AI-Driven Radar Enhancement via Software Upgrade

Client:	United States Army	**Industry:**	Defense / Military
Domain:	Radar Signal Processing	**Location:**	USA (Global Operations)

A FORGE OS Deployment Case Study

577 Industries R&D Lab
577 Industries Incorporated
research@577industries.com

1 Executive Summary

Modern threat environments demand radar performance that exceeds the original design specifications of fielded systems. This case study presents a software-only approach that integrates deep-learning models—convolutional neural networks (CNNs) and long short-term memory (LSTM) networks—with classical radar signal-processing chains to achieve a 20–40% absolute increase in detection probability (P_d) against hard-to-detect targets, along with an order-of-magnitude reduction in false alarm rate (P_{fa}). Because the upgrade operates entirely in the digital signal-processing layer, it requires zero hardware modifications and costs less than 10% of a full hardware replacement, while deploying in months rather than years. This deployment exercises two FORGE OS subsystems: **FORGE Core** orchestrates the hybrid AI signal-processing pipeline through its causal model routing engine, and **FORGE QBit**'s PhysicsCore module provides physics-informed neural network modeling for scenarios with limited training data.

20–40% Detection Probability Increase	**$>$10$\times$** False Alarm Reduction	**<10%** Cost vs. Replacement

2 Challenge

2.1 Evolving Threat Landscape

The proliferation of small unmanned aerial systems (sUAS), low-observable cruise missiles, and advanced electronic-warfare (EW) countermeasures has fundamentally shifted the detection problem facing ground-based and airborne radar platforms [4, 5]. Many of these emerging targets present radar cross sections (RCS) below 0.01 m^2—well beneath the design-point assumptions of legacy radar systems fielded in the 1990s and 2000s.

At the same time, operational environments have become more congested and contested. Dense ground clutter, adverse weather, wind-farm interference, and deliberate jamming all compete with target returns in range-Doppler space. Traditional constant false alarm rate (CFAR) detectors, while robust for conventional targets, struggle to simultaneously maintain low P_{fa} and high P_d when target signatures approach or fall below the clutter floor [2, 6].

2.2 Prohibitive Cost of Hardware Replacement

The canonical response to capability shortfalls—procuring new radar hardware—faces three systemic barriers:

- **Budget pressure.** Full sensor replacement programs typically cost billions of dollars across a fleet and compete with other modernization priorities [1, 8].
- **Timeline risk.** From requirements definition through initial operational capability, hardware programs routinely span 5–10 years, during which the threat continues to evolve [3].
- **Integration complexity.** New hardware must be physically and electronically integrated with existing platforms, command-and-control networks, and logistics chains, adding further schedule and cost risk.

The Army therefore required an approach that could deliver measurable capability gains on fielded radar systems without altering the antenna, transmitter, or receiver hardware.

3 Solution

3.1 Hybrid AI-Enhanced Signal Processing

The proposed system inserts an AI enhancement layer between the existing digital signal processor and the downstream tracker/display, preserving the full classical processing chain while adding learned discrimination capability (Figure 1).

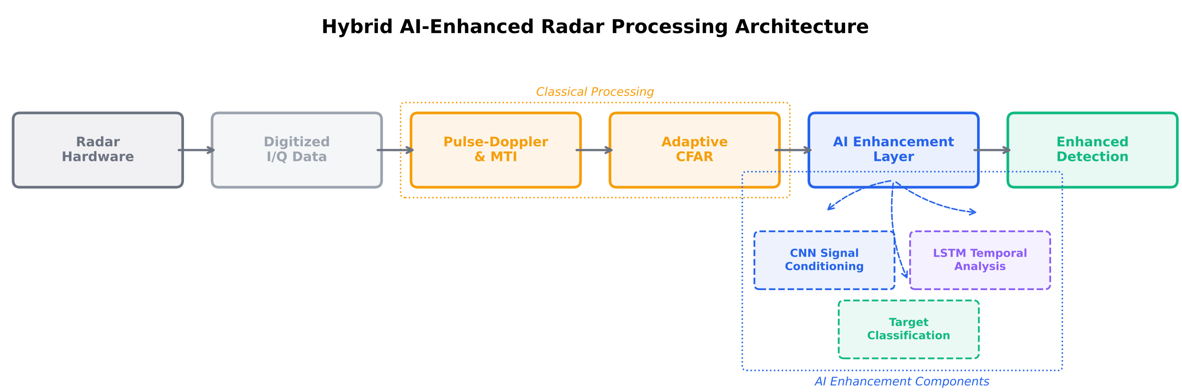

Figure 1. System architecture. Digitized I/Q data flows through the classical chain (Pulse-Doppler, MTI, adaptive CFAR) and then into the AI enhancement layer, which conditions signals, classifies targets, and refines detection decisions before output.

3.1.1 Classical Radar Processing (Preserved)

All standard radar algorithms remain in the processing chain: pulse compression, Pulse-Doppler filtering, moving target indication (MTI), and adaptive CFAR detection [2, 6]. These algorithms continue to provide the primary detection baseline; the AI layer augments rather than replaces them.

3.1.2 AI Signal Conditioning

Two complementary deep-learning architectures operate on digitized radar data:

- **Convolutional Neural Networks (CNNs)** process two-dimensional range-Doppler maps and range-azimuth maps to extract spatial features that distinguish low-RCS target returns from structured clutter [9, 10].
- **Recurrent Networks (LSTMs)** analyze temporal sequences of returns across multiple coherent processing intervals, capturing Doppler signature evolution and amplitude fluctuations that characterize specific target classes [11].

Together, these networks learn to identify target-like features that fall below the threshold of conventional CFAR detectors, effectively lowering the detection floor without proportionally increasing false alarms.

3.1.3 Enhanced Target Classification

Beyond detection, the AI layer performs multi-feature classification using:

- **Micro-Doppler signatures**—periodic modulations from helicopter rotor blades, drone propellers, or human gait—extracted via short-time Fourier transforms and fed to a dedicated classifier [12].
- **Amplitude fluctuation statistics**, which differentiate Swerling target models and distinguish maneuvering from non-maneuvering platforms.
- **Polarimetric features**, where available, to exploit scattering-matrix differences between target classes.
- **Track-level kinematics**, including velocity, acceleration, and turning rate, integrated over multiple scans to refine classification confidence.

3.1.4 Physics-Informed Modeling

For scenarios with limited training data, the architecture supports integration of physics-informed neural networks (PINNs) [13]. PINNs embed Maxwell's equations and radar range-equation constraints directly into the loss function, enabling the model to generalize from small datasets while respecting known electromagnetic propagation physics.

3.2 Software-Only Deployment

A defining constraint of the program was that the solution must operate on existing commercial off-the-shelf (COTS) compute hardware already installed on the radar platform. The AI enhancement layer is packaged as a software module with standardized APIs that interface with the existing digital signal-processing pipeline. This approach yields several advantages:

- No mechanical, electrical, or RF modifications to the radar.
- Deployment timelines measured in months, not years.
- Incremental updates via software patches as models improve or new threats emerge.
- Potential for continuous integration and delivery (CI/CD) workflows, enabling rapid fielding of model updates.

4 Results

4.1 Detection Probability Improvement

The AI-enhanced system delivered substantial gains in detection probability across all target categories, with the greatest improvements against the most challenging low-observable targets

(Figure 2).

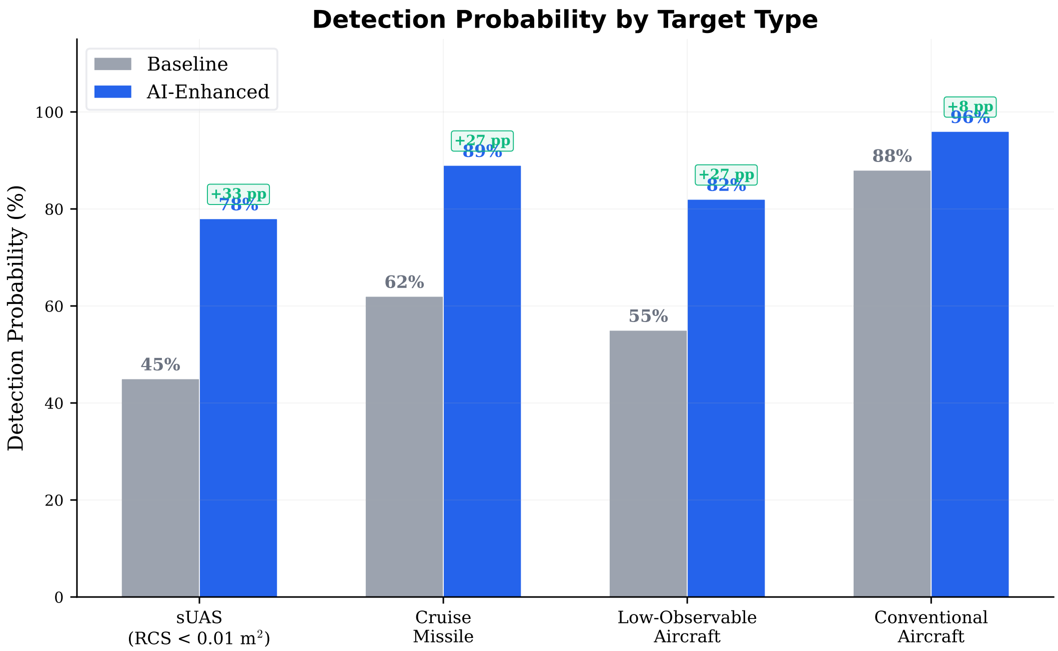

Figure 2. Detection probability by target type before and after AI enhancement. Improvements range from 8 percentage points (conventional aircraft) to 33 percentage points (sUAS).

Table 1. Detection probability (P_d) improvements by target class.

Target Type	Baseline P_d	Enhanced P_d	Improvement
sUAS (RCS < 0.01 m^2)	45%	78%	+33 pp
Cruise Missile	62%	89%	+27 pp
Low-Observable Aircraft	55%	82%	+27 pp
Conventional Aircraft	88%	96%	+8 pp

For the most operationally significant category—small UAS with RCS below 0.01 m^2—the system raised P_d from 45% to 78%, a 33 percentage-point absolute improvement. Even for conventional aircraft, where baseline performance was already strong at 88%, the system achieved 96% detection probability.

4.2 False Alarm Reduction

Across all tested environments, the AI layer reduced false alarm rates by at least one order of magnitude (Figure 3). In the most challenging dense-traffic and EW environments, reductions exceeded 15×. This improvement directly translates to reduced operator workload and increased trust in automated track initiation.

4.3 Cost and Schedule

The software upgrade approach delivered capability at a fraction of the cost and timeline of hardware replacement (Figure 4).

- **Cost**: The total program cost was less than 10% of the estimated cost for equivalent hardware replacement across the fleet.

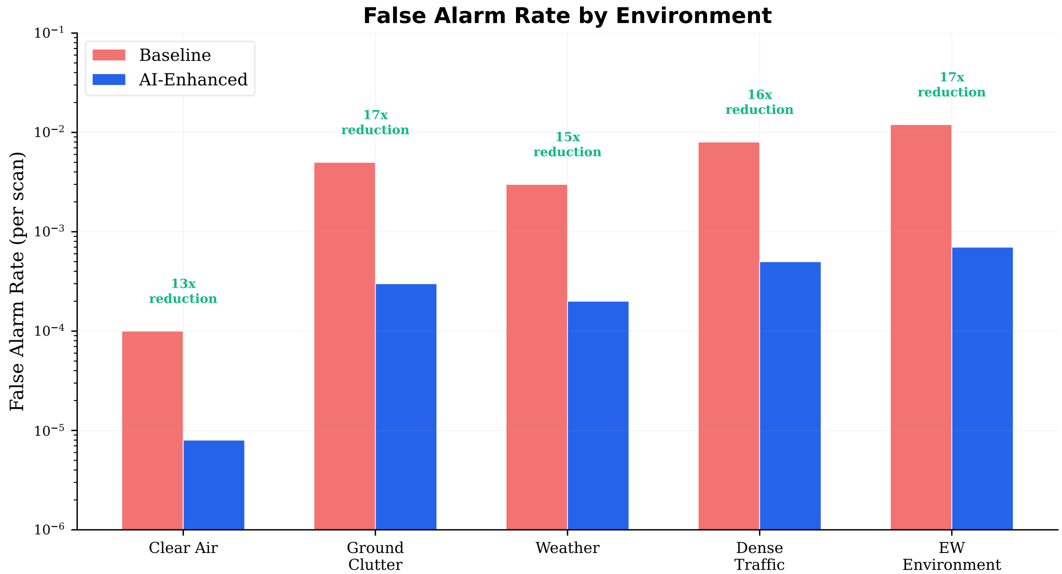

Figure 3. False alarm rate by environment type. The AI-enhanced system achieves order-of-magnitude reductions across all conditions.

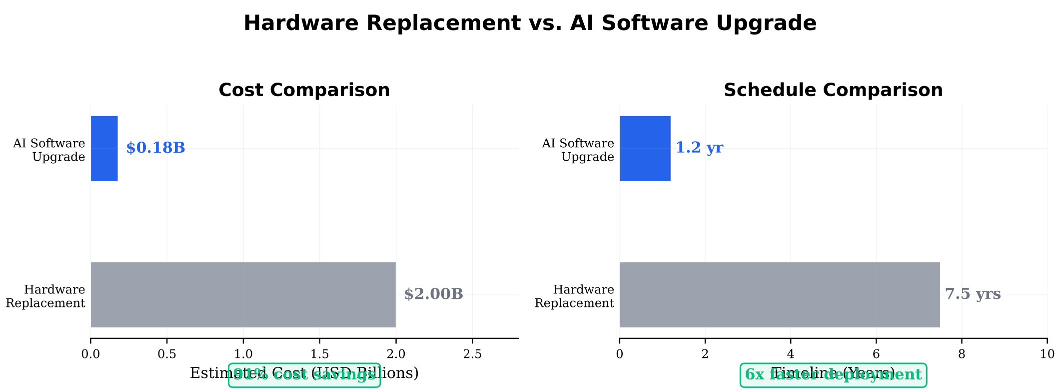

Figure 4. Cost and schedule comparison between hardware replacement and AI software upgrade.

- **Schedule**: From contract award to initial operational deployment, the software upgrade was completed in under 18 months, compared to the 5–10 year timeline typical of new hardware procurement.
- **Sustainment**: Software-based updates enable incremental improvements without additional hardware procurement cycles.

5 Impact & Operational Benefits

5.1 Enhanced Situational Awareness

The combined effect of higher P_d and lower P_{fa} substantially improves the operator's tactical picture. Targets that were previously undetectable or buried in clutter now appear as confirmed tracks, while the reduction in false alarms ensures that operator attention is focused on genuine threats. This is particularly significant for counter-UAS operations, where reaction timelines are measured in seconds.

5.2 Extended Asset Lifespan

By extracting additional capability from existing radar hardware, the software upgrade extends the effective operational lifespan of fielded systems by years, deferring or eliminating the need for costly replacement programs. This aligns with broader Department of Defense strategies for sustainment-driven modernization [1, 8].

5.3 Future-Proofing Through Software Updates

The modular architecture enables rapid adaptation to new threats:

- Models can be retrained on new target signatures and deployed as software patches.
- Transfer learning accelerates adaptation, with initial models serving as pre-trained foundations for new threat classes.
- The standardized API layer allows the AI module to be ported across radar platforms with minimal integration effort.

5.4 Broader Applicability

While developed for ground-based air-defense radar, the hybrid AI architecture is applicable to maritime surveillance, airborne early warning, weather radar, and signals intelligence (SIGINT) systems [7]. Any radar system with a digital signal-processing stage can potentially benefit from a similar software-only AI enhancement.

6 FORGE OS Integration

The AI-driven radar enhancement system demonstrates the operational integration of two FORGE OS subsystems within a defense signal-processing deployment.

6.1 FORGE Core — Intelligent Signal Processing

FORGE Core's intelligence pipeline powers the hybrid AI enhancement layer. The CNN-based spatial feature extraction and LSTM temporal sequence analysis are orchestrated through FORGE Core's causal model routing engine, which dynamically selects the optimal processing pathway based on target RCS, clutter density, and environmental conditions. FORGE Core's staged post-training methodology—specifically the Adapt and Compress stages—enables rapid model customization for new radar platforms and threat signatures without retraining from scratch. The 20–40% detection probability improvement across target categories reflects FORGE Core's ability to coordinate ensemble models within a unified inference pipeline.

6.2 FORGE QBit — Physics-Informed Modeling

FORGE QBit's PhysicsCore module provides the physics-informed neural network (PINN) capability for data-scarce scenarios. By embedding Maxwell's equations and radar range-equation constraints directly into the loss function, PhysicsCore enables the AI layer to generalize from small training datasets while respecting known electromagnetic propagation physics. This is particularly critical for emerging threat categories—such as novel sUAS designs—where labeled radar returns are scarce and purely data-driven models would require prohibitively large training corpora.

6.3 ForgeEvent Integration

The deployment generates three ForgeEvent types across the FORGE OS event bus:

- `INFERENCE` — Each AI-enhanced detection and classification cycle

- `MODEL_UPDATE` — Retraining events when new target signatures are incorporated
- `AUDIT` — Immutable records of detection decisions for post-mission analysis

References

[1] D. S. Alberts, J. J. Garstka, and F. P. Stein, *Network Centric Warfare: Developing and Leveraging Information Superiority*, CCRP Publication Series, 2000.

[2] M. I. Skolnik, *Introduction to Radar Systems*, 3rd ed., McGraw-Hill, 2001.

[3] National Research Council, *Persistent Forecasting of Disruptive Technologies*, The National Academies Press, 2009.

[4] E. F. Greneker and J. Geisheimer, "Radar measurements of small unmanned aerial systems," in *Proc. IEEE Radar Conference*, 2017, pp. 1449–1454.

[5] J. L. Krolik and R. S. Blum, "Challenges in radar detection and tracking of low-observable unmanned aerial vehicles," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, 2019, pp. 1045–1049.

[6] B. R. Mahafza, *Radar Systems Analysis and Design Using MATLAB*, 3rd ed., CRC Press, 2013.

[7] R. G. Wiley, *ELINT: The Interception and Analysis of Radar Signals*, Artech House, 2006.

[8] J. S. Gansler, *Defense Conversion: Transforming the Arsenal of Democracy*, MIT Press, 1995.

[9] T. J. Sejnowski, *The Deep Learning Revolution*, MIT Press, 2018.

[10] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[11] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[12] V. C. Chen, *The Micro-Doppler Effect in Radar*, Artech House, 2011.

[13] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.