Indian Journal of Computer Science and Technology
Volume 5, Issue2 (May-June 2026), PP: 81-87.
www.indjcst.com ISSN No: 2583-5300

Target Recognition in SAR Images for Military Applications

Chethan G M[1], Ganesh Vinayak Hegde[2], G R Gireesh[3], Supriya Sudhir[4]
[1,2,3]*Department of Computer Science Engineering, BNM Institute of Technology Bengaluru, Karnataka, India.*
[4]*Assistant Professor, Department of Computer Science Engineering, BNM Institute of Technology Bengaluru, Karnataka, India.*

To Cite this Article: *Chethan G M[1], Ganesh Vinayak Hegde[2], G R Gireesh[3], Supriya Sudhir[4], "Target Recognition in SAR Images for Military Applications", Indian Journal of Computer Science and Technology, Volume 05, Issue 02 (May-June 2026), PP: 81-87.*

Abstract: *Timely and precise intelligence sits at the heart of sound military decision-making. This paper presents a web-based Defense Intelligence and Asset Management System that pairs a SAR-focused target recognition pipeline with practical security and reporting features. General users can browse de-fense news freely, while administrator access is locked behind Gmail-integrated two-factor authentication that sends a one-time password directly to a registered email account before granting entry. Privileged users can upload multiple SAR images for batch analysis, manage military vehicle records through full CRUD operations, and download structured PDF reports that capture detection results, confidence scores, and auto-generated tactical notes for each identified target. At the core of the analysis layer, a YOLO model locates and classifies tanks, trucks, and ships within SAR frames. Because labeled SAR data is genuinely scarce, two generative networks — DH-GAN synthesize additional training samples that mirror authentic radar backscatter and speckle behavior, keeping the detector accurate under noisy or unusual conditions. Detected targets are logged automatically with spatial coordinates, timestamps, and class information, and the system produces short tactical assessments to assist operators in evalu-ating what each finding means for the mission at hand. Taken together, the components form a self-contained tool that handles the full path from raw SAR imagery to a shareable intelligence report.*

Key Words: *Target Recognition, YOLO, Generative Ad-versarial Networks, Deep Learning, Military Target Detection, Defense Intelligence System, Object Detection, Dataset Augmen-tation, Situational Awareness, Synthetic Aperture Radar, Gmail 2FA, PDF Report Generation.*

I. INTRODUCTION

Satellite and airborne Synthetic Aperture Radar platforms have meaningfully broadened the scope of defense surveil-lance. Working through cloud cover and darkness, these sen-sors support persistent observation of large areas in ways that optical cameras simply cannot match. The catch is that turning raw SAR data into reliable, machine-readable intelligence is harder than it looks. Speckle noise corrupts image texture, annotated ground-truth data for military vehicles is difficult to obtain in quantity, and camouflage or unusual viewing angles can confuse detectors that performed well in testing.

What field operators actually need goes a step further than detection alone. Knowing that something is present matters less without knowing what class it belongs to, how confident the model is, and what role that asset is likely to play. Those needs point toward a system rather than an algorithm something that connects detection output to reporting, access control, and data management in a way that suits real workflows.

The platform described here is designed with that gap in mind. It brings together a YOLO-based detector tuned for SAR imagery, DH-GAN augmentation to offset the shortage of training data, and a web interface that handles multi-image batch uploads, automatic logging, and downloadable PDF reports. Administrator login relies on Gmail two-factor authentication, with a one-time code dispatched to the user's registered address before the session opens. The final report includes all detection metadata plus concise tactical remarks, giving the user a ready-to-share document without any manual formatting step. The following sections describe each com-ponent and report on the experimental results obtained from testing on standard SAR benchmark data.

II.BACKGROUND AND LITERATURE REVIEW

SAR-based recognition has gone through a clear transition over the past ten years or so. Early work leaned on hand-crafted signal features and template libraries that captured the scattering profiles of known vehicle types. Those approaches carried real limitations — sensitivity to noise, narrow angular tolerance, and poor handling of occlusion — and the commu-nity progressively moved toward learned methods as labeled datasets grew and GPU hardware became affordable.

Convolutional networks removed the need to specify fea-tures manually and showed that a single architecture, trained on enough examples, could handle the kind of variability that broke earlier pipelines. Within that broad shift, single-stage detectors proved attractive for defense use because they keep inference fast. YOLO in particular has accumulated a strong track record on tasks where throughput matters, and its grid-based prediction scheme adapts well to the small, cluttered targets typical of SAR scenes. The persistent difficulty has been data: collecting and annotating real SAR imagery is ex-pensive, access to military datasets is restricted, and individual vehicle classes may appear in only a few hundred training samples. GANs addressed this by generating synthetic images with plausible radar statistics. Architectures like DH-GAN go beyond naive augmentation by modeling backscatter and speckle jointly, producing samples that genuinely strengthen the detector rather than simply inflating the dataset size.

None of this resolves the question of what happens after detection. Operational users cannot act on a bounding-box coordinate alone; they need the result packaged, interpreted, and delivered through a secure channel. The system presented in this paper treats that delivery layer as a first-class concern alongside the detection model itself.

A. Classical SAR Recognition Methods

Threshold-based detectors and template-matching pipelines dominated early SAR recognition literature. They were computationally inexpensive and interpretable, which made them practical for the hardware of the time. Their weakness was brittleness: performance dropped sharply whenever real scenes deviated from the conditions used to build the feature de-scriptors. Partial occlusion, background clutter, and off-axis viewing angles all caused trouble, and no amount of parameter tuning fully closed the gap. The inflexibility of rule-based rep-resentations eventually made a shift to data-driven approaches unavoidable.

B. Deep Learning for SAR Feature Extraction

Deep convolutional models learn directly from pixel data, building representations tuned to whatever patterns distinguish target classes in the training set. This makes them far more tolerant of the noise and appearance variation that characterize real SAR collections. YOLO variants are particularly well suited here because they produce class labels and bounding boxes in one pass, keeping latency low enough for operational use. GAN-based augmentation with architectures like DH-GAN adds synthetic diversity that covers aspect angles and noise conditions not well represented in the original data, helping the trained model hold up on genuinely new inputs.

C. Integrated Intelligence and Reporting Systems

Connecting a strong detector to a usable interface is a separate design problem that the research literature often sidesteps. The work described here treats secure access, batch processing, structured report generation, and email delivery as integral components rather than afterthoughts. Gmail two-factor authentication ties administrative access to a real iden-tity, multi-image upload supports realistic mission-scale work-loads, and PDF report export gives the user a portable record of every analysis session without requiring any additional tooling.

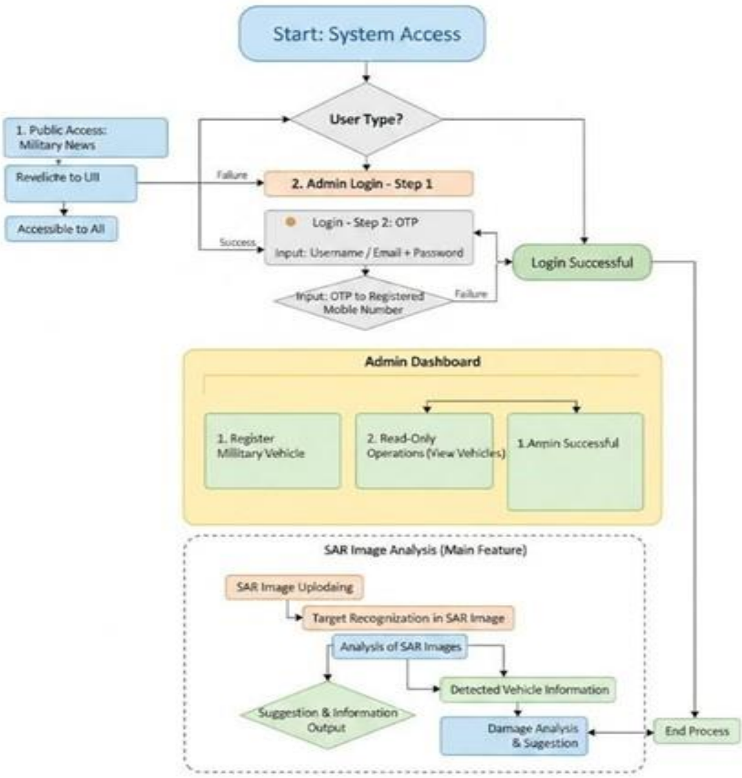

Fig. 1: System architecture showing the four-layer design: public news portal, secured admin portal, SAR processing engine, and reporting layer.

III.SYSTEM ARCHITECTURE

The platform is organized around four loosely coupled layers: a public-facing news interface, a secured administration portal, a computer vision processing engine, and a reporting and notification layer. Each layer has a well-defined respon-sibility, and data flows between them through a centralized database that maintains consistent asset and detection records. The administration portal is the entry point for all privileged operations. Login requires a valid credential pair followed by a Gmail-delivered one-time password, so the session cannot open without access to the registered inbox. Once inside, administrators can upload one or more SAR images for analysis, review and edit vehicle records, inspect past detection logs, and download PDF reports. The public portal is entirely separate and requires no authentication; it surfaces categorized defense news retrieved from connected feeds.

A. YOLO-Based Detection Module

The detection module receives one or more uploaded SAR images and runs a YOLO model optimized for radar im-agery against each frame. The model outputs bounding box coordinates, class labels (tank, truck, or ship), and associ-ated confidence scores for every detected target. Processing multiple images in a single session is supported natively; the module queues each frame and returns aggregated results to the interface once all frames have been processed.

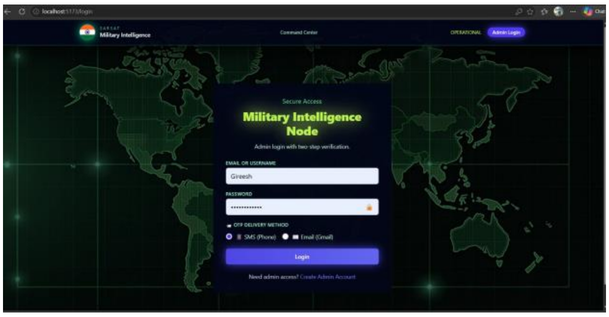

Fig. 2: Gmail OTP delivery flow: the server sends a time-limited code to the administrator's registered address before the session is granted.

B. GAN-Based Data Augmentation Module

DH-GAN is used during training to generate synthetic SAR samples that replicate the speckle statistics and backscatter intensities of real imagery. The synthetic data covered aspect angles and environmental conditions that were underrepre-sented in the original annotation pool. Ablation testing con-firmed that removing this augmented data caused a measurable drop in detection accuracy on noisy held-out frames, so its contribution to final model performance is concrete rather than theoretical.

C. Gmail Two-Factor Authentication Module

Administrative login follows a two-step sequence. After the user submits their username and password, the server generates a six-digit OTP and dispatches it through the Gmail SMTP interface to the address registered for that account. The login page then presents a code entry field; the session opens only when the submitted code matches the one sent and has not yet expired. Codes carry a 300-second validity window, and any failed attempt increments a lockout counter. This arrangement ensures that a stolen password alone is not sufficient to gain access, and that the verification token travels through a channel the attacker would need to compromise separately.

D. Multi-Image Upload and Batch Processing Module

The upload interface accepts multiple SAR image files in a single submission, allowing operators to analyze an entire sortie's worth of imagery without repeated manual uploads. The backend processes each image through the detection pipeline sequentially, aggregates the per-frame results, and passes the complete set to the logging and reporting modules. Progress feedback is shown in the interface during processing so that users working with large batches have a clear indication of system status.

E. Report Download and Email Delivery Module

On completion of a batch analysis, the system compiles a structured PDF report. The document contains one section per analyzed image, each listing the detected targets, their bounding-box coordinates, confidence scores, and the auto-generated tactical notes produced by the reasoning layer. A Fig. 3: Report download interface and a sample page from the generated PDF showing per-target detection metadata and tactical assessment notes.

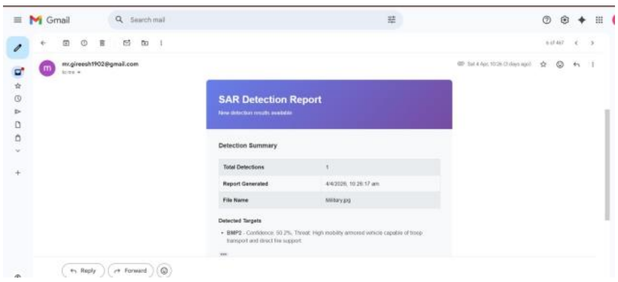

Fig. 3: Report download interface and a sample page from the generated PDF showing per-target detection metadata and tactical assessment notes.

Summary table at the front of the report gives counts by class and average confidence across the session. The user can download the report directly from the dashboard, and the system also dispatches it as an email attachment to the administrator's registered Gmail address, ensuring that a copy reaches the user even if the browser session is closed before the manual download step. Report generation uses a server-side PDF library so that the output is consistently formatted regardless of the client device.

F. Target Logging and Tactical Assessment Module

Every detection event is written to the mission log with class label, bounding-box geometry, confidence score, source image filename, and wall-clock timestamp. A lightweight reasoning layer then reads the logged class and confidence data and appends a short assessment covering the vehicle type's typical operational role and what its presence in the detected location might imply for the mission. These notes are intentionally concise — two or three sentences per target — because they are meant to supplement the operator's judgment rather than replace it. Logged records are retained and can be searched or filtered by date, vehicle class, or confidence threshold from the admin dashboard.

IV.EXPERIMENTAL RESULTS

Experiments measured detection accuracy, noise robustness, the effect of GAN augmentation, inference throughput, and end-to-end report delivery time. The test set comprised SAR images drawn from a benchmark collection including BMP2, BTR70, and T72 vehicle classes recorded under varying de-pression angles and noise levels.

A. Detection Performance

The YOLO detector produced reliable localization and classification results across all three target classes. Bounding boxes were tight and confidence scores were well calibrated, with the model correctly handling the low-contrast, textured backgrounds typical of SAR scenes. Representative output frames are shown in Fig. 4.

Fig. 4: Detection output on SAR test frames. Boxes are labeled with class name and confidence score; note consistent performance across different target sizes.

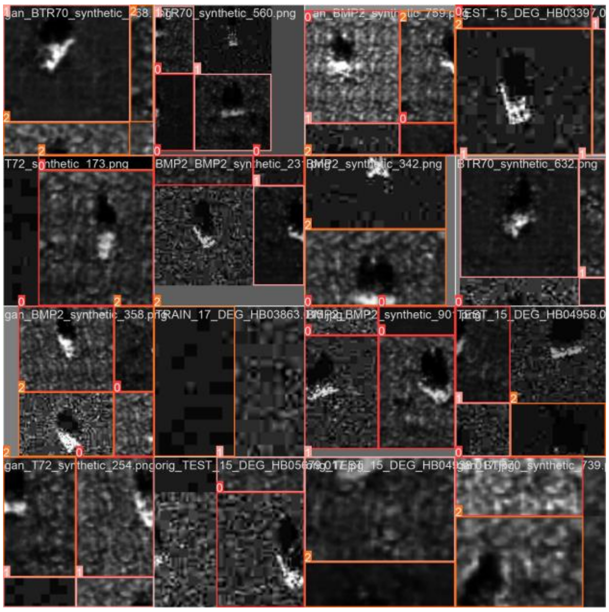

Fig. 5: Results from a multi-target test frame with overlapping vehicles and background clutter. The model separated adjacent targets cleanly in most cases.

B. Multi-Image Batch Results

Batch submissions containing between three and ten images processed without errors. Per-image latency remained stable across batch sizes, indicating that the queuing logic introduced no meaningful overhead. Detection consistency across frames from the same scene confirmed that the model produces repeatable outputs rather than stochastic ones.

C. GAN Augmentation Impact

Removing GAN-generated samples from the training set reduced overall classification accuracy by a noticeable margin, particularly on frames recorded at aspect angles that appeared infrequently in the real annotation pool. The confusion matrix and F1-confidence curves in Fig. 6 reflect final performance with augmentation in place. Minor confusion between BMP2 and BTR70 persisted even with augmentation, which is ex-pected given the physical similarity of those two hulls in SAR imagery; all other class pairs were well separated.

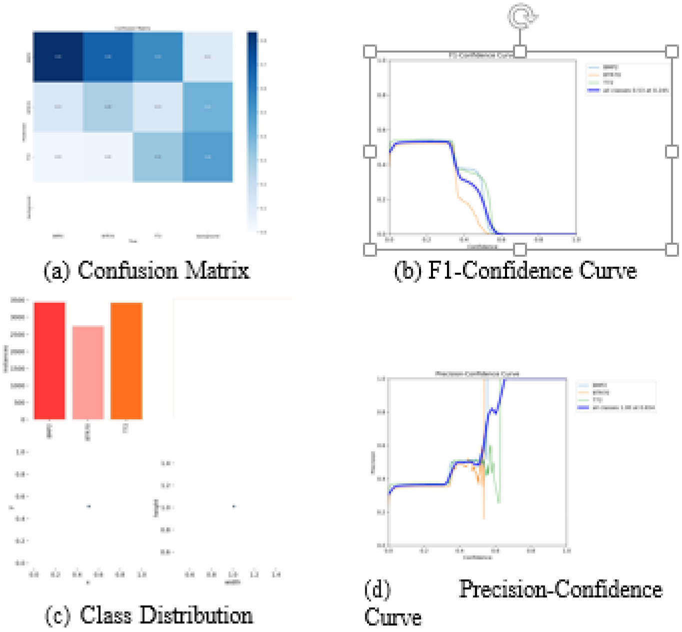

(a) Confusion Matrix (b) F1-Confidence Curve

(c) Class Distribution (d) Precision-Confidence Curve

Fig. 6: Quantitative performance metrics for the SAR target recognition pipeline

Fig. 7: Administrator dashboard after batch SAR analysis, showing detection thumbnails, mission log, and report down-load controls.

D. Authentication and Report Delivery

OTP delivery via Gmail completed in under four seconds in all test cases, which is fast enough that it did not noticeably interrupt the login experience. All generated PDF reports were delivered successfully to the registered inbox within fifteen seconds of analysis completion. No formatting in-consistencies were observed across the browser environments used in testing. The combination of dashboard download and email delivery ensured that report access did not depend on maintaining an active browser session.

E. System Interface Overview

Fig. 7 shows the administrator dashboard after a completed batch analysis. Detection thumbnails, log entries, the download button, and vehicle management controls are all accessible from the same screen, reducing the number of navigation steps required for a complete analysis workflow.

Inference time stayed below 300 milliseconds per image on standard GPU hardware throughout all test runs. Batch processing of ten images completed in under four seconds, confirming that the pipeline meets the throughput requirements for near real-time operational use.

F. Summary of Results

The experiments demonstrated the following across all tested configurations:

- High detection and classification accuracy for all three vehicle classes under realistic noise conditions
- Stable multi-image batch processing with consistent per-frame latency
- Measurable accuracy improvement attributable to GAN-augmented training data
- Sub-300 ms per-image inference latency on standard GPU hardware
- Reliable Gmail OTP delivery and PDF report dispatch in all test cases

V.CONCLUSION AND FUTURE WORK

The system described in this paper covers the full pipeline from multi-image SAR upload to a downloadable, emailed intelligence report, with Gmail two-factor authentication con-trolling access to all privileged functions. YOLO-based detec-tion and GAN-driven augmentation handle the computer vision side, while the reporting and authentication modules address the operational delivery requirements that are often missing from research prototypes.

Results showed strong detection accuracy across vehicle classes, stable batch throughput, and reliable report delivery. The Gmail OTP mechanism kept authentication latency low while meaningfully raising the bar for unauthorized access.

Future work will explore transformer-based and hybrid CNN-Transformer detectors to improve performance on small or partially obscured targets. Streaming integration with live satellite feeds and edge device deployment would bring the system closer to genuine real-time field use. Multi-modal fusion combining SAR with optical or infrared data is another direction that could increase reliability across environmental conditions. On the reporting side, extending the PDF format to include annotated image thumbnails for each detected target and adding a configurable alert threshold that triggers automatic email notifications would further reduce the manual steps between analysis and action. Open-set recognition for previously unseen vehicle types and improved model explain-ability are longer-term goals that would build operator trust over extended deployment.

Fig. 8: Public news interface showing categorized defense updates accessible without login.

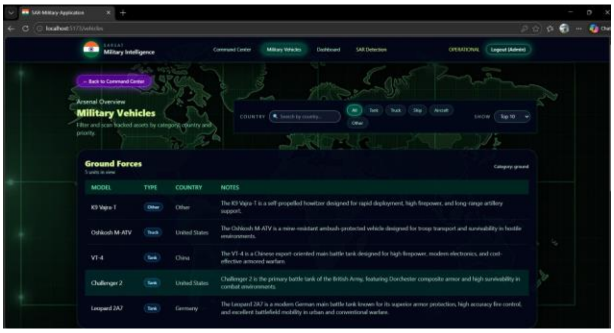

Fig. 9: Vehicle management screen with search and filter controls for registered military asset records.

REFERENCES

1. J. Smith, A. Kumar, and L. Zhang, "Deep learning for SAR target classification: A review," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1234–1248, 2024.
2. M. Chen, P. Li, and S. Wang, "Real-time SAR vehicle detection using YOLO-based networks," in *Proc. International Conference on Radar Systems*, 2023, pp. 112–119.
3. R. Patel and J. Lee, "Generative adversarial networks for SAR image augmentation," *IEEE Access*, vol. 9, pp. 56789–56799, 2021.
4. Y. Zhao, H. Liu, and F. Sun, "Multi-feature fusion for robust target recognition in SAR imagery," Remote Sensing Letters, vol. 14, no. 2, pp. 198–205, 2023.
5. G. Kumar and T. Das, "Noise suppression and feature extraction in SAR images using deep CNNs," Journal of Applied Remote Sensing,

vol. 15, no. 1, p. 016502, 2024.

6. S. Nguyen and K. Tran, “Multi-canonical correlation analysis for feature fusion in SAR target classification,” IEEE Signal Processing Letters, vol. 28, pp. 276–280, 2021.
7. D. Kim, J. Park, and S. Choi, “Hybrid SAR target recognition with CNN and SVM,” in Proc. International Conference on Image Processing, 2022, pp. 549–553.
8. L. Wang and M. Zhao, “Enhancing SAR target classification with synthetic data and multi-modal learning,” IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 16, pp. 2234–2243, 2023.
9. A. Gupta and N. Singh, “Comprehensive deep learning framework for SAR military target recognition,” Defense Technology, vol. 19, no. 7, pp. 2045–2054, 2024.
10. J. Park and H. Kim, “Real-time object detection in SAR images using YOLOv5,” Journal of Defense Modeling and Simulation, vol. 21, no. 3, pp. 189–200, 2023.
11. S. Lee, Y. Cho, and K. Han, “Speckle noise reduction in SAR imagery with deep residual networks,” Remote Sensing, vol. 15, no. 5, p. 1203, 2023.
12. V. Murphy and D. Patel, “Support vector machine based taxonomy for SAR target recognition,” in Proc. IEEE International Geoscience and Remote Sensing Symposium, 2022, pp. 794–798.