



(Estd. 1870)

CONTESTING THE COGNITIVE-CYBER- KINETIC KILL CHAIN

A CROSS-DOMAIN DOCTRINE FOR
AGENTIC CYBER OPERATIONS



COGNITIVE
Shape Perception
Influence Decision



CYBER
Infiltrate | Disrupt
Degrade | Deceive



AGENTIC
Autonomous Agents
Adaptive Execution



KINETIC
Precision Effects
Achieve Objectives



RESILIENCE
Detect | Deter
Defend | Recover

CONTEST – DISRUPT – DEGRADE – DENY – DEFEAT



Lieutenant Colonel Nishank Sharma
Dr Surya Prakash
Dr Hemprasad Patil

About the Occasional Paper

This paper argues that contemporary cyber campaigns increasingly begin not with malicious code but with cognitive manipulation, and that cyber defence must, therefore, evolve from a network-centric to a cognitive-centric posture. Drawing on Proactive Response and Hostile Action Analysis Regime (PRAHAAR)—an indigenous, experimental cross-domain framework developed jointly at Indian Institute of Technology-Indore and Military College of Telecommunication Engineering, Mhow—it treats disinformation, social engineering, and network intrusion as sequential phases of one continuous campaign—the Cognitive-Cyber-Kinetic Kill Chain.

The paper advances two findings. The ‘Authenticity Trap’ holds that intelligence pipelines, which discard unverified material as noise, also discard the signal that most reliably precedes an attack; in a hundred-day live evaluation, this signal led hostile activity by three to four days. The second holds that a defender who seeds an adversary’s path with believable deception can invert the attacker’s cost advantage; in controlled trials, adaptive deception cut the simulated breach rate roughly threefold.

Situating the framework against how the major cyber powers—China, Russia, the United States, the United Kingdom, Germany, and Pakistan—fight across both domains, it evaluates the operational, doctrinal, and command implications, including human oversight, legal mandate, and sovereign control. It concludes that a state, which learns to read the cognitive contest deceive an adversary within its own networks and keep machine-speed autonomy bounded by human command, will hold a decisive advantage in conflicts already being fought before the first packet is sent.

Introduction: The Contest before the First Packet

Modern cyber campaigns are designed to deny the defender any early warning. State-sponsored Advanced Persistent Threat (APT) groups no longer begin at the network boundary; they begin at what practitioners call Layer 8—the human cognitive layer.¹ Long before a malicious packet reaches a firewall, the adversary has already begun shaping the information environment through synthetic media, inauthentic activity, and amplified disinformation, leading the public to believe falsehoods as truth.² The sequence is consistent: from the 2016 Internet Research Agency operations to recent state-linked activity; narrative manipulation appears first, then social engineering delivers the intrusion, and technical exploitation completes the chain.³ India has experienced this significantly. During the 2025 confrontation following Operation Sindoor, the adversary's cognitive operations—including Artificial Intelligence (AI)-tailored propaganda, deepfakes, and demoralising messaging targeting military families—ran in parallel with cyber probing. These were not separate contests; rather they were phases of a single campaign.

The key argument follows: future cyber campaigns begin not with malicious code, but, with cognitive manipulation, and hence, cyber defence must evolve from a network-centric model to a cognitive-centric one that integrates information operations, social engineering, and technical defence systems into a single, doctrine-driven framework.

Defensive architecture remains partitioned along the lines the adversary intends to exploit.⁴ Signature-based intrusion detection mechanisms inspect network traffic alone; now Large Language Model (LLM)-driven penetration tools operate strictly within the exploitation domain.⁵ Multi-Agent Reinforcement

Learning (MARL) environments scrutinise this network but, never reason about cognitive shift.⁶ No instrument in the Security Operations Centre (SOC) today connects a population-level shift in sentiment to the spear-phishing wave that follows days later—denying commanders their most valuable early warning, with direct bearing on national cyber-intelligence doctrine and the command relationships that must govern autonomous systems acting at machine speed.⁷

Why Today's Models Arrive Too Late

The usual reference point is the cyber kill chain formulated at Lockheed Martin, running from reconnaissance to action on the objective.⁸ Its weakness is that it begins too late: by the time reconnaissance is detectable on the wire, the adversary has already spent weeks shaping the cognitive terrain. MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) precisely catalogues the technical tactics and techniques of the adversary, but begins at the point of intrusion, remaining largely silent on the cognitive preparation that precedes it.⁹ LLM-driven tools such as PentestGPT reason effectively within this exploitation domain, yet they have no visibility of the population being cognitively primed for their lure. Reinforcement-learning environments such as Cyber Operations Research Gym (CybORG) model attacker–defender interactions within the network in detail, but remain confined entirely to the network environment. The cognitive-warfare literature describes influence richly but conversely rarely connects it in measurable time to the intrusion that follows. The gap is structural: cognitive warfare has been studied chiefly by strategists concerned with perception, cyber intrusion chiefly by engineers concerned with code, and the institutions built on them rarely share a common picture or clock. The unified kill chain, which stitches reconnaissance, intrusion, and action together, functions firmly within the technical

domain.¹⁰ The space between these bodies of work is the space the Cognitive-Cyber-Kinetic Kill Chain (CKKC) occupies, connecting a measurable shift in the information environment to the technical attack it foreshadows. The paper first highlights how the major powers fight, then focusses on the models that describe them.

How Major Powers Fight in the Cognitive-Cyber Domain: A Comparative Baseline

Before any framework claims novelty, it must be measured against how the principal cyber powers fight. The states that India must plan against do not treat information operations and network intrusion as separate trades—they run them as one campaign, on one clock. The differences lie in how openly this fusion is acknowledged, which institution owns it, and whether the posture is offensive, defensive, or both. Table 1 sets out the full comparison and Figure 1 plots it; refer to both for amplification of the discussion that follows.




Figure 1: Indicative Analytic Positioning of National Cognitive-Cyber Posture

Source: Created by author

China's 'Three Warfares'—public-opinion, psychological, and legal warfare, formalised for the People's Liberation Army (PLA) in 2003—and its 'Cognitive Domain Operations' concept treat the mind, not the network, as a target, waged continuously in peacetime.¹¹ The 2024 reorganisation replaced the Strategic Support Forces with dedicated information support, aerospace, and cyberspace forces, which has sharpened this further.¹² For India, this is not something abstract: the doctrine shaping opinion on Taiwan shapes it along the Himalayan frontier.

Russia frames the same idea as 'Information Confrontation', codified in its 2016 Information Security Doctrine and divided into technical and psychological components.¹³ Russia's Main Directorate's (GRU's) 'Disruptive Playbook' in Ukraine—destructive cyber operations timed for psychological effect, executed by APT44/Sandworm—is the clearest open-source confirmation that a capable adversary already runs the cross-domain campaign that CKKC describes.¹⁴

The United States (US), a leading technological power, reaffirmed the concepts of 'Defend Forward' and 'Persistent engagement' in its 2023 Defense Cyber Strategy and maintains influence and cyber instruments institutionally separate.¹⁵ The United Kingdom explicitly recognises this convergence. Its National Cyber Force's 'Doctrine of Cognitive Effect' uses cyber means to undermine an adversary's will—an openly cognitive but offensive approach. The Cognitive-CKKC applies the same principle to defence.¹⁶ Germany's Cyber and Information Domain Service, elevated in 2024 to a fourth military service, unifies cyber, electronic warfare, psychological operations, and intelligence under one command. This is the organisational model this paper recommends—though constitutional limits confine it largely to wartime.¹⁷

Pakistan offers the most relevant case for India. Islamabad's effort, exemplified through APT36/Transparent Tribe, compensates for modest tooling with tight cross-domain integration around live flashpoints—Kashmir, Pahalgam, Operation Sindoor.¹⁸ Open-source telemetry from the 2025 episode constitutes this paper's most significant data point: anomalous spear-phishing targeting Indian government and defence networks preceded the 22 Apr trigger by several days—the precursor lead that the CKKC is designed to capture.¹⁹ A network-centric defence would have had no reason to act on it; a cognitive-centric defence would have. Israel's Unit 8200 and Stuxnet remain the canonical proof that a cyber operation can produce kinetic effect²⁰; North Korea's Lazarus Group shows the opposite emphasis—formidable technical capability turned chiefly to revenue—a reminder that the CKKC targets geopolitically driven campaigns, not every raid.

Power Organisation or	Doctrine: Cognitive Cyber Fused how and are	Signature Illustrative Operations or	Posture and Relevance to India
China: PLA; Information Support and Cyberspace Force	'Three Warfares' plus cognitive domain operations; perception as continuous war domain.	Taiwan Himalayan influence campaigns; sustained cyber-espionage.	Most integrated doctrine; India's highest-priority adversary.
Russia: GRU (APT28, Sandworm); <i>Federalnaya Sluzhba Bezopasnosti</i> (Fed-eral Security Service); <i>Sluzhba Vnesheyny Razvedki</i> (Foreign Intelligence Service)	'Information Confrontation', technical plus psychological components (2016 doctrine).	GRU 'Disruptive playbook', Ukraine; NotPetya; wiper attacks timed to kinetic action.	Cross-domain by design; open-source proof the CKKC campaign is real.

Power Organisation or	Doctrine: Cognitive Cyber Fused how and are	Signature Illustrative Operations or	Posture and Relevance to India
The US: US Cyber Command; National Security Agency	Technically dominant; 'Defend forward', 'Persistent engagement'. Cognitive/influence held separately.	Hunt-forward operations against adversary infrastructure.	Strong execution, partitioned from the cognitive precursor.
The United Kingdom: National Cyber Force	Explicit 'Doctrine of cognitive effect': cyber means to shape adversary perception and will.	Covert operations degrading adversary cognition.	Openly cognitive but offensive; CKKC turns the insight to defence.
Germany: Cyber and Information Domain Service	Cyber, electronic warfare, psychological-operations, intelligence under one service.	Hybrid-threat defence; North Atlantic Treaty Organization (NATO) exercises (Locked Shields).	Best organisational model; constrained to a wartime mandate.
Pakistan: APT36 /Transparent Tribe	Modest tooling, tight cross-domain integration around live flashpoints.	Pahalgam/Sindoor-themed lures (2025); phishing detected days before the trigger.	Most India-relevant case; real-world proof of the precursor lead.
Israel: Unit 8200	Cyber tied to physical effect; world-class tradecraft.	Stuxnet-kinetic outcome from a cyber operation.	Canonical demonstration of the chain's sixth phase.

Power Organisation or	Doctrine: Cognitive Cyber Fused	how and are	Signature Illustrative Operations or	Posture and Relevance to India
North Korea: Lazarus Group	Technical capability oriented revenue, cognitive shaping.	to not	Financial heists; ransomware (e.g. WannaCry).	Weakest cognitive precursor; not the CKKC's target case.
India (today): Defence Cyber Agency; Computer Emergency Response Team-In (CERT-In)	Capable improving, network-organised; fused doctrine.	and but no	National incident response; tri-service cyber focus.	The position this paper argues India should change (Figure 1 arrow).

Table 1: Major Cyber Powers and their Cognitive and Technical Domains.

Source: Created by author

Two conclusions emerge. First, integration is the clear direction of travel for every major power, with China, Russia, and Pakistan already positioned high on the integration axis. Second, despite possessing considerable technical capabilities, India remains unnecessarily constrained by a network-centric approach and lacks a doctrine that links the cognitive precursors of an attack with its technical execution. The arrow in Figure 1 represents the doctrinal shift advocated in this paper.

The Anatomy of Cross-Domain Campaign

The CKKC extends the conventional model backwards into the pre-kinetic period in which a campaign is conceived, and forwards into the physical consequences a cyber operation is designed to

enable (Figure 2). It comprises six phases governed by a single overarching intent.




Figure 2: The Six-Phase Cognitive-Cyber-Kinetic Kill Chain

Source: Adapted from the Proactive Response and Hostile Action Analysis Regime (PRAHAAR) framework

The campaign opens with preparation of a quiet build-up of infrastructure and cross-correlation of economic, military, political, and diplomatic indicators to locate a region under strain. The campaign then proceeds towards cognitive dominance, where coordinated narrative manipulation, botnet amplification and synthetic media prime the population through a flood of unfiltered information. As authentic news struggles against the noise, the gap between what is verified and what merely circulates becomes the opening through which the target is captured. Target intelligence and reconnaissance follow, while initial access is secured by AI agents running spear-phishing at machine speed—a single click admits the intruder; lateral movement then carries malware through the network, including operational technology or supervisory control; and data acquisition segments governing physical infrastructure. The chain closes with a kinetic effect—

damage in the cyber domain enabling action in the physical world, from a stock-exchange outage to the metro services being halted for days—thus, narrowing an opponent’s options before any conventional move is contemplated.

What Cyber-Kinetic Kill Chain Does That Other Models Do Not

The comparison with other models shows what the CKKC contributes that the existing toolkit does not. It is not a new attack technique, rather a reframing that closes a measurable gap. Table 2 sets the framework against the models a SOC and an information-warfare cell use; Figure 3 shows the same point as coverage across the six phases.

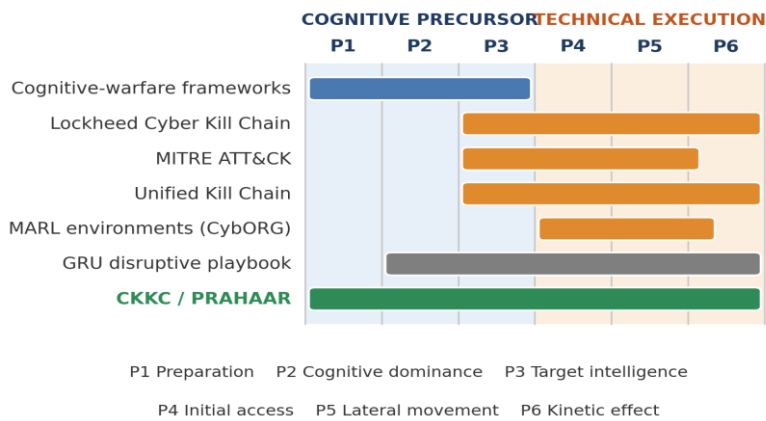


Figure 3: Only Cyber-Kinetic Kill Chain (in green) Spans the Entire Chain

Source: Authors’ assessment

Every widely used technical model begins after the decisive phase—the Lockheed Cyber Kill Chain, MITRE ATT&CK, the Unified Kill Chain and reinforcement-learning environments—all open at or beyond the point of access. The

cognitive-warfare literature, including the NATO's own conceptualisation, has the mirror-image limitation.²¹ In this gap, the framework places three things, none of which is an algorithm. First, is a single clock: the testable claim that the divergence between verified and unverified information streams can be measured and leads the intrusion by a usable number of days. This lead, which can be quantified and acted upon, is the contribution. Second, is a defensive inversion: a deception logic that grows stronger as the adversary grows more capable, thus, reversing cyber defence's two-decade cost asymmetry. Third, is the bounded autonomy: a command structure fixing machine speed below the network layer and human judgement above it? The first buys warning, the second economy, and the third control.

Model	The Start and the End	Cognitive Precursor or Single Clock	Defensive Early Warning and Gap it Leaves
Lockheed Cyber Kill Chain	Reconnaissance to action	No cognitive phase; no cross-domain clock.	Detects intrusion, not the preparation before it.
MITRE ATT&CK	Initial access to lateral movement	No; technical technique catalogue only.	Superb post-intrusion taxonomy; silent before access.
Unified Kill Chain	Recon objectives to	No; extends the technical chain only.	Longer chain; still no pre-kinetic warning.
MARL environments (CybORG)	In-network only	No; the network is the entire world.	Rich attacker and defender modelling; blind to the population.

Model	The Start and the End	Cognitive Precursor or Single Clock	Defensive Early Warning and Gap it Leaves
Cognitive-warfare frameworks	Influence/perception	Cognitive only; no measured link to intrusion.	Describes the precursor; no bridge to the attack.
GRU disruptive playbook (adversary)	Cognitive to kinetic	Yes; fused, but as offensive tradecraft.	Proof the campaign is real; not a defence.
CKKC or PRAHAAR	Preparation to kinetic effect	Yes; one measurable clock, both domains.	Days of warning, adaptive deception, bounded autonomy.

Table 2: Cyber-Kinetic Kill Chain against Established Models and Reference Adversary

Source: Created by Author

The First Finding: The Authenticity Trap

The first finding concerns intelligence, and it indicts a discipline rather than a tool. Open-source intelligence pipelines are built on a verification-first principle: information that cannot be authenticated is discarded as noise.²² This approach is sound for conventional intelligence analysis but can be actively misleading when applied to the detection of cognitive operations. When an adversary deliberately surfaces any unverified content as the opening move, the divergence between the official narrative and the unverified stream is not noise—it is the attack itself, taking shape in plain sight. This is the authenticity trap: when an attack is imminent, the discarded data stream carries more information

than it does under peacetime conditions, yet the filter removes precisely the material that matters most.

The value of an insight also lies in its timing. Cross-correlation between open-source patterns and hostile activity is well documented, and social-sentiment analysis has a record of anticipating geopolitical events.²³ Tested against real crises, the verified channels read calm while the unverified stream climbs, and a warning fires days before any visible escalation. During a 100-day live evaluation, the signal crossed its alert threshold roughly three to four days before the breach as highlighted in Figure 4. Also, during the 2025 Operation Sindoor episode, anomalous spear-phishing against Indian networks was detected on or about 17 Apr, several days before the 22 Apr trigger.²⁴




Figure 4: The Authenticity Trap and Timing Problem

Source: Adapted from the PRAHAAR framework

Crucially, the signal has been validated outside the laboratory. SAKSHI, the regional-risk engine that answers where an adversary is likely to move, sustained a detection score of $F1=0.84$ across a continuous 100-day live evaluation spanning

several real theatres. It provides evidence that the read was stable on live open-source feeds, not merely on curated data. Tested against real historical crises, this contrast is even sharper: a verified-only pipeline never once raised the alarm (its recall was effectively zero), while the same inputs read through the unverified stream flagged every episode, days ahead. The ‘Authenticity Trap’ is, therefore, not just a theoretical worry but a measured failure mode—the same inputs, with and without the discarded stream, separate a usable precursor signal from none.

The signal becomes useful only once measured and reduced to two instruments answering the two questions a commander asks—where and when. The first one scores which region and sector are under geopolitical and institutional strain.²⁵ The second reads the information environment in parallel with verified sources against the unfiltered stream—expressing the gap as a divergence index that rises whenever a population is primed. Neither suffices alone; the window opens only when a stressed region is simultaneously under active manipulation. Therefore, this coincidence is the framework’s single most actionable warning as showcased in Table 3.

Signal	Question	What it Reveals	Why it Matters to the Commander
Regional-Risk (where)	WHERE	Which region/sector is under geopolitical and institutional strain.	Identifies terrain an adversary will move against, well before any packet is sent.
Divergence (when)	WHEN	Gap between official narratives and the unverified stream.	Crosses its alert threshold three to four days before the strike.

Signal	Question	What it Reveals	Why it Matters to the Commander
Both together (the window)	THE WINDOW	The attack window opens only when a stressed region is under active manipulation.	Neither signal alone suffices; their coincidence is the single most actionable warning.

Table 3: The Two-Signal Lock at the Heart of Cyber-Kinetic Kill Chain

Source: Created by author

The same signature is visible closer to home. When examination paper leaks emerge, official channels may remain silent for days while unverified information proliferates—an instructive illustration of the pattern, though not an act of hostile statecraft. Even state-of-the-art filters cannot reliably distinguish a coherent, personally relevant spear-phishing message from a legitimate one. Defence, therefore, cannot depend solely on staff recognising a hostile message; it must detect the campaign before that message is crafted—upstream, where the earliest warning signals lie.

The Spear-Phishing Multiplier

The bridge between the cognitive and technical phases is the spear-phishing wave. What matters strategically is not the absolute rate at which an AI-generated lure evades a filter, but the multiplier effect: enriching a payload with cross-domain cognitive context made it, on average, nearly nine times more likely to evade detection than the same agent’s unenriched output.²⁶ The value lies in transferring intelligence from the cognitive layer to the social-engineering layer, thereby, producing a message tailored to what its recipient fears—and why.

The multiplier is remarkably consistent. Measured across five filter families—from rule-based SpamAssassin through the adaptive filters of Gmail and Proofpoint to Microsoft’s Defender for Office 365 with its detonation chamber—cognitive enrichment lifted the evasion rate by between 7.2 and 11.3 times, a geometric mean of 8.8x with a 95 per cent confidence interval of 7.0 to 10.9x (Table 4). The enriched payloads themselves cleared filters about 41 per cent of the time, squarely within the range the published literature reports for AI-written phishing. The claim is not a record bypass rate but a consistent multiplicative effect that holds across filter generations. The strategic consequence is stark: across a campaign of roughly 70 targeted messages, even that bypass rate makes at least one successful breach a near-certainty—on the order of 99 per cent. Filters cannot improve fast enough; interdiction must occur upstream before the enrichment chain closes.

Filter family	Baseline	Enriched	Lift
SpamAssassin (rule-based)	0.06	0.68	11.3x
Gmail (Bayesian/ML)	0.05	0.36	7.2x
Proofpoint (behavioural)	0.04	0.31	7.8x
MS Defender O365 (detonation)	0.03	0.28	9.3x
Custom Bayesian (in-house)	0.05	0.44	8.8x
Mean	0.05	0.41	8.8x

Table 4: Filter-Evasion Lift from Cognitive Enrichment, Anchored to Published Phishing-LLM Literature

Source: Adapted from the PRAHAAR framework

The Second Finding: Inverting the Cost Asymmetry

The second finding belongs to the defender, thus overturning one of cyber security's most demoralising truths. Rather than blocking every move, the defending agent plants believable decoy nodes, credentials, and documents, which are modelled on documented APT behaviour throughout the network. An intruder acting on planted intelligence and triggering a high-confidence alarm is expelled and wastes effort, pursuing a dead end. The defender then raises alert levels on adjacent nodes and repositions decoys along the intruder's most frequently probed routes, while the intruder learns to avoid exposed traps. Through self-play, both continually adapt to each other and move towards an equilibrium.²⁷

Deception inverts the conventional rule that a more capable attacker imposes greater cost on the defender. The more aggressively an intruder pursues its objective, the more decoys it touches and the faster it exposes itself. The effect was measured over 3,000 engagements in each condition (Figure 5). Against the same self-modifying attacker, a defender with no deception was breached 45 per cent of the time; static decoys cut that to 27 per cent; and a defender that co-evolved its decoys against the adapting attacker cut it to 15 per cent—a threefold reduction. Over the same conditions, detection rose from 54 to 85 per cent, the time taken to evict an intruder more than halved, and the rate at which an evicted attacker regained a foothold fell from 40 to 16 per cent.²⁸ A static decoy set helped; an adaptive one helped considerably more.




Figure 5: Effectiveness of Co-Evolving Deception

Source: Adapted from the PRAHAAR framework

The decisive property is that the trap is planted information, not a software flaw. A better-resourced adversary is no less likely to act on a convincing decoy, because deception operates on the decision rather than the code. This is the rare defence that does not degrade as the adversary’s tooling improves. Decoys cut that to 27 per cent; and a defender that co-evolved its decoys against the adapting attacker cut it to 15 per cent, marking a threefold reduction. Over the same conditions, detection rose from 54 to 85 per cent, the time taken to evict an intruder more than halved, and the rate at which an evicted attacker regained a foothold fell from 40 to 16 per cent.²⁹ A static decoy set helped; an adaptive one helped considerably more.

The Cyber-Kinetic Kill Chain in Action: Three Scenarios

Three short scenarios make the argument concrete, contrasting a network-centric and a cognitive-centric defence against the same campaign; outcomes appear in Table 5, timing in Figure 6.




Figure 6: Timing and Early Warning

Source: Created by author

Scenario 1: The Substation and the Rumour. A fabricated story of contaminated water, amplified by inauthentic accounts, seeds grievance in a city under communal strain. Officials ignore it for four days while the unverified stream climbs; on day five, a spear-phishing wave tuned to that grievance reaches a state utility, and one click admits an intruder toward the substation network. A network-centric centre sees nothing until the intrusion; a cognitive-centric defence sees the divergence cross its threshold on day one and hardens the utility's perimeter before the lure is written.

Scenario 2: Operation Sindoor, Replayed. Following the Apr 2025 Pahalgam attack, a western adversary-linked actor ran exactly the same sequence: disinformation and impersonation, then spear-phishing built from the crisis, then attempted intrusion and data theft. Anomalous phishing began around 17 Apr—days before the trigger—with malicious documents proliferating by 24 Apr. Analysed through the CKKC, this is one chain whose cognitive phase announced the technical one; read through a network-centric lens, two separate teams handled two separate problems,

and the precursor was discarded. The claim is simply that the first reading was available, in time, from open sources.

Scenario 3: The Honeyed Headquarters. A well-resourced intruder breaches a defence headquarters despite good hygiene—the familiar asymmetry, where the attacker need succeed once and the defender must be right everywhere. A deception-seeded network, salted with believable decoy credentials and hosts, inverts this—the harder the intruder presses, the faster it betrays itself, since the trap is planted information rather than a software flaw. The intruder is detected, evicted, and left having spent effort on nothing.

Dimension	Network-Centric Defence	Cognitive-Centric Defence
First detection point	At the intrusion (P4).	In the cognitive phase (P2), as the population is primed.
Warning lead	Effectively none; reactive.	Roughly three to four days; proactive.
Decision options	Containment under time pressure.	Harden, pre-position, interdict before the lure is written.
Cost imposed on attacker	Low; needs only to succeed once.	High and rising—deception turns into aggression into exposure.

Dimension	Network-Centric Defence	Cognitive-Centric Defence
Likely outcome	Breach contained late, damage done.	Breach pre-empted or detected early, at the attacker's expense.

Table 5: Comparison of Network-Centric Defence and Cognitive-Centric Defence

Source: Created by author

The Strategic Dividends: Why India Should Adopt This Model

Decisions of this nature are driven by strategic dividends, not architectural considerations. Table 6 identifies six such dividends, none of which can be fully realised while information operations, social engineering, and network defence remain institutionally separated. These dividends include decision advantage—replacing reactive responses with deliberate action; cost inversion—whereby deception exposes stronger attackers while reducing the cost of sustained defence; sovereign control—ensuring that such a sensitive capability is developed domestically; doctrinal unification—providing the information-warfare cell and SOC with a common operating picture and decision cycle; bounded autonomy—combining machine speed below the network layer with human authority above it; and deterrence by denial—whereby an adversary that expects its preparations to be detected early, confronted with decoys, and penalised rather than rewarded can no longer surprise India at minimal cost.

Dividend	Mechanism	Strategic Payoff	Primary Indian beneficiary
Decision advantage	Divergence signal leads intrusion by 3-4 days.	Reactive scramble becomes deliberate decision.	Defence Cyber Agency; CERT-In; theatre commands
Cost inversion	Adaptive deception turns effort into exposure.	Defence scales favourably against a stronger adversary.	Critical information infrastructure operators
Sovereign control	Indigenous build; air-gapped simulation.	Capability held nationally, not procured abroad.	Ministry of Defence; National Security Council Secretariat
Doctrinal unification	One clock across IW cells and SOCs.	Ends the partition adversaries exploit.	Tri-service and national cyber architecture
Bounded autonomy	Machine speed below network layer; human above.	Machine-speed defence without losing command.	Operational commanders; legal authority
Deterrence by denial	Precursor read + decoys + cost inversion.	Removes the adversary's easy, unattributed win.	National deterrence posture

Table 6: Strategic Dividends of Cognitive-Centric Posture for India

Source: Created by author

Together, these six dividends answer the policymaker's central question—not 'How does it work?' but 'What does it deliver?' The answer is early warning, economy, sovereignty, coherence, control, and deterrence.

Keeping the Commander in Command

A framework that can simulate offensive operations at machine speed raises an unavoidable question of control, sharpened by the proliferation of uncensored AI models. PRAHAAR treats a strict human-in-the-loop requirement as a strict condition of deployment, dividing command across three tiers (Figure 7).




Figure 7: Three-Tier Human-in-the-Loop command

Source: Adapted from PRAHAAR framework.

At the machine-speed tier, automated processes handle anomaly detection and decoy placement at millisecond latency, where human intervention would be neither feasible nor useful. At the human-machine tier, escalation from passive deception to active isolation requires a human to authorise the move. At the highest tier, kinetic or legal escalation remains significantly a matter of human authority. These maps onto familiar roles: the SOC owns machine speed, the watch officer owns escalation to active isolation, and the commander with legal authority alone sanctions kinetic or legal consequence. A composite risk score—weighing intelligence confidence, target susceptibility, network

exposure, and organisational criticality—triages alerts so only high-risk cases demand a human in real time. The principle should be doctrinal: machine speed below the network layer, human judgement above it.

Three conditions make the tiering credible. First, rules of engagement must be written, not assumed. Second, the audit must be adversary-proof—every automated action logged under a tamper-evident cryptographic record. Third, authority must be named—the operations centre owning machine speed, the watch officer owning escalation, and the command-and-legal authority owning kinetic or legal consequence. Command of machines faster than any human is a standing responsibility, not a problem solved once.

Implications for India's Cyber Defence Doctrine

The framework is a lens, but its consequences bear directly on India's cyber defence. One caveat bounds the claim: it is built for geopolitically driven campaigns, not a purely financial ransomware raid, which throws a far weaker signal. India's architecture has matured rapid—the Defence Cyber Agency provides tri-service focus, CERT-In coordinates national response—but remains organised around the network. India now needs less new agencies than a new division of labour among those that exist. Five propositions follow.

- **Move Intelligence Upstream.** Divergence between verified and unverified narratives must be a weighted intelligence indicator, not noise discarded. National threat intelligence should be reorganised to capture and act on it; analysts must be rewarded for reading the stream current practice instructs them to delete.

- **Build Deception into Defensive Doctrine.**

Active deception is the most efficient way to draw an adversary into expending resources against nothing. Deception ranges, refreshed through cyber war-gaming, should become standing features of critical-infrastructure defence, not an experimental novelty.

- **Hold Cross-domain Capability as Sovereign Asset.**

This capability carries obvious dual-use risk and should be developed and held nationally rather than procured abroad—with controls built in from the outset: these include air-gapped offensive simulation, ethics approval for live payload generation, an explicit legal mandate, tamper-evident audit, and responsible disclosure of filter-evasion findings.³⁰

- **Codify Human Command Before Fielding Autonomy.**

Any operational system must be governed by the tiered structure above, underwritten by a clear legal mandate and rigorous audit—designed in from the outset, not retrofitted after the first serious incident.

- **Learn from the Integration, Not the Aggression.**

China's Three Warfares, Russia's information confrontation, and Pakistan's blended operations are instructive, but not as templates for imitation. These are the evidence that cross-domain campaigning is the prevailing form of the contest. India's task is to build—indigenously, within a clear legal and ethical mandate—an early-warning read of the pre-kinetic contest, fused to an autonomous but human-commanded defence.

Conclusion: Doctrine for Pre-Kinetic Contest

The contest described here is not a future possibility; it is already unfolding—days before the first malicious packet is sent and within an information environment that most defensive architectures are not designed to observe. The argument is not that India should acquire a particular tool, but that it must reconsider the layer at which it expects the contest to begin. A defence organised around the network perimeter encounters the adversary only after the cognitive environment has been shaped and the lure crafted. By contrast, a defence focused on cognitive precursors engages the threat at its earliest stage, when intervention is the least costly and the warning period is longest. The two findings show this is measurable advantage, not aspiration—a precursor signal leading hostile activity by several days, and a deception logic that grows stronger as the adversary grows more capable—neither available to a force that still treats information operations, social engineering, and network intrusion as separate problems.

The strategic takeaway is a single line of doctrine. A state that learns to read the contest in the cognitive domain, to deceive the adversary inside its own networks, and to keep the commander firmly in command of machines that act faster than any human can—and that builds and holds this capability as its own—will enter every future confrontation already ahead in the phase that decides the rest.

India's next confrontation in cyberspace will not begin with a hostile packet. It will reveal itself days earlier through a manipulated information environment that current defensive architectures are not designed to detect. India's adversaries already operate in this manner. The question, therefore, is no longer whether the contest has entered the cognitive domain—it has—

but whether India will develop the ability to detect and counter it there first. The required capability is technically feasible, governable and capable of sovereign development. The cost of delay will be measured in the critical days of warning forfeited each time an adversarial campaign unfolds unseen.

Endnotes

¹ Daniel Ayzenshteyn, Roy Weiss, and Yisroel Mirsky, “Cloak, Honey, Trap: Proactive Defenses Against Large Language Model Agents”, *34th USENIX Security Symposium*, 2025, accessed 11 Jul 2026, <https://www.usenix.org/conference/usenixsecurity25/presentation/ayzenshteyn>

² Thomas Rid, *Active Measures: The Secret History of Disinformation and Political Warfare* (New York: Farrar, Straus and Giroux, 2020); RA Chesney and DK Citron, “Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security”, *California Law Review* 107, no. 6 (2019): 1753-1820, accessed 12 Jul 2026, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3213954

³ Renne DiResta et al, “The Tactics and Tropes of the Internet Research Agency”, *New Knowledge White Paper*, 2018, accessed 12 Jul 2026, <https://dataspace.princeton.edu/handle/88435/dsp01fb494c31z>; Gerardo Spagnuolo, “The Cognitive Battlefield of Hybrid Warfare”, *NATO Defense College Foundation*, accessed 12 Jul 2026, <https://www.natofoundation.org/food/the-cognitive-battlefield-of-hybrid-warfare/>

⁴ ENISA, “Threat Landscape 2024”, *European Union Agency for Cybersecurity*, 2024, accessed 13 Jul 2026, <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024>

⁵ Gelei Deng et al, “PentestGPT: Evaluating and Harnessing Large Language Models for Automated Penetration Testing”, *34th USENIX Security Symposium* 2024, accessed 13 Jul 2026, <https://www.usenix.org/conference/usenixsecurity24/presentation/deng>

⁶ Maxwell Standen et al, “CybORG: A Gym for the Development of Autonomous Cyber Agents”, *IJCAI-21 1st International Workshop on Adaptive Cyber Defense*, 2021, accessed 13 Jul 2026, <https://arxiv.org/abs/2108.09118>

⁷ Major General BK Sharma and Major General Sanjeev Chowdhry, “Cognitive Warfare by China and India's Response”, *USI Monograph No. 3*, 2026, accessed 13 Jul 2026, https://usiofindia.org/pdf/Monograph_no_3_organized.pdf

⁸ Eric M Hutchins, Michael J Cloppert, and Rohan M Amin, “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains”, *Lockheed Martin Corporation technical report*, 2011, accessed 14 Jul 2026, <https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf>

⁹ MITRE ATT&CK Framework, “Enterprise Matrix v14.0”, *MITRE Corporation*, 2024, accessed 14 Jul 2026, <https://attack.mitre.org/>

¹⁰ Paul Pols, “The Unified Kill Chain: Designing a Unified Kill Chain for Analysing, Comparing and Defending against Cyber Attacks”, *Cyber Security Academy*, 2017, accessed 14 Jul 2026, <https://www.unifiedkillchain.com/assets/The-Unified-Kill-Chain.pdf>

¹¹ ES Cochran, “China’s ‘Three Warfares’: People’s Liberation Army Influence Operations”, *International Bulletin of Political Psychology* 20, no. 3; accessed 14 Jul 2026, <https://commons.erau.edu/ibpp/vol20/iss3/1/>; US DoD, “Military and Security Developments Involving the People’s Republic of China”, *Annual report of US Department of Defense*, 2024, accessed 14 Jul 2026, <https://perma.cc/RJT3-SZPR>

¹² Ibid.

¹³ Gavin Wilde, “Cyber Operations in Ukraine: Russia’s Unmet Expectations”, *Carnegie Endowment for International Peace*, Dec 2022, accessed 15 Jul 2026, <https://carnegieendowment.org/research/2022/12/cyber-operations-in-ukraine-russias-unmet-expectations?lang=en>

¹⁴ Dan Black and Gabby Roncone, “The GRU’s Disruptive Playbook”, *Google Cloud Blog*, 2023, accessed 14 Jul 2026, <https://cloud.google.com/blog/topics/threat-intelligence/gru-disruptive-playbook>

¹⁵ US DoD, “2023 DoD Cyber Strategy Summary”, *US Department of Defense*, 2023, accessed 14 Jul 2026, https://media.defense.gov/2023/Sep/12/2003299076/-1/-1/1/2023_DOD_Cyber_Strategy_Summary.PDF

¹⁶ “Responsible Cyber Power in Practice”, *National Cyber Force*, 2023, accessed 15 Jul 2026, <https://www.gov.uk/government/publications/responsible-cyber-power-in-practice/responsible-cyber-power-in-practice-html>

¹⁷ “Cyber and information space-a new domain”, *Bundeswehr*, accessed 15 Jul 2026, <https://www.bundeswehr.de/en/organization/the-cyber-and-information-domain-service>; Nicholas Fiorenza, “Germany Restructures Bundeswehr”, *Janes*, 2024, accessed 15 Jul 2026,

<https://www.janes.com/defence-intelligence-insights/defence-news/c4isr/germany-restructures-bundeswehr>

¹⁸ Rhishav Kanjilal, “Advisory: Pahalgam Attack themed decoys used by APT36 to target the Indian Government”, *Seqrite Blog*, Apr 2025, accessed 15 Jul 2026, <https://www.seqrite.com/blog/advisory-pahalgam-attack-themed-decoys-used-by-apt36-to-target-the-indian-government/>; “Dark Web Profile: APT36”, *SOCRadar*, 2025, accessed 15 Jul 2026, <https://socradar.io/blog/dark-web-profile-apt36/>

¹⁹ Prakhar Paliwal, Atul Kabra, and Manjesh Kumar Hanawal, “Cyber Warfare During Operation Sindoor: Malware Campaign Analysis and Detection Framework”, *21st International Conference on Information Systems Security*, 2025, accessed 16 Jul 2026, <https://arxiv.org/abs/2510.04118>; Seqrite, “Cyber Threat Landscape Analysis: India-Pakistan Conflict Dynamics”, *Seqrite*, 2025, accessed 16 Jul 2026, <https://www.seqrite.com/cyber-threat-landscape-analysis-india-pakistan-conflict-dynamics/>

²⁰ Kim Zetter, *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon* (New York: Crown, 2014).

²¹ Francois du Cluzel, “Cognitive Warfare”, *NATO Allied Command Transformation/Innovation Hub*, 2020, accessed 16 Jul 2026, https://innovationhub-act.org/wp-content/uploads/2023/12/20210122_CW-Final.pdf

²² S Glassman and M Krueger, “The Interoperability of Open Source Intelligence with Big Data Analytics”, *Intelligence and National Security* 36, no. 3 (2021): 372-389.

²³ Nanda Rani, Bikash Saha, and Sandeep Kumar Shukla, “Leveraging OSINT in Advanced Proactive Cybersecurity”, *IEEE Access*, Aug 2025, accessed 16 Jul 2026, <https://ieeexplore.ieee.org/stamp.jsp?arnumber=11143131>; Neelam Tripathi and D Srinivasa Rao, “Deep Learning Sentiment Analysis of News and Social Media in Geopolitical Crises”, *International Information and Engineering Technology Association*, Aug 2025, accessed 17 Jul 2026, <https://www.iieta.org/journals/isi/paper/10.18280/isi.300825>

²⁴ Ibid, Paliwal, Kabra, and Hanawal, “Cyber Warfare During Operation Sindoor”, Seqrite, “Cyber Threat Landscape Analysis”.

²⁵ Nishank Sharma et al, “SAKSHI: Agentic AI for Regional Intelligence using Structured Open-Source Intelligence”, *IEEE Explore*, Feb 2026, accessed 17 Jul 2026, <https://ieeexplore.ieee.org/abstract/document/11545466>

²⁶ Fredrik Heidig et al, “Devising and Detecting Phishing: Large Language Models vs. Smaller Human Models”, *IEEE Access* 12 (2024): 1-14, accessed 17

Jul 2026, <https://arxiv.org/abs/2308.12287>; Julian Hazell, “Spear Phishing with Large Language Models”, *University of Oxford working paper*, 2023, accessed 17 Jul 2026, <https://arxiv.org/abs/2305.06972>; John Seymour and Philip Tully, “Weaponizing Data Science for Social Engineering: Automated E2E Spear Phishing on Twitter”, *DEF CON*, 2016, accessed 17 Jul 2026, <https://blackhat.com/docs/us-16/materials/us-16-Seymour-Tully-Weaponizing-Data-Science-For-Social-Engineering-Automated-E2E-Spear-Phishing-On-Twitter-wp.pdf>

²⁷ John Schulman et al, “Proximal Policy Optimization Algorithms”, *arXiv*, 2017, accessed 17 Jul 2026, <https://arxiv.org/abs/1707.06347>; Richard S Sutton and Andrew Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (Cambridge, MA: MIT Press, 2018); Nicolas Papernot et al, “The Limitations of Deep Learning in Adversarial Settings”, *IEEE EuroSec&P*, 2016, accessed 18 Jul 2026, <https://arxiv.org/abs/1511.07528>

²⁸ Quanyan Zhu and Tao Li, “Agentic AI for Cyber Resilience: A New Security Paradigm and Its System-Theoretic Foundations”, *arXiv*, 2025, accessed 18 Jul 2026, <https://arxiv.org/html/2512.22883v1>

²⁹ Ibid.

³⁰ Sharma and Chowdhry, “Cognitive Warfare”; Charles Cleveland, Daniel Brookes, and Davi Maxwell, “Cognitive Warfare: An Allied Blueprint and a Pentagon Opportunity”, *Small Wars Journal*, 16 Jan 2026, accessed 18 Jul 2026, <https://smallwarsjournal.com/2026/01/16/cognitive-warfare/>

About the Authors



Lieutenant Colonel Nishank Sharma is an Indian Army Corps of Signals officer and a cyber security researcher specialising in Artificial Intelligence (AI)-enabled cyber defence and agentic autonomous systems. He is. Currently pursuing an MTech on Agentic AI and Cyber Defence. He holds Google and National Cyber Crime Training Centre certifications and contributes regularly to military journals.



Dr Surya Prakash is Professor of Computer Science and Engineering at Indian Institute of Technology (IIT), Indore, and served as the Head of Department from 2017 to 2020. He has studied at IIT Kanpur, IIT Madras, and University Institute of Engineering and Technology, Kanpur. His research focuses on computer vision, image processing, machine learning, deep learning, pattern recognition, and biometric security.



Dr Hemprasad Patil is associated with the Military College of Telecommunication Engineering, Mhow, where he specialises in AI-enabled cyber defence. His research spans machine learning, deep learning, and agentic AI for military operations. He mentors student officers and advances indigenous defence capabilities.

About the United Service Institution of India (USI)

The USI was founded in 1870 by a soldier scholar, Colonel (later Major General) Sir Charles MacGregor 'For the furtherance of interest and knowledge in the Art, Science and Literature of National Security in general and Defence Services, in particular'. It commenced publishing its Journal in 1871. The USI also publishes reports of its members and research scholars as books, monographs, and occasional papers (pertaining to security matters). The present Director General is Vice Admiral Sanjay J Singh, SYSM, PVSM, AVSM, NM, PhD (Retd).



United Service Institution of India (USI)
Rao Tula Ram Marg, Opposite Signals Enclave,
New Delhi-110057

Tele: 2086 2316/ Fax: 2086 2315,
E-mail: dircl@usiofindia.org