NEXUS-HC: A Framework for Technology Foresight in Strategic Communications

Yukai Zeng
Singapore Ministry of Defence Communications Organisation (MCO), Singapore

zeng_yukai@defence.gov.sg

Abstract: Strategic communications teams increasingly operate in an information environment shaped less by what people say to each other and more by what platforms choose to show, boost, or hide. Automation, synthetic media, and faster technology cycles mean influence activity can scale quickly, while decision-makers must still choose what to monitor, test, and prepare for. Many organisations track technology "hype" informally through news and vendor claims, but this often fails to produce clear priorities or explainable trade-offs. This paper introduces NEXUS-HC (Navigational Examination and Understanding System for Hype Cycle Analysis), a decision-oriented framework that adapts the familiar hype-cycle idea for contested, next-generation information environments. NEXUS-HC standardises technology assessment through a lightweight cross-impact rubric to compare likely effects, and links assessment to action through five portfolio postures: Monitor, Pilot, Adopt, Govern/Harden, or Counter. The framework separates (a) maturity and diffusion signals, (b) disruption mechanisms relevant to influence and integrity, and (c) adversarial leverage and governance readiness. We provide practical guidance for running the workflow in a small StratCom team, recording evidence notes and confidence, and setting reassessment triggers. A short three-technology illustration and light-touch retrospective validation—synthetic media, algorithmic Information-Operations (IO), and AI avatars—demonstrate how technologies associated with synthetic saturation, algorithmic contestation, and reality mediation can still lead to different actions depending on leverage and governance. The contribution is a repeatable approach that improves transparency (how judgements were made), comparability (across technologies and time), and briefing value (clearer decisions for leaders). NEXUS-HC is designed to complement, not replace, deeper intelligence assessments and technical evaluation. It is intended primarily for StratCom teams, with cyber or technical specialists contributing when required. It provides a common language for briefings, helps teams avoid over-reacting to short-lived attention spikes, and highlights when governance should start early because standards, labelling practices, and organisational controls often lag behind capability diffusion. When applied on a quarterly cadence, NEXUS-HC produces an auditable decision record that captures evidence, confidence, and triggers for reassessment—supporting learning across cycles and reducing reliance on individual judgement.

Keywords: Strategic communications, Horizon scanning; Technology assessment, Synthetic media, Algorithmic IO, AI avatars; Decision support

1. Introduction

Strategic communications (StratCom) has traditionally focused on crafting messages for human audiences. That remains important, but it is no longer sufficient. Today, what people see is strongly shaped by ranking systems, recommendation engines and automated accounts that amplify or suppress content at scale (Bolt and Lange-Ionatamishvili, 2026).

This shift changes the operational question for StratCom leaders (Chairman of the Joint Chiefs of Staff, 2022). It is not only "what story should we tell?" but also "what system will deliver, distort or drown out that story?". New capabilities such as synthetic media and algorithmically-optimised influence techniques can reduce the cost of manipulation, increase speed and widen reach, while also increasing ambiguity and undermining trust (Wardle and Derakhshan, 2017; Tabassi, 2023).

Technology anticipation is often informal: teams track trends through headlines or vendor pitches, but struggle to compare technologies, explain trade-offs and translate scanning into decisions. NEXUS-HC addresses this by combining maturity signals with disruption mechanisms and governance readiness, producing an action posture that can be briefed and revisited (UK Ministry of Defence, 2021).

2. Background

Hype-cycle models are widely used because they offer an intuitive vocabulary for expectation dynamics. However, stage placement can be opaque and different technologies follow different trajectories, which limits their use for high-stakes decisions (Dedehayir and Steinert, 2016).

Readiness frameworks such as Technology Readiness Levels (TRLs) provide disciplined evidence gates for engineering and acquisition, but they were not designed to capture adversarial leverage, platform incentives or the institutional burden of verification in contested information environments (U.S. Department of Defense, 2005; Defense Acquisition University, n.d.; Mankins, 2009; U.S. Government Accountability Office, 2020).

For StratCom, the practical need is a repeatable way to (a) track maturity and diffusion (adoption and spread) signals, (b) translate technical affordances into mission-relevant disruption, and (c) decide whether to Monitor, Pilot, Adopt, Govern/Harden, or Counter a technology. Across current approaches, the persistent weakness is prioritisation: many methods detect and describe emerging technologies well, but still struggle to convert noisy signals into defensible, decision-relevant recommendations that make trade-offs explicit and can be revisited over time (Australian DSTG, 2016).

3. The NEXUS-HC Framework

Here, the Navigational Examination and Understanding System for Hype Cycle Analysis (NEXUS-HC) is introduced. NEXUS-HC is designed as a navigational decision-support framework that helps organisations prioritise emerging technologies under conditions of uncertainty and adversarial pressure. It adapts the hype-cycle idea into a transparent decision-support rubric for StratCom and information-warfare organisations operating in contested, platform-shaped information environments. Rather than predict exact futures, it structures how a small team can (a) translate an environment scan into comparable disruption dimensions, (b) stage technologies using observable "maturity and diffusion (adoption/spread)" signals, and (c) derive repeatable portfolio actions (Monitor, Pilot, Adopt, Govern/Harden, Counter) supported by evidence notes, confidence, and reassessment triggers (Gordon and Hayward, 1968). Figure 1 provides an overview of NEXUS-HC. The intended primary user is a small StratCom team rather than a full cyber threat-intelligence cell.

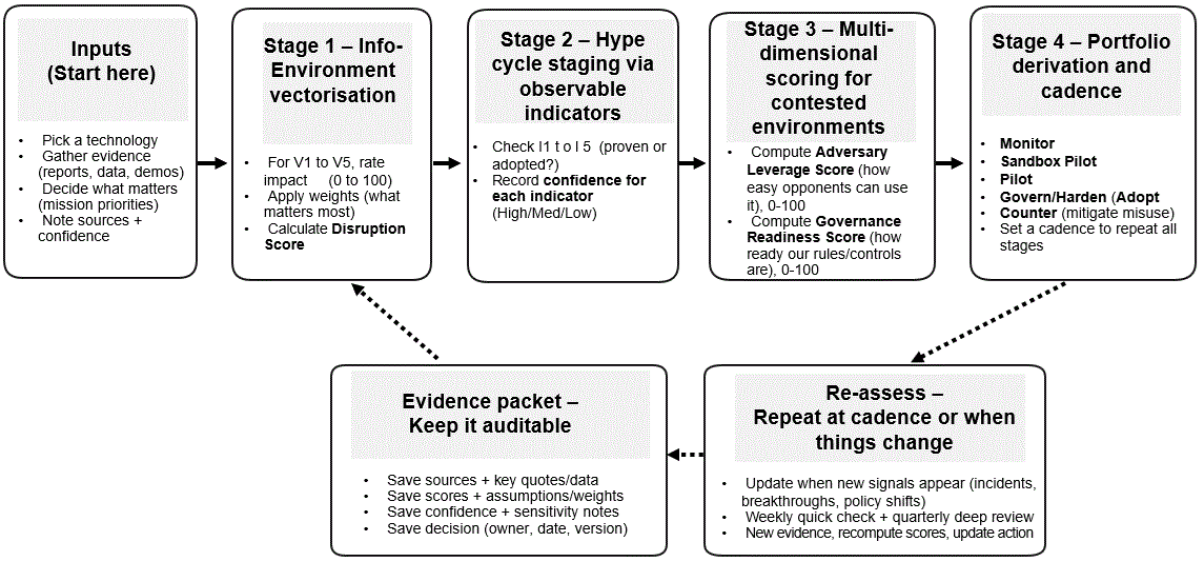

Figure 1: NEXUS-HC framework overview

Stage 1 — Information-environment vectorisation. NEXUS-HC begins with an explicit environmental baseline (here, NATO's Next-Gen Information Environment framing), but the same workflow can be re-run using other baselines (e.g., national strategies, regional architectures, or threat-intelligence scans) to compare how priorities change under different assumptions. The baseline is operationalised into five vectors (Table 1) that represent distinct ways technologies can reshape information dynamics.

Table 1: Next Gen information environment vectors

Vector	Definition
V1: Reality Mediation	Shift from direct information consumption to AI-curated immersive experiences; agents and avatars as reality gatekeepers
V2: Authority Shift	Movement of epistemic authority from institutions (media, academia, government) to technical systems (algorithms, AI assistants, verification infrastructure)
V3: Algorithmic Contestation	Campaigns and counter-campaigns targeting computational systems (recommendation algorithms, training data, model behaviour) rather than solely human audiences
V4: Synthetic Saturation	Scale and indistinguishability of synthetic content (deepfakes, generated text, synthetic audiences) overwhelming detection and verification capacity
V5: Human-Machine Convergence	Neuro-data collection and affective/behavioural modelling enabling finer-grained manipulation; blurring boundaries between human cognition and machine augmentation

Vector scoring is intentionally lightweight. For each technology and each vector, assessors give a Magnitude score, *M* (0–100) that represents the expected size of the technology's impact on that specific vector for the mission setting being analysed (i.e., the organisation's target audiences, platforms, time horizon, and threat model). Next, assessors proceed to apply two small adjustments: Directness, *D* (0.5–1.0) for how directly the technology enables the vector's disruption pathway (e.g., reality mediation, authority shift, algorithmic contestation, synthetic saturation, or human–machine convergence), and Timeframe, *T* (1.0 near-term; 0.75 mid-term; 0.5 long-term) for how soon that vector impact is expected to materialise at operationally relevant scale within that mission setting. Here, $i \in \{1,\ldots,5\}$ indexes the five Next-Gen information environment vectors (V1–V5) defined in Table 1. Let $VectorScore_i$ denote the score for vector V_i. In other words:

$$VectorScore_i = M_i \times D_i \times T_i$$

The overall *Disruption Score* is then a weighted average across the five vectors, where w_i is the organisational weight assigned to V_i:

$$\text{DisruptionScore} = \Sigma_{i=1}^{5}\, w_i \cdot VectorScore_i, \quad \text{with } \Sigma_{i=1}^{5}\, w_i = 1$$

Weights w_i can be set uniformly by default, but should ideally reflect mission priorities; NEXUS-HC makes these weights explicit and auditable, and teams can elicit them quickly using the Analytic Hierarchy Process (AHP) (Saaty, 1980; Saaty, 2008).

Stage 2 — Hype-cycle staging via observable indicators. Instead of relying on vendor-defined, proprietary stage assignments (e.g., the Gartner Hype Cycle), NEXUS-HC stages technologies using the same familiar hype-cycle stages but grounds placement in five observable indicators of adoption, institutionalisation, and adversarial use (Table 2): I1 Attention, I2 Investment, I3 Deployment, I4 Standardisation, and I5 Adversarial Adoption. Each indicator family is assessed on a simple ordinal scale (e.g., Low/Medium/High) with a brief evidence note (and, where useful, an evidence-quality tier), and the combined pattern is used to justify stage placement. Assessors compile "observable signals" (e.g., media/conference volume, funding/procurement signals, production deployments and incidents, standards activity, and documented misuse) and select the hype-cycle stage most consistent with the balance of evidence across the five indicator families (i.e., the stage supported by most indicators and the strongest evidence). Thresholds are transparent starting points: teams are expected to calibrate them to their language/platform ecosystem and record any changes to the default thresholds (and why) for auditability. Where indicators conflict, NEXUS-HC treats disagreement as uncertainty rather than forcing false precision; evidence quality tiers and a confidence level (or a simple stage probability distribution) can be recorded alongside the stage assignment.

Table 2: Five observable indicators of adoption, institutionalisation, and adversarial use

Indicator Family	Measurement Approach*
I1: Attention	Volume and volatility of technical and mainstream media coverage; Google Trends; conference session counts; social media mentions
I2: Investment	Venture funding rounds and amounts; government procurement/contract awards; corporate R&D allocations
I3: Deployment	Evidence of operational use (pilots, production integrations, documented incidents); user adoption metrics where available
I4: Standardisation	Emergence of standards bodies, technical specifications, regulatory frameworks, certification programs, audit methodologies
I5: Adversarial Adoption	Observed misuse in influence operations; red-team findings; threat-intelligence reports; vulnerability disclosures

* Note: I1–I5 are recorded on an ordinal Low/Medium/High scale with brief evidence notes; teams may also record evidence quality and confidence.

Stage 3 — Multi-dimensional scoring for contested environments. Beyond disruption, NEXUS-HC adds two dimensions that often diverge from "maturity" in adversarial settings. *Adversarial Leverage* is scored on a 0–100 scale and is computed as the sum of four 0–25 component scores: (a) access barriers (how hard it is for an adversary to obtain or access the capability), (b) deployment simplicity (how easy it is to operationalise at scale), (c) attribution difficulty (how hard it is to detect, trace, or prove responsibility), and (d) impact asymmetry (how large the effect is relative to the attacker's cost or risk). Because adversaries differ in resources and risk tolerance, teams may score distinct attacker archetypes (e.g., state-linked, criminal, hacktivist) and report either a mission-weighted score, scenario-specific scores, or a conservative min–max range. *Governance Readiness* is

also scored on a 0–100 scale and is computed as the sum of four 0–25 component scores covering (a) technical countermeasures, (b) standards/protocols, (c) regulatory/policy coverage, and (d) organisational capacity to implement controls—so that the overall score reflects the combined availability and maturity of mitigation mechanisms.

Stage 4 — Portfolio action derivation and cadence. NEXUS-HC converts the combined picture into action categories (Table 3) using an all-encompassing two-layer mapping over (a) the indicator profile across I1–I5 recorded on a Low/Medium/High scale and (b) the three scores (*Disruption*, *Adversarial Leverage*, and *Governance Readiness*). First, NEXUS-HC assigns a capability posture—Monitor, Pilot, Adopt, or Govern/Harden—based on *Disruption* and the diffusion-related indicators (especially I2 Investment, I3 Deployment, and I4 Standardisation) together with *Adversarial Leverage* and *Governance Readiness*. Second, it applies a threat overlay: Counter is added whenever I5 Adversarial Adoption and/or *Adversarial Leverage* indicates credible hostile use or high exploitability, particularly when *Governance Readiness* is weak (UK Ministry of Defence, 2023). This two-layer structure ensures that all possible combinations of indicator ratings and scores map to a clear decision posture, while still allowing combined actions (e.g., Pilot+Counter or Govern+Counter) when the capability and threat pictures diverge. Where indicators conflict, NEXUS-HC records uncertainty (evidence quality and confidence tiers) and defaults to the higher-urgency posture when the strongest credible signals justify it.

For cyber-relevant technologies, NEXUS-HC is not intended to replace established cyber risk management processes such as NIST RMF or ISO/IEC 27005. Rather, it gives a StratCom team an upstream prioritisation layer and a common briefing language. Where a technology interacts with attack surfaces or platform control systems, I3 Deployment and I5 Adversarial Adoption can draw on threat-informed inputs such as ATT&CK-mapped TTP reporting, incident reporting, and formal advisories, while Governance Readiness can be cross-checked against existing organisational cyber risk processes (Joint Task Force, 2018; ISO/IEC, 2022; MITRE, 2025).

To reduce scoring burden for small teams, NEXUS-HC can also be applied in a light-touch calibration mode. The aim is consistency rather than false precision: assessors can anchor 0–100 scores around illustrative bands of 20 (low), 50 (medium), and 80 (high); keep Directness at 1.0 unless the pathway is clearly indirect; and use Timeframe only to distinguish near-term, mid-term, and longer-term salience. A filled example for synthetic media is provided in Appendix A.

Dynamic reassessment is built in. Teams record explicit reassessment triggers (e.g., major deployments, regulatory actions, observed attacks, technical breakthroughs) and use scheduled reviews matched to action category (e.g., quarterly for Monitor, monthly for Counter, and event-driven for Pilot and Govern). Automated monitoring of I1–I5 can flag material shifts for expert re-evaluation. To address "leapfrogging" (rapid transitions over weeks), NEXUS-HC recommends a two-speed cadence: (a) a lightweight fast-lane weekly scan that updates only high-velocity indicators (I1 Attention, I2 Investment, and I3 Deployment/misuse signals) and (b) a quarterly full-lane review that refreshes slower-moving baselines (especially I4 Standardisation and I5 Adversarial Adoption) and re-runs weighting/sensitivity checks. The result is a short assessment card per technology that preserves an auditable record of why priorities changed across cycles.

Table 3: Portfolio action categories and trigger conditions

Action category	Complete trigger logic (using I1–I5 on Low/Medium/High scale and 0–100 Scores)	Organisational response
A. Monitor (capability posture)	Assign Monitor if *Disruption* <60 and I3 Deployment is Low and I2 Investment is not High (I1 Attention may be High).	Periodic review; lightweight tracking; knowledge building.
B1. Sandbox Pilot (Sandboxed Discovery)	Else Assign if *Disruption* > 80 and *Governance Readiness* < 50.	Isolated 'safe-to-fail' environment; no live audiences or production systems; objective is governance-building (evidence notes, standards/procedures (I4), countermeasures) with rapid feedback into Stage 3 scoring. Graduate to standard Pilot once Readiness ≥ 50.
B2. Pilot (capability posture)	Else assign Pilot if (I3 is Medium/High or I2 is High) and *Disruption* is 60–80 (or *Disruption* <60 with a defined use case), and *Governance Readiness* is not critically low (≥50).	Controlled experiments; capability assessment; gap analysis.

Action category	Complete trigger logic (using I1–I5 on Low/Medium/High scale and 0–100 Scores)	Organisational response
C. Adopt (capability posture – Resilient integration)	Else assign Adopt if I3 Deployment is High and *Governance Readiness* is High (≥75), and the disruption profile is stable (no recent step-change in *Disruption* Score drivers or diffusion indicators.)	Full operational integration; formal training programs; lifecycle maintenance; transition from "project" to business-as-usual, with conditional reassessment triggers.
D. Govern/Harden (capability posture)	Else assign Govern/Harden (all remaining cases), e.g., *Disruption* >70, or I4 Standardisation is Medium/High, or *Governance Readiness* <60 while I2/I3 indicate emerging adoption.	Standards development; policy formulation; capability acquisition; defensive investment.
E. Counter (threat overlay; may be combined)	Add Counter in addition to the capability posture if any holds: I5 *Adversarial Adoption* is High; or *Adversarial Leverage* >75; or (I5 is Medium and *Governance Readiness* <60).	Active defence; attribution investment; defensive capability development; intelligence collection.

* Note: Apply the capability posture rules (A–D) in order, then add the Counter overlay (E) when triggered. If indicators are mixed or evidence quality varies, record uncertainty (evidence-quality tiers and confidence) and default to the higher-urgency posture supported by the strongest credible signals.

4. Worked Example Determining Action Postures

Table 4 summarises a worked example using three technologies—synthetic media, algorithmic information operations (Algorithmic IO), and AI avatars—to show how NEXUS-HC translates evidence and scores into portfolio actions.

The example is intended to make the action logic inspectable rather than to forecast outcomes. Assessors score disruption vectors (V1–V5) to compute the *Disruption Score*, rate the five indicator families (I1–I5) on a Low/Medium/High scale based on observable signals, and then interpret the combined profile alongside *Adversarial Leverage* and *Governance Readiness* to assign a capability posture (Monitor/Pilot/Adopt/Govern) and, where warranted, a Counter overlay (Bolt and Lange-Ionatamishvili, 2026).

Synthetic media (including deepfakes and other AI-generated or AI-altered audio, images and video) receives a high Disruption score because it scales convincing content production and accelerates synthetic saturation, while also weakening trust and provenance (Table 1, V4). In Table 4, the indicator profile (I1–I3 High; I4 Medium; I5 High) and an elevated Adversarial Leverage score suggest both broad diffusion and active hostile uptake; Governance Readiness is improving but uneven, supporting a combined Govern/Harden + Counter posture focused on coordinated controls, labelling practices, monitoring and resilience measures (Doss et al., 2023; National Security Agency, Federal Bureau of Investigation and Cybersecurity and Infrastructure Security Agency, 2024; European Commission, 2025).

Algorithmic IO targets the computational systems that allocate attention—such as ranking, recommendation and moderation workflows—rather than relying only on direct persuasion. Its high *Disruption* Score reflects system-level leverage over amplification and visibility (Table 1, V3). Table 4 shows a Medium indicator profile (I1–I3 Medium; I4 Low; I5 Medium) paired with high *Adversarial Leverage* and low *Governance Readiness*, indicating exploitability even before strong standardisation or institutional controls emerge. Under NEXUS-HC, this combination supports a Govern/Harden posture, with Counter measures as a closely linked workstream: the priority is to close governance gaps early (policies, monitoring, platform engagement, and organisational controls) while hardening measurement and operational processes against manipulation. Accordingly, the recommended posture emphasises governance-first mitigation—standards and controls where feasible, monitoring and response playbooks, and targeted countermeasures for coordinated manipulation—rather than treating detection alone as sufficient (Bolt and Lange-Ionatamishvili, 2026; Tabassi, 2023).

AI avatars—represented here as interactive digital humans or virtual influencers—score primarily on reality mediation because they insert synthetic spokespeople into quasi-social interactions rather than only modifying isolated artefacts. Current research suggests that audiences can form one-sided audience attachments to virtual influencers comparable to those formed with human influencers, while engagement is shaped by human-likeness and perceived authenticity (Stein, Breves and Anders, 2024; Ju, Kim and Im, 2024). In NEXUS-HC terms, this pattern suggests rising salience and operational potential, but still uneven deployment, standards, and safeguards. For a StratCom team, that combination supports a Pilot posture, paired with governance safeguards around disclosure, provenance, consent, and impersonation risk.

Table 4: Worked example—synthetic media, algorithmic IO and AI avatars produce different action postures under NEXUS-HC

Technology	*Disruption* Score	Indicators (I1–I5)	*Adversarial Leverage* Score	*Governance Readiness* Score	Action
Synthetic media — Synthetic saturation	80	I1: High I2: High I3: High I4: Mid I5: High	75	58	Govern/ Harden +Counter
	High-end-to-end synthetic content at scale, undermining trust/provenance and saturating verification capacity	High attention; mature tools; active misuse	Low-cost, widely accessible tooling that's easy to operationalise and hard to attribute, with outsized impact	Improving, but uneven (policy + platform practice varies)	
Algorithmic IO — Algorithmic contestation	88	I1: Mid I2: Mid I3: Mid I4: Low I5: Mid	82	32	Govern/ Harden
	Directly targets the attention allocation machinery (ranking/recommendation/moderation), amplifying influence effects system-wide	Growing attention; feasible tactics	Highly scalable exploitation of platform dynamics with low barriers, low attribution, and strong impact asymmetry	Low–medium (limited controls and measurement)	
AI avatars — Reality mediation	72	I1: High I2: Mid I3: Mid I4: Low I5: Low	58	55	Pilot
	Synthetic spokespersons and interactive digital humans mediate audience attention and social presence, shifting engagement from direct human communication to avatar-led interaction.	High public visibility; growing vendor ecosystem; meaningful but still uneven deployment.	Moderate barriers to entry, scalable persona creation, and some impersonation potential, but less mature hostile use than synthetic media.	Partial safeguards exist, but disclosure, provenance, consent, and authenticity norms remain uneven across platforms and use cases.	

5. Retrospective Validation Illustration

As a retrospective validation, the three technologies in Table 4 can be compared against their observed trajectories over roughly the last 5–8 years. The purpose is not statistical proof, but to test whether NEXUS-HC would have generated action postures broadly consistent with what later became operationally salient.

On this reading, synthetic media is a plausible Govern/Harden + Counter case: the literature shows falling barriers to realistic manipulations and substantial public vulnerability, while governance responses have accelerated but remain uneven (Doss et al., 2023; European Commission, 2025). Algorithmic IO is a plausible Govern/Harden case because manipulation increasingly targets the systems allocating visibility, and even modest changes in collective behaviour can alter downstream recommendation outcomes (Fredheim and Pamment, 2025; Matias, 2023). AI avatars are a plausible Pilot case: virtual influencers and other digital humans already attract engagement and one-sided audience attachment, but operational norms for disclosure,

authenticity, and impersonation remain unsettled (Stein, Breves and Anders, 2024; Ju, Kim and Im, 2024). Taken together, these cases suggest that NEXUS-HC is most useful not when it predicts exact dates, but when it separates maturity from leverage and governance in a way that produces defensible early action.

6. Discussion

NEXUS-HC is intended to make technology prioritisation easier to justify and easier to revisit. It reduces a common failure mode: treating novelty or maturity as a proxy for urgency. In contested environments, urgency can arise early when adversaries can combine low-cost tools into effective workflows, while governance and standards lag behind diffusion. This complements related work on actor interventions in narrative evolution, which likewise emphasises timing, pathways, and intervention points rather than technology novelty alone (Ng, Zeng and Ponmani, 2026).

The framework also supports balanced mitigation. Counter does not mean only technical detection; it can include resilience measures. Evidence for prebunking and inoculation suggests that technique-based interventions can improve resistance to misinformation, which can complement technical and governance controls (Roozenbeek et al., 2022; Lewandowsky and van der Linden, 2021).

In contested information environments, organisations may encounter a catch-22 where a technology is highly disruptive yet poorly governed: they cannot safely pilot it without controls, yet cannot design effective controls without practical learning. NEXUS-HC therefore distinguishes between a standard Pilot and a Sandbox Pilot (Sandboxed Discovery) for 'leapfrogging' technologies. A Sandbox Pilot is permitted when *Disruption* > 80 and *Governance Readiness* < 50, but it is conducted under strict constraints (e.g., isolated environments, no interaction with live audiences, and no production recommendation engines) and exists primarily to generate the evidence required to raise *Governance Readiness*. Pilot objectives in this state are explicitly governance-building—producing evidence notes for attribution/traceability, drafting internal standards and procedures (contributing to I4 Standardisation), and validating technical countermeasures. Scoring becomes recursive: findings from the Sandbox Pilot feed rapidly back into Stage 3 (e.g., updating *Adversarial Leverage* and *Disruption* mechanisms), while *Governance Readiness* is updated as controls are tested and formalised. Once *Governance Readiness* crosses the 50 threshold, the technology graduates from Sandbox Pilot to a standard Pilot posture; a Counter overlay may still be triggered in parallel if *Adversarial Leverage* is high.

To address potential drift between organisational priorities and the evolving information environment, NEXUS-HC treats vector weights (w_i) as periodically reviewable parameters rather than fixed constants. Organisations should establish a standing review cadence (e.g., quarterly or aligned to major platform or policy shifts) to reassess whether existing weights continue to reflect mission risk. Trigger-based reviews are also recommended following significant ecosystem changes (such as major platform access restrictions or sudden adversarial capability shifts). Where material divergence is identified, weights may be adjusted through a documented governance process, with changes logged alongside technology score updates. This adaptive weighting mechanism preserves comparability over time while ensuring the *Disruption* Score remains aligned with current strategic priorities in contested information environments.

Limitations remain. Scoring involves judgement, signals are uneven across technologies, and some sources are grey literature. In addition, some indicator evidence may be unavailable, delayed, or systematically obscured (e.g., low-visibility deployments, restricted reporting, or capabilities developed in closed settings), which can bias assessments toward what is most observable rather than what is most consequential. NEXUS-HC should be used as a structured aid to judgement, not a mechanical ranking, and teams should document confidence and reasons for disagreement.

7. Conclusion

This paper presented NEXUS-HC, a hype-cycle-informed decision rubric for anticipating emerging technologies in next-generation information environments. By separating maturity signals from disruption mechanisms and governance readiness, it produces clearer portfolio actions and an auditable record of why priorities change over time.

Future work can extend the framework with empirical evaluation (e.g., retrospective back-testing of stage assignments), richer indicator libraries for specific technology families, and integration with measurement pipelines that track adoption and incident signals. These extensions would further strengthen the ability of StratCom organisations to manage the algorithmic burden of proof as new capabilities emerge.

Ethics Declaration: This paper is conceptual and did not involve human participants, personal data collection or activities requiring ethical clearance.

AI Declaration: AI tools were used to assist with language refinement and editing. The author reviewed all changes and is responsible for the final content.

References

Australian Defence Science and Technology Group (2016) Horizon Scan of Emerging Technologies and Trends for ADF Combat Service Support 2016. DST-Group-GD-0946. Edinburgh, SA: Department of Defence, Australian Government.

Bolt, N. and Lange-Ionatamishvili, E. (2026) The NextGen Information Environment, NATO Strategic Communications Centre of Excellence. Available at: https://stratcomcoe.org/publications/the-nextgen-information-environment/339 (Accessed 21 February 2026).

Chairman of the Joint Chiefs of Staff (2022) Joint Publication 3-04: Information in Joint Operations. Washington, DC.

Dedehayir, O. and Steinert, M. (2016) 'The hype cycle model: A review and future directions', Technological Forecasting and Social Change, 108, pp. 28–41. https://doi.org/10.1016/j.techfore.2016.04.005.

Defense Acquisition University (n.d.) Technology Readiness (resource page). Available at: https://www.dau.edu/index.php/cop/stm/resources/technology-readiness (Accessed 21 February 2026).

Doss, C., Mondschein, J., Shu, D. et al. (2023) 'Deepfakes and scientific knowledge dissemination', Scientific Reports, 13, 13429. https://doi.org/10.1038/s41598-023-39944-3.

European Commission (2025) Commission launches work on a code of practice on marking and labelling AI-generated content, Shaping Europe's digital future (press release, 5 November). Available at: https://digital-strategy.ec.europa.eu/en/news/commission-launches-work-code-practice-marking-and-labelling-ai-generated-content (Accessed 21 February 2026).

Fredheim, R. and Pamment, J. (2025) 'Assessing the risks and opportunities posed by AI-enhanced influence operations on social media', Place Branding and Public Diplomacy, 21, pp. 319–326. https://doi.org/10.1057/s41254-023-00322-5.

Gordon, T.J. and Hayward, H. (1968) 'Initial experiments with the cross-impact matrix method of forecasting', Futures, 1(2), pp. 100–116. https://doi.org/10.1016/S0016-3287(68)80003-5.

ISO/IEC (2022) ISO/IEC 27005:2022 Information security, cybersecurity and privacy protection — Guidance on managing information security risks. Available at: https://www.iso.org/standard/80585.html (Accessed 12 April 2026).

Joint Task Force (2018) Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. NIST Special Publication 800-37 Rev. 2. Gaithersburg, MD: National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-37r2.

Ju, N., Kim, T. and Im, H. (2024) 'Fake human but real influencer: the interplay of authenticity and humanlikeness in virtual influencer communication?', Fashion and Textiles, 11, 16. https://doi.org/10.1186/s40691-024-00380-0.

Lewandowsky, S. and van der Linden, S. (2021) 'Countering misinformation and fake news through inoculation and prebunking: A review of recent evidence and recommendations', European Review of Social Psychology, 32(2), pp. 348–384. https://doi.org/10.1080/10463283.2021.1876983.

Mankins, J.C. (2009) 'Technology readiness assessments: A retrospective', Acta Astronautica, 65(9–10), pp. 1216–1223. https://doi.org/10.1016/j.actaastro.2009.03.058.

Matias, J.N. (2023) 'Influencing recommendation algorithms to reduce the spread of unreliable news by encouraging humans to fact-check articles, in a field experiment', Scientific Reports, 13, 11715. https://doi.org/10.1038/s41598-023-38277-5.

MITRE (2025) MITRE ATT&CK. Available at: https://attack.mitre.org/ (Accessed 12 April 2026).

National Security Agency, Federal Bureau of Investigation and Cybersecurity and Infrastructure Security Agency (2024) Deploying AI Systems Securely, Joint Cybersecurity Information Sheet. Available at: https://www.cisa.gov/news-events/alerts/2024/04/15/joint-guidance-deploying-ai-systems-securely (Accessed 21 February 2026).

Ng, L.H.X., Zeng, Y. and Ponmani, M. (2026) 'The Architects of Narrative Evolution: Actor Interventions Across the SAGES Framework in Information Campaigns'. Available at: https://scholarspace.manoa.hawaii.edu/items/2a06600f-e299-4b94-9cbb-d45b89181bf7 (Accessed 16 April 2026).

Roozenbeek, J., Traberg, C.S. and van der Linden, S. (2022) 'Technique-based inoculation against real-world misinformation', Royal Society Open Science, 9(5), 211719. https://doi.org/10.1098/rsos.211719.

Saaty, T.L. (1980) The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation. New York: McGraw-Hill.

Saaty, T.L. (2008) 'Decision making with the analytic hierarchy process', International Journal of Services Sciences, 1(1), pp. 83–98. https://doi.org/10.1504/IJSSCI.2008.017590.

Stein, J.-P., Breves, P.L. and Anders, N. (2024) 'Parasocial interactions with real and virtual influencers: The role of perceived similarity and human-likeness', New Media & Society, 26(6), pp. 3433–3453. https://doi.org/10.1177/14614448221102900.

Tabassi, E. (2023) Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1.

U.S. Department of Defense (2005) Technology Readiness Assessment (TRA) Deskbook, May. Available at: https://spacese.spacegrant.org/SEModules/Technology%20Mods/TRA-TRL_2005%20DoD.pdf (Accessed 21 February 2026).

U.S. Government Accountability Office (2020) Technology Readiness Assessment Guide: Best Practices for Evaluating the Readiness of Technology for Use in Acquisition Programs and Projects (GAO-20-48G). Available at: https://www.gao.gov/products/gao-20-48g (Accessed 21 February 2026).

UK Ministry of Defence (2021) The Orchestration of Military Strategic Effects. Available at: https://www.gov.uk/government/publications/the-orchestration-of-military-strategic-effects-guide (Accessed 21 February 2026).

UK Ministry of Defence (2023) Allied Joint Doctrine for Information Operations (AJP-10.1), Edition A Version 1. Available at: https://www.gov.uk/government/publications/allied-joint-doctrine-for-information-operations-ajp-101 (Accessed 21 February 2026).

Wardle, C. and Derakhshan, H. (2017) Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making. Strasbourg: Council of Europe.

Appendix A: Light-touch Assessment Card (Synthetic Media)

Table A1 provides a filled example of the light-touch calibration approach for synthetic media. It is intentionally illustrative rather than normative, and is included to show how a small StratCom team can document scores, evidence notes and the resulting posture without over-engineering the process.

Table A1: Filled light-touch assessment card for synthetic media

Field	Illustrative value	Rationale / evidence note
Mission setting	Public-facing StratCom monitoring across major social and video platforms	Assumes a team concerned with audience exposure, trust effects and response options.
Dominant vector	V4: Synthetic Saturation	Synthetic media most directly increases scale and indistinguishability of manipulated content.
Magnitude (0–100)	80	High potential effect on trust, verification workload and content volume.
Directness (0.5–1.0)	1.0	The technology directly produces the artefact that creates the disruption.
Timeframe	1.0 (near-term)	Operational salience is already present in current platforms and workflows.
I1 Attention	High	Sustained public, policy and platform attention.
I2 Investment	High	Mature tooling ecosystem and continued vendor/platform investment.
I3 Deployment	High	Widespread legitimate and malicious use.
I4 Standardisation	Medium	Labelling, provenance and policy responses are advancing, but unevenly.
I5 Adversarial Adoption	High	Active misuse and deception risks are already evident.
Adversarial Leverage	75	Low-cost tooling, scale, attribution difficulty and outsized effect.
Governance Readiness	58	Improving, but uneven across organisations and platforms.
Recommended posture	Govern/Harden + Counter	Governance and resilience measures are needed alongside active defensive monitoring.