IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.01

# Sentinel Vision: A Hub-and-Spoke Federated Vision-Language Framework with Differential Privacy for Privacy-Preserving Autonomous Drone Intelligence

**Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N**
Department of Information Science and Engineering,
Vemana Institute of Technology (Autonomous), Bengaluru, Karnataka, India
**Dr.Rajanna M**
Professor and Head & Department of Information Science and Engineering,
Vemana Institute of Technology (Autonomous), Bengaluru, Karnataka, India
https://orcid.org/0009-0006-4297-7014
rajannam318@gmail.com https://ror.org/00ha14p11

**Publication History**
Manuscript Reference No: IJIRIS/RS/Vol.12/Issue07/OCIS10080
Research Article | Open Access | Double-Blind Peer-Reviewed| Article ID: IJIRAE/RS/Vol.12/Issue07/OCIS10080
Received:02,September 2026,Revised:28,September 2026,Accepted:03,October 2026,Published Online:30,October2026
https://www.ijiris.com/volumes/Vol12/iss-07/01.OCIS10080.pdf
**Article Citation:**Dr.Rajanna,Mohammed,Pavan,Dheena,Vishruth(2026), Sentinel Vision: A Hub-and-Spoke Federated Vision-Language Framework with Differential Privacy for Privacy-Preserving Autonomous Drone Intelligence. IJIRIS: International Journal of Innovative Research in Information Security, Volume 12, Issue 07 of 2026 pages 338-343
**Doi:>** https://doi.org/10.26562/ijiris.2026.v1207.01 **BibTeX Key: Dr.Rajanna@2026Sentinel**
IJIRIS papers should be cited as IJIRIS(International Journal of Innovative Research in Information Security, AM Publications, India 2026, ISSN 2349-7017, https://doi.org/10.26562/ijiris.2026.v1207.01 The journal's official abbreviation is IJIRIS. **ORCiD:** https://orcid.org/0009-0004-9398-7488 About the License: Copyright©2026 copyright by the authors. This article is an open access and license under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/)

**Abstract:** Autonomous drones deployed for disaster response, surveillance, and remote-area monitoring face high communication latency, unreliable connectivity, data-privacy risk, and limited onboard computer. This paper presents Sentinel Vision, a hub-and-spoke federated learning framework in which each drone (an edge node) adapts a Vision-Language Model (VLM), SmolVLM-Instruct running in bfloat16, on its own locally retained imagery and transmits only a weight delta ($\Delta W$) to a Central Aggregator Server, which combines the deltas with Federated Averaging (FedAvg) into a global model. Parameter-Efficient Fine-Tuning keeps the core vision and language layers frozen and trains only the lm_head; Target Masking (prompt tokens set to −100) restricts the loss to the safety-protocol label. Before transmission, calibrated Gaussian noise is injected into the clipped $\Delta W$ matrix, providing a Differential Privacy defense against gradient-leakage and reconstruction attacks. The task is Supervised Federated Learning that maps novel visual disaster inputs, such as building demolitions, to strict protocol syntax ([PROTOCOL UPDATE] Level 2 Hazard Detected). A working prototype was implemented and validated on Apple Silicon (MPS). In the holdout evaluation, Drone 2 ingested the FedAvg payload and correctly classified a completely unseen holdout image, demonstrating cross-modal visual generalization without prompt engineering or a centralized RAG database.

**Index Terms:** federated learning, federated averaging, differential privacy, vision-language models, parameter-efficient fine-tuning, edge AI, drone intelligence, hub-and-spoke architecture

## I. INTRODUCTION

Autonomous drones are increasingly proposed for disaster management, surveillance, search-and-rescue, border monitoring, and environmental sensing. These missions require drones to interpret visual data and issue decisions in environments where bandwidth is scarce and the imagery itself is sensitive. Streaming raw imagery to a central server increases latency, bandwidth cost, and the attack surface for interception, particularly for imagery captured over populated or contested areas. Individual drones are also constrained in compute and battery capacity, which limits how much of a model can be trained on the airframe. Sentinel Vision addresses these constraints jointly with a hub-and-spoke federated learning design. Each drone keeps its captured imagery on-device, performs local supervised training of a small trainable component of a VLM, and sends only a privacy-protected weight delta to a Central Aggregator Server. The server aggregates the deltas into a global model and returns it to the drones. Raw imagery never leaves a drone.

**The contributions of this paper are:**

- A hub-and-spoke federated architecture in which edge nodes compute updates locally and transmit them to a Central Aggregator Server that maintains the global model.
- A FedAvg-based payload of weight deltas ($\Delta W$) rather than raw per-step gradients (FedSGD), for network stability and convergence.
- A Differential Privacy pipeline that clips $\Delta W$ and injects calibrated Gaussian noise before transmission, to defend against gradient leakage.
- A parameter-efficient VLM adaptation (SmolVLM-Instruct, bfloat16, lm_head only) with Target Masking, applied to a Supervised Federated Learning hazard-protocol task.

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.01

- A working prototype validated on Apple Silicon (MPS), with a holdout evaluation showing cross-modal visual generalization.

## II. RELATED WORK

### A. Federated Learning Foundations

McMahan et al. introduced Federated Averaging, in which clients train locally for several steps and share model updates with a coordinating server, reducing communication cost relative to transmitting raw data [1]. Sentinel Vision adopts this mechanism with each drone acting as a client. Bonawitz et al. describe a production-scale, centrally coordinated federated learning system [2], which matches the hub-and-spoke topology used here.

### B. Privacy Risk and Differential Privacy

Zhu et al. show that gradients exchanged during federated training can be inverted to reconstruct approximations of the original training inputs (deep leakage from gradients) [3]. Encrypting an update protects it in transit but not from inversion by a party that can decrypt it, including an honest-but-curious aggregator. Differential privacy bounds what an output can reveal about any individual record [4]. Abadi et al. introduced gradient clipping with Gaussian noise for training with differential privacy [5], and McMahan et al. applied user-level differential privacy to federated averaging [6]. Sentinel Vision follows this clip-and-noise approach on the transmitted weight delta.

### C. Vision-Language Models for Edge Use

SmolVLM is a family of compact multimodal models designed for efficient inference on constrained hardware [7]. Its small size and the option of training only a small output head make it suitable for per-drone adaptation, where transmitting full-model updates would be impractical. Table I summarizes how each work is used here.

**Table I-** Literature Foundations and Their Role in Sentinel Vision

| Work | Contribution | Role here |
|---|---|---|
| FedAvg [1] | Local multi-step training; share updates, not raw data | Drone-level training and $\Delta W$ payload |
| FL at scale [2] | Centrally coordinated FL system design | Hub-and-spoke topology |
| Deep leakage [3] | Gradients can be inverted toward inputs | Motivates the differential-privacy mechanism |
| Differential privacy | Clipping and Gaussian noise, including for FedAvg | Noise injection into $\Delta W$ |
| SmolVLM [7] | Compact vision-language models | Base model for on-drone adaptation |

## III. RESEARCH GAP AND CONTRIBUTIONS

Federated learning, parameter-efficient fine-tuning, and differential privacy are each documented in the literature. The gap addressed here is a specified, working integration for drone fleets: which parameters a drone trains, what exactly it transmits, how that payload is protected against reconstruction, and whether the resulting global model transfers knowledge across drones to visual inputs that none of them saw. Sentinel Vision supplies this integration and demonstrates it end to end on a prototype.

## IV. SYSTEM ARCHITECTURE

### A. Hub-and-Spoke Topology

The system has two tiers. The spokes are edge nodes (drones), each with a camera, onboard processor, and a local copy of the model. The hub is a Central Aggregator Server that holds the global model, receives privacy-protected weight deltas from the drones, aggregates them, and redistributes the updated global model. Drones do not exchange updates with one another; all collaboration is mediated by the hub. Raw imagery stays on each drone.

### B. Edge Node: VLM Inference and Local Training

Each drone runs SmolVLM-Instruct in bfloat16 for inference and local training. The core vision encoder and language layers are completely frozen. Only the lm_head, the final projection from hidden states to vocabulary logits, is trained. This Parameter-Efficient Fine-Tuning approach means a round's update is confined to one weight matrix, which keeps the payload small, limits the compute and memory needed on the drone, and preserves the pretrained visual and linguistic representations.

### C. Central Aggregator Server

The server maintains the global lm_head weights. In each round it collects the noised weight deltas from participating drones, computes their weighted average (FedAvg), applies it to the global weights, and returns the updated model to the drones.

## V. METHODOLOGY

### A. Supervised Federated Learning Task and Local Data

The task is Supervised Federated Learning of a hazard-response protocol. Each drone holds locally captured disaster imagery paired with a target response. The model learns to map novel visual disaster inputs, for example building demolitions, to strict military and safety syntax: [PROTOCOL UPDATE] Level 2 Hazard Detected. A fixed instruction prompt accompanies each image, and the target is the protocol string. Imagery remains on the drone that captured it.

### B. Target Masking

Training examples are tokenized as a prompt segment (instruction and image tokens) followed by the target label. Prompt tokens are assigned the label −100, which the cross-entropy loss ignores. The loss therefore optimizes exclusively on the tokens of the targeted safety-protocol label, so the model is not trained to reproduce the prompt and the update reflects only the protocol mapping.

### C. Local Training and the Weight-Delta Payload

At round $t$, each drone $k$ receives the global weights $W^t$ and trains the lm_head locally for $E$ epochs on its $n_k$ examples, obtaining $W_k^{t+1}$.

E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.01

It then forms the weight delta and transmits $\Delta W_k$, not the per-step gradients. This is Federated Averaging (FedAvg) rather than FedSGD. Averaging multi-step deltas, as opposed to single-step gradients, reduces the number of communication rounds, tolerates uneven drone connectivity, and improves convergence stability across non-identical local datasets [1].

$$\Delta W_k = W_k^{t+1} - W^t \quad (1)$$

### D. Differential Privacy on ΔW

Before transmission, each drone applies a two-step mechanism directly to its $\Delta W$ matrix:

1) Clipping. The delta is bounded to a maximum $L_2$ norm $C$:

$$\Delta W_k' = \Delta W_k / \max(1, ||\Delta W_k||_2 / C) \quad (2)$$

This limits the influence of any single drone, and of any single training sample within it.

2) Calibrated Gaussian noise. Noise is injected into the clipped matrix:

$$\Delta \tilde{W}_k = \Delta W_k' + N(0, \sigma^2 C^2 I) \quad (3)$$

Where the noise multiplier $\sigma$ is calibrated to the clipping bound $C$ and the target privacy budget ($\varepsilon$, $\delta$). Only the noised matrix leaves the drone. This directly targets the reconstruction attacks of Zhu et al. [3]. An adversary who intercepts the payload, or an honest-but-curious aggregator, observes a clipped, noised multi-step delta confined to the lm_head rather than an exact per-batch gradient, and the injected noise masks the contribution of individual training samples. Because the noise is zero-mean and independent across drones, it averages out in aggregation, preserving the global signal as the number of participating drones grows. Encryption of the transport channel may still be layered on top, but privacy against reconstruction rests on the differential-privacy mechanism, not on encryption.

### E. Global Aggregation (FedAvg)

The server computes

$$W^{t+1} = W^t + \Sigma_k (n_k / n) \Delta \tilde{W}_k, \quad n = \Sigma_k n_k \quad (4)$$

and redistributes $W^{t+1}$. Each drone loads the updated lm_head and continues local operation, so a drone benefits from patterns learned on other drones' data without ever receiving their imagery.

## VI. IMPLEMENTATION

The prototype implements the full pipeline: SmolVLM-Instruct loaded in bfloat16; freezing of all vision and language parameters with the lm_head trainable; Target Masking with −100 on prompt tokens; local supervised training on each simulated drone; $\Delta W$ computation; clipping and Gaussian noise injection; FedAvg aggregation at the Central Aggregator Server; and redistribution of the global model. All experiments run on Apple Silicon using the Metal Performance Shaders (MPS) backend. Table II lists the configuration.

**Table II** - Prototype Configuration

| Component | Implementation |
|---|---|
| Base model | SmolVLM-Instruct, bfloat16 |
| Trainable | lm_head only; vision and language layers frozen |
| Loss | Cross-entropy, prompt tokens masked (−100) |
| Payload | Weight delta $\Delta W$, FedAvg |
| Privacy | $L_2$ clipping ($C$) plus Gaussian noise ($\sigma$) on $\Delta W$ |
| Topology | Drones (spokes) to Central Aggregator Server (hub) |
| Hardware | Apple Silicon, MPS backend |
| Task | Hazard imagery → [PROTOCOL UPDATE] Level 2 Hazard Detected |

The prototype uses 5 drones, each holding 4 local training examples (20 in total). Each drone trains the lm_head locally for $E = 12$ epochs at a learning rate of $10^{-3}$. Aggregation is performed for a single communication round. Differential privacy uses a clipping bound $C = 1.0$ and a noise multiplier $\sigma = 1.0$, so the Gaussian noise added to each clipped $\Delta W$ has standard deviation $\sigma C = 1.0$.

## VII. EVALUATION METHODOLOGY

### A. Functional Validation

The first check is that the pipeline runs end to end: local training reduces loss on the masked target, $\Delta W$ is produced, clipped and noised, the server aggregates, and drones load the global model.

### B.Holdout Evaluation for Cross-Modal Visual Generalization

The central test withholds an image from every drone's training data. After aggregation, a drone that has not trained on that image or its scene type is given the holdout image together with the standard instruction prompt. The test passes if the model outputs the correct protocol label. No prompt engineering, in-context examples, or centralized Retrieval-Augmented Generation (RAG) database is used, so any correct output must come from the weights delivered by the FedAvg payload.

### C. Privacy Evaluation

The privacy mechanism is characterized by its clipping bound, noise multiplier, and the privacy budget ($\varepsilon$, $\delta$) obtained from a privacy accountant, and can be probed empirically with gradient-inversion attacks on the transmitted payload.

## VIII. SECURITY AND PRIVACY CONSIDERATIONS

Table III relates each threat to the corresponding response and its residual consideration. The hub-and-spoke design concentrates trust and availability in the server. Differential privacy reduces what the server can learn from any one drone, and drones continue inference locally if the server is unreachable.

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.01

**Table III** - Threat Model and Responses

| Threat | Response | Residual |
|---|---|---|
| Interception in transit | Only a clipped, noised $\Delta W$ is sent; transport encryption may be added | Metadata (timing, drone identity) is not hidden by noise |
| Reconstruction (deep leakage) | Clipping plus calibrated Gaussian noise on $\Delta W$ | Strength depends on $\sigma$, $C$, and number of rounds |
| Curious aggregator | Server only receives noised deltas | Aggregator must be access-controlled |
| Raw-data exposure | Imagery never leaves the drone | A captured drone still holds local data; encrypt at rest |
| Hub failure | Drones keep using the last global model | Aggregation pauses until hub is reachable |

## IX. EXPERIMENTAL SETUP

Experiments were run on Apple Silicon with PyTorch on the MPS backend. Drones are simulated as independent clients, each with its own local dataset of disaster imagery paired with the protocol target. The Central Aggregator Server runs FedAvg over the clipped, noised $\Delta W$ matrices. The holdout image is excluded from all drones' training data. Model, precision (bfloat16), trainable component (lm_head), and loss masking (−100 on prompt tokens) are identical across all drones.

## X. RESULTS

### A. Prototype Status

A working prototype of the full pipeline has been implemented and validated on Apple Silicon (MPS). Local training, weight-delta extraction, differential-privacy noise injection, FedAvg aggregation at the Central Aggregator Server, and global-model redistribution each execute end to end.

### B. Holdout Evaluation

The holdout evaluation tested whether knowledge learned across drone's transfers to an input none of them saw. Drone 2 ingested the FedAvg payload from the aggregator and was shown a completely unseen holdout image of a disaster scene (a building demolition). It correctly classified the image and produced the target protocol output, [PROTOCOL UPDATE] Level 2 Hazard Detected. This was achieved without prompt engineering and without a centralized RAG database; the only change between the pre-aggregation and post-aggregation drone was the aggregated lm_head weights delivered by the FedAvg payload (Table IV).

**Table IV**- Validation Outcomes

| Test | Condition | Outcome |
|---|---|---|
| Pipeline execution | Local training, DP noise, FedAvg, redistribution on MPS | Completed |
| Holdout classification | Drone 2, unseen image, FedAvg payload only, no prompt engineering, no RAG | Correct protocol output |

This result demonstrates cross-modal visual generalization: the lm_head, trained under Target Masking on a small set of disaster imagery and aggregated across drones, maps the frozen VLM's visual representation of a new disaster scene to the correct protocol syntax. Because the vision and language layers are frozen, the generalization comes from the pretrained visual representations combined with the federated lm_head update.

## XI. DISCUSSION

The design shows that a drone can acquire a hazard-response behavior from the experience of other drones while exchanging nothing but a small, noised weight matrix. Three choices carry this result. Training only the lm_head keeps the payload small and the pretrained representations intact. Target Masking focuses learning on the protocol label. A FedAvg overweight delta gives a stable aggregation that carries the learned mapping to a drone that never saw the holdout scene. The move from a peer-to-peer network to a hub-and-spoke topology simplifies the system: there is no reconciliation of conflicting model versions across partitions, and aggregation is performed in one place under one rule. The cost is dependence on server reachability for updates; local inference remains available without it.

## XII. LIMITATIONS

The prototype is validated on a single Apple Silicon (MPS) platform with simulated drones, not on flight hardware. The holdout evaluation demonstrates generalization to an unseen image and should be extended to a larger and more varied set of holdout scenes, hazard levels, and negative (non-hazard) inputs to measure false-positive behavior. The strength of the differential-privacy guarantee depends on the chosen $\sigma$, $C$, and number of rounds and should be reported as ($\varepsilon$, $\delta$) alongside empirical gradient-inversion tests. Field deployment will additionally require evaluation of battery draw, link reliability to the aggregator, and on-device encryption at rest.

## XIII. FUTURE WORK

Planned work includes evaluating on a broader holdout set with quantitative accuracy; reporting the privacy budget and running gradient-inversion attacks against the noised payload; studying the utility cost of different noise multipliers; scaling to more drones with heterogeneous local data; porting the pipeline to embedded flight hardware; and extending the protocol vocabulary to multiple hazard levels.

## XIV. CONCLUSION

Sentinel Vision presents and validates a hub-and-spoke federated learning framework for privacy-preserving drone intelligence. Drones adapt SmolVLM-Instruct in bfloat16 by training only the lm_head with Target Masking, transmit FedAvg weight deltas protected by clipping and calibrated Gaussian noise, and receive a global model from a Central Aggregator Server.

A working prototype implemented on Apple Silicon (MPS) completed the full pipeline, and in the holdout evaluation Drone 2, using the FedAvg payload alone, correctly classified a completely unseen image with the output [PROTOCOL UPDATE] Level 2 Hazard Detected, without prompt engineering or centralized RAG. These results show that federated, differentially private, parameter-efficient training can transfer hazard-response behavior across drones without sharing raw imagery.

## ACKNOWLEDGMENT

The authors thank Dr.Rajanna M, Professor and Head, Department of Information Science and Engineering, Vemana Institute of Technology, for his guidance and coordination during the project. The authors also thank Dr.Vijayasimha Reddy B G, Principal, and the teaching and non-teaching staff of the department for their support.

## AUTHOR CONTRIBUTION STATEMENT

**Conceptualization: Dr.Rajanna M**
**Literature Review and Methodology design:**Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Software:** Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Validation:** Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Formal Analysis:** Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Investigation: Dr.Rajanna M**
**Resources:** Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Data Curation:** Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Writing original draft preparation:** Mohammed Rayyan, Pavan Choudhary, S Dheena, Vishruth Reddy N
**Writing review and Editing: Dr.Rajanna M**
**Visualization: Dr.Rajanna M**
**Supervision: Dr.Rajanna M**
**Conflict of interest:** The authors declare no conflicts of interest.
**Data availability statement:** Data supporting these findings are available within the article, at https://doi.org/10.26562/ijiris.2026.v1207.01, or upon request.
**Supplementary materials:** The supplementary materials are available at https://doi.org/10.26562/ijiris.2026.v1207.01
**Publisher's note:** AM Publications, India. IJIRIS (International Journal of Innovative Research in Information Security) Journals stays neutral with regard to jurisdictional claims in published maps and institutional affiliations. All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.
https://ampublications-india.com/

## REFERENCES

1. Vedant, Shivani, Ajay, Vikas (2024). Performance Analysis of Share market by Using ML Technique. International Journal of Innovative Research in Information Security (IJIRIS), Vol.10, Issue 01, 17-20 https://doi.org/10.26562/ijiris.2023.v1001.03
2. H.B.McMahan,E.Moore,D.Ramage,S.Hampson, and B.guerayArcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017.
3. K.Bonawitz et al., "Towards federated learning at scale: System design," in Proc. 2nd SysML Conf., 2019.
4. Kashish,Chhavi,Arya,Ishrat,Geetanjali,Dr.Shivani(2025). Proposed Framework for Validation Requirements Using Semantic and Functional Models. International Journal of Innovative Research in Information Security, Volume 11, Issue 01, Pages 08-13. https://doi.org/10.26562/ijiris.2025.v1101.03
5. L.Zhu, Z.Liu, and S.Han, "Deep leakage from gradients," in Advances in Neural Information Processing Systems (NeurIPS), vol. 32, 2019.
6. C.Dwork and A.Roth, "The algorithmic foundations of differential privacy," Foundations and Trends in Theoretical Computer Science, vol. 9, nos. 3–4, pp. 211–407, 2014. https://doi.org/10.1561/0400000042
7. M.Abadi et al., "Deep learning with differential privacy," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2016. https://doi.org/10.1145/2976749.2978318
8. Nishant,Shivani,Vikas,Ajay (2024). Review on Digital Camoufalge Pattern Generation system using Artificial Intelligence. International Journal of Innovative Research in Information Security (IJIRIS), Vol.10, Issue 01, 21-26 https://doi.org/10.26562/ijiris.2023.v1001.04
9. Mohammad,Shivani,Vikas,Ajay (2024).Faculty Feedback Mining system using Sentiment Analysis. International Journal of Innovative Research in Information Security (IJIRIS), Vol.10, Issue 01, 27-30 https://doi.org/10.26562/ijiris.2023.v1001.05
10. H.B.McMahan, D.Ramage, K.Talwar, and L. Zhang, "Learning differentially private recurrent language models," in Proc. Int. Conf. Learning Representations (ICLR), 2018.
11. A.Marafioti et al., "SmolVLM: Redefining small and efficient multimodal models," arXiv:2504.05299, 2025.
12. Mohammad,Shivani,Vikas,Ajay (2024). Presentation Investigation of Divide Marketplace by Using Machine Learning System.International Journal of Innovative Research in Information Security (IJIRIS),Vol.10,Issue 01,31-35 https://doi.org/10.26562/ijiris.2023.v1001.06
13. Vihas,Pankaj,Ayush (2024).Professional Analysis of chatbots on Facebook Messenger using Artificial Intelligence. International Journal of Innovative Research in Information Security (IJIRIS),Vol.10,Issue 01,36-41 https://doi.org/10.26562/ijiris.2023.v1001.07