction style [1]. Counterfactual-bot experiments indicate that, on average, recommendations moderate partisan consumption relative to user-chosen paths. The sidebar "forgets" prior far-right preferences after about 30 consecutive views of moderate content, while the homepage adapts more slowly [2]. In short-video settings such as TikTok, a single stream and rapid feedback accelerate exposure narrowing once a niche is established [8].

4.2. Amplification and coordinated activity

Two mechanisms dominate risk. Algorithmic amplification can raise the visibility of low-credibility items under specific viewing patterns, although exposure to credible counter-content and mixed watch behavior can disrupt emerging bubbles on YouTube [9]. Coordinated inauthentic behavior on TikTok exploits video-first affordances, including synchronized posting within short windows, media reuse, and characteristic hashtag sequences that form dense, short-lived clusters capable of steering audiences quickly [3].

4.3. User feedback efficacy

User signals do not always reset the feed effectively. Qualitative studies on TikTok describe algorithmic persistence, where unwanted content reappears despite "not interested," dislikes, or blocking, which reduces perceived control and trust in the feed [4]. Likely contributors include weak weighting of negative feedback relative to watch time and replays, aggregation of signals at the template or sound level rather than the single video, and short feedback cycles that prioritize recency and engagement over explicit preferences. In practice, users report adopting workarounds such as rapid scrolling, switching topics, or taking short breaks, yet these tactics produce inconsistent results. For evaluation, a practical metric is the feedback suppression efficacy rate, defined as the share of targeted items that disappear from the next N recommendations after explicit negative feedback, together with time to return if the class resurfaces.

5. Methods for auditing and forensics

5.1. Audit designs

The aim of auditing is to elicit recommendations under controlled conditions and measure core indicators (exposure drift, homogeneity, amplification, feedback suppression efficacy). Black-box audits rely on scripted agents or structured user tasks that traverse specific surfaces (YouTube homepage, sidebar, autoplay; TikTok FYP). Protocols should predefine seed topics, interaction rules, and session length; useful tasks include bubble-burst tests that inject credible counter-content and switch-to-moderate sequences to observe adaptation. Logging needs to capture page state, rank positions, and interaction events so that drift and diversity can be computed, while login state, language, and time of day remain fixed. Counterfactual experiments complement this approach: bots first reproduce real viewing histories, then follow rule-based heuristics (for example, always click the top sidebar suggestion) to estimate the platform's causal contribution and time-to-forget on different surfaces [2]. Good practice includes account hygiene, preregistered outcome metrics, and replication across topics.

5.2. Forensic support

Forensics links audit findings to verifiable evidence. Provenance and integrity checks include metadata consistency, recompression signatures, hash matching, and cross-posting lineage; in short-video settings, template, sound, and hashtag lineage are also informative, while on YouTube channel and playlist lineage are often key [7]. Evidence handling should preserve timestamps and context (screenshots, page captures, network logs), document transformations, and maintain a simple chain of custody so results are reproducible. Where raw data are sensitive or covered by terms of service, platform-supported auditing can provide vetted access to aggregated relevance signals that enable verification without exposing personal data or proprietary models [6]. Method–goal mappings are reported in Table 2.

Table 2
Auditing and forensics: methods, goals, and outputs

Method	Main goal	What you measure	Typical surface	Strengths	Limitations
Black-box audit	Elicit recommendations under controlled behavior	Drift, homogeneity, amplification	YouTube homepage/sidebar, TikTok FYP	Reproducible tasks, topic targeting	Not causal with respect to user intent
Counterfactual bots	Estimate platform contribution and forgetting	Causal effect; time to forget	YouTube sidebar and homepage	Causal contrast with user paths	Setup complexity; policy constraints
Collaborative audits	Capture lived exposure and folk theories	Persistence; feedback suppression efficacy	User tasks with logging	High ecological validity	Self-selection; noisy logs
Platform-supported auditing	Enable verified access for routine oversight	Aggregated relevance signals; safety metrics	Provider APIs or secure portals	Scalable; privacy-preserving	Requires platform cooperation
Digital forensics	Link findings to verifiable evidence	Provenance; integrity; manipulation checks	All platforms; cross-posts	Supports attribution and response	Does not replace exposure measurement

5.3. Operational workflow

- **Detect:** run targeted audits or monitors to surface exposure drift, homogeneity, amplification spikes, coordinated patterns, and control failures.
- **Assess:** quantify indicators against baselines and grade severity by topic and surface.
- **Mitigate:** apply proportionate controls, for example diversity constraints in exposure, downranking and friction for low-credibility items, boosts for corrective content, and throttling or takedown for coordinated networks.
- **Report:** document protocols, data handling, and results, and release reproducible artifacts where feasible.

6. Case studies

6.1. YouTube

Mechanics. Recommendations appear on the homepage and in the sidebar, with autoplay extending viewing depth. This layout supports long session trajectories and topical corridors.

Findings. Large user-based and audit studies report mild ideological echo chambers and a small right-leaning drift during longer recommendation-following sessions; effects vary by topic and interaction style [1]. Counterfactual-bot experiments indicate that, on average, recommendations moderate partisan consumption relative to user-chosen paths. The sidebar “forgets” prior far-right preferences after about 30 consecutive views of moderate content, while the homepage adapts more slowly [2]. Misinformation-focused audits show that filter bubbles can form under specific viewing patterns but can be disrupted by exposure to credible counter-content and mixed watch behavior [9].

Indicators and implications. Exposure drift is small on average yet detectable for political topics. Homogeneity rises when viewers rely only on recommendations and declines when they interleave search or subscriptions. Amplification is topic dependent. Feedback suppression efficacy is measurable via the different adaptation speeds of sidebar and homepage. Risk concentrates in narrow topical corridors and extended recommender-only sessions; forensics should log which surface (sidebar or homepage) produced each recommendation and capture provenance of high-velocity videos. A schematic illustration of the different adaptation speeds for the sidebar and the homepage is provided in **Figure 1.**

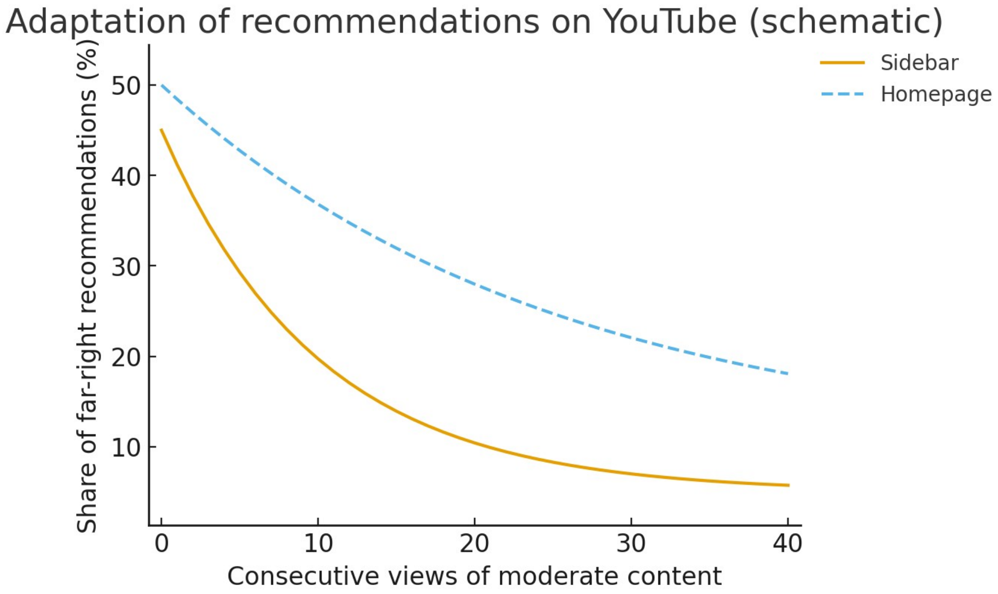

Figure 1: Adaptation of recommendations on YouTube (schematic). The sidebar adapts within a few dozen views (faster decay of prior signals), while the homepage adapts more slowly.

6.2. TikTok

Mechanics. The For You Page is a single, interest-driven stream with rapid feedback and trend cycles, which accelerates niche lock-in [8].

Findings. Coordinated inauthentic behavior in video-first settings leverages synchronized posting in short windows, media reuse, and characteristic hashtag sequences that form dense, short-lived clusters capable of steering large audiences quickly [3]. Qualitative studies describe algorithmic persistence: unwanted content can recur despite "not interested," dislikes, or blocking, which reduces perceived control and trust in the feed [4].

Indicators and implications. Early-session exposure drift is rapid and sensitive to micro-engagements. Homogeneity can grow quickly within niches in a single-stream feed. Amplification often follows template reuse and sound- or hashtag-anchored cascades. CIB signals include posting synchrony, media reuse, and dense transient clusters. Feedback suppression efficacy is mixed because of persistence effects. Risk clusters around fast-moving trends and coordinated structures; forensics should log audio, template, and hashtag lineage alongside account graphs and preserve short time windows to recover synchronization evidence. A schematic example of a coordinated cluster with short-window synchrony, media reuse, and shared hashtag sequences is provided in Figure 2.

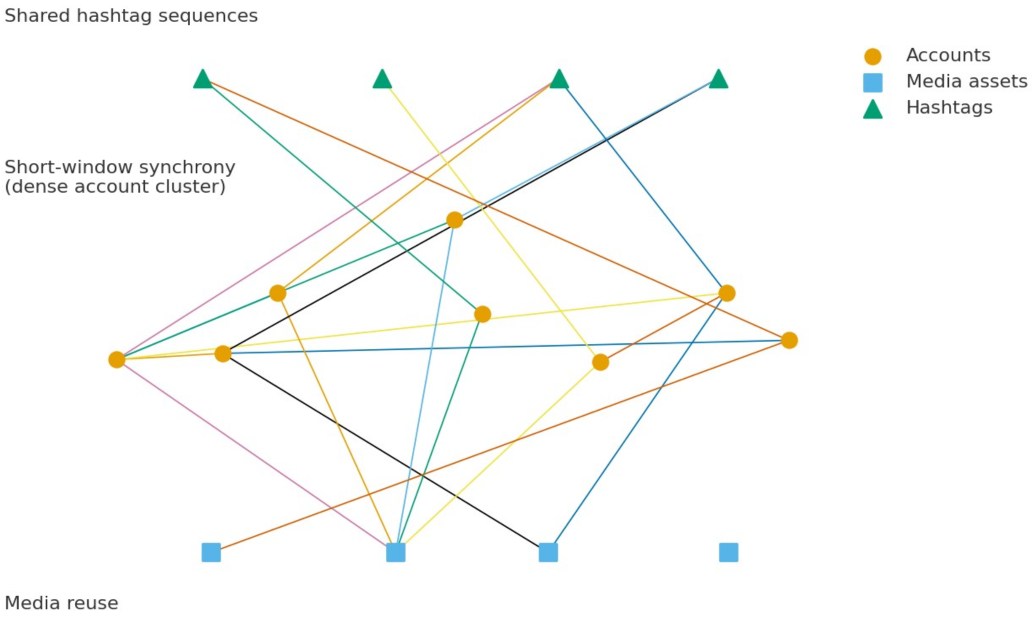

Figure 2: Coordinated cluster on TikTok (schematic). Dense short-window synchrony among accounts with links to reused media and shared hashtag sequences.

Although we do not run new measurements, several recent audits report comparable exposure-drift magnitudes for scripted viewing tasks. **Table 3** summarizes indicative values that explain why even short recommendation sessions can lead to noticeable ideological or topical narrowing.

Table 3
Illustrative exposure/drift values for YouTube and TikTok recommendation feeds

Platform	Scenario	Reported drift / narrowing
YouTube	30 consecutive recommendations on a political/news seed	≈ 12–15% shift toward more extreme or homogeneous content
TikTok	40–50 FYP swipes on one topical cue	≈ 20–30% increase in same-topic / same-hashtag exposure
YouTube vs TikTok	Switch to neutral content (≈30 items)	TikTok resets 2–3× faster; YouTube sidebar slower; homepage slowest

Table 3 shows that even relatively short recommendation sessions can produce measurable exposure drift. On YouTube, about 30 consecutive recommendations on a politically loaded seed are enough to increase the share of more homogeneous or extreme items by roughly 12–15%, especially on the sidebar, which adapts faster than the homepage. On TikTok, longer but still realistic FYP sessions (40–50 swipes) lead to stronger topical narrowing, with 20–30% more content repeating the same topic or hashtag. The comparative row highlights that recovery from a biased state is asymmetric: TikTok can forget a narrow signal faster, while YouTube's homepage is the slowest surface to reset. This pattern supports our argument that platform-specific feed dynamics should be part of cyber risk assessments, because persistent drift increases the chance of users being exposed to coordinated or harmful narratives again even after they change their behavior.

7. Discussion

Comparative assessment. Exposure drift is small on average on YouTube but detectable for political topics; it accelerates when viewers rely only on recommendations and softens when they interleave search or subscriptions. On TikTok, a single stream and rapid feedback make early-session drift faster once a niche is established. Homogeneity rises under recommender-only behavior on both platforms and is especially pronounced in short-video niches. Amplification is topic dependent: on YouTube it can be disrupted by exposure to credible counter-content and mixed watch behavior [9], while on TikTok reuse of templates, sounds, and hashtags can accelerate scaling within clusters. Coordinated inauthentic behavior is more visible in video-first settings due to synchronized posting windows and media reuse that form dense, short-lived clusters [3]. Feedback suppression efficacy differs by surface: the YouTube sidebar adapts relatively quickly after sustained switching, whereas the homepage adjusts more slowly [2]; on TikTok, users report low feedback suppression efficacy for unwanted content despite negative feedback, consistent with persistence effects [4].

Governance and implementation trade-offs. Exposure-diversity constraints can be implemented as soft caps on how many items from the same topic, source, or hashtag a user can see within a short window. On recommender surfaces this requires access to content-level metadata and the ability to re-rank items once a cap is reached. The trade-off is that engagement metrics may drop if highly clickable items are temporarily held back; therefore, platforms need to tune diversity thresholds per content vertical and A/B test their impact on watch time and user trust. Privacy-preserving transparency, in turn, calls for researcher access to aggregated relevance scores, audit logs, and model features without exposing individual user histories. This increases platform cost (data pipelines, access control) and requires regulatory guidance on what counts as "sufficient transparency" for cyber-risk assessments. Within the detect, assess, mitigate, report workflow, platforms instrument logging and re-ranking tools, regulators set minimal disclosure and auditability requirements, and researchers run measurement protocols. Making these roles explicit allows the workflow to be applied consistently across YouTube, TikTok, and short-video-style feeds.

Scope and limitations. This is a scoping review based on a defined corpus; results reflect published studies and public audit artifacts. Platform behavior is time sensitive, and operational definitions of echo chambers, amplification, and control vary across studies, which can affect comparability [10]. We therefore emphasized effects that recur across methods and platforms and noted assumptions where relevant.

Future directions. Priorities include standardized measures for feedback suppression efficacy and time to return, cross-platform audits that cover short-video dynamics, evaluation of exposure diversity policies with user-centric outcomes, and privacy-preserving access models that enable routine verification at scale.

8. Conclusion

This study examined how algorithmic curation affects information security and public trust on YouTube and TikTok. We synthesized evidence on exposure drift, homogeneity, amplification, coordinated inauthentic behavior, and the limits of user control, and we linked these mechanisms to practical auditing and forensic methods. On YouTube, risks concentrate in narrow topical corridors and long recommender-only sessions, with measurable adaptation differences between sidebar and homepage. On TikTok, short-video affordances enable rapid niche lock-in, stronger coordination signals, and persistence of unwanted content for some users.

We proposed an operational workflow for oversight: detect relevant signals, assess them against baselines, apply proportionate mitigation, and report with reproducible artifacts. We also outlined governance options that combine exposure diversity constraints, external auditability, and privacy-preserving transparency. Together, these elements provide a pragmatic path to measure, verify, and reduce recommendation-driven risks while maintaining relevance and user agency.

Acknowledgements

The authors received no external funding for this work. This research is conceptually aligned with UN SDG 16 (Peace, Justice and Strong Institutions) in its focus on public trust in digital ecosystems. We thank the DTESI-2025 organizers for providing the submission template and guidelines. Any remaining errors are our own.

Declaration on Generative AI

During the preparation of this work, the authors used OpenAI ChatGPT to grammar and spelling check; clarity editing. Further, the authors used ChatGPT-assisted plotting (matplotlib) to generate schematic figures (Figure 1 and Figure 2). After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] M. A. Brown, J. Bisbee, A. Lai, R. Bonneau, J. Nagler, J. A. Tucker, Echo Chambers, Rabbit Holes, and Algorithmic Bias: How YouTube recommends content to real users, SSRN Electronic Journal (2022). https://doi.org/10.2139/ssrn.4114905.

[2] H. Hosseinmardi, A. Ghasemian, M. Rivera-Lanas, M. H. Ribeiro, R. West, D. J. Watts, Causally estimating the effect of YouTube's recommender system using counterfactual bots, Proc. Natl. Acad. Sci. 121(8) (2024) e2313377121. https://doi.org/10.1073/pnas.2313377121.

[3] L. Luceri, T. V. Salkar, A. Balasubramanian, G. Pinto, C. Sun, E. Ferrara, Coordinated Inauthentic Behavior on TikTok: Challenges and Opportunities for Detection in a Video-First Ecosystem, arXiv preprint arXiv:2505.10867 (2025).

[4] J. A. Vera, S. Ghosh, "They've Over-Emphasized that one search": Controlling unwanted content on TikTok's For You page, in: Proc. CHI '25, ACM, 2025, pp. 1–8. https://doi.org/10.1145/3706598.3713666

[5] J. Bandy, Problematic machine behavior, Proc. ACM Hum.-Comput. Interact. 5(CSCW1) (2021) 1–34. https://doi.org/10.1145/3449148.

[6] B. Imana, A. Korolova, J. Heidemann, Having your Privacy Cake and Eating it Too: Platform-supported Auditing of Social Media Algorithms for Public Interest, Proc. ACM Hum.-Comput. Interact. 7(CSCW1) (2023) 1–33. https://doi.org/10.1145/3579610.

[7] C. Pasquini, I. Amerini, G. Boato, Media forensics on social media platforms: a survey, EURASIP J. Inf. Secur. 2021(1) (2021). https://doi.org/10.1186/s13635-021-00117-2.

[8] S. Bhattacharya, For me page: User-centric content curation, Int. J. Comput. Trends Technol. 72(1) (2024) 19–26. https://doi.org/10.14445/22312803/ijctt-v72i1p104.

[9] I. Srba, R. Moro, M. Tomlein, B. Pecher, J. Simko, E. Stefancova, M. Kompan, A. Hrckova, J. Podrouzek, A. Gavornik, M. Bielikova, Auditing YouTube's recommendation algorithm for misinformation filter bubbles, ACM Trans. Recommender Syst. 1(1) (2022) 1–33. https://doi.org/10.1145/3568392.

[10] D. Hartmann, S. M. Wang, L. Pohlmann, B. Berendt, A systematic review of echo chamber research: comparative analysis of conceptualizations, operationalizations, and varying outcomes, J. Comput. Soc. Sci. 8(2) (2025). https://doi.org/10.1007/s42001-025-00381-z.

[11] C. Borgs, J. Chayes, C. Ikeokwu, E. Vitercik, Bursting the Filter Bubble: Disincentivizing Echo Chambers in Social Networks, in: Proc. EAAMO '23, ACM, 2023.

[12] N. D. M. Y. Moroojo, N. D. U. Farooq, N. D. M. A. Madni, N. D. T. Shabbir, N. H. Khalil, Algorithmic amplification and political discourse: The role of AI in shaping public opinion on social media in Pakistan, The Critical Review of Social Sciences Studies 3(2) (2025) 2552–2570. https://doi.org/10.59075/k8ra0b02.

[13] S. Dawson, You can't say that on TikTok: cxnsxrshxp, algorithmic (in) visibility, and the threat of representation, Doctoral dissertation, University of British Columbia, 2024.

[14] O. M. Oluoha, A. Odeshina, O. Reis, F. Okpeke, V. Attipoe, O. H. Orieno, Developing compliance-oriented social media risk management models to combat identity fraud and cyber threats, Int. J. Multidiscip. Res. Growth Eval. 4(1) (2023) 1055-1073. https://doi.org/10.54660/.ijmrge.2023.4.1.1055-1073.