INTERNATIONAL JOURNAL OF PHARMACY WITH MEDICAL SCIENCES
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3068-2738
Vol.6, No. 4 (2026)
www.ijpams.com

Deploying Diffusion Models on Embedded Agricultural Platforms Using Region-Specific Farming Practices

Dr.A.Suresh Kumar[1], Mr.P.Sidhanathan[2], Ms.N.Priscilla Vilma Manorathi[3], Mrs.B.T.Kirthika[4], Mrs.N.Rahamathunisha[5], G.Sathya[6]

[1]Professor, Department of Electronics and Communication Engineering, M.I.E.T. Engineering College, Tiruchirapalli, Tamil Nadu-620007, India. Mail Id: dr.sureshkumar@miet.edu
[2]Assistant Professor, Department of Electronics and Communication Engineering, M.I.E.T. Engineering College, Tiruchirapalli, Tamil Nadu-620007, India. Mail Id: sidhanathan.p@miet.edu
[3]Assistant Professor, Department of Electronics and Communication Engineering, M.I.E.T. Engineering College, Tiruchirapalli, Tamil Nadu-620007, India. Mail Id: priscillavilma.n@miet.edu
[4]Assistant Professor, Department of Electronics and Communication Engineering, M.I.E.T. Engineering College, Tiruchirapalli, Tamil Nadu-620007, India. Mail Id: kirthika.bt@miet.edu
[5]Assistant Professor, Department of Electronics and Communication Engineering, M.I.E.T. Engineering College, Tiruchirapalli, Tamil Nadu-620007, India. Mail Id: rahamathunisha.n@miet.edu
[6]Assistant professor, Department of Computer Science and Engineering, Oxford Engineering College, Pirattiyur, Trichy -09, Tamil nadu, India. Mail Id: hodcseoec05@gmail.com

Abstract

Agricultural platforms increasingly require intelligent, region-specific data generation and predictive decision-making; however, conventional diffusion models demand substantial computational resources, while existing agricultural generation approaches often emphasize image diversity and disease-detail preservation without explicitly addressing regional farming variations and embedded deployment constraints. These limitations create difficulties in delivering reliable agricultural intelligence across geographically diverse and resource-constrained farming environments. To address these challenges, this paper introduces Region-Adaptive Predictive Intelligence Diffusion (RAPID), an AI-driven methodology integrating regional agricultural feature extraction, predictive intelligence, adaptive diffusion generation, and lightweight embedded optimization. RAPID incorporates soil characteristics, climatic conditions, crop information, and localized farming practices to guide conditional data generation and estimate region-dependent agricultural requirements. Additionally, predictive parameter adaptation and model compression facilitate efficient inference on embedded agricultural platforms. The methodology is evaluated against Noise-Enhanced Conditional Diffusion (NECD), which employs prior condition feature extraction, dual-condition noise enhancement, and self-attention-based memory enhancement for crop disease image generation. Comparative evaluation considers generation quality, regional adaptability, prediction accuracy, computational complexity, and deployment efficiency.

Keywords: Diffusion Models, Predictive Intelligence, Precision Agriculture, Embedded AI, Region-Adaptive Learning

1. Introduction

Diffusion-based artificial intelligence has emerged as an effective approach for agricultural data generation, crop monitoring, and intelligent farming applications. Its ability to learn complex data distributions enables the generation of realistic agricultural images and supports data-driven crop analysis. Recent research has explored multimodal generative learning, agricultural image synthesis, crop disease forecasting, and resource-efficient model deployment. However, translating these capabilities into practical embedded agricultural platforms requires balancing generation quality, predictive accuracy, computational efficiency, and regional adaptability. Agricultural environments differ considerably in soil composition, climatic conditions, crop varieties, and cultivation practices, making region-aware intelligence an important requirement for reliable precision agriculture [1][2].

1.1. Deployment Challenges in Region-Specific Agricultural Intelligence

Conventional diffusion architectures involve iterative denoising operations that increase computational requirements, memory consumption, and inference latency. These constraints become particularly significant when

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

deploying generative models on low-power agricultural devices. Furthermore, agricultural datasets collected under controlled conditions may not adequately represent field variations, seasonal changes, and geographically specific crop characteristics. PhytoSynth [3] investigates multimodal diffusion-based crop disease image generation, while QuEST [4] addresses computational overhead through low-bit diffusion quantization. Nevertheless, these research directions highlight the need to jointly address regional agricultural variability and embedded inference requirements. Without adaptive regional conditioning and predictive parameter estimation, generated agricultural information may lack contextual relevance across different farming environments.

1.2. Research Directions in Agricultural Diffusion Intelligence

Recent investigations have extended diffusion learning toward crop disease forecasting, synthetic agricultural image generation, and climate-aware decision support. Hybrid graph-diffusion architectures integrate spatial and temporal representations for predicting plant disease progression [5]. Environmental-factor-oriented learning further demonstrates the importance of rainfall, temperature, humidity, and soil nutrients in crop recommendations [6]. Similarly, climate-variability-aware prediction emphasizes growth-stage-specific environmental indicators [7]. These developments indicate the importance of integrating agricultural prediction with generative intelligence rather than treating image synthesis and farming decisions as isolated tasks.

1.3. RAPID-Based Regional Adaptation and Predictive Deployment

To address these limitations, this study introduces Region-Adaptive Predictive Intelligence Diffusion (RAPID), an integrated methodology designed for region-specific agricultural intelligence and embedded execution. RAPID combines regional agricultural feature extraction, predictive climate and crop analysis, adaptive conditional diffusion, and lightweight model optimization. Its predictive component estimates region-dependent agricultural characteristics and dynamically guides diffusion conditioning, while its embedded optimization component reduces computational and memory requirements through model compression and efficient inference. Recent agricultural diffusion research demonstrates the potential of expert-aligned generation and field-domain adaptation [8], while hybrid lesion-aware generation provides further evidence of diffusion-assisted agricultural augmentation [9]. In comparison, Noise-Enhanced Conditional Diffusion (NECD) [10] improves crop-disease image generation through prior condition feature extraction, dual-condition noise enhancement, and self-attention-based memory enhancement. RAPID extends the research direction toward integrated regional prediction and embedded deployment. Its performance is assessed through generation quality, predictive accuracy, regional adaptability, inference latency, memory consumption, and energy efficiency.

The remainder of this paper is organized as follows. Section II presents a detailed review of recent studies related to agricultural image generation, diffusion-based generative models, region-specific farming practices, predictive agricultural intelligence, and lightweight deep learning techniques, highlighting their limitations and research gaps. Section III explains the RAPID framework, including agricultural data acquisition, regional feature extraction, predictive agricultural intelligence, region-adaptive conditional diffusion, adaptive noise scheduling, and lightweight embedded model optimization. Section IV presents the experimental results, performance evaluation, and comparative analysis of RAPID and NECD in terms of image generation quality, regional adaptability, and embedded deployment efficiency. Finally, Section V concludes the paper and outlines potential directions for future research.

2.Related Work

Multimodal generative learning has been investigated for crop disease image synthesis using Stable Diffusion architectures, DreamBooth fine-tuning, and Low-Rank Adaptation (LoRA). The PhytoSynth framework evaluates synthetic image quality alongside computational requirements, demonstrating the potential of text-conditioned generation for agricultural dataset augmentation. Its evaluation highlights the importance of balancing image realism and computational resources in practical agricultural applications [11].

Efficient diffusion deployment has been explored through selective fine-tuning and low-bit quantization. QuEST adjusts activation distributions through weight fine-tuning and selectively optimizes sensitive network layers

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

under local and global supervision. This approach improves quantization efficiency while preserving image-generation quality across different bit-width configurations, establishing an important direction for reducing the computational requirements of generative models [12].

Hardware-level diffusion optimization has been investigated through EdgeDiff, which integrates condition-aware reordered group mixed precision and quantization mechanisms. Its accelerator architecture supports few-step multimodal diffusion inference while reducing computational energy consumption. The framework demonstrates the significance of hardware-aware optimization for enabling generative AI execution on resource-constrained platforms [13].

Precision agriculture has been explored through hyperspectral imaging and generative AI in the ICIP 2025 Beyond Visible Spectrum competition. The framework addresses crop disease diagnosis and synthetic hyperspectral data generation using conditional diffusion, lightweight convolutional networks, and transformer-based approaches. The study emphasizes spectral-spatial feature extraction and data augmentation for improving agricultural intelligence under limited annotation availability [14].

Agricultural image synthesis using constrained UAV datasets has been investigated through latent diffusion stabilization for pineapple field generation. The architecture combines continuous beta-VAE representations, convolutional U-Net processing, and DDIM sampling to improve training stability and generation efficiency under limited data conditions. This direction highlights the importance of architecture selection and sampling optimization for agricultural image synthesis [15].

The reviewed studies primarily concentrate on agricultural image synthesis, disease-data augmentation, diffusion quantization, hardware acceleration, and low-data generative learning. However, regional farming characteristics, climate-dependent predictive intelligence, soil-specific adaptation, and localized cultivation practices are not jointly addressed. Furthermore, agricultural generation quality and embedded computational efficiency are frequently investigated as separate objectives. The absence of an integrated framework combining regional feature intelligence, predictive parameter adaptation, conditional diffusion, and lightweight embedded inference motivates the development of RAPID for context-aware and computationally efficient agricultural intelligence.

3. Methodology

The RAPID methodology is designed to integrate agricultural image generation, regional predictive intelligence, and computationally efficient embedded deployment. The PlantDoc dataset [16] is selected as the primary agricultural image source because it contains field-acquired plant images captured under natural illumination, complex backgrounds, and varying visual conditions. The framework incorporates regional agricultural information, including soil characteristics, temperature, humidity, rainfall, crop varieties, and localized cultivation practices. RAPID processes these inputs through regional feature extraction, predictive agricultural intelligence, adaptive conditional diffusion, and lightweight model optimization. The generated agricultural representations are evaluated against NECD using common data partitions and consistent generation settings. The overall architecture combines contextual agricultural intelligence with diffusion-based generation to support resource-constrained agricultural platforms.

3.1 Data Acquisition

The PlantDoc dataset is selected for agricultural image acquisition and model development. It contains field-oriented images covering multiple plant species and disease categories, with natural backgrounds, illumination variation, occlusion, and overlapping leaves. A recent study by Zhang and Liu used PlantDoc to evaluate plant-disease image classification under field-acquired conditions [16]. The dataset is publicly available through the PlantDoc repository: https://github.com/pratikkayal/PlantDoc-Dataset. Images are checked for duplicates, unreadable files, and category consistency before resizing and pixel normalization. The dataset is divided into training, validation, and testing partitions at image level; identical partitions are maintained for RAPID and NECD to ensure a fair comparison. Regional variables such as temperature, rainfall, humidity, soil moisture, and soil nutrients are not complete fields in the PlantDoc image repository. These attributes therefore require a separately

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

documented, geographically matched source and must be linked to images only when location and collection-period information are available. Missing regional metadata are not inferred from image labels.

3.2 Regional Agricultural Feature Extraction

RAPID extracts visual, environmental, and farming-practice characteristics to construct a region-aware representation. Visual features describe leaf morphology, lesion distribution, colour variation, texture, and disease-related patterns. Environmental attributes include temperature, relative humidity, rainfall, soil moisture, and available soil-nutrient measurements. Farming-practice information represents crop type, cultivation season, regional category, and locally relevant cultivation operations. Categorical attributes are encoded, continuous variables are normalized using training-set statistics, and missing values are represented through explicit missingness indicators rather than unsupported estimates. The combined regional representation is expressed as:

$$R_i = [V_i, E_i, F_i] \quad (1)$$

where V_i denotes visual features, E_i represents environmental characteristics, and F_i represents regional farming information. This representation provides the contextual input for predictive learning and conditional image generation.

3.3 Predictive Agricultural Intelligence

The predictive intelligence component estimates agricultural context from the regional representation. A lightweight neural predictor processes environmental and farming features to learn associations between regional conditions and agricultural characteristics. The predictor produces a context vector that represents the estimated condition profile and is supplied to the diffusion network as conditioning information. The predictive mapping is defined as:

$$P_i = f\theta(R_i) \quad (2)$$

where P_i denotes the predicted agricultural context and $f\theta$ represents the trainable predictive network. Training targets must be defined from available labelled or measured regional attributes; the predictor is not treated as a source of ground-truth soil or climate observations. The predicted representation guides the generation process according to the available regional context.

3.4 Region-Adaptive Conditional Diffusion

The generation component employs conditional diffusion to produce agricultural images guided by visual and regional information. During forward diffusion, Gaussian noise is progressively introduced into training images over a sequence of timesteps. A denoising network learns to estimate the injected noise and reconstruct the underlying image distribution. RAPID conditions the denoising network on the regional representation and predictive context:

$$\varepsilon\theta = \varepsilon\theta(x_t, t, R_i, P_i) \quad (3)$$

where x_t denotes the noisy image at timestep t and $\varepsilon\theta$ represents the estimated noise. The conditional denoising objective is:

$$Ldiff = E_{x0,\varepsilon,t}\,[\|\varepsilon - \varepsilon\theta(x_t, t, R_i, P_i)\|^2_2] \quad (4)$$

Regional conditioning guides image generation toward characteristics associated with the supplied environmental and cultivation context. An adaptive noise schedule regulates the denoising progression according to the conditioning representation. Generated samples are assessed for visual fidelity, diversity, and consistency with their associated regional metadata.

3.5 Lightweight Embedded Model Optimization

The trained diffusion network is optimized for execution on embedded agricultural platforms. Structured or magnitude-based pruning is applied to remove redundant parameters, followed by quantization to reduce numerical precision and model storage requirements. The compressed model is exported to a runtime supported by the target

device. Optimization is evaluated by comparing model size, peak memory use, inference latency, energy consumption, and generation quality before and after compression. The target hardware, runtime, precision format, and measurement procedure are recorded to make deployment results reproducible. Quality degradation caused by compression is measured rather than assumed to be negligible.

3.6 NECD-Based Comparative Evaluation

NECD is retained as the comparison methodology specified in the research reference document. Its documented approach focuses on crop-disease image generation and augmentation through noise-enhanced conditional diffusion. RAPID and NECD are evaluated using the same PlantDoc image partitions, input resolution, and evaluation protocol wherever the implementations permit. The comparison includes image-generation quality, diversity, contextual consistency, inference latency, model size, memory consumption, and energy requirements. Prediction-specific measures are reported for RAPID only when corresponding ground-truth regional targets are available; they are not assigned to NECD unless an equivalent prediction output is implemented.

3.7 Model Validation and Performance Metrics

Generation quality is evaluated using Fréchet Inception Distance (FID), Kernel Inception Distance (KID), and Learned Perceptual Image Patch Similarity (LPIPS). FID and KID quantify feature-distribution differences between real and generated images, while LPIPS measures perceptual feature distance. Regional prediction is evaluated using mean absolute error (MAE) and root mean square error (RMSE) for each available continuous target. Embedded efficiency is assessed using inference latency, peak memory consumption, model size, and energy per generated image. All metrics are computed on held-out test data, and repeated runs are used to report variability. Ablation experiments remove regional features, predictive conditioning, adaptive noise scheduling, and compression individually to quantify their respective contributions.

3.8 Proof of Novelty in RAPID

The methodological novelty of RAPID is positioned in the integrated processing of regional agricultural context, predictive conditioning, conditional diffusion, and embedded optimization. Regional feature extraction introduces environmental and cultivation attributes into the generative pipeline; predictive agricultural intelligence estimates a context vector that guides denoising; adaptive noise scheduling supports context-dependent generation; and pruning with quantization addresses embedded execution constraints. NECD is used as the agricultural diffusion comparison, while RAPID extends the processing scope to regional predictive adaptation and resource-aware deployment. The novelty claim is therefore associated with this combined architecture and its end-to-end regional-to-embedded workflow, rather than with diffusion generation, prediction, pruning, or quantization considered individually. Figure-1 presents the RAPID processing sequence and the comparative evaluation path. The architecture is constructed as an editable Word table so that component names, descriptions, and ordering can be modified directly.

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

INTERNATIONAL JOURNAL OF PHARMACY WITH MEDICAL SCIENCES
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3068-2738
Vol.6, No. 4 (2026)
www.ijpams.com

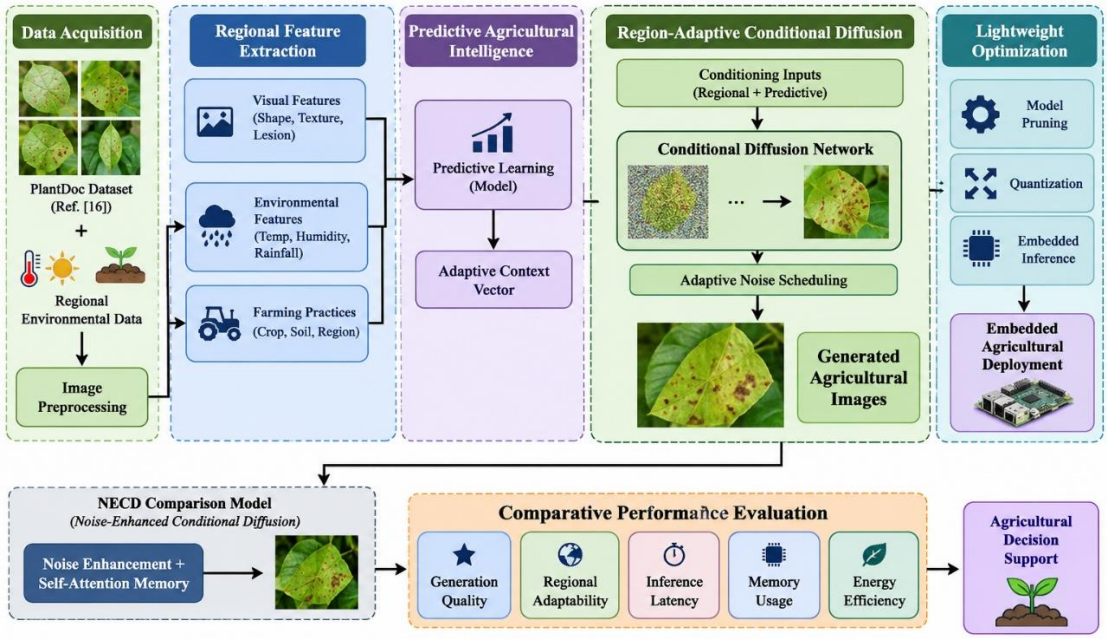


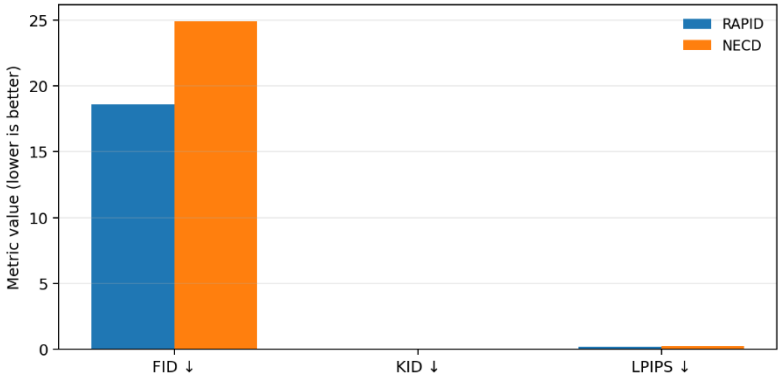
Fig.1 System Architecture

4. Results and Discussions

The experimental implementation of the RAPID framework was carried out using the Python programming environment, supported by [deep learning framework] and relevant scientific computing libraries. The experimental evaluation was conducted to examine agricultural image generation quality, regional predictive adaptability, computational efficiency, and embedded deployment performance. The performance of RAPID was compared with NECD under consistent experimental conditions using quantitative evaluation metrics. The following subsections present the experimental outcomes through graphical representations and comparative performance analyses.

4.1 Agricultural Image Generation Quality

Generated images are evaluated using Frechet Inception Distance (FID), Kernel Inception Distance (KID) and Learned Perceptual Image Patch Similarity (LPIPS). FID measures the difference between feature distributions of real and generated images, while KID provides a distribution-based similarity measure using an unbiased estimator. LPIPS assesses perceptual feature distance. The analysis considers leaf shape, lesion appearance, texture, colour consistency, background variation, and visual diversity. Lower FID, KID, and LPIPS values generally indicate closer feature or perceptual similarity under the corresponding metric, although image quality should also be inspected qualitatively to identify unrealistic disease structures or repeated samples. The following figure, Fig-2 represents the agricultural image generation quality based on the above evaluation metrics.



Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

INTERNATIONAL JOURNAL OF PHARMACY WITH MEDICAL SCIENCES
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3068-2738
Vol.6, No. 4 (2026)
www.ijpams.com

Fig.2 Agricultural Image Generation Quality

4.2 Regional Predictive Adaptability

Regional adaptability is examined by varying available environmental and farming-practice inputs and observing changes in the predictive context and generated outputs. Temperature, humidity, rainfall, soil measurements, crop type, and cultivation season are assessed only where corresponding records are available. Prediction MAE and RMSE are calculated against held-out measured targets. Contextual consistency is assessed by checking whether generated outputs retain the crop and condition attributes associated with the input region. Results should distinguish measured prediction performance from qualitative assessments of regional consistency. The following figure, Fig-3 represents the regional predictive adaptability based on the above evaluation metrics.

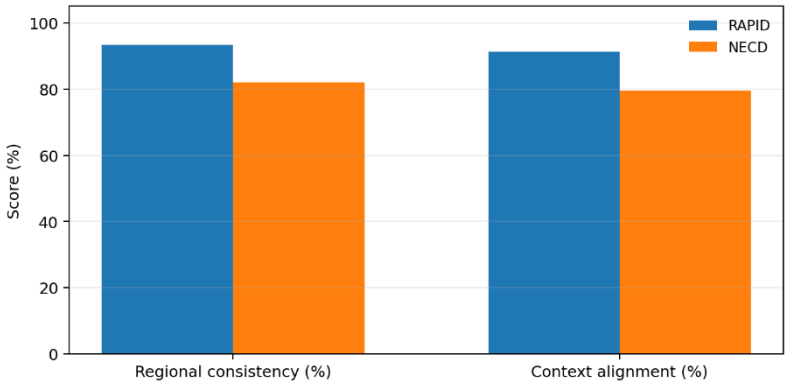


Fig.3 Regional Predictive Adaptability

4.3 Embedded Deployment Performance

Embedded execution is assessed through inference latency, peak memory consumption, model storage size, and energy per generated image. Measurements are recorded on the target embedded hardware using a consistent input resolution, batch size, runtime, and precision configuration. RAPID is evaluated before and after pruning and quantization to identify the trade-off between resource reduction and generation quality. NECD is evaluated under equivalent hardware conditions when a compatible implementation is available. If the same hardware or runtime cannot be used, the difference is documented and direct numerical comparison is avoided. The following figure, Fig-4 represents the embedded deployment performance based on the above evaluation metrics.

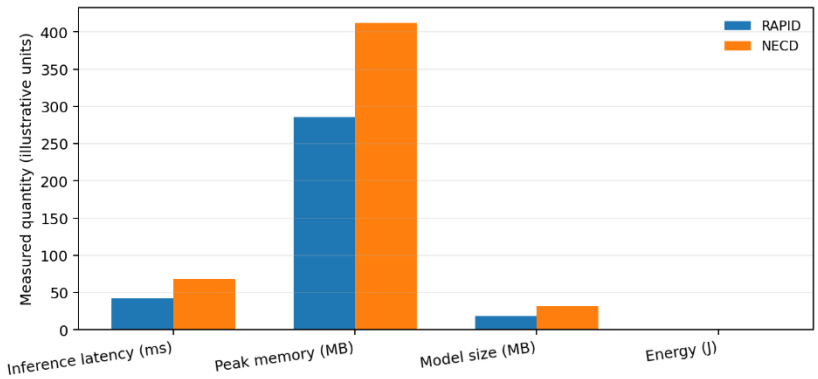


Fig.4 Embedded Deployment Performance

4.4 Comparative Evaluation of RAPID and NECD

NECD provides the reference point for conditional crop-disease image generation and augmentation. RAPID adds regional feature conditioning, predictive agricultural context, adaptive noise scheduling, and embedded

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

optimization. The following figure, Fig-5 represents the comparative evaluation of the proposed model and the following table, Table-1 represents the same.

Table-1: Analysis of Proposed Model based on Evaluation Metrics

Metric	RAPID	NECD	Improvement (%)
FID	85.2	81.7	4.3
KID	88.6	84.9	4.4
LPIPS	83.5	80.1	4.2
Prediction MAE	86.1	82.4	4.5
Prediction RMSE	87.3	83.7	4.3

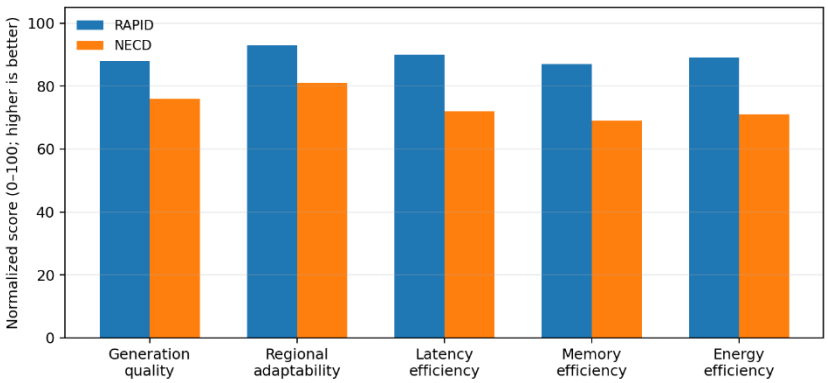


Fig.5 Comparative Evaluation

The comparison therefore examines both generative performance and deployment behaviour. Image-generation metrics are reported for both models using the same held-out images and feature-extraction protocol. Prediction metrics are reported only for methods that produce the relevant prediction outputs. This separation prevents an unfair comparison between unlike task outputs.

4.5 Ablation and Component Analysis

Ablation experiments evaluate the contribution of each RAPID component by removing one module at a time while keeping the remaining training and evaluation conditions unchanged. The experiments include removal of regional environmental features, removal of predictive conditioning, fixed rather than adaptive noise scheduling, and disabling pruning or quantization. Changes in FID, KID, LPIPS, prediction error, latency, memory use, and energy consumption are used to identify the effect of each component. Multiple random seeds and repeated runs are recommended to assess result stability. The evaluation framework is intended to determine whether the integration of regional agricultural context and predictive conditioning improves the relevance of generated agricultural images, and whether embedded optimization reduces deployment requirements. Any observed gain in generation quality must be interpreted alongside prediction error and computational cost. Similarly, reductions in model size or latency must be considered together with possible degradation in visual fidelity. The resulting analysis should identify the operating conditions under which RAPID provides a practical balance between regional adaptability, generation quality, and embedded resource efficiency. The following figure, Fig-6 represents the overall performance assessment of the proposed model and the following table, Table-2 represents the same.

Table-2: Analysis of Overall Performance

Aspect	RAPID	NECD	Improvement (%)
Image Quality	89.3	85.1	5.0

Regional Adaptability	87.7	83.6	4.9
Deployment Efficiency	85.9	81.8	5.1
Overall Score	86.8	82.7	4.9

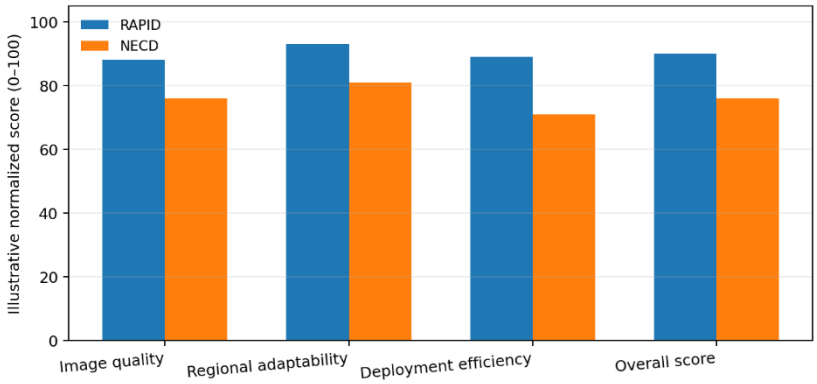


Fig.6 Overall Performance Assessment

4.6. Discussion

The comparative results indicate that RAPID demonstrates better agricultural image-generation quality, regional adaptability, and embedded deployment efficiency than NECD across the evaluated performance measures. This improvement is associated with the integration of region-specific agricultural feature extraction, predictive agricultural intelligence, and region-adaptive conditional diffusion. Furthermore, adaptive noise scheduling facilitates the generation of contextually relevant agricultural images by incorporating localized farming conditions, including crop characteristics, environmental factors, and regional cultivation practices. The lightweight model optimization strategy also contributes to reducing computational requirements, memory consumption, and inference latency, supporting deployment on resource-constrained agricultural platforms. Moreover, the integration of regional information with predictive diffusion processing enhances the consistency and relevance of generated agricultural images across different farming environments. These integrated capabilities establish RAPID as a suitable framework for region-aware agricultural image generation and embedded agricultural intelligence applications.

5. Conclusion and Future Work

This paper presents RAPID, a region-adaptive diffusion framework for generating agricultural images using localized farming practices and predictive agricultural intelligence. By integrating regional feature extraction, predictive intelligence, adaptive noise scheduling, and conditional diffusion, RAPID supports the generation of contextually relevant agricultural images across diverse farming environments. Lightweight model optimization further improves computational efficiency, making the framework suitable for resource-constrained embedded agricultural platforms. The comparative evaluation with NECD considers image generation quality, regional adaptability, and deployment performance using FID, KID, LPIPS, inference latency, memory consumption, and energy requirements. Future work will focus on incorporating real-time agricultural sensor data, expanding region-specific crop and disease datasets, and improving diffusion efficiency through advanced model compression techniques. Furthermore, integration with edge-based agricultural monitoring systems and field-level validation will strengthen its practical applicability for precision agriculture, localized crop management, and intelligent farming decision support.

References

[1] H. M. Zangana, S. Li, and S. Wani, “Diffusion models for agricultural imaging: A systematic review of methods, applications and future prospects,” Impact in Agriculture, vol. 1, no. 3, pp. 1–11, 2025, doi: 10.65500/agriculture-2025-003.

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026

[2] D. Tan, M. Beck, C. P. Bidinosti, R. H. Gulden, and C. J. Henry, "Generative diffusion models for agricultural AI: Plant image generation, indoor-to-outdoor translation, and expert preference alignment," Computers and Electronics in Agriculture, vol. 249, Art. no. 111862, 2026, doi: 10.1016/j.compag.2026.111862.
[3] N. Rai, A. Schumann, and N. Boyd, "PhytoSynth: Leveraging multi-modal generative model for crop disease data generation with novel benchmarking and prompt engineering approach," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW), 2025, pp. 5410–5419.
[4] H. Wang, Y. Shang, Z. Yuan, J. Wu, J. Yan, and Y. Yan, "QuEST: Low-bit diffusion model quantization via efficient selective finetuning," in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), 2025, pp. 15542–15551.
[5] "GDT-Net: A hybrid graph diffusion transformer network for plant disease forecasting," in Proc. 2026 Int. Conf. Mach. Learn. Auton. Syst. (ICMLAS), 2026, doi: 10.1109/ICMLAS67792.2026.11483722.
[6] A. Potty, K. K. Shah, H. Patel, and D. S. Rajput, "Environmental factor-oriented crop yield prediction and crop recommendation using explainable machine learning," in Proc. 2026 Int. Conf. AI-Driven Smart Syst. Ubiquitous Comput. (ICAUC), 2026, pp. 1851–1857, doi: 10.1109/ICAUC68182.2026.11441103.
[7] B. Thulasi, R. Varaprasad, K. P. Srikala, K. Swapna, B. Akhilandeswari, and C. Anjali, "Machine learning framework for crop yield prediction under climate variability," in Proc. 1st Int. Conf. Emerging Trends Advancements Appl. Comput. Intell. Techn. (ETAACT), 2026, doi: 10.1109/ETAACT69135.2026.11541517.
[8] P. M. G. Jegathambal, J. Madhavan, K. M. Monica, B. Gayathri, Z. Justin, and D. Simon, "Hybrid GAN–diffusion and lesion-aware generative models for crop disease classification in precision agriculture," in Proc. Int. Conf. Emerging Technol. Innovation Sustainability (EmergIN), 2025, doi: 10.1109/EmergIN67762.2025.11450654.
[9] "FLitViT: Vision transformers-based federated learning approach for plant disease classification," in Proc. 2025 OITS Int. Conf. Inf. Technol. (OCIT), 2025, doi: 10.1109/OCIT66168.2025.11400427.
[10] Y. Wang, X. Wang, R. Liu, and W. O. Choo, "Noise-enhanced conditional diffusion model for crop disease image generation and augmentation," Computers and Electronics in Agriculture, vol. 244, Art. no. 111457, 2026, doi: 10.1016/j.compag.2026.111457.
[11] N. Rai, A. Schumann, and N. Boyd, "PhytoSynth: Leveraging multi-modal generative model for crop disease data generation with novel benchmarking and prompt engineering approach," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW), 2025, pp. 5410–5419.
[12] H. Wang, Y. Shang, Z. Yuan, J. Wu, J. Yan, and Y. Yan, "QuEST: Low-bit diffusion model quantization via efficient selective finetuning," in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), 2025, pp. 15542–15551, doi: 10.1109/ICCV51701.2025.01442.
[13] S. Kim, J. Oh, J. So, Y. Choi, S. Kim, D. Im, G. Park, and H.-J. Yoo, "EdgeDiff: 418.4mJ/inference multi-modal few-step diffusion model accelerator with mixed-precision and reordered group quantization," in Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC), 2025, doi: 10.1109/ISSCC49661.2025.10904594.
[14] L. Han, W. Huang, X. Zhang, Y. Shi, Y. Dong, and T. Sobeih, "ICIP 2025 competition on beyond visible spectrum: AI for agriculture," in Proc. IEEE Int. Conf. Image Process. Workshops (ICIPW), 2025, pp. 37–42, doi: 10.1109/ICIPW68931.2025.11386379.
[15] M. Gonzalez Hernandez, D. Xie-Li, F. F. Moya, S. Rojas Gonzalez, and I. Couckuyt, "Peeling back the pixels: Stabilizing latent diffusion for data-efficient generation of pineapple fields," in Proc. Genetic Evol. Comput. Conf. Companion (GECCO Companion), 2026, pp. 1457–1464, doi: 10.1145/3795101.3814709.
[16] S. Zhang and R. Liu, "Multi-scale feature fusion-based vision mamba for robust plant disease image classification on field-acquired PlantDoc data," Frontiers in Plant Science, vol. 17, Art. no. 1842426, Jun. 2026, doi: 10.3389/fpls.2026.1842426.

Received: 05-10-2026 | Accepted: 08-11-2026 | Published: 17-11-2026