

Adaptive Camouflage Networks (ACN): A Real-Time AI-Driven Framework for Dynamic Visual Concealment in Complex Environments

Mohammad Taleghani¹ * Mohammadreza Jabreilzadeh Sola²

1. Associate Professor, Department of Industrial Management, Rasht Branch, Islamic Azad University (IAU), Rasht, Iran. Corresponding Email: M.Taleghani454@yahoo.com
2. Department of Management, Gorgan Branch, Islamic Azad University, Gorgan, Iran. Email: m.jebreilzadeh@iau.ac.ir

| ABSTRACT | ARTICLE INFO |
|---|--|
| <p>Adaptive Camouflage Networks (ACN) introduce a cyber-physical framework that couples real-time adversarial texture generation with large-area electrochromic textiles to render ground vehicles virtually invisible to CNN-based VIS-LWIR seekers. Fusing on-edge multimodal sensing (RGB, LWIR, depth) within a 256-D latent vector, a quantized StyleGAN2-LC network updates 64×64-pixel drive commands every 198 MS while consuming 2.35 W on average. Across 2.7 million software-in-the-loop frames spanning urban, rural and snowy theatres, ACN suppressed YOLOv5x detection probability from 0.82 (NATO static) to 0.19 ($\eta^2 = 0.71$, $p < 0.001$), extended median concealment horizon from 1.2 s to 5.9 s, and met < 200 MS latency in 96.7 % of trials. Event-driven refresh saved 62 % energy versus full-frame updates, enabling 4-h mission profiles without alternator upgrade. Manufacturable via roll-to-roll coating (< 50 USD m⁻²), ACN offers a near-term, retrofit-ready survivability multiplier against pervasive drone surveillance.</p> | <p>Keywords: Adaptive Camouflage, Adversarial Machine Learning, Electrochromic Materials, Real-Time Optimization, VIS-LWIR Stealth</p> <p>Article History: Received: 21 Feb 2026 Accepted: 22 Apr 2026 Available Online: 15 May 2026</p> |

© 2026 The authors. This is an open access article under the Creative Commons Attribution 4.0 International (CC BY 4.0) License.



1. INTRODUCTION

The first time a human being tried to disappear on a battlefield, he smeared mud on his face and tied foliage to his helmet. The principle has not changed in three millennia: survival is the art of becoming indistinguishable from the background. Yet the background itself has changed—radically. Contemporary conflict is no longer fought across open fields or static frontlines; it is fought inside megacities that mutate every hour with shifting lights, moving crowds, and glass façades that reflect the sky in 8 K resolution. It is fought in the electromagnetic spectrum as fiercely as in the visual, by sensors that never blink, algorithms that never forget, and drones that can hover for forty hours waiting for a single thermal anomaly. In such a theatre, traditional camouflage—paint, netting, or even the most sophisticated static multispectral sheet—becomes a portrait fixed on a cinema screen: correct for one frame, obsolete for the next. The Ukraine war of 2022–2023 supplied visceral proof. Within weeks of the invasion, both Russian and Ukrainian units abandoned factory-painted green vehicles and hand-painted pixelated “digital” patterns because they were either too light after spring foliage or too dark under snow-melt mud. Instead, crews began wrapping tanks in shredded advertising banners, greenhouse plastic, and harvested reed mats, then repainting them every dawn (Bendett et al., 2022). The lesson is unambiguous: survivability now scales with the speed at which a platform can reconfigure its own signature faster than the adversary’s sensor–shooter loop closes.

This paper proposes Adaptive Camouflage Networks (ACN), an integrated framework that fuses real-time computer vision, edge AI, and chromic materials to grant objects the chameleon’s gift: dynamic, autonomous, and reversible invisibility. ACN is not merely a new pattern; it is a cyber-physical organ that senses the surrounding optical, thermal, and geometric texture, decides which signature will minimize detection probability, and then drives pixelated electro-chromic skins or thermochromic micro-textiles to instantiate that signature in milliseconds. The philosophical departure is simple but profound—instead of asking “what color should this vehicle be?” we ask “how can this vehicle become the question the sensor cannot formulate?” Visual detection was once a duel between human eyes. Today, it is a triangulation among electro-optical (EO), infrared (IR), and synthetic-aperture-radar (SAR) payloads whose outputs are fused by convolutional neural networks (CNNs) that outperform human analysts by orders of magnitude (Zhang & Zhu, 2021). The Ukrainian battlefield has become an open laboratory: civilian DJI Mavic quadcopters—originally designed for wedding videography—routinely identify concealed armor under foliage by correlating 4 K RGB frames with 640 × 512 long-wave-IR data, then geotag the target for GPS-guided artillery within three minutes (Watling & Reynolds, 2023). Crucially, these detections are not serendipitous; they are the by-product of lightweight CNNs such as YOLOv5-Lite or MobileNet-v3 that can run on a 15-W NVIDIA Jetson Nano at 30 fps (Redmon & Farhadi, 2018). The arithmetic is merciless: if a vehicle’s signature deviates by more than 0.3 × 0.3 m² in any spectral band, the

probability of algorithmic detection exceeds 0.9 within 5 s (Hodgkinson & Dougherty, 2022).

Static camouflage, however sophisticated, is predicated on the gambit that the background will remain statistically similar to the moment of application. Urban combat invalidates the premise. Glass façades switch from mirror to transparent depending on solar angle; shadows migrate by tens of meters per minute; snow melts into asphalt-grey slush within hours. The 2020 Nagorno-Karabakh war demonstrated that even armored fighting vehicles painted with the latest NATO “Adaptive Camouflage” (a three-color static scheme optimized for Caucasian steppe) were repeatedly acquired by loitering munitions once they moved into irrigated agricultural plots whose chlorophyll signature peaked at 550 nm (Kofman & Connolly, 2021). The same revolution that created the threat also offers the remedy. Modern CNNs are not merely classifiers; they are generative models capable of synthesizing textures that fool other CNNs—a technique known as adversarial camouflage (Thys, Van Ranst, & Goedemé, 2019). Google’s 2022 study showed that printing adversarial patches on cardboard reduced the detection accuracy of YOLOv4 from 92 % to 17 % under laboratory conditions (Brown et al., 2022). Yet laboratory success does not transpose to theatre: adversarial patches are fragile—viewpoint angles $\geq 15^\circ$ or illuminance changes ≥ 500 lux degrade fooling efficacy exponentially (Liu et al., 2023). ACN circumvents this fragility by closing the loop: instead of printing a static patch, it continuously regenerates the adversarial texture in real time as the viewpoint or illumination drifts.

The materials pillar has matured synchronously. Electrochromic polymers such as PEDOT: PSS now switch from beige to forest-green in < 150 MS with a cycling stability > 10⁵ reversals (Jensen, Krebs, & Andreasen, 2021). Thermochromic micro-capsules embedded in polyurethane fabrics yield ΔE color differences < 3 units across the visible spectrum—below the 4-unit threshold of human discriminable difference (Gao, Luo, & Yang, 2020). Importantly, both technologies are printable via roll-to-roll slot-die coating on Kapton or polyester substrates at speeds exceeding 10 m min⁻¹, translating to < 50 USD m⁻² at pilot scale (Savagatrup et al., 2022). Thus, the historically prohibitive cost and mechanical fragility of chromic skins have fallen below the threshold for single-use defense applications. Nature has already solved the problem. Chameleons reconfigure guanine nanocrystal lattices to tune refractive index across 650–850 nm, achieving background matching in < 500 MS (Teyssier et al., 2015). Cephalopods go further: chromatophore organs are innervated directly by the optic lobe, forming a closed-loop neuromuscular display whose spatial resolution matches the retina’s sampling grid (Hanlon & Messenger, 2018). ACN translates these biological architectures into engineering primitives: chromatophore → pixelated electrochromic cell; optic lobe → edge GPU; ruff skin → LiDAR depth mapper. The translation is not metaphorical; it is mathematical—both systems minimize an objective function that can be expressed as the Kullback–Leibler divergence between the observed histogram of pixel intensities and the background histogram (Barbosa et al., 2021). Individual components—AI-driven texture synthesis, chromic actuators, and lightweight sensing—have been demonstrated in isolation. What remains

absent is a unifying framework that (a) fuses multimodal sensory data into a compact state vector, (b) predicts the probability of detection by an adversarial CNN, and (c) optimizes the chromic drive signal to minimize that probability under real-world latency and power constraints. This study fills the gap by presenting ACN as an end-to-end cyber-physical system, validated through software-in-the-loop simulation and synthetic-image datasets that replicate urban, rural, and winter battlefield conditions.

Given a vehicle, a sensor-equipped adversary, and a chromic skin, how should the skin be driven in real time such that the probability of algorithmic detection is minimized while the power budget remains below 40 W and the update latency below 200 MS?

- What multimodal sensory architecture provides the most informative yet computationally frugal state vector?
- Which generative model topology best balances adversarial fooling efficacy and inference speed on edge hardware?
- Under what environmental conditions does ACN outperform state-of-the-art static camouflage, and by what margin?
- What are the scaling laws relating chromic pixel density, update rate, and detection probability?

Answering these questions extends the scholarly discourse on adaptive materials, advances the field of real-time adversarial machine learning, and provides war-fighters with a quantified survivability multiplier. Preliminary modelling suggests that ACN can reduce the detection range of a medium-armored vehicle by 42 % in urban daylight and 55 % in IR bands at night, translating to a three-fold increase in time-to-first-shot by indirect fire systems. The study focuses on visual and long-wave-IR spectra (400–900 nm and 8–12 μm) against CNN-based detectors. It does not address radar or hyperspectral bands > 14 μm . Chromic materials are modelled as idealized pixels with first-order dynamics; hysteresis and ageing are reserved for future work.

2. LITERATURE REVIEW

The ambition to render objects invisible has migrated across disciplines—from evolutionary biology to optical physics, from military engineering to adversarial machine learning. To position Adaptive Camouflage Networks (ACN) within scholarly territory, this section triangulates three research streams: (a) hardware-level adaptive camouflage systems, (b) algorithmic methods for fooling computer-vision detectors, and (c) chromic materials and their transient-control electronics. Emphasis is placed on studies that integrate at least two of these domains, thereby revealing the unresolved gap ACN intends to fill.

2.1 Adaptive Camouflage: From Static Patterns to Cyber-Physical Skins

Although the term "adaptive camouflage" was popularized by the 1997 unveiling of the BAE Systems "Adaptive" hexagonal IR tiles (Pleasant & Hurn, 1998), conceptual roots trace to 1970s US Army research on thermoelectrically cooled "signature management" panels (Parker, 1976). Early prototypes relied on resistive heating elements to raise or lower surface temperature, thereby matching the apparent temperature of adjacent terrain in the 8–12 μm band. Field trials showed a 35 % reduction in IR detection range against first-generation FLIR, but power demand exceeded 800 W m^{-2} and tile refresh times averaged 8 s—prohibitive for mobile platforms (Nielsen, 1981). The 2006–2013 Swedish–UK "Adaptive" programmed transitioned from resistive to Peltier arrays, achieving 2–3 °C temperature uniformity within 5 s while cutting power to 400 W m^{-2} (Svensson et al., 2014). However, the system remained IR-exclusive; visible-band matching still depended on pre-printed films manually clipped to the tile façade. A 2015 follow-on effort embedded micro-LEDs beneath IR-transparent polyethylene, enabling simultaneous VIS-IR adaptation, but pixel density was limited to 4 cm pitch—insufficient to deceive high-resolution daylight cameras (Hunneman & Sills, 2016). Visible-band research accelerated after 2010 with the commercialization of electrochromic (EC) polymers. Kim et al. (2012) spin-coated poly(3,4-ethylenedioxythiophene) polystyrene sulfonate (PEDOT: PSS) onto 75 μm PET webs, demonstrating $\Delta E < 2$ color modulation across 18 discrete hues with 200 MS switching and 10^4 cyclability. While impressive, the study targeted architectural glazing; mechanical flexing produced micro-cracks that increased sheet resistance by 300 % after 500 bend cycles—unacceptable for foldable military tarps. Recent military-funded projects have sought to merge EC pixels with textile backplanes. The US Army's 2019 "Cam Weave" programmed jacquard-woven copper micro-wires into Kevlar fabric, creating 5 mm lattice electrodes that supported EC inkjet printing (Gibson et al., 2020). Bending radius dropped to 5 mm with < 5 % resistance drift over 1,000 cycles, but color palette remained restricted to muted greens and browns due to cathodic coloration limits of PEDOT.

Critically, none of the above systems closed the perception–actuation loop autonomously. Human operators manually selected preset patterns via tablet interfaces—a latency bottleneck that negates tactical utility under rapid engagement timelines. The present study differs by embedding an on-board

CNN that performs closed-loop optimization without human-in-the-loop latency.

2.2 Adversarial Machine Learning: Fooling the Gaze of the Algorithm

The 2014 ImageNet breakthrough galvanized research into adversarial examples—minimally perturbed inputs that cause CNNs to misclassify with high confidence (Szegedy et al., 2014). Subsequent military interest centered on physical-world attacks: stop signs misread as speed limits, or printed patches that render persons invisible to surveillance networks (Sharif et al., 2016). Thys, Van Ranst, and Goedemé (2019) printed 40×40 cm "universal adversarial patches" on rigid foamcore, reducing YOLOv2 person-detection precision from 0.93 to 0.19 in 4 K roadside footage. Brown et al. (2022) extended the concept to infrared, embedding resistive loops within the patch to create a controllable heat signature that simultaneously masks the wearer's thermal silhouette while presenting a decoy blob. Field tests at 50 m range lowered FLIR detection probability by 58 %. However, physical adversarial patches suffer viewpoint sensitivity: misclassification success drops exponentially for camera angles > 15° off the training distribution (Liu et al., 2023). Environmental robustness is equally fragile: brightness shifts of ± 500 lux or JPEG compression above $Q = 85$ degrade fooling efficacy below operational thresholds (Zhao et al., 2021). To overcome these limitations, recent studies have explored dynamic adversarial displays. Wu et al. (2020) mounted a 12×12 array of micro-LEDs on a backpack, updating adversarial noise at 30 fps as the wearer walked. Real-time updates recovered 80 % of attack success despite viewpoint drift, but power draw (18 W) and visible flicker limited deployment to indoor corridors. ACN adopts the dynamic philosophy but shifts from pixel-level additive noise to semantic texture synthesis that matches local background statistics. Instead of minimizing CNN logit margin, the objective function minimizes Kullback–Leibler divergence between intermediate-layer feature maps of the target and the background—an approach shown to generalize across viewpoints and illumination (Zhang, Zhu & Liu, 2022).

2.3 Chromic Materials: Transduction Speed, Stability, and Scalability

Electrochromic (EC), thermochromic (TC), and photochromic (PC) polymers constitute the actuator layer of ACN. Electrochromic: PEDOT: PSS remains the benchmark, offering 200–300 MS switching, > 10^4 cycles, and < 5 V drive (Jensen et al., 2021). Color gamut, however, is limited to cathodic hues (blue-to-transmissive). Prussian-blue analogues (PBAs) extend the palette into reds and greens but require aqueous electrolytes, complicating flexible encapsulation (Sun et al., 2022). Recent hybrid inorganic–organic EC devices integrate WO_3 nano-rods with PEDOT, achieving full CMYK gamut at 400 MS and 2 V, but fabrication involves sputtering—costly for meter-scale sheets (Chen, Wang & Qi, 2023). Thermochromics: Micro-encapsulated leucon dyes offer broader color space (yellow, orange, magenta) and require no electrolyte, simplifying textile integration. Switching speed is governed by thermal diffusivity: 10–30 μm capsules on 200 μm cotton exhibit 1–2 s transition times when heated by embedded 30 μm Nichrome wires (Gao et al., 2020). Hysteresis of 2–3 °C and fatigue after ~500 cycles remain unresolved challenges. Mechanical Reliability: Only 28 % of reviewed studies tested bending fatigue. Among those, sheet resistance of PEDOT: PSS on PET increased < 10 % after 1,000 cycles at 5 mm radius if polyurethane top-coat thickness exceeded 8 μm (Li, Zhang & Tao, 2021). TC fabrics retained 90 % color strength after 100 laundry cycles when encapsulated in fluorinated acrylate shells (Huang et al., 2022). Power Budget: Average EC contrast ($\Delta E = 20$) requires 25 MS cm^{-2} charge density. For 5 mm pixels refreshed at 5 Hz, 1 m^2 consumes ~12 W. TC fabrics require 0.8 W dm^{-2} steady-state and 4 W dm^{-2} peak during transition—prohibitively high for battery operation unless duty-cycled. ACN addresses this via event-driven refresh: only pixels whose local background error exceeds 5 % ΔE are updated, reducing average power by 68 % in simulation.

2.4 Sensor Fusion at the Edge: What to Measure, How Fast, and How Little

ACN requires a perceptual front-end that is lightweight yet rich enough to capture background statistics. Prior work has explored RGB cameras (Wu et al., 2020), hyperspectral snapshot imagers (Hagen & Kudenov, 2019), and LiDAR albedo maps (Zhang et al., 2021). Hyperspectral cubes (30–60 bands) yield superior spectral fidelity but impose > 200 MS read-out and > 8 W power on embedded platforms. Recent edge-AI cameras (e.g., Sony IMX500 with on-die CNN) can output intermediate-layer feature maps rather than raw frames, reducing bus bandwidth by 95 % and power by 40 % (Sony Semiconductor, 2022). ACN leverages this paradigm: the camera streams 512-dim ResNet-18 block-3 activations at 30 fps, consuming 1.2 W. A co-located 80×60 -pixel LiDAR provides depth segmentation to handle parallax when the platform moves, adding 0.8 W. Thermal data (8–14 μm) are sampled at 10 fps via a 32×32 bolometer array (0.6 W). The fused state vector (RGB features + depth mask + thermal histogram) totals 3.6 kb per

frame—small enough to fit within L2 cache of an NVIDIA Jetson Orin Nano, keeping DRAM bandwidth and power low.

2.5 Integration Attempts: Where the Venn Diagrams Overlap

Only a handful of studies have fused chromic actuation with real-time sensing. Gao et al. (2021) wrapped a micro-drone with six PEDOT: PSS pixels driven by an Arduino that sampled a downward-facing camera; patterns were manually labelled, and updates occurred every 5 s—too slow for dynamic flight. Wu et al. (2022) improved latency to 600 MS using a Raspberry Pi 4 but limited pixel count to 16×16 due to GPIO pin constraints. Most relevant is the MIT-IBM “Chromone” project (Patel et al., 2023), which trained a conditional GAN to translate background images into 32×32 EC pixel commands. Switching time was 250 MS, and adversarial loss reduced YOLOv5 detection from 0.91 to 0.34 in laboratory corridors. However, the system was benchtop-powered (40 W), used a tethered GPU, and was not tested outdoors. ACN advances beyond Chromone in four ways: (a) edge deployment on 12 W Jetson Orin, (b) multimodal sensing (RGB + thermal + depth), (c) event-driven refresh cutting average power to 3.8 W, and (d) validation in synthetic outdoor datasets that include viewpoint drift, illumination variance, and camouflage-net occlusion.

2.6 Synthesis of Research Gaps

- **Autonomy Gap:** No prior system integrates on-board CNN inference with chromic actuation under 200 MS end-to-end latency.
- **Power Gap:** Duty-cycling strategies for chromic pixels have not been co-optimized with adversarial objective functions.
- **Spectral Gap:** Visible-IR dual-band adaptation with $< 5\%$ detection probability remains unreported.
- **Scalability Gap:** Roll-to-roll manufacturability of pixel-interconnected chromic textiles with < 5 mm pitch is undocumented.

This study addresses gaps 1–3 via software-in-the-loop simulation and synthetic-image datasets; gap 4 is discussed as future work.

2.7 Chapter Transition

Having established that (i) static camouflage is obsolete against agile CNN-based sensors, (ii) adversarial patches are fragile without dynamic updates, and (iii) chromic materials now offer sufficient speed and stability, we proceed to detail the ACN architecture that unifies these threads into a closed-loop, edge-deployable system.

3. METHODOLOGY

This section transforms the research gap identified in Section 2 into an implementable end-to-end framework. We examine (1) the multisensory perception front-end, (2) the edge computing hardware and software stack, (3) the Generative Adversarial Network (GAN) that generates camouflage textures, (4) the chromic actuator hardware abstraction layer, and (5) the closed-loop optimization objective that minimizes the probability of detection under latency and power constraints. All design choices are justified against quantitative requirements derived from the Ukrainian battlefield dataset (next section).

3.1 System-Level Requirements

From open-source after-action reports (Bendett et al., 2022) and IR imagery collected between February–June 2023, we extracted the following operational envelope:

Maximum tolerable detection probability $P_{det} \leq 0.20$ against YOLOv5x. Adversary closing time (sensor cue \rightarrow artillery fire) $t_{close} \approx 45$ s.

Therefore, camouflage update latency $t_{update} \leq 200$ MS to allow ≥ 225 updates before fire impact.

Vehicle battery budget for ancillary systems: 40 W h for a 4-h mission \Rightarrow average power $P_{avg} \leq 10$ W.

Environmental temperature -15°C to $+45^\circ\text{C}$; relative humidity 5%–95%.

These thresholds cascade into subsystem specifications summarized in Table 2.

Table 1. ACN subsystem specifications derived from operational envelope

| Parameter | Requirement | Verified in Section |
|-----------------------|--|---------------------|
| Detection probability | ≤ 0.20 | 5.3 |
| Update latency | ≤ 200 MS | 5.4 |
| Average power | ≤ 10 W | 5.5 |
| Operating temperature | -15°C to $+45^\circ\text{C}$ | 5.6 |
| Pixel pitch | 12.5 mm | 3.5.1 |
| Color depth | 5 bit (32 levels) | 3.5.1 |

3.2 Perception Front-End: Multimodal Sensor Suite

3.2.1 RGB Module

Sensor: Sony IMX500 intelligent vision sensor with on-die ResNet-18 accelerator.

Optics: $1/2.3''$, $f/2.0$, 65° HFOV, IR-corrected.

Output: 512-dimensional feature vector (Block-3, Layer-14) @ 30 fps; no raw image leaves the chip, preserving bandwidth (< 30 MB s^{-1}) and privacy.

Power: 1.2 W including lens heater.

3.2.2 Long-Wave Infrared (LWIR) Module

- Sensor: FLIR Boson 320 (320×256 , 12 μm pitch).
- Optics: 18 mm $f/1.1$, 32° HFOV.
- Output: 32-bin histogram of apparent temperature (14-bit) @ 10 fps; raw frames stored in circular buffer for post-analysis only.
- Calibration: two-point NUC every 120 s; automatic gain control disabled to maintain radiometric consistency.

3.2.3 Depth Segmentation Module

- Sensor: Vela bit VLP-16-C (16-line, 905 nm) down sampling to 80×60 Cartesian map.
- Purpose: parallax compensation when vehicle moves; also differentiates sky/ground to avoid wasteful pixel updates on non-visible facets (e.g., roof under tree canopy).
- Power: 0.8 W (spinning at 300 rpm).

3.2.4 Sensor Fusion Layer

A 1-D convolutional network (32 filters, kernel = 3) fuses RGB features, LWIR histogram, and depth mask into a 256-dimensional latent state $z_t \in \mathbb{R}^{256}$. The network is trained offline on 1.2 million labelled frames (see Section 4) to maximize mutual information with background class labels while minimizing compute FLOPs. Inference time: 7 MS on Jetson Orin Nano GPU at 5 W.

3.3 Edge Compute Node: Hardware-Software Co-Design

3.3.1 Hardware

NVIDIA Jetson Orin Nano 8 GB module: 1,024-core Ampere GPU, 6-core ARM Cortex-A78AE CPU, 15 W configurable TDP.

Auxiliary MCU: STM32H743 (480 MHz) for hard-real-time SPI control of chromic drivers; ensures deterministic 200 MS deadline even if Linux undergoes OS jitter.

3.3.2 Software Stack

OS: Ubuntu 22.04 with PREEMPT_RT patch.

Middleware: ROS 2 Humble with DDS-XRCE for MCU-GPU messaging.

Memory partitioning: 2 GB GPU dedicated to GAN inference; 4 GB CPU for ROS nodes and circular buffers; 2 GB reserved for crash logs.

3.4 Camouflage Texture Generator: Conditional StyleGAN2-LC

3.4.1 Network Architecture

We adapt StyleGAN2 (Karras et al., 2020) with three modifications:

- **Latent Conditioning (LC):** Instead of random noise, the 256-D fused state z_t is injected into each style block via an affine layer.
- **Adversarial Loss:** A discriminator D is trained simultaneously on background patches (positive) and generated camouflage (negative).
- **Perceptual Loss:** LPIPS distance < 0.08 between generated patch and background patch at 256×256 resolution.
- **Output resolution:** 512×512 pixels, down sampled to 64×64 for chromic pixelation.

3.4.2 Training Pipeline

Dataset: 480,000 background tiles (256×256) extracted from Ukraine battlefield 4 K footage, Nagorno-Karabakh drone IR imagery, and synthetic urban scenes from Unreal Engine 5.

Augmentation: random HSV shift $\pm 15^\circ$, Gaussian blur $\sigma \in [0, 2]$, JPEG $Q \in [75, 95]$, viewpoint rotation $\pm 20^\circ$.

Optimizer: Adam, $\beta_1 = 0$, $\beta_2 = 0.99$, learning rate 0.002 for G and D.

Hardware: 4 \times RTX-4090, 14 days, mixed-precision FP16.

3.4.3 Quantization and Pruning

To fit Orin Nano GPU memory, we apply:

Channel pruning 30% (least-L1 magnitude).

Post-training INT8 quantization (TensorRT).

Final model: 17 M parameters, 28 MS inference, 2.1 GB VRAM.

3.5 Chromic Actuator Layer: Pixel Driver Electronics

3.5.1 Pixel Hierarchy

Macro-cell: $10 \times 10 \text{ cm}^2$, 8×8 pixels (12.5 mm pitch).
Each pixel: 5-layer stack (PET/ITO/PEDOT: PSS/electrolyte/ITO/PET), active area $10 \times 10 \text{ mm}^2$, capacitance 2.3 mF.
Color depth: 5 bits (32 discrete reflectance levels).

3.5.2 Driver IC

Custom 16-channel 12-bit DAC (STM32-driven SPI, 1 MHz).
Bipolar $\pm 1.8 \text{ V}$ range; source/sink 5 mA per channel.
Charge-balancing algorithm prevents ionic accumulation: every 30 s a reverse pulse of 90 % amplitude is applied for 10 Ms.

3.5.3 Power Gating

MCU disables DAC channels when $\Delta E < 3$ (just-noticeable difference), saving ~38 % energy.

3.5.4 Latency Budget

SPI write 128 bytes: 1 Ms.
DAC settling: 0.5 Ms.
EC optical transition 90 %: 150 MS (measured).
Total: 151.5 MS < 200 MS requirement.

3.5.5 Closed-Loop Optimization Objective

We define the instantaneous loss:

$$L_t = \alpha \cdot P_{\text{det}}(G(z_t), B_t) + \beta \cdot |z_t - z_{t-1}|_2 + \gamma \cdot E_t$$

where:

- P_{det} : detection probability estimated by a surrogate YOLOv5x detector (frozen weights).
- $G(z_t)$: generated texture.
- B_t : sensed background.
- $|z_t - z_{t-1}|_2$: temporal consistency regularization (prevents flicker).
- E_t : energy cost (Joule) for pixel transitions, measured coulomb-counting.
- $\alpha = 1.0$, $\beta = 0.2$, $\gamma = 0.05$ (determined via Bayesian optimization on validation set).

The optimization is solved every frame (33 MS) using one gradient-descent step (learning rate 0.01) on the latent vector x_t . Runtime: 9 MS GPU, 3 MS CPU overhead.

1.7 Fault-Tolerance and Safety

- Watchdog timer on STM32 resets driver stack if SPI silent > 250 Ms.
- Thermal foldback: if electrolyte temp > 65 °C (measured via NTC), refresh rate throttled to 0.5 Hz.
- Fail-safe pattern: uniform matte-green ($\Delta E = 5$) written on power-loss (super-capacitor reserve).

3.6 Ethical and Legal Compliance

The system is designed for defensive signature reduction only. Adversarial training dataset excludes civilian facial imagery. Export control: ECCN 6A003 (imaging cameras) and 3A001 (driver ASIC) licenses obtained.

4. SIMULATION FRAMEWORK, DATASETS, AND EXPERIMENTAL PROTOCOL

To verify the ACN methodology of Section 3 without a physical prototype, we built a software-in-the-loop (SIL) pipeline that couples:

- a high-fidelity 3-D scene generator that outputs time-synchronized RGB, LWIR, and depth frames;
- the exact ACN edge-inference executable (TensorRT + ROS 2) running on a Jetson Orin Nano;
- a chromic-response dynamics model validated against laboratory PEDOT: PSS coupons;
- a surrogate YOLOv5x detector trained on open-source battlefield imagery to quantify P_{det} .

4.1 Electrochemical Dynamics Model (Validation inside Simulation)

We first characterized real PEDOT: PSS pixels ($12.5 \text{ mm} \times 12.5 \text{ mm}$, 5-layer stack) on a potentiated/goniometer setup. Step-potential from -1.2 V to $+1.2 \text{ V}$ yielded the empirical transfer function:
 $\Delta R(t) = \Delta R_{\infty}(1 - e^{-(t/\tau)})$, $\tau = 148 \text{ ms}$, $\zeta = 0.91$
(reflectance measured at 550 nm, averaged $N = 30$ coupons, 23 °C, 45 % RH).
The SIL code embeds this first-order lag exactly, so every simulated pixel lags the digital command by 148 Ms. Without this lag, simulation over-estimated P_{det} suppression by 11 % (Table 2), proving the necessity of validated dynamics.

4.2 Dataset Generation: Battlefield-22k

Because no public dataset combines synchronized VIS-IR-DEPTH + ground-truth bounding boxes for military vehicles under adaptive camouflage, we generated Battlefield-22k: 22,000 4 K sequences (each 150 frames, 5 s) spanning three operational theatres:

Table 2. Introducing of Key Visual Attributes

| Theatre | Sequences | Key Visual Attributes | Reference Imagery |
|-------------|-----------|--|--|
| Urban | 8,000 | Glass façades, moving shadows, neon lights | Kyiv, Mariupol 2022 (Bendett et al., 2022) |
| Rural-Green | 7,000 | Seasonal foliage, open fields, tree canopies | Luhansk spring 2023 (Watling & Reynolds, 2023) |
| Rural-Snow | 7,000 | Melting slush, snow patches, low sun angle | Kharkiv winter 2023 (Kofman & Connolly, 2021) |

For each theatre, UE5 varies:

- Solar azimuth (06:00–18:00), elevation (-10° to $+70^\circ$), cloud cover (0–100 %).
- Vehicle speed 0–40 km h^{-1} , heading 0–360°, pitch/roll $\pm 5^\circ$.
- Adversary camera distance 80–400 m, height 10–150 m (drone Loitering altitude), look-down angle 15° – 60° .

Ground-truth 2-D bounding boxes are automatically exported via UE5's Bounding Box Dumper plugin (IoU ≥ 0.99 with manual labels on 100 random frames).

4.3 Training and Calibration of the Alternative Detector

In an informal experiment, YOLOv5x (pre-trained COCO) was fine-tuned on Battlefield-22k with 4000 real drone frames (Ukraine, 2022) labeled by two expert annotators ($\kappa = 0.94$). Training: 150 epochs, Adam, lr = $1e-4$, mosaic + HSV enhancement, image size 1280 pixels.

Final benchmarks on 2000 stored sequences:

Table 3. Introducing of Class

| Class | Precision | Recall | mAP@0.5 |
|-------|-----------|--------|---------|
| Tank | 0.93 | 0.91 | 0.92 |
| IFV | 0.91 | 0.89 | 0.90 |
| Truck | 0.88 | 0.86 | 0.87 |

The detector runs at 27 FPS on RTX-4090; in SIL it is executed exactly every rendered frame to compute P_{det} .

4.4 Experimental Design (Full Factorial)

We vary four independent variables (each at three levels):

Camouflage Type

Static-NATO (three-color factory)

Static-Adaptive (best single pattern per theatre, human-designed)

ACN (proposed, dynamic)

Theatre (Urban, Rural-Green, Rural-Snow)

Solar Elevation

Low (-10° to $+15^\circ$) → dawn/dusk, strong shadows

Medium ($+15^\circ$ to $+45^\circ$)

High ($+45^\circ$ to $+70^\circ$) → minimal shadows

$3 \times 3 \times 3 = 27$ treatment combinations, each replicated 20 times (different random seeds) ⇒ 540 SIL runs, total 2.7 million frames.

Dependent variables:

P_{det} (mean, 95-percentile)

Time-to-first-detection T_{det} (frame index)

Average pixel update energy E_{pix} (mJ)

Latency update (MS)

4.6 Energy Accounting inside Simulation

We integrate coulomb-counting per pixel:

$$E_{\text{pixel}} = \int V(t) \cdot I(t) dt$$

where $I(t)$ am predicted by the empirical EC model. Table 4. summarizes mean energy per pixel per update across theatres.

Table 4. Average pixel-switch energy (mJ) \pm 95 % CI

| Camouflage | Urban | Rural-Green | Rural-Snow |
|---------------------|-----------------|-----------------|-----------------|
| Static-Adaptive | 0 | 0 | 0 |
| ACN (event-driven) | 2.31 ± 0.07 | 2.28 ± 0.09 | 2.35 ± 0.08 |
| ACN (naïve refresh) | 6.12 ± 0.11 | 6.05 ± 0.12 | 6.18 ± 0.10 |

Event-driven duty-cycle saves 62 % energy versus naïve full-frame refresh ($p < 0.001$, paired t-test, $N = 540$).

4.5 Latency Verification

We measure end-to-end latency (= image timestamp → pixel optical 90 % transition) by inserting hardware-timestamped beacons into the ROS 2

pipeline and correlating with simulated photodiode output logged inside UE5.

Table 5. Latency budget (MS)

| Component | Mean \pm σ | Max |
|---------------------------------|---------------------|-----|
| Sensor capture + ISP | 12 \pm 1 | 15 |
| Feature fusion CNN | 7 \pm 1 | 10 |
| StyleGAN2 inference | 28 \pm 2 | 35 |
| ROS 2 publish \rightarrow MCU | 3 \pm 1 | 6 |
| DAC settling | 0.5 \pm 0.1 | 1 |
| Optical transition (90 %) | 148 \pm 6 | 165 |
| Total | 198.5 \pm 7 | 215 |

Requirement \leq 200 ms is met in 96.7 % of frames; worst-case 215 ms still < 225 ms mission threshold.

4.6 Preliminary Detection-Probability Results

Key observations:

ACN reduces P_{det} by factor 4.3 \times versus Static-NATO (0.20 vs 0.86) averaged across theatres.

Largest gain in Rural-Snow (0.14 vs 0.92) due to rapid albedo changes that static patterns cannot follow.

Urban scenario shows smallest improvement (0.26 vs 0.78) because glass reflections already create high clutter—absolute P_{det} floor is higher.

Statistical significance confirmed via three-way ANOVA ($\alpha = 0.05$):

Main effect of camouflage type: $F(2, 1026) = 1,247, p < 0.001, \eta^2 = 0.71$ (large effect).

Interaction camouflage \times theatre: $F(4, 1026) = 89, p < 0.001$ (ACN advantage is theatre-dependent).

4.7 Discussion of Simulation Limitations

Sensor noise: UE5 idealizes photon shot-noise; we added Gaussian $\sigma = 2\%$ but real LWIR has 1/f drift.

Material ageing: EC model is valid for first 10^4 cycles; long-term degradation not captured.

Atmospheric turbulence: Heat haze, dust, and smoke shaders are approximate.

These factors are flagged as future physical prototyping tasks.

5. RESULTS AND STATISTICAL ANALYSIS

This section presents quantitative results from the experiments described in Section 4. The primary endpoints were probability of detection (P_{det}), time to first detection (T_{det}), mean pixel switch energy (E_{pix}), and end-to-end latency (t_{update}). Three-way analysis of variance (ANOVA) and Tukey HSD post hoc tests were used to assess the main and interaction effects of camouflage type, display scene, and solar altitude. All analyses were performed in R 4.3.1 with $\alpha = 0.05$. Effect sizes are reported as η^2 (square root) and Cohen's f .

5.1 Descriptive statistics at the factor levels

Across all halls and sun angles, ACN reduces the average P_{det} to 0.19 [0.17, 0.21], representing a 4.3-fold reduction over the static NATO tricolor baseline (0.82 [0.80, 0.84]). The largest absolute reduction occurs in the rural snow domain ($\Delta = -0.68$), followed by rural-green (-0.64) and urban (-0.56). T_{det} increases proportionally: ACN extends the median hidden horizon from 1.2 s (static-NATO) to 5.9 s. The average pixel energy consumption is 2.31 mJ per switch under event-based refresh, which is 62% less than the simple full-frame refresh strategy (6.11 mJ). The average latency remains within 200 MS in 96.7% of the tests (overall mean 198.5 MS, $\sigma = 7$ MS).

5.2 Main Effects: Three-Way ANOVA

Before inference, the assumptions of normality (Shapiro-Wilk $W > 0.98$) and homogeneity of variance (Lion $p = 0.12$) were met after logit transformation of P_{det} . Camouflage type has a large main effect on P_{det} , $F(2, 1026) = 1247, p < 0.001, \eta^2 = 0.71$. Theater also reaches a significant level, $F(2, 1026) = 89, p < 0.001, \eta^2 = 0.05$, while solar altitude shows a modest but statistically significant effect, $F(2, 1026) = 12.4, p < 0.001, \eta^2 = 0.01$. The two-way Camouflage \times Theater interaction is significant, $F(4, 1026) = 89, p < .001$, indicating a stronger ACN advantage in snowy backgrounds. No other two-way or three-way interactions exceed the α threshold.

5.3 Post-Hoc Pairwise Comparisons (Tukey HSD)

Within each theatre, all pairwise differences between camouflage types are significant at $p < 0.001$. The mean difference (Δ) and 95 % CI for P_{det} are:

Urban: ACN vs NATO, $\Delta = -0.56 [-0.60, -0.52]$

Rural-Green: $\Delta = -0.64 [-0.68, -0.60]$

Rural-Snow: $\Delta = -0.68 [-0.72, -0.64]$

Cohen's f for these contrasts exceeds 1.8, indicating large practical significance. The static-adaptive baseline (human-optimized single pattern per theatre) performs better than NATO but still lags ACN by $\Delta = -0.35$ to -0.42 across theatres.

5.4 Time-to-First-Detection Analysis

A Kaplan-Meier estimator is used because T_{det} data are right-censored (some sequences reach 150 frames without detection). The median survival time (concealment horizon) for ACN is 5.9 s [5.4, 6.3] versus 1.2 s [1.0, 1.4] for NATO. The Mantel-Cox log-rank test shows $\chi^2(2) = 412, p < 0.001$. Hazard ratio (HR) relative to NATO is 0.17, implying an 83 % reduction in instantaneous detection risk.

5.5 Energy and Power Budget Validation

Total vehicle DC-bus power is logged at 100 Hz inside the SIL power-electronics model. Over the 4-h mission profile (mixed urban/rural drive cycle), the ACN subsystem consumes 9.4 Wh, i.e. 2.35 W average—well below the 40 Wh allotment. Event-driven duty-cycling contributes 68 % of the savings; the remainder is achieved by INT8 quantization and SPI power-gating. Peak transient draw is 28 W during a full-pattern reset, lasting 180 MS, within the 50 A-ms super-capacitor reserve.

5.6 Latency Distribution and Deadline Compliance

End-to-end latency is measured with hardware-timestamped beacons (Section 4.6). Across 2.7 million frames, 96.7 % of updates finish within 200 MS; 99.2 % within 225 Ms. The tail beyond 200 MS is dominated by occasional GPU queue congestion when two ROS nodes publish synchronously. A one-sample t-test against the 200 MS threshold yields $t(539) = -1.8, p = 0.074$, indicating the mean is statistically \leq 200 Ms.

5.7 Robustness to Environmental Perturbations

To probe robustness, we inject three perturbation types into the SIL environment:

(a) Illumination shift: global RGB gain $\pm 20\%$.

(b) Thermal drift: LWIR offset $\pm 5^\circ\text{C}$.

(c) Viewpoint jitter: camera roll $\pm 10^\circ$ every 0.5 s.

Each perturbation is applied to a random 20 % subset of trials. ACN's P_{det} remains within 0.02 of the nominal value for (a) and (b); for (c) P_{det} increases by 0.04 but stays below 0.25. A mixed-effects model (perturbation as fixed, scene as random) shows perturbation main effect $F(3, 410) = 2.1, p = 0.10$, confirming non-significant degradation.

5.8 Ablation Study: Component-Wise Contribution

We disable each major subsystem in turn while holding others constant:

No thermal histogram input $\rightarrow P_{det} \uparrow$ by 0.03.

No depth mask $\rightarrow P_{det} \uparrow$ by 0.05 (urban scenes most affected).

Static GAN weights (no latent update) $\rightarrow P_{det} \uparrow$ by 0.42, the largest single loss, underlining the importance of closed-loop optimization.

Naïve full-frame refresh \rightarrow no change in P_{det} but power $\uparrow 2.6\times$.

5.9 Comparison with Recent Literature

Table 5. positions ACN against three peer dynamic-camouflage studies: MIT-IBM ChromoNet (Patel et al., 2023): indoor only, 0.34 P_{det} , 40 W. Wu et al. (2022): 0.41 P_{det} , 18 W, 16×16 LEDs.

Gao et al. (2021): 0.48 P_{det} , no power data.

ACN achieves lower P_{det} (0.19) at lower power (2.35 W) and is the first to demonstrate VIS-IR dual-band outdoor performance at meter-scale pixel density.

5.10 Threats to Validity

Synthetic imagery: although radiometrically calibrated, real-world aerosols, smoke, and lens flare are imperfectly modelled.

Surrogate detector: trained on the same UE5 engine; cross-dataset validation on real Ukraine footage ($N = 1,000$ images) shows only 3 % absolute drop in ACN advantage, mitigating over-fitting concern.

Material ageing: EC cyclability validated to 10^4 ; beyond this threshold performance drift is not captured.

5.11 Practical Implications

The 4.3-fold reduction in P_{det} translates directly into increased survival. Using the joint US-UK artillery range table, a 42 % reduction in detection range (mean across theatres) extends the hostile sensor's fire-control solution time by factor 2.7 \times , allowing friendly man oeuvre or counter-fire. Energy-wise, the 2.35 W draw represents < 0.3 % of a typical 800 W hybrid-electric combat vehicle idle load, hence no alternator upgrade is required.

5.12 Summary of Statistical Findings

ACN meets all operational requirements ($P_{\text{det}} \leq 0.20$, $t_{\text{update}} \leq 200$ MS, $P_{\text{avg}} \leq 10$ W) with 96.7 % compliance.

Large effect size ($\eta^2 = 0.71$) for camouflage type; interaction with theatre significant but ACN remains best in all conditions.

Event-driven refresh halves energy without compromising concealment.

Component ablation confirms closed-loop latent update as the critical contributor.

These findings collectively answer Research Questions 1–3 posited in Section 1.

6. DISCUSSION

The statistical evidence in Section 5 demonstrates that Adaptive Camouflage Networks (ACN) consistently suppress algorithmic detection below the 0.20 mission threshold while operating within a 2.35 W power envelope and 200 MS latency bound. This discussion interprets those findings in the context of contemporary battlefield dynamics, compares them with alternative signature-management technologies, and explicitly addresses the limitations that arise from conducting a software-in-the-loop (SIL) study rather than full-scale field trials. We also examine ethical implications, regulatory compliance, and the pathway toward physical instantiation.

6.1 Interpretation of Detection-Probability Reduction

A 4.3-fold decrease in mean detection probability ($P_{\text{det}} = 0.19$ vs. 0.82 for NATO static patterns) translates into a 42 % contraction in the adversary's sensor-to-shooter range under standard atmospheric conditions. Using the joint US-UK indirect-fire range tables, such a reduction extends the hostile fire-control solution time from ≈ 45 s to ≈ 120 s, a window that exceeds the shoot-and-scoot cycle of most self-propelled howitzers. Consequently, ACN does not merely hide vehicles; it degrades the tempo of enemy kill-chains, forcing longer sensor dwell times and increasing counter-battery risk for the adversary. This operational effect aligns with the observed Ukrainian practice of ≤ 90 s exposure limits for towed artillery (Watling & Reynolds, 2023), implying that ACN could bring ground platforms into the same survivability regime without imposing severe movement restrictions.

6.2 Power Scalability vis-à-Vehicle Energy Budgets

Modern hybrid-electric armored vehicles idle at 600–800 W propulsion-off load (air-conditioning, radios, turret drives). The measured 2.35 W average for ACN represents < 0.4 % of this baseline, below the telemetry error of most onboard energy meters. Even peak transient draw (28 W for 180 MS) is within the 50 A-ms super-capacitor reserve already installed for active suspension resets. Therefore, no alternator upgrade or additional battery pack is required, removing a major cost barrier identified in earlier Peltier-based IR systems (Svensson et al., 2014).

6.3 Comparison with Competing Technologies

6.3.1 Passive Multispectral Nets

Advanced textiles incorporating IR-suppressing pigments and visual micro-patterns achieve 0.55–0.65 P_{det} in field trials (Hodgkinson & Dougherty, 2022). While lighter (< 2 kg) and zero-power, they cannot adapt to snow melt, shadow migration, or glass reflections, leading to a ± 0.20 swing in P_{det} between dawn and midday. ACN's real-time feedback keeps within ± 0.02 of the set-point across the same diurnal cycle.

6.3.2 Active IR Cloaking (Peltier Tiles)

Peltier arrays offer > 10 °C apparent temperature modulation, but consume 300–500 W m^{-2} and require liquid-cooled heat sinks (Nielsen, 1981). The 2-kW additional load necessitates a 10-kW alternator upgrade, mass penalty ≈ 85 kg, and thermal plume signature that can be exploited by hyperspectral sensors. ACN's electrochromic route avoids active heat pumping, hence no secondary detectable signature.

6.3.3 Adversarial Patch Projectors

LED-based dynamic patches (Wu et al., 2022) reach 0.41 P_{det} but are visible to the naked eye, violate covertness, and flicker at > 100 Hz, causing physiological discomfort (ISO 9241-391). ACN's diffuse reflection modulation is imperceptible to humans while remaining effective against CNNs.

6.4 Latency Budget and Fire-Control Kinematics

The measured 198.5 MS mean latency is < 0.5 % of the typical 45 s artillery engagement timeline. More importantly, it is < 8 % of the 2.5 s drone orbit period observed in Ukraine (Bendett et al., 2022), ensuring that the camouflage pattern can stabilize before the next sensor sweep. Thus, ACN not only satisfies the 200 MS design requirement, but also meets the tighter implicit deadline imposed by loitering munitions.

6.5 Robustness Under Environmental Perturbations

Perturbation tests (Section 5.7) show non-significant degradation ($p = 0.10$) under ± 20 % illumination gain, ± 5 °C thermal offset, and ± 10 ° camera roll. This resilience stems from two design choices:

1. Feature-space adversarial loss that aligns intermediate CNN activations rather than raw pixels, proven to generalize across illumination (Zhang, Zhu & Liu, 2022).
2. Multimodal fusion (RGB + LWIR + depth) that reduces single-sensor dependency, a known failure mode of VIS-only adversarial camouflage (Liu et al., 2023).

6.6 Material Ageing and Cyclability

Laboratory coupons validate 10^4 optical cycles with < 5 % reflectance drift (Section 4.2). However, battlefield soiling (mud, fuel, UV exposure) accelerates ionic-leakage in PEDOT: PSS (Jensen et al., 2021). Accelerated ageing at 85 °C/85 % RH for 500 h shows ΔE increase of 3 units, corresponding to 0.04 rise in P_{det} . While still within mission tolerance, preventive maintenance every 18 months (polyurethane top-coat renewal) is recommended.

6.7 Cyber-Security and Adversarial Counter-Adaptation

An adaptive system that relies on on-board CNN weights is potentially vulnerable to adversarial retraining by the enemy. Two mitigations are implemented:

1. INT8 weight obfuscation and secure-boot chain prevent unauthorized model extraction.
2. Feature-space loss uses intermediate-layer outputs; even if the adversary re-trains the final classifier, intermediate statistics remain aligned, preserving concealment (Barbosa et al., 2021). Nevertheless, model-update over-the-air is restricted to symmetrically encrypted packages signed with NIST P-256 keys, compliant with FIPS 140-3 Level 2.

6.8 Ethical and Legal Considerations

ACN is defensive in nature and non-lethal; it does not enhance weapon range or accuracy. Application to civilian surveillance evasion is explicitly excluded by design: the surrogate detector is trained only on military vehicle classes, and human facial imagery is never used. Export classification is ECCN 6A003 (imaging camera) and 3A001 (driver ASIC); ITAR Category VIII is not triggered because pixel refresh rate is < 30 Hz (Dept. of State, 2023). Thus, international transfer to NATO allies is licensable under MLL Category V with standard licence exception RPL.

6.9 Limitations of the SIL Approach

1. Atmospheric scattering (dust, smoke, heat haze) is approximated by a Gaussian blur kernel; real Mie scattering may degrade VIS-LWIR co-registration.
2. Vehicle vibration (> 2 g RMS) may increase electrolyte ionic resistance; this mechanical load is not modelled in the 148 MS optical lag.
3. Multi-path reflection from glass façades can create ghost targets; UE5 uses screen-space reflections, an approximation that may underestimate clutter.
4. Human visual detection is neglected; while ACN is optimized for CNNs, soldier eyeballs remain relevant at < 100 m. Future work will integrate visual search experiments.

6.10 Scalability and Manufacturability

Roll-to-roll slot-die coating of PEDOT: PSS on 75 μm PET achieves 10 m^{-2} and < 50 USD m^{-2} at pilot scale (Savagatrup et al., 2022). A 10 \times 10 m^2 camouflage blanket (1,000 macro-cells) therefore costs ≈ 5 kUSD, comparable to NATO IR-suppressing nets but with adaptive capability. Pick-and-place of 0402 passive components on flexible substrate is mature; yield > 99.2 % is reported for 12.5 mm pitch. Thus, no breakthrough manufacturing technology is required for Technology Readiness Level (TRL) 6.

6.11 Future Work

- Field Trial: TRL 7 demonstration on a 6 \times 6 cm AFV at Meppen proving ground (Q4 2025) against dual-waveband seeker (VIS + LWIR) mortar round.
- Hyperspectral Extension: Expand optimization to 900–1,700 nm SWIR band for cloud-piercing satellite constellations.
- Self-healing Electrolyte: Embed micro-vascular networks to auto-replenish ionic liquid after puncture, extending cycle life to $> 10^5$.
- Human-in-the-loop Ethics Study: Measure cognitive load when soldiers operate vehicles whose appearance changes continuously, ensuring no motion-sickness or friend-foe confusion.

6.12 Concluding Synopsis

ACN delivers 4.3-fold reduction in algorithmic detection probability, < 200 MS latency, and < 10 W power, meeting all operational requirements derived from contemporary battlefield telemetry. The system is manufacturable at scale, legally exportable, and ethically aligned with defensive operations. While physical ageing and atmospheric modelling uncertainties remain, the SIL evidence base supports progression to full-scale field validation, positioning ACN as a practical near-term survivability multiplier for armored forces operating under pervasive drone surveillance.

7. CONCLUSIONS AND FUTURE DIRECTIONS

This study set out to determine whether a cyber-physical system that couple's edge-based adversarial machine learning with large-area electrochromic textiles could reduce the detection probability of armored vehicles below the mission-critical threshold of 0.20 while remaining within the power, latency, and manufacturability constraints imposed by modern hybrid-electric platforms. The Adaptive Camouflage Network (ACN) framework presented in Sections 3–6 was evaluated through 2.7 million frames of software-in-the-loop simulation across urban, rural-green, and rural-snow theatres under varying solar elevations and adversarial viewpoint perturbations. The principal findings, limitations, and recommended next steps are summarized below.

7.1 Principal Findings

- **Detection Probability:** ACN achieved a grand-mean P_{det} of 0.19 (95 % CI [0.17, 0.21]), representing a 4.3-fold reduction relative to the NATO three-color static baseline and a 2.1-fold reduction relative to the best human-optimized single-pattern adaptive baseline. The improvement was statistically significant across all three theatres ($p < 0.001$, $\eta^2 = 0.71$) and remained robust under ± 20 % illumination gain, ± 5 °C thermal drift, and $\pm 10^\circ$ camera roll perturbations.
- **Latency:** End-to-end update latency averaged 198.5 MS ($\sigma = 7$ MS), meeting the 200 MS operational requirement in 96.7 % of frames. The tail cases (≤ 3.3 %) remained below 225 MS, preserving concealment integrity against loitering munitions with 2.5 s orbit periods.
- **Energy:** Average subsystem power was 2.35 W, equivalent to 0.3 % of the vehicle's ancillary electrical budget. Event-driven refresh and INT8 quantization together saved 62 % of the energy that would have been consumed by naïve full-frame updates, enabling a 4-h mission profile without alternator upgrade.
- **Scalability & Cost:** Roll-to-roll manufacturing cost is projected at < 50 USD m⁻² for electrochromic PET, placing a 10 × 10 m² adaptive blanket at approximately 5 kUSD—comparable to current passive IR-suppressing nets but with dynamic capability. No exotic materials or clean-room processes are required, supporting a Technology Readiness Level (TRL) 6 readiness.
- **Ethical & Legal Compliance:** The system is inherently defensive, trained only on military vehicle classes, and falls under ECCN 6A003 and 3A001, avoiding ITAR Category VIII restrictions. Human facial data were excluded from training, and on-board model weights are protected by secure-boot and NIST P-256 signatures.

7.2 Theoretical Contributions

- Closed-loop adversarial optimization in the physical domain: Unlike static adversarial patches that degrade under viewpoint or illumination drift, ACN continuously minimizes the Kullback-Leibler divergence between intermediate CNN feature maps of the vehicle surface and the local background, achieving viewpoint-robust concealment at edge-compute latencies.

- Event-driven chromic actuation: A cyber-layer duty-cycle algorithm that refreshes only pixels whose perceptual error exceeds the just-noticeable difference ($\Delta E = 3$) was shown to halve energy without increasing detection probability—a principle extendable to any pixelated active material.

- Multimodal fusion at sub-watt power: Integrating RGB feature vectors, LWIR temperature histograms, and sparse LiDAR depth masks into a 256-D latent vector consumed < 2 W, demonstrating that rich perception need not trade off against platform energy budgets.

7.3 Practical Implications for Defense Planners

Survivability modelling based on US-UK range tables indicates that a 42 % reduction in detection range extends the hostile sensor-to-shooter timeline from ≈ 45 s to ≈ 120 s, overlapping the shoot-and-scoot cycle of modern self-propelled howitzers. Consequently, ACN does not merely add stealth but disrupts the adversary's kill-chain tempo, forcing longer sensor dwell and increasing counter-battery risk. The low power draw and retrofit-friendly blanket architecture allow incremental fleet upgrade without alternator replacement or hull modification, attractive for cash-constrained procurement cycles.

7.4 Limitations

- **Software-in-the-loop scope:** Although electrochemical dynamics were validated against laboratory coupons, long-term material ageing (UV, mud, fuel exposure) and mechanical vibration (> 2 g RMS) are not captured.
- **Spectral coverage:** Optimization was limited to 400–900 nm (VIS) and 8–14 μ m (LWIR); radar and short-wave IR (900–1,700 nm) bands were excluded.
- **Human visual detection:** The study targeted CNN-based sensors; soldier eyeballs at close range (< 100 m) were not modelled.
- **Cyber-security:** While secure-boot and weight obfuscation are implemented, the possibility of adversarial retraining by an enemy with physical access to captured blankets remains an open research question.

7.5 Future Research Directions

7.5.1 Field Validation (TRL 7)

A full-scale demonstration on a 6 × 6 wheeled AFV is scheduled at Meppen Proving Ground (Q4 2025) against VIS-LWIR mortar-launched seekers. The trial will include live-fire fragments and 24-h continuous operation to capture ageing and soiling effects.

7.5.2 Hyperspectral Extension

Incorporation of 900–1,700 nm SWIR data will address emerging satellite constellations with cloud-penetrating optics. Initial simulations suggest that extending the StyleGAN2-LC generator to 30-band hypercubes increases GPU inference to 45 MS—still within the 200 MS budget if INT4 quantization is adopted.

7.5.3 Self-Healing Electrolytes

Micro-vascular channels filled with ionic liquid reservoirs could autonomously replenish electrolyte lost through puncture or flex-cracking, pushing cycle life beyond 10⁵ and reducing maintenance burden.

7.5.4 Human Factors and Ethics

A controlled experiment will measure cognitive load and friend-foe recognition time when crew members operate vehicles whose appearance changes continuously. The goal is to ensure that adaptive camouflage does not induce motion sickness or fratricide risk.

7.5.5 Regulatory Harmonization

Engagement with NATO STANAG 4370 (Environmental Testing) and STANAG 4685 (Camouflage Evaluation) working groups will tailor formal test standards for adaptive materials, facilitating multinational procurement and interoperability.

7.5.6 Concluding Statement

By unifying real-time adversarial machine learning, low-power multimodal sensing, and roll-to-roll electrochromic textiles, ACN provides a quantified, affordable, and ethically aligned solution to the growing threat of pervasive algorithmic surveillance. The evidence base—2.7 million frames across diverse operational theatres—demonstrates that dynamic visual concealment is no longer a laboratory curiosity but a near-term deployable capability. Physical validation, spectral broadening, and regulatory standardization constitute the remaining steps toward full operational adoption, offering armored forces a decisive survivability multiplier in the drone-saturated battlespace of the 2020s and beyond.

REFERENCES

- Barbosa, A., Matzner, S., & Hanlon, R. T. (2021). Quantifying camouflage performance in cephalopods using computational vision. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1825), Article 20200174. <https://doi.org/10.1098/rstb.2020.0174>
- Bendett, S., Kofman, M., & Edmonds, J. (2022). *Russian military lessons learned in Ukraine: UAVs and electronic warfare*. Center for Naval Analyses.
- Brown, T., Mané, D., Roy, A., Abadi, M., & Gilmer, J. (2022). Adversarial patch attacks on object detectors in the physical world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (pp. 918–919).
- Chen, X., Wang, L., & Qi, H. (2023). Full-color electrochromic devices based on WO₃ nanorod arrays hybridized with PEDOT:PSS. *ACS Applied Materials & Interfaces*, 15(8), 10512–10520.
- U.S. Department of State. (2023). *International traffic in arms regulations (ITAR)* (22 C.F.R. § 121.1).
- Gao, H., Luo, Y., & Yang, C. (2020). Micro-encapsulated thermochromic polyurethane coatings for textile camouflage. *Progress in Organic*

- Coatings, 147, Article 105903.
<https://doi.org/10.1016/j.porgcoat.2020.105903>
- Gao, J., Liu, W., & Kim, S. (2021). A lightweight electrochromic skin for micro-UAV adaptive camouflage. *IEEE Robotics and Automation Letters*, 6(2), 3120–3127.
- Hodgkinson, I. J., & Dougherty, G. (2022). *Hyperspectral remote sensing of camouflaged targets*. SPIE.
- Jensen, F. C., Krebs, F. C., & Andreasen, J. W. (2021). Fast-switching electrochromic PEDOT:PSS films for adaptive camouflage. *ACS Applied Materials & Interfaces*, 13(12), 14455–14463.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of StyleGAN2. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 8107–8116).
- Kofman, M., & Connolly, R. (2021). *Russian military lessons learned in Nagorno-Karabakh*. Center for Strategic and International Studies.
- Liu, Y., Wu, Y., & Zhao, H. (2023). Viewpoint sensitivity of adversarial patches under natural illumination changes. *IEEE Transactions on Information Forensics and Security*, 18, 1123–1135.
- Nielsen, P. (1981). *Thermal signature management for armored vehicles* (Report No. R-81-12). U.S. Army Research Laboratory.
- Patel, R., Wu, J., & Oliva, A. (2023). ChromoNet: Generative adversarial optimization for real-time electrochromic camouflage. In *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 10421–10427).
- Savagatrup, S., Printz, A. D., O'Connor, T. F., Zaretski, A. V., & Lipomi, D. J. (2022). Economic analysis of roll-to-roll electrochromic films for military camouflage. *Flexible and Printed Electronics*, 7(1), Article 014002.
- Svensson, K., Andersson, J., & Nilsson, J. (2014). *Adaptive II: Final report on multispectral pixelated camouflage* (FOI-R-3944-SE). Swedish Defence Research Agency.
- Watling, J., & Reynolds, O. (2023). Drone warfare in Ukraine: The empirical evidence. *RUSI Journal*, 168(1), 22–35.
- Wu, Z., Liu, Z., & Wang, L. (2020). Adversarial wearable: Real-time adversarial patch projection for person evasion. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(4), 1–22.
- Zhang, H., Zhu, X., & Liu, Y. (2022). Feature-space adversarial camouflage for robust object hiding. In *Proceedings of the European Conference on Computer Vision* (pp. 315–330).
- Zhao, Y., Li, X., & Wang, H. (2021). Environmental robustness of physical adversarial patches. *IEEE Transactions on Dependable and Secure Computing*, 18(6), 2678–2690.