Generative AI Technologies for Electronic Warfare Applications

Dr. Mohiuddin "Mohin" Ahmed, P.E.

mohin@hrl.com

HRL Laboratories, LLC

Malibu, CA

AOC International Symposium & Convention

National Harbor, MD

December 10, 2025

HRL | LABORATORIES

Generative AI for EMS: From Adaptive to Generative EW for Dominance by 2035

- **EW Evolution: f**rom Adaptive (library-based reaction) to Cognitive (learning-based classification) and now to Generative (creation of novel effects).
  - Unlike traditional AI (classifies signals), Generative AI creates new data distributions.
  - synthesizing novel waveforms, jamming strategies, and protocols that have never been seen before.
- **The Strategic Necessity:** Adversaries use SDRs and Cognitive Radars that change signatures faster than we can update libraries
  - Cannot “classify” a signal that you haven’t seen before
- **The 2035 Strategic Goal:** from systems that "look up" answers to systems that "invent" solutions in microseconds.
- **Key Enabler:** Convergence of massive compute at the edge and new architectures like Transformers and Diffusion Models, moving us from "fighting the last war's emitter" to "countering the next second's pulse."

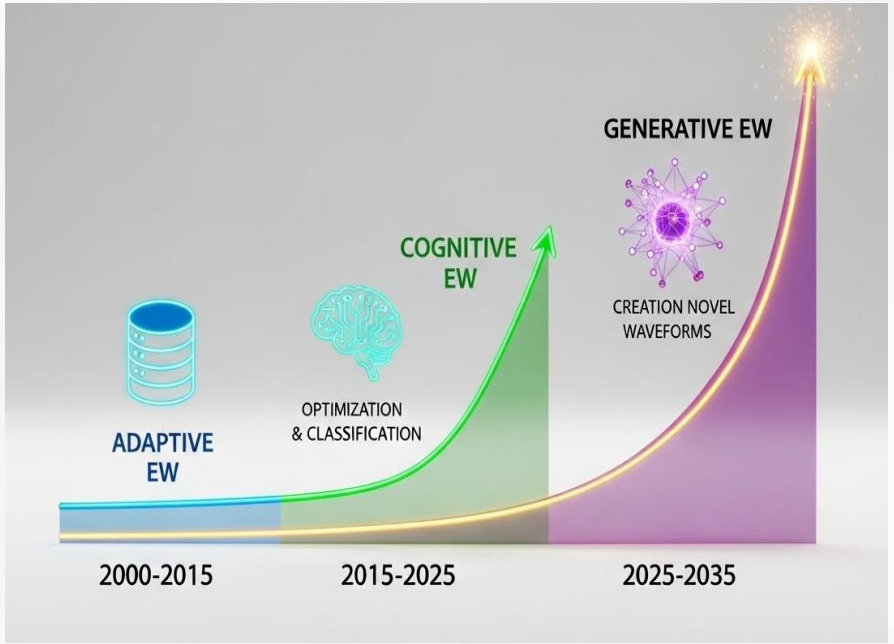

From Adaptive to Generative: The Strategic Paradigm Shift

The Strategic Imperative: Move Beyond the "Library Problem" in a Software-Defined Conflict

- **The Core Bottleneck:** Current EW systems rely on static databases (Pulse Descriptor Word (PDW) libraries for radar and Protocol Dictionaries for communications). *This approach collapses when adversary is software-defined!*
  - E.g. Agile Multi-Function Radars (MFRs) generate novel waveforms that simply do not exist in any threat library
  - Novel comms can use proprietary, non-standard waveforms (e.g., custom 5G slices) where no pre-recorded demodulation key or "signature" exists. Novel LPI/LPD techniques keep comm signals below noise floor
  - “Hyper-agility” - threats change their signature per pulse (radar) or per packet (comms) so library-based system is always one update cycle behind
- **The 2035 Vision:** Must move from Recognition (matching a known signal) to Characterization (understanding an unknown signal).
  - Predict the next pulse
  - Reverse-engineer the protocol grammar on the fly to spoof or jam comms intelligently.

From Adaptive to Generative: The Strategic Paradigm Shift

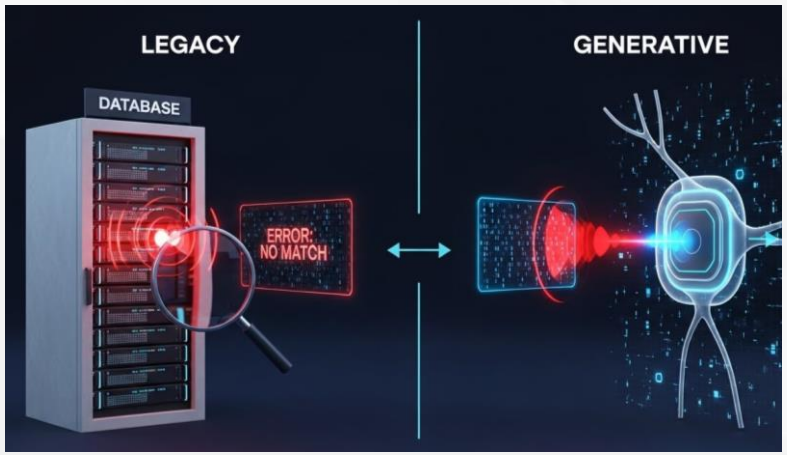

Gen AI Use Cases: Asymmetric Advantages in Near-Peer Conflict

1. Zero-Day Jamming & Protocol Spoofing (Waveform Synthesis)

– Problem: Novel LPI waveforms & unknown proprietary drone protocols.

– Solution (GANs): Analyzes signal distribution; learns the "grammar" of the threat.

– Effect: Synthesizes "smart noise" (Radar) or valid spoofed packets (Comms) in real-time—no library required.

2. Cognitive Prediction (Spectrum Forecasting)

– Problem: Reactive jamming is too slow for hyper-agile frequency hoppers.

– Solution (Transformers): Treats pulses/hops like words; predicts the next move.

– Effect: Anticipatory Jamming—striking the target frequency before the enemy arrives.

3. The Training Data Revolution (Synthetic Data)

– Problem: "Data Starvation"—lack of labeled data for classified/rare threats.

– Solution (Diffusion Models): "Hallucinates" millions of high-fidelity synthetic IQ spectrograms.

– Effect: Trains robust classifiers on Day 0 without needing real-world intercepts.

GenAI Primer (1): Transformers & Attention — Predicting the Next Move

- **Core Concept:** Transformers rely on a mechanism called Self-Attention.
  - Instead of reading data sequentially (left-to-right), the model looks at the entire sequence at once, assigning a "weight" or importance to every element's relationship with every other element.
- **The EW Translation: R**eplace text "tokens" with Pulse Descriptor Words (PDWs)—Frequency, Pulse Width, and PRI, other parameters as relevant. For comms, tokenization is I/Q symbols, or packet bytes (headers, sync bits)
- **Operational Application:**
  - Deinterleaving: Disentangles mixed pulse trains from multiple emitters in dense environments.
  - Protocol Inference: Learns the "grammar" of unknown protocols; predicts next header byte or symbol without a spec sheet.
  - Sequence Prediction: By analyzing the "grammar" of a radar's previous 1,000 pulses, or comm waveform's emissions, the model predicts the parameters of the next pulse, enabling jamming capabilities that are predictive rather than reactive.

Transformers: The "Attention" Mechanism Applied to Pulse Trains and Waveforms

LLM's and Gen AI Implementation

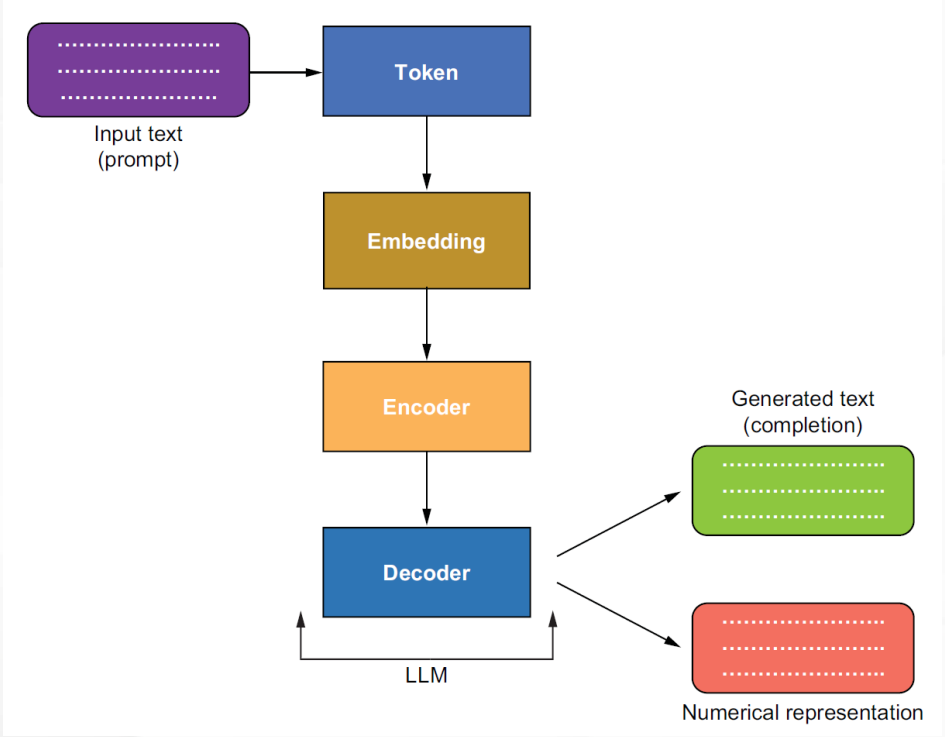

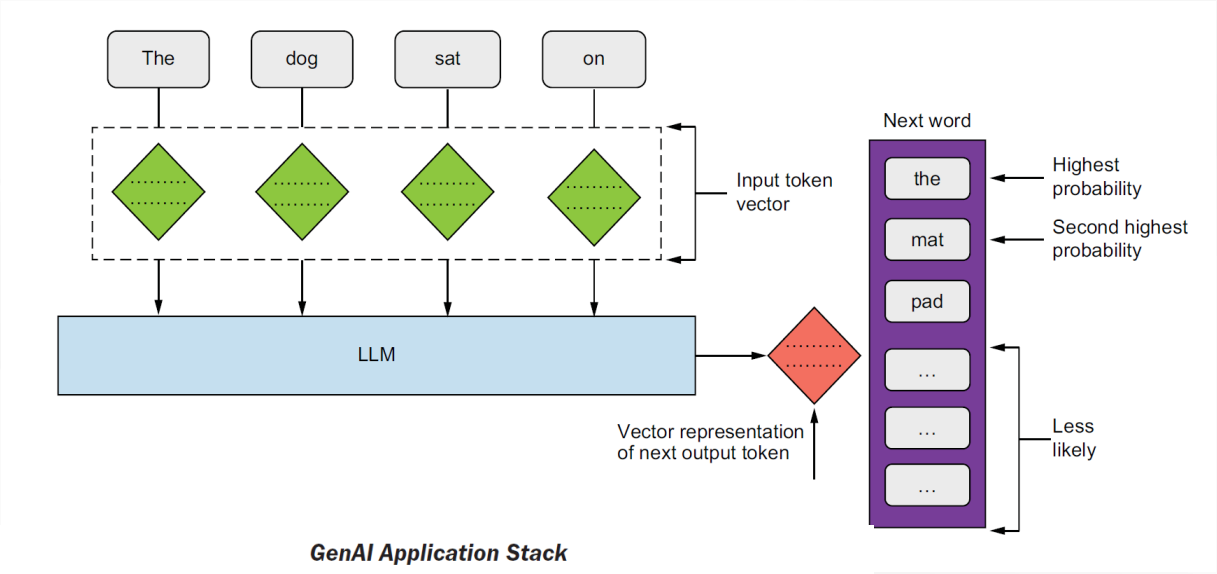

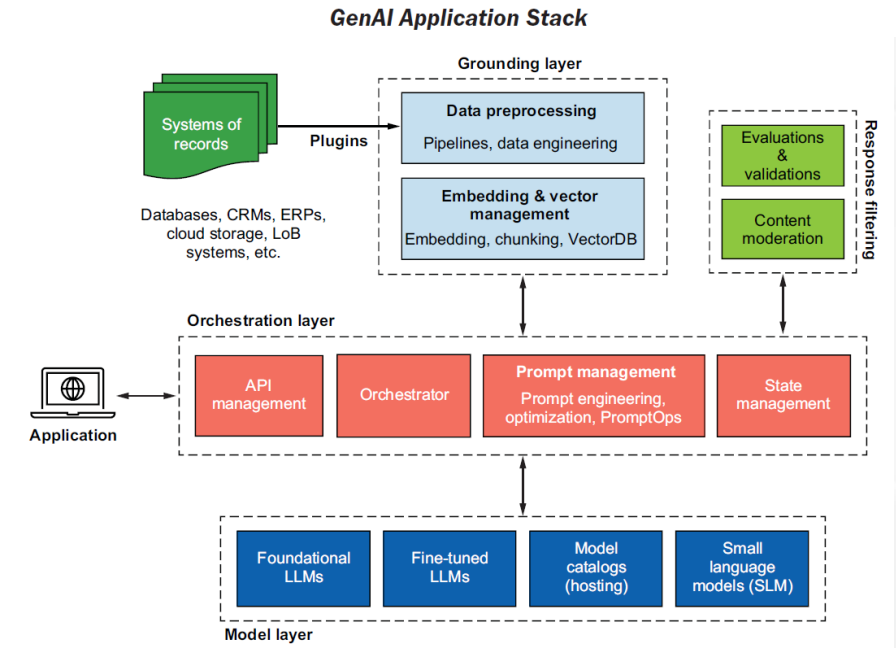

Courtesy: Bahree, et al

GenAI Primer (2): Generative Adversarial Networks (GAN) — The Ultimate Cat & Mouse Game

- **Core Concept:** Automating the jammer-radar / covert comms arms race – Generators vs. Discriminators.
  - Consists of two dueling neural networks: **Generator** tries to create fake data (the "forger"), and **Discriminator** tries to detect it (the "police"). They improve by competing against each other.
- **The EW Translation:**
  - Radar application: Generator = The Jammer. Goal is to synthesize waveform that looks like valid skin return. Discriminator = The Enemy Radar. Goal is to filter out jamming signals using CFAR (Constant False Alarm Rate) logic.
  - Comms application: Generator creates “deepfake packets” e.g. handshake/auth headers to spoof secure network. Discriminators train against such scenarios
- **Operational Application:** Train GAN against surrogate of the enemy radar or comm EA system
  - Over thousands of iterations, Generator learns to create "Smart Noise" or deceptive targets or spook packets that are mathematically optimized to bypass the specific filters of radars or comm systems.

GANs for EW: Automating the SIGINT, ELINT, COMINT Arms Race

GenAI Primer (3): Diffusion Models — Recovering Signal from the Noise Floor

- **Core Concept:** Diffusion models (e.g. Stable Diffusion) work by destroying training data with noise (Forward Process) and then learning to reverse that process to recover the clean data (Reverse Process).
- **The EW Translation:** Signal Recovery
  - Radar: Treat RF spectrograms (Time-Frequency plots) as images. Recover LPI echoes buried in clutter via diffusion model trained on "clean" emitter signatures
  - Comms: Denoise direct sequence spread spectrum signals
- Operational Outcome:
  - Diffusion model reconstructs coherent signal structure by extracting faint signals from background noise where traditional energy detectors would fail.
  - Detection sensitivity beyond physical limits of energy detectors

Diffusion: "Hallucinating" the Signal Back into Existence

Morphing Technology (1): RF as Language — Tokenizing Pulse Descriptor Words (PDWs)

- **The Translation Layer:** Large Language Models (LLMs) ingest text tokens. For EW, we "tokenize" the RF spectrum
- **The Vocabulary:**
  - Word = Pulse. Each pulse is defined by a vector: {Frequency, Pulse Width (PW), Amplitude, Time of Arrival (TOA)}.
  - Sentence = Pulse Train. A sequence of these pulses forms a "sentence" that describes the radar's behavior over time.
  - For comms: demod or packet bytes for ({Header, Payload, Sync}).
- **The Generative Task:** Auto-Complete
  - Input: sequence of 50 pulses or packets into a Transformer Model (e.g. Signal Former)
  - Output: Model "auto-completes" the sequence, predicting the next likely PDW (e.g., the next frequency hop or protocol command) with high probability
  - Result: We jam the future, not the present (e.g. the frequency before the radar ev

Tokenization: Turning Waves into Words
If signals are a language, we can auto-complete the sentence

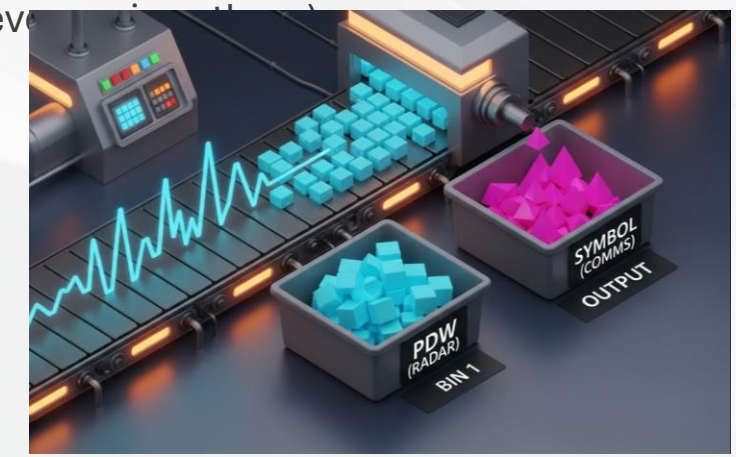

Morphing Technology (2): RF as Vision — Applying Computer Vision to I/Q Spectrograms

- **The Visual Analogy:** Complex RF signals (like modern 5G or LPI radar) can be converted In-Phase/Quadrature (I/Q) data into Spectrograms (Time vs. Frequency plots), effectively turning a radio problem into a computer vision problem.
- **Adapting the Architecture:**
  - In visual domain, can use Vision Transformers (ViTs) or Diffusion Models (like those used for AI art) to process the "image" of the RF spectrum.
  - Scenario: Adversary signal is partially jammed or corrupted by noise (occluded).
  - Solution: Mask corrupted section of the spectrogram. Generative model uses "Image In-Painting" to hallucinate the missing pieces of the signal based on the surrounding context, reconstructing full enemy command structure from fragments.
- Benefit: Allows use of commercial AI tools (e.g. Stable Diffusion) for RF problems

Spectograms: Treating Spectrum as an Image Generation Problem

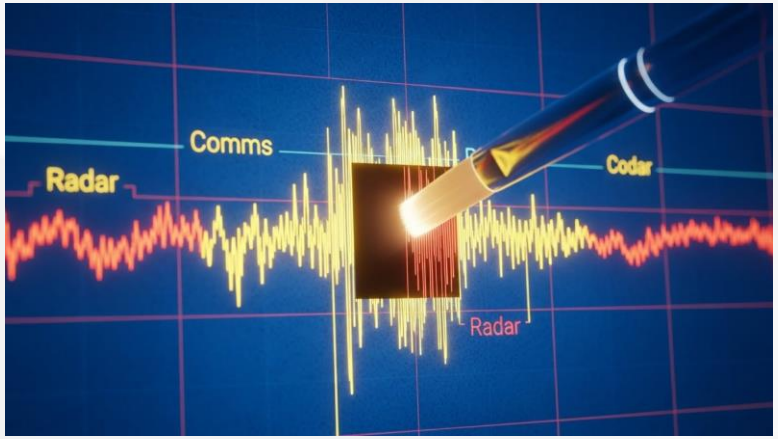

Operational Scenarios: Predictive Dominance

- **Scenario A: “Transcendental” Jamming**
  - Digital Radio Frequency Memory (DRFM) jammers are reactive (receives a pulse, delays and retransmits it). Latency reveals it is a fake against advanced "leading-edge" tracking filters.
  - GenAI Solution: GAN predicts the next microsecond of adversary waveform (trained on radar’s micro-doppler signature)
  - Effect: jamming pulse arrive simultaneously with skin return, or even before peak of the real pulse (predictive lead), thus adversary cannot distinguish the jammer from the target using timing analysis. We achieve "coherent confusion" where the jammer lives inside the victim's processing loop, not chasing it.
- **Scenario B: Anticipatory Link Severing (Comms)**
  - *Action:* Transformer predicts the next frequency hop in a pseudo-random sequence.
  - *Effect:* The jammer is waiting at the target frequency *before* the enemy drone attempts to transmit. Link lost instantly.

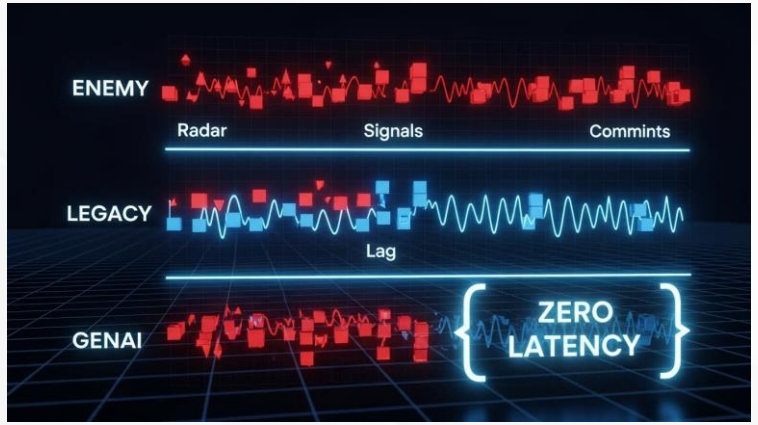

HRL | LABORATORIES

Operational End-State: Zero-Latency Effects

One Critical Enabler: Synthetic Data Generation for Rare & Classified Threat Training

- **The Bottleneck:** Cannot train reliable AI classifiers for enemy emitters (e.g., a specific Mode-5 seeking radar) because we lack sufficient intercepted data. Cannot train a Deep Learning model on 5 samples.
  - Rare threats: we have zeros samples of “wartime modes” for adversary radars
  - Elusive comms: adversary drones use proprietary encrypted links rarely seen
- **The GenAI Solution (Diffusion Models):**
  - Apply Denoising Diffusion Probabilistic Models (DDPM)—the same tech behind image generators—to RF spectrograms.
  - Process: We feed the model a handful of "Golden Intercepts" (rare, high-quality captures). The model learns the statistical manifold of that signal structure.
- **The Result:** Generate 100,000 synthetic variations of that threat. Vary the SNR, multipath fading, and antenna scan rates in the synthetic data.
- **Benefit:** Move from "few-shot" learning to massive supervised learning, training robust classifiers that work on Day 1 of the conflict without needing months of collection.

Overcome a Roadblock: the "Data Starvation" Problem in Electronic Warfare

GenAI in Action for Synthetic Data Generation

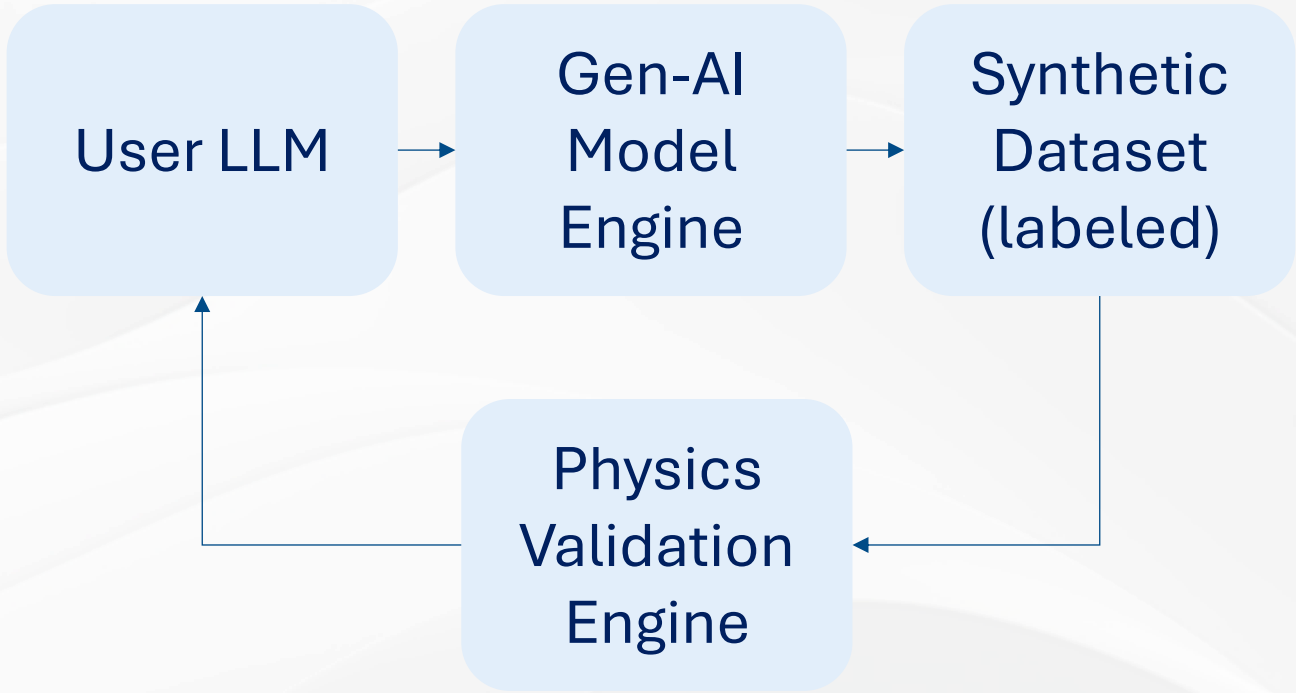

Implementation Challenges: Hallucinations, Latency, and SWaP-C at the Edge

1. **The Hallucination Risk (Fratricide):**
   - Risk: A generative jammer might "hallucinate" a waveform matching a friendly comms protocol, inadvertently jamming own network.
   - Mitigation: Guardrail Models. A lightweight, separate classifier (discriminator) sits at the output to veto any waveform that correlates with friendly crypto/waveforms before transmission.
2. **The Latency Penalty:**
   - Risk: Large Transformers (LLMs) are slow. Generating a token might take 50ms, but a radar pulse width is 1µs.
   - Mitigation: Knowledge Distillation. We train a massive "Teacher" model offline (in the cloud) and distill it into a tiny "Student" model (quantized to 4-bit) that runs on edge FPGAs or NPUs (like NVIDIA Jetson).
3. **SWaP-C (Size, Weight, Power, and Cost):**
   - Reality: We cannot put a data center on a loitering munition.
   - Research Area: Neuromorphic Computing and FPGA accelarators. Using spiking neural networks to run GenAI inference with milliwatts of power instead of kilowatts.

The "Trust Gap“: Why GenAI Isn't on the Jet… Yet

A Roadmap to Cognitive Dominance & Research Priorities

- **The Strategic Horizon:**
  - Today (2025): Offline GenAI to write code for Software Defined Radios (SDRs) and generate synthetic training data.
  - Near Term (2028): Human-on-the-Loop GenAI suggests jamming courses of action (COAs); EW Officers approve them.
  - Far Term (2035): Autonomous Generative EW. "AI-Defined Radios" that rewrite their own FPGA bitstreams in real-time to counter physics-compliant but unknown threats.
- **Call to Action / Research Priorities:**
  - Invest in RF Foundation Models (like "RadioGPT") pre-trained on massive unclassified spectrum data.
  - Develop Verifiable AI for mission-critical kill chains (solving the hallucination trust problem).
  - Focus on Edge-Native Architectures—moving the "brain" from the cloud to the sensor.

The Path to 2035: From "Software-Defined" to "AI-Defined" Warfare