AI-Driven Spectrum Monitoring and Awareness for Tactical Units

Jack R. Driscoll and Jason L. Lind
jack@groundbreaker.solutions & jason@groundbreaker.solutions
26 March 2025

We are working on topic 1 (Detect, separate from background noise, and characterize signals of interest, including identifying the communications protocol for signals within the RF communications band), our solution operates from 10MHz-6GHz. The proposed system aims to provide dismounted warfighters with on-the-spot awareness of the radio frequency (RF) environment by integrating advanced artificial intelligence (AI) and a portable software-defined radio (SDR). At present, tactical units often lack immediate spectrum intelligence because signals are recorded and sent back to specialized teams for analysis. By the time feedback returns, it can be too late for effective decision-making. This solution replaces that delay with a near-real-time capability that fits into the warfighter's standard equipment load, offering a live view of the signals in the environment as well as their estimated locations. The system harnesses deep learning for rapid signal classification, identifies various communication protocols (such as Wi-Fi, Bluetooth, LTE, and push-to-talk radios), and delivers user-friendly data overlays through the Android Tactical Assault Kit (ATAK).

One of the central innovations is applying modern neural network architectures to raw IQ or spectrally processed data. Deep learning has shown the ability to detect and classify signals in milliseconds, even when operating in crowded or noisy conditions. However, successful classification depends heavily on the breadth and quality of the training data, which must represent realistic military and commercial waveforms—some potentially noisy or intentionally jamming. By combining public and custom datasets, the proposed system becomes more robust in real-world deployments. Once trained, the neural network performs the classification task efficiently in low size, weight, and power (SWaP) environments. This efficiency is critical for a wearable or handheld system, which must run continuously with minimal battery drain and without burdening an operator who is already carrying essential combat gear.

Beyond simply identifying different signal types, the system also attempts real-time localization of the emitter. The proposed method leverages principles of electronic support, such as measuring variations in received signals or phase differences across multiple antennas. In settings where multipath reflections are significant, machine learning further refines location estimates by using patterns of signal propagation. The result is an intuitive display on the ATAK device that marks an emitter's approximate location on a map or provides directional bearings if exact positioning cannot be calculated.

Ease of use for non-expert operators is a driving principle behind the system's design. Warfighters already rely on ATAK for situational awareness, mapping, and friendly-force tracking. By presenting spectrum alerts through the familiar ATAK interface, the new functionality avoids complicated stand-alone displays. Operators can see icons indicating detected signals along with descriptions of the signal type, all supported by a confidence score. In practice, a single operator on patrol can continue normal tasks but glance down only when the device buzzes or signals the presence of a new or suspicious emitter. The aim is to ensure minimal cognitive load: known and friendly signals can be grayed out or ignored by default, allowing the soldier to focus on potentially hostile activity.

In addition to immediate battlefield benefits, the proposed solution creates a record of each signal's activity over time, forming a pattern-of-life database that identifies recurring emissions or anomalous transmissions. This feature can prove essential when a unit remains in the same operational area for extended periods and needs to detect daily rhythms in adversary radio usage. For example, if an enemy force transmits short bursts at specific hours each morning, the system logs and highlights that pattern, giving intelligence analysts or platoon leaders actionable evidence for planning.

From a hardware standpoint, the system envisions using commercial off-the-shelf multichannel SDRs, particularly DeepSig's AIR-T and Ettus X310, which support simultaneous channels for communication and scanning. One channel can maintain regular voice or data communication, while another is dedicated to the AI-driven detection process. Because the ultimate goal is to remain as close as possible in size and weight to existing fielded radios, a major Phase I focus is proving that these capabilities do not unduly compromise battery life, thermal limits, or the radio's primary function. The integration would preserve current waveforms and encryption standards, ensuring that ongoing voice transmissions remain unaltered.

Phase I is designed to confirm feasibility through a six-month series of tasks and deliverables. The plan begins with defining the system architecture and carefully selecting suitable SDR hardware, antennas, and data processing modules. That phase will include a survey of potential training datasets and initial tests of signal classification algorithms on recorded samples. The next step develops and refines the AI models, training convolutional neural networks to recognize a core set of signals commonly encountered in military or dual-use scenarios. By the midpoint of Phase I, the research team intends to demonstrate lab-based real-time classification of a handful of signal classes and produce a preliminary direction-finding simulation that estimates bearing or position.

Once the classification and localization algorithms reach sufficient maturity, they will be integrated into a prototype with an Android device running ATAK. This phase includes implementing a customized ATAK plugin that communicates bidirectionally with the SDR to retrieve raw samples or partially processed data. The plugin will receive classification results from the AI model, then display alerts on a tactical map. A short field test in a controlled environment is planned to show that a user can power on the system, walk around, and see localized signals in real time without requiring specialized RF expertise. Based on results from these tests, the design can be revised to correct issues with false alarms, missed detections, or unwieldy user interface elements.

During the final segment of Phase I, the development team will systematically measure performance against key benchmarks that include classification accuracy, time to detection, and localization error. If results fall short, the project schedule includes time for algorithmic optimizations, alternative antenna placements, or additional training data collection. A comprehensive technical report will outline these findings and detail how the system is expected to evolve in Phase II, where a more robust prototype can be tested under realistic operational conditions, potentially with a designated test unit or a specialized training environment. The report also covers interoperability considerations, such as whether the SDR's power usage impacts battery performance and how easily the technology can transition to a fully militarized radio chassis.

Beyond the military, the proposal anticipates that first responders and commercial entities could benefit from an AI-driven spectrum monitoring system. Firefighters or police could be alerted to unauthorized transmissions or interference sources at a disaster site. Cell providers, regulators, and private companies might deploy the technology to locate sources of illegal broadcasting or system interference in crowded urban centers. The software-defined nature of this platform makes it relatively easy to adapt: re-training the AI model on publicly licensed waveforms or adjusting the frequency range are straightforward modifications once the core detection engine is in place. A single, scalable architecture could address multiple market segments with only minor alterations to reflect the distinct frequencies, waveforms, and user interface preferences in each domain.

Ultimately, the main innovation of this approach lies in embedding cutting-edge AI signal classification into a portable, soldier-friendly form factor. Past signal detection techniques often required large computers, advanced expertise, and post-mission analysis to identify hostile emissions. By making those insights available instantly, right where the operator stands, the system aims to improve force protection and tactical decision-making. Whether the challenge is identifying a hidden jammer or recognizing a stray push-to-talk transmission in an urban environment, real-time intelligence means a squad can maneuver confidently, disrupt enemy communications, or report suspicious signals up the chain of command without waiting hours for offline analysis.

Concluding the proposed Phase I plan, the expected outcome is a demonstration of core feasibility. The team will have shown that the hardware, firmware, and AI algorithms can operate at sufficient speed and accuracy to be viable for tactical units. This initial demonstration also paves the way for a more advanced Phase II prototype, including higher levels of integration into standard-issue radios and refined modules for localization. If successful, the system will enter a transitional period in which military units test it in field conditions, provide feedback on usability, and shape final requirements for a potential program of record. Dual-use pathways will then encourage deployment into public safety or commercial domains. By introducing advanced AI for real-time spectrum awareness in a practical form, the proposal represents a significant step toward bridging the long-standing gap between tactical warfighters and the elusive complexity of the electromagnetic spectrum.

Performance of Artificially Degraded Data

If the CNN is trained on simulated or "clean" signals, accuracy is degraded from 95% to less than 50%!
We are training the CNN using a large variety of signals collected from a wide variety of hardware.

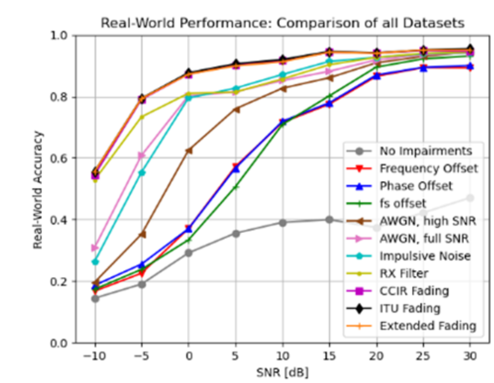

Figure 8. Comparison of the real-world performance for all eleven training datasets

"Https://Arxiv.Org/Pdf/2206.12967." Accessed March 20, 2025. https://arxiv.org/pdf/2206.12967.