


Policy Brief

Series Information:
EPINOVA Policy Brief
Series on Strategic
Competition,
Geopolitics, and
Emerging Security

**Recommended
Citation:**
Wu, S. (2026). The
Machine Audience of
War: AI-Generated
Video, Narrative
Occupancy, and
Machine-Mediated
Information
Competition in the
U.S.–Israel–Iran
Conflict (Policy Brief
No. EPINOVA–2026–
PB–71). Global AI
Governance and Policy
Research Center,
EPINOVA LLC.
<https://doi.org/10.67037/epinova.pb.2026.071>

Disclaimer:
The analysis in this
publication is based on
publicly available
information, official
statements, media
reporting, and the
author’s independent
analysis. It does not
represent the official
position of any
government. This
publication is intended
solely for research and
policy discussion and
does not constitute
legal, military,
operational, or
sanctions compliance
advice.



GLOBAL AI
GOVERNANCE
RESEARCH CENTER



The Machine Audience of War:

AI-Generated Video, Narrative Occupancy, and Machine-Mediated Information Competition in the U.S.–Israel–Iran Conflict

Author: Shaoyuan Wu
Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC
Date: September 7, 2026

Core Judgment

AI-generated video in the U.S.–Israel–Iran conflict should be assessed not only by its ability to persuade or deceive human audiences. As search systems, retrieval pipelines, large language models, and AI agents increasingly mediate digital information, synthetic content may remain retrievable beyond its period of direct human attention. This brief introduces **narrative occupancy**, the persistent representation of a claim or association across machine-accessible information objects, and **machine addressability**, the degree to which such information can be discovered, parsed, and retrieved by automated systems. Together, these concepts extend information-warfare analysis from human attention toward machine-mediated information competition.

Key Judgments

- **Generative AI is lowering the cost and accelerating the production of wartime narratives.** Familiar objectives such as mobilization, deterrence, delegitimization, polarization, and political advantage remain, but events, relationships, symbolic coalitions, and imagined battlefield outcomes can now be constructed visually with greater speed and scale.
- **The U.S.–Israel–Iran conflict illustrates several overlapping uses of synthetic political video.** Iranian state-linked dissemination illustrates synthetic mobilization; President Donald Trump’s Kharg Island post illustrates synthetic battlefield representation; and Prime Minister Benjamin Netanyahu’s campaign video illustrates relational fabrication. These are analytical lenses rather than a comprehensive taxonomy.
- **Human attention is an incomplete measure of information availability.** Large-scale automated crawling and retrieval are already features of the contemporary web, but machine accessibility should not be confused with retrieval, acceptance, or influence (Weiss et al., 2026). Cloudflare reports that more than half of the traffic it observes is now non-human, while 52 percent of identified crawler requests are associated with AI training and more than 36 percent with mixed-use crawlers.
- **The policy problem therefore extends beyond deepfake detection.** The central integrity challenge is whether machine-mediated systems preserve distinctions among an event, a depiction, an allegation, a denial, and independently verified evidence.

Policy Brief

Executive Summary

Generative AI has become unusually visible in the information environment surrounding the 2026 U.S.–Israel–Iran conflict. Brookings found that more than 5,000 X Community Notes had referenced AI-generated content since the conflict began, while emphasizing that this is a proxy for contested material rather than a direct measure of how much synthetic content actually circulated. CSIS separately documented Iranian-linked use of AI-generated imagery, misleading battlefield claims, deepfakes, and coordinated narratives within the conflict’s information environment (Jensen et al., 2026; Wirtschafter, 2026).

These developments reinforce the established concern that generative AI expands the production capacity of information operations. They also raise a different question. Search systems, retrieval tools, language models, and AI agents increasingly mediate information before it reaches an end user. The relevant distinction is therefore not simply between true and false content, but between human attention and machine accessibility.

This brief develops narrative occupancy to describe the persistent representation of a claim or association across machine-accessible information objects. Three cases illustrate different forms of synthetic political representation: an Iranian video compressed historically separate grievances into a common retaliatory frame; Trump’s Kharg Island post visually represented an attack for which Reuters found no contemporaneous supporting evidence; and Netanyahu’s AI-generated campaign video placed politically distinct actors into a fictitious common interaction (Arguello, 2026; Stewart & Lawder, 2026; ToI Staff & Tress, 2026).

The cases do not demonstrate that these materials were intentionally designed for AI systems or that machine retrieval altered downstream decisions. They support a narrower policy question: whether synthetic political claims can remain accessible through machine-mediated environments after the original artifact loses direct visibility, and whether automated systems preserve their evidentiary status when retrieving and summarizing them.

Why This Matters

Information warfare has traditionally been assessed through human exposure, persuasion, and behavior. Machine-mediated information systems complicate that model because content can now be crawled, indexed, retrieved, and summarized without a user visiting the original source.

Cloudflare reports that automated consumption has become a large component of contemporary web traffic and that content is increasingly crawled and used without corresponding referral traffic to originating sites. These data describe the broader architecture of machine consumption rather than political influence, but they establish the environment within which machine-mediated information competition can occur (Weiss et al., 2026).

The policy question is therefore not only whether synthetic content persuades people. It is whether automated systems preserve the difference between what happened, what was claimed, what was depicted, and what was generated.

Policy Brief

Analytic Note and Case Selection

This brief is an exploratory, mechanism-oriented policy assessment. Machine audience is used as analytical shorthand for heterogeneous automated systems, including crawlers, search indexes, retrieval systems, language models, recommendation systems, and agents, that perform different functions and should not be treated as a single technical category.

The three cases were purposively selected because they illustrate different uses of synthetic representation within the broader wartime political-information environment. They are illustrative rather than representative, sequential, or mutually exclusive. The Iranian case is based on documented circulation by Iranian state broadcaster IRIB and an IRGC-affiliated Telegram channel; this establishes dissemination but does not independently establish who produced the video. The Trump case separates the original social-media claim, the contemporaneous absence of corroboration, a subsequent U.S. official statement that Kharg Island had not been targeted, and Vice President JD Vance's later description of the post as a message to Iran. The Netanyahu case concerns domestic electoral communication that incorporates actors and imagery from the ongoing conflict and is included as an example of wartime political communication rather than a direct military information operation.

Source publication dates are retained as reported. The analytical cutoff for this brief is September 6, 2026, 12:00 ET. Where EPINOVA analytical days are used, they follow U.S. Eastern Time from 12:00 ET to 12:00 ET the following day. Publication dates are not automatically converted into event dates when a precise event time is unavailable.

1. From Deepfakes to Synthetic Political Representation

Research on deepfakes has examined authenticity, deception, uncertainty, trust, social context, and the political consequences of manipulated audiovisual evidence (Chesney & Citron, 2019; Paris & Donovan, 2019; Vaccari & Chadwick, 2020). The contribution of this brief is therefore not to replace that literature, but to extend the problem toward machine-mediated circulation and derivative information objects.

The U.S.–Israel–Iran conflict illustrates why that extension matters. Brookings found a sharp increase in X Community Notes referencing AI-generated content after the conflict began, while cautioning that the measure captures contested material rather than proving that every flagged item was generated by AI. CSIS documented a wider environment in which AI-generated imagery, deepfakes, false claims of battlefield success, and coordinated political narratives circulated alongside conventional information operations (Jensen et al., 2026; Wirtschafter, 2026).

Not all strategically relevant synthetic media, however, depends on factual deception. Netanyahu's September 4 campaign video was labeled as AI-generated and depicted New York City Mayor Zohran Mamdani, Turkish President Recep Tayyip Erdoğan, Iranian Supreme Leader Mojtaba Khamenei, and Palestinian Authority President Mahmoud Abbas participating in a fictitious conversation about Netanyahu's electoral defeat. It then displayed imagined election results showing Netanyahu winning, after which the depicted participants canceled their celebration (ToI Staff & Tress, 2026).

Because the video was explicitly labeled as AI-generated, its representational mechanism did not depend on presenting the conversation as authentic. It instead converted a political association into a simulated interaction.

Policy Brief

This brief describes that mechanism as relational fabrication: **the synthetic depiction of distinct actors as participating in a shared interaction or coordinated relationship that is not established by independent evidence.**

The distinction matters because demonstrating that the interaction is fictitious does not necessarily remove the association created by the representation. This does not establish a measurable political effect; it identifies a mechanism through which synthetic political media can communicate association even when its artificial character is openly acknowledged.

2. Three Uses of AI-Generated Video in the Conflict

The cases examined here illustrate three overlapping analytical functions: **synthetic mobilization, synthetic battlefield representation, and relational fabrication.** They are not a unified taxonomy, and a single synthetic artifact may perform several functions simultaneously.

Table 1. Three Analytical Uses of AI-Generated Video in the U.S.–Israel–Iran Conflict

Case	Synthetic Representation	Analytical Lens	Primary Compression
Iranian state-linked video ending with “One vengeance for all”	Historical grievances and contemporary conflicts linked to Iranian retaliation	Synthetic mobilization	Time, geography, grievance, identity
Trump Kharg Island video	Destruction of a strategic Iranian energy hub without contemporaneous corroboration of such an attack	Synthetic battlefield representation	Capability, intention, event, effect
Netanyahu campaign video	Politically distinct actors depicted within a common anti-Netanyahu interaction	Relational fabrication	Actors, motivations, alignment

Note: The categories are analytical rather than mutually exclusive. “Compression” refers to the reduction of temporal, relational, or political complexity into a simplified visual representation. The categories identify representational functions, not demonstrated audience effects.

Sources: Author’s analysis based on Arguello (2026), Reuters (2026), Stewart and Lawder (2026), and ToI Staff and Tress (2026).

2.1 Synthetic Mobilization

In March 2026, Iranian state broadcaster IRIB and an IRGC-affiliated Telegram channel circulated an AI-generated video that ended with the slogan “One vengeance for all.” The video connected Native American dispossession, Hiroshima and Nagasaki, Vietnam, recent Middle Eastern conflicts, Qassem Soleimani, Ali Khamenei, and other symbols before depicting an Iranian missile striking the Statue of Liberty (Arguello, 2026). The available reporting establishes dissemination through state-linked channels but does not independently establish the production chain.

Its analytical significance lies in the compression of geographically and historically separate grievances into a common retaliatory frame. This brief describes the mechanism as synthetic narrative compression: the use of generative media to place otherwise distinct events, actors, and grievances within a simplified common representation.

Policy Brief

The resulting artifact creates symbolic alignment rather than evidence of an actual political coalition. This is consistent with a broader systemic-warfare framework in which operational and political effects are shaped through interpretation and information competition rather than produced automatically by material events (Wu, 2026c).

2.2 Synthetic Battlefield Representation

Trump's Kharg Island post illustrates a different mechanism. On August 31, Reuters reported that Trump had posted a video depicting Kharg Island under attack and assessed the clip as most likely synthetically generated. Reuters found no other contemporaneous evidence that the Iranian energy hub was under attack (Stewart & Lawder, 2026).

A separate Reuters report subsequently stated that a U.S. official said the military had not targeted Kharg Island and quoted Vance describing Trump's post as a message to Iran rather than an announcement of military action (Reuters, 2026).

One evidence-generating pathway begins with a physical event and moves through observation and verification toward interpretation. Synthetic representation permits a different sequence in which a strategic claim is first given visual form and then circulated for interpretation.

The second pathway does not establish that viewers accepted the representation as fact. Its significance lies in the ability of generative media to visualize a threatened, imagined, or potential outcome without waiting for that outcome to produce physical evidence.

This extends earlier work on battlefield observability, in which recoverable artifacts and independently verified imagery can be transformed through documentation, validation, and amplification into informational and political leverage (Wu, 2026d). Synthetic media creates a different informational object whose evidentiary status must remain distinct from physical verification.

2.3 Relational Fabrication

Netanyahu's September 4 video shifts the object of synthetic representation from battlefield outcomes to political relationships. It places Mamdani, Erdoğan, Khamenei, and Abbas within a fictitious shared conversation and then depicts them reacting negatively to an imagined Netanyahu election victory (ToI Staff & Tress, 2026).

The distinction is important. The video establishes neither common coordination nor identical political interests among the figures depicted. Instead, it converts an existing campaign argument about common opposition to Netanyahu into a synthetic social scene.

The same broad message had appeared earlier in Likud campaign material. The Times of Israel reported that a previous campaign billboard used a similar grouping, although Abbas was replaced by Hezbollah leader Naim Qassem (ToI Staff & Tress, 2026).

Generative AI therefore need not create a new political proposition. It can make an existing association more concrete by converting an abstract political claim into a simulated interaction.

Policy Brief

3. Narrative Occupancy

Conventional measures such as views, shares, engagement, and sentiment capture forms of human exposure but do not fully describe how long a claim remains represented across digital information objects.

This brief defines **narrative occupancy** as: **the persistent representation of a specific claim, interpretation, or actor relationship across machine-accessible information objects over time.**

The concept is distinct from persuasion and narrative dominance. A claim can remain retrievable while being rejected, corrected, mocked, or accompanied by competing interpretations.

To avoid reducing occupancy to simple online existence, three observational dimensions are important. **Temporal persistence** asks whether the proposition remains retrievable across multiple observation points. **Source lineage** distinguishes independent sources from derivative reporting, reposts, or summaries originating from the same artifact. **Stance and epistemic status** distinguish assertion, quotation, denial, satire, correction, and verification.

This last distinction is essential. “A and B coordinated an action” and “there is no evidence that A and B coordinated an action” contain similar entities but do not express the same proposition. Simple entity co-occurrence should therefore not be treated as narrative occupancy of the same claim.

Synthetic content can nevertheless generate derivative information objects, including reporting, screenshots, transcripts, commentary, fact-checks, and machine-generated summaries, that preserve some elements of the originating representation. Contemporary research on manipulated audiovisual content already emphasizes that meaning is shaped through circulation and social context rather than by the original artifact alone (Paris & Donovan, 2019; Vaccari & Chadwick, 2020).

Correction can therefore produce different outcomes. If systems preserve negation, provenance, and source dependence, additional corrective material may improve epistemic accuracy. If those qualifiers are lost during retrieval or summarization, derivative references may preserve an erroneous association or create the appearance of multiple corroborating sources.

Narrative occupancy is therefore an observable property of information persistence, not evidence of persuasion or propaganda success.

4. Machine Addressability

This brief defines **machine addressability** as **the degree to which an information object can be discovered, parsed, and retrieved by automated systems.** Attribution accuracy, preservation of evidentiary status, incorporation into model output, and downstream action are separate questions.

Accessibility does not guarantee indexing; indexing does not guarantee retrieval; and retrieval does not imply acceptance or action.

Policy Brief

This boundary matters because large-scale crawling does not establish influence. Cloudflare’s 2026 data demonstrate substantial automated consumption of web content, but crawler requests do not show whether a particular item entered a model, appeared in retrieval, affected an answer, or shaped subsequent behavior (Weiss et al., 2026). Cloudflare also distinguishes AI-training crawlers from mixed-use systems combining search, agent use, and training, reinforcing the need to separate machine-access stages rather than treat them as a single process.

The broader expansion of large-scale data aggregation, indexing, and automated processing may further reinforce these conditions by increasing the number of machine-accessible representations and the pathways through which information can be rediscovered across sources. In this sense, contemporary data infrastructures can strengthen the conditions for narrative occupancy without themselves establishing influence.

The concept also differs from Generative Engine Optimization (GEO), which concerns deliberate efforts to improve the visibility of content in generative-engine responses (Aggarwal et al., 2023). GEO formalizes optimization strategies for increasing representation in generative search outputs.

Machine addressability is descriptive rather than optimization-oriented: it concerns whether information is technically and semantically available for machine discovery and retrieval, regardless of whether anyone has attempted to optimize it.

Deliberate optimization of wartime synthetic media for machine retrieval is not established in the cases examined here, but it represents a plausible emerging incentive.

5. The Machine Audience

The term **machine audience** describes the growing role of non-human systems as intermediate consumers and processors of political information. It does not imply that crawlers, search indexes, recommendation systems, retrieval-augmented generation systems, language models, and autonomous agents perform the same function or interpret information as humans do.

Traditional communication moves from actor to message to human audience. Platform-mediated communication inserts ranking and recommendation systems. Machine-mediated environments add further stages of retrieval, classification, synthesis, and, in some cases, automated action before information reaches a human user or another system.

CSIS has examined how AI agents can support information operations by identifying sources, gathering and revisiting information, detecting changes, summarizing material, and improving the timing and targeting of information activity (Yayboke, 2025).

The broader implication runs in both directions. Machines can function as tools used by information operators, but they can also operate as intermediaries processing information produced by other actors.

Machine audience is therefore best understood as analytical shorthand for non-human information consumption and mediation rather than as a single technical category.

Figure 1 summarizes how machine addressability, human attention, machine-mediated processing, and epistemic-status preservation interact within the information environment.

Policy Brief




Figure 1. Human and Machine Pathways of Synthetic Information Persistence

Note: Machine addressability does not imply retrieval, acceptance, or downstream influence. The figure identifies a conditional pathway through which synthetic information may remain retrievable after direct human engagement declines. Epistemic compression occurs when provenance, qualification, or source dependence is lost during subsequent processing.

Sources: Author’s conceptual synthesis based on Paris and Donovan (2019), Vaccari and Chadwick (2020), Yayboke (2025), Weiss et al. (2026), Wu (2026a), and Wu (2026d).

6. Human Disengagement and Machine Persistence

A further possibility is that human engagement with synthetic content may diverge from its technical availability to machines.

Research on synthetic political media shows that manipulated content can generate uncertainty and reduce trust as well as deceive, but current evidence does not establish a broad population-level aversion to AI-generated political video (Paris & Donovan, 2019; Vaccari & Chadwick, 2020).

The relevant proposition is therefore conditional: **declining human engagement does not necessarily produce declining machine retrievability.**

A user may quickly dismiss an artificial-looking video, while automated systems may continue to encounter its description, transcript, screenshots, metadata, or secondary reporting. The original visual artifact may therefore lose direct human attention while derivative information objects remain available for search or retrieval.

The relevant possibility is a divergence between human attention and machine accessibility. Whether continued retrievability produces meaningful downstream consequences remains to be established.

Policy Brief

7. From Observability Competition to Synthetic Presence

Earlier analysis of the U.S.–Israel–Iran conflict identified growing competition over evidence, attribution, and observability as the conflict evolved. By Days 44–50, evidence production, attribution clarity, and narrative competition had become increasingly prominent within the conflict environment (Wu, 2026a).

Related analysis showed how recoverable battlefield artifacts can support documentation, attribution, validation, and subsequent informational leverage (Wu, 2026d).

A material evidence pathway begins with a physical event and proceeds through observation, recoverable evidence, verification, and interpretation. Synthetic media creates a different pathway in which a narrative intention can first be converted into an informational artifact and then circulated for interpretation without requiring a corresponding physical event.

The two pathways must remain evidentially distinct. Synthetic artifacts do not verify the events they depict; corroborated imagery, recoverable evidence, technical data, and reliable independent reporting retain different evidentiary status.

The risk emerges when that status is not preserved during derivative processing. A source may accurately report that a political leader posted an AI-generated depiction, while a later summary retains the depicted event but drops the qualification. Multiple articles may also derive from the same original artifact and therefore should not automatically be treated as independent corroboration.

This brief calls such a failure epistemic compression: **the loss of distinctions among event, claim, depiction, simulation, allegation, denial, and verification during repeated representation or summarization.**

The risk is conditional. Where provenance, negation, and evidentiary status are preserved, derivative reporting may instead improve accuracy. The policy question is therefore whether machine-mediated systems can reliably maintain those distinctions.

8. Implications for MCCM and Conflict Monitoring

The Multi-Domain Conflict and Coupling Model (MCCM) v2.3.4 structures conflict monitoring through 35 mechanisms across ten analytical layers, including Narrative Warfare. Previous MCCM analysis identified narrative polarization and information amplification among the largest persistent risk-oriented contributors across the first 159 analytical days of the conflict (Wu, 2026b).

Generative AI does not create the distinction between narrative production and amplification; propaganda has always been capable of manufacturing, staging, or selectively presenting representations. What generative AI changes is the cost, speed, and scale at which politically usable pseudo-events, imagined outcomes, and synthetic relationships can be produced independently of observed events.

The current evidence does not justify adding another quantified MCCM indicator. A more disciplined near-term approach is to treat synthetic-content production as an observed input to existing narrative-amplification and information-distortion mechanisms while maintaining a separate provenance and retrieval record.

Policy Brief

Table 2. Minimal Monitoring Record for Synthetic-Content Tracking

Record Field	Analytical Purpose
Original artifact and publication date	Fix the object of analysis
Claim, depiction, denial, or verification status	Preserve epistemic qualification
Original source and derivative lineage	Avoid duplicate corroboration
Retrieval date, system, and query	Support reproducibility
Qualification preserved in output?	Detect epistemic-status loss

This approach preserves MCCM model stability while creating an empirical lane through which narrative occupancy and machine addressability can be monitored and, if repeated observations justify it, later formalized.

9. Policy Implications

- **Monitor machine-mediated visibility.** Information-environment monitoring should complement measures of reach, engagement, and sentiment with limited retrieval-oriented testing. The purpose is not to measure “machine belief,” but to determine whether automated systems preserve source attribution, evidentiary qualification, correction, and source dependence when handling contested claims.
- **Preserve epistemic status.** Terms such as claimed, depicted, AI-generated, denied, and independently verified are substantive information. A system that converts “an AI-generated video depicted Kharg Island under attack” into “Kharg Island was attacked” has changed the evidentiary status of the underlying claim.
- **Extend provenance through derivative circulation.** The Coalition for Content Provenance and Authenticity (C2PA) provides standards for recording assertions about digital-content provenance within digitally signed and tamper-evident claims. Its specifications allow provenance assertions to describe creation, transformation, ingredients, and algorithmically generated media, but explicitly do not treat validation of provenance assertions as a judgment that the represented factual content is true (C2PA, 2026). The harder problem arises when content is transformed into screenshots, transcripts, quotations, news reports, or machine-generated summaries, where the relationship to the originating asset and its provenance information may weaken or disappear.
- **Test source dependence under narrative saturation.** Crisis-oriented AI evaluation should examine whether systems recognize when numerous derivative references originate from one synthetic artifact rather than treating repetition as independent corroboration. During conflict, information volume can grow much faster than independently verified evidence.

Policy Brief

10. Limitations

This assessment relies on three purposively selected cases. They illustrate different forms of synthetic political representation but cannot establish the prevalence of those mechanisms across the U.S.–Israel–Iran conflict or contemporary warfare more broadly. Nor do the cases demonstrate agent-directed intent or actual machine retrieval. The changing information architecture makes such pathways plausible; it does not demonstrate their operation in these specific cases.

The core concepts remain exploratory. Narrative occupancy is not yet a validated quantitative measure, while machine addressability does not establish machine influence. Future observation would need to distinguish temporal persistence, independent source lineages, stance, retrieval frequency, preservation of qualification, and downstream use. Cloudflare’s data provide background on automated web activity rather than evidence that any particular wartime narrative shaped AI outputs (Weiss et al., 2026).

Finally, neither broad human disengagement from synthetic media nor downstream political effects are established here. The individual videos may support framing, mobilization, or signaling without measurably changing public attitudes, electoral behavior, or strategic decisions. The brief therefore identifies mechanisms and testable risks, not demonstrated causal effects.

Conclusion

Generative AI is becoming embedded in wartime political communication, but its significance extends beyond the familiar problem of convincing fakes. The U.S.–Israel–Iran conflict shows synthetic media being used to compress grievances, visualize unverified battlefield outcomes, and simulate political relationships.

These artifacts circulate within an information environment increasingly mediated by search, retrieval, language models, and agents. Human attention remains central, but direct consumption is no longer the only route through which an information object can remain technically accessible.

Narrative occupancy identifies the persistent representation of claims and relationships across machine-accessible information objects. Machine addressability identifies whether those objects can be discovered, parsed, and retrieved. Neither establishes influence, nor do the cases demonstrate deliberate machine-targeted propaganda.

Together, they identify a testable possibility: synthetic wartime claims may remain retrievable after direct human attention declines, while their downstream consequences remain to be established.

The policy challenge is therefore not only to detect synthetic media. It is to preserve, across increasingly machine-mediated information systems, the distinction between what occurred, what was claimed, what was depicted, and what was generated.

Policy Brief

References

- Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2023). *GEO: Generative engine optimization*. arXiv. <https://doi.org/10.48550/arXiv.2311.09735>
- Arguello, A. V. (2026, March 25). *Iran posts AI video showing missile striking Statue of Liberty*. *The Algemeiner*. <https://www.algemeiner.com/2026/03/25/iran-posts-ai-video-showing-missile-striking-statue-liberty/>
- Chesney, R., & Citron, D. K. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6), 1753–1820. <https://www.californialawreview.org/print/deep-fakes-a-looming-challenge-for-privacy-democracy-and-national-security>
- Coalition for Content Provenance and Authenticity. (2026). *C2PA specifications, version 2.4*. https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html
- Jensen, B., Vacca, N., & Macias, J. M., III. (2026, May 5). *How to lose an information war in 10 days*. Center for Strategic and International Studies. <https://www.csis.org/analysis/how-lose-information-war-10-days>
- Paris, B., & Donovan, J. (2019, September 18). *Deepfakes and cheap fakes: The manipulation of audio and visual evidence*. Data & Society Research Institute. <https://datasociety.net/research-library/deepfakes-and-cheap-fakes-the-manipulation-of-audio-and-visual-evidence/>
- Reuters. (2026, August 31). *Trump was sending message to Iran with Kharg Island post, says Vance*. <https://www.reuters.com/world/middle-east/trump-was-sending-message-iran-with-kharg-island-post-says-vance-2026-08-31/>
- Stewart, P., & Lawder, D. (2026, August 31). *Trump says Iran's Kharg is being "blown to smithereens," but gives no details*. Reuters. <https://www.reuters.com/world/middle-east/us-iran-exchange-fire-flare-up-bessent-signals-more-sanctions-2026-08-31/>
- ToI Staff, & Tress, L. (2026, September 5). *In AI-made Likud ad, Mamdani, Khamenei and Erdogan prematurely celebrate Netanyahu's defeat*. *The Times of Israel*. <https://www.timesofisrael.com/in-ai-made-likud-ad-mamdani-khamenei-and-erdogan-celebrate-netanyahus-expected-defeat/>
- Vaccari, C., & Chadwick, A. (2020). Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society*, 6(1). <https://doi.org/10.1177/2056305120903408>
- Weiss, A., Albertson, Z., & Lanfear, E. (2026, July 1). *Content Independence Day, one year on: Building the business model for the agentic Internet*. Cloudflare. <https://blog.cloudflare.com/agentic-internet-bot-report/>
- Wirtschafter, V. (2026, April 8). *Generative AI as a weapon of war in Iran*. Brookings Institution. <https://www.brookings.edu/articles/generative-ai-as-a-weapon-of-war-in-iran/>
- Wu, S. (2026a). *Escalation without collapse: High-pressure systemic equilibrium in the U.S.–Israel–Iran conflict, Days 1–50* (Policy Brief No. EPINOVA–2026–PB–35). Global AI Governance and Policy Research Center, EPINOVA LLC. <https://doi.org/10.5281/zenodo.19645873>
- Wu, S. (2026b). *The evolving structure of the U.S.–Iran–Israel conflict: MCEA and NMF analysis of analytical days 1–159 using MCCM v2.3.4* (Policy Brief No. EPINOVA–2026–PB–66). Global AI Governance and Policy Research Center, EPINOVA LLC. <https://doi.org/10.67037/epinova.pb.2026.066>
- Wu, S. (2026c). *Systemic warfare in the networked age: Operational systems, information competition, and cumulative pressure* (EPINOVA Working Paper No. EPINOVA–WP–F–2026–07). Global AI Governance and Policy Research Center, EPINOVA LLC. <https://doi.org/10.5281/zenodo.19078936>
- Wu, S. (2026d). *What cannot be recovered cannot be leveraged: Debris, evidence, and power in the Iran battlefield* (Policy Brief No. EPINOVA–2026–PB–24). Global AI Governance and Policy Research Center, EPINOVA LLC. <https://doi.org/10.5281/zenodo.19432715>
- Yayboke, E. (2025, September 19). *Channeling Augustus: On agentic offensive information operations*. Center for Strategic and International Studies. <https://www.csis.org/analysis/channeling-augustus-agentic-offensive-information-operations>