ISSN(O): 2320-5407 | ISSN(P): 3107-4928

International Journal of Advanced Research

Publisher's Name: Jana Publication and Research LLP

www.journalijar.com

REVIEWER'S REPORT

Manuscript No.: IJAR-58680

Title: LEAKAGE-AWARE LATENT-DIMENSION SELECTION FOR INTRUSION DETECTION IN AGRICULTURAL IOT NETWORKS.

Recommendation:

Accept as it is .....

Accept after minor revision..... .....

Accept after major revision ..... yes .....

Do not accept (*Reasons below*) .....

Rating	Excel.	Good	Fair	Poor
Originality			Y	
Techn. Quality			Y	
Clarity			Y	
Significance			Y	

Reviewer's ID: JPR-Dr.shaweta Sachdeva

Detailed Reviewer's Report

1. The novelty should be articulated more explicitly. The manuscript combines CNN-Lite, IWHO-Lite, and Random Forest, but the authors should clearly distinguish the proposed contribution from existing CNN-based IDS, feature-selection, and metaheuristic approaches.
2. A stronger ablation study is required. In particular, the authors should compare: CNN-Lite + RF without latent selection, CNN-Lite + conventional feature-selection methods, CNN-Lite + alternative optimizers, and the proposed IWHO-Lite approach under identical experimental conditions.
3. The selected-latent versus full-latent comparison is not a controlled comparison. The proposed model uses 300 RF trees, whereas the reference uses 150 trees, and CICIoT2023 also uses different data limits. Therefore, the reported differences cannot be attributed specifically to IWHO-Lite. An equal-budget comparison should be added.
4. Only three random seeds are used. Although the manuscript correctly acknowledges this limitation, additional independent runs would provide more reliable estimates of variability and confidence intervals.
5. Statistical significance testing should be considered. Mean ± standard deviation across three seeds is useful descriptively, but statistical tests or confidence intervals would help establish whether differences between selected and full representations are meaningful.
6. The choice of the sparsity coefficient $\lambda s = 0.05$ requires justification. A sensitivity analysis should examine how different values of λs affect the selected number of dimensions, F1-score, and overall objective.
7. The choice of 128 latent dimensions should be justified. The authors should explain why 128 was selected and ideally evaluate alternative latent sizes such as 32, 64, 256, etc., to determine whether the observed compression is dependent on the initial representation size.
8. The optimizer comparison is insufficient. Since IWHO-Lite is central to the proposed methodology, comparisons with simpler and established feature-selection strategies and other metaheuristic optimizers would help demonstrate whether the proposed search mechanism

ISSN(O): 2320-5407 | ISSN(P): 3107-4928
International Journal of Advanced Research
Publisher's Name: Jana Publication and Research LLP
www.journalijar.com

REVIEWER'S REPORT

provides a meaningful advantage. The manuscript itself appropriately notes that superiority over other optimizers cannot currently be claimed.

9. The Farm-Flow results require deeper investigation. The reported F1-score is only 0.4724, with precision of 0.3143 and specificity of 0.3167, despite very high recall. This indicates a substantial false-positive problem. The manuscript should discuss possible reasons and provide a confusion matrix or class-wise error analysis.
10. The near-perfect CICIoT2023 performance should be investigated carefully. An F1-score of 0.9953 is unusually high. The authors should provide additional analysis of class distribution, attack categories, potential dataset artifacts, and whether the splitting strategy adequately prevents related traffic patterns from appearing across partitions.
11. Cross-dataset generalization is not demonstrated. Although three datasets are used, the experiments appear to evaluate each dataset independently. Cross-dataset testing—for example, training on one dataset and testing on another where feature compatibility permits—would provide stronger evidence of robustness.
12. The leakage-aware methodology is a positive aspect, but more details are needed for reproducibility. The authors should provide the exact preprocessing pipeline, partitioning procedure, randomization procedure, and implementation details so that other researchers can reproduce the leakage controls.
13. The treatment of contradictory-label observations deserves additional explanation. The manuscript reports contradictory labels in Farm-Flow and UNSW-NB15 and retains some of them in the final test population. The authors should explain quantitatively how these observations affect the reported metrics and whether results change when they are excluded.
14. The computational efficiency claims should remain carefully qualified. The manuscript correctly states that latency and serialized model size do not establish hardware-level efficiency. If the paper emphasizes edge/gateway deployment, experiments on an actual constrained edge device would substantially strengthen the practical contribution.
15. Memory consumption should be reported. Since the motivation includes edge deployment, reporting peak RAM usage during inference and model loading would provide a more complete assessment of deployment feasibility.
16. Energy consumption is missing. For an agricultural IoT/edge application, energy usage can be an important deployment constraint. If hardware experiments are not feasible, the authors should explicitly frame energy efficiency as future work rather than implying complete resource efficiency.
17. The calibration results need better interpretation. Logistic calibration actually worsens both Brier score and ECE for all three datasets. The manuscript acknowledges this, but the authors should explain why calibration is retained and whether alternative calibration methods such as isotonic regression or Platt scaling variants were considered.
18. The fixed F1-based threshold may not reflect practical agricultural cybersecurity requirements. In real deployments, false negatives and false positives may have very different consequences. A cost-sensitive threshold analysis would make the practical implications stronger.
19. The paper should report confusion matrices for all three datasets. Aggregate F1, precision, recall, and specificity are useful, but confusion matrices would allow readers to understand the actual distribution of false alarms and missed attacks.
20. The discussion could be made more concise. Several methodological qualifications and limitations are repeated in the Results, Discussion, and Conclusion sections. Reducing repetition would improve readability and allow greater emphasis on the key scientific findings.
21. The related-work section should include a more systematic comparison table. A table containing recent IDS methods, datasets, classifiers, feature/latent selection strategies, reported metrics, and limitations would make the research gap much clearer.

ISSN(O): 2320-5407 | ISSN(P): 3107-4928
International Journal of Advanced Research
Publisher's Name: Jana Publication and Research LLP
www.journalijar.com

REVIEWER'S REPORT

22. The manuscript should clarify whether IWHO-Lite is genuinely an adaptation of IWHO or a new binary search algorithm. The mathematical relationship between the original IWHO and the proposed binary implementation should be explained more rigorously.
23. The computational complexity of IWHO-Lite should be analyzed. The paper reports population size, iterations, restarts, and multi-fidelity evaluation, but a formal or empirical analysis of optimization cost would help determine whether the dimensionality reduction actually reduces the total computational burden.
24. The reproducibility of the implementation should be improved. Providing source code, configuration files, exact package versions, hardware specifications, and dataset preprocessing scripts would considerably strengthen the work.
25. The conclusion is appropriately cautious, but the practical contribution should be stated more precisely. The current evidence supports latent-space sparsification and performance-parsimony analysis, rather than demonstrating a deployable agricultural IDS. The conclusion should consistently maintain this distinction.

Overall recommendation

Major Revision

The manuscript presents an interesting and potentially useful approach to latent-dimension selection for IoT intrusion detection, and its leakage-aware experimental design is a notable strength. However, controlled ablation studies, stronger optimizer comparisons, more extensive repeated experiments, justification of hyperparameters, deeper analysis of the unusually high CICIoT2023 performance, and hardware/resource evaluation are needed before the methodological and practical claims can be considered sufficiently established.