



Diminishing Returns of Algorithmic Complexity: A Cost-Benefit Evaluation of UAV Image Enhancement for Precision Weed Detection in Rice Farming

Muhamad Syaiful Huda*

Institut Teknologi Sepuluh Nopember,
Indonesia

Budi Setiyono

Institut Teknologi Sepuluh Nopember,
Indonesia

***Corresponding author:**

Budi Setiyono, Institut Teknologi Sepuluh
Nopember, Indonesia.

✉ budisetiyono@its.ac.id

Article Info:

Article history:

Received: July 30, 2026

Revised: September 12, 2026

Accepted: October 03, 2026

Available Online: October 13, 2026

Keywords:

Precision Agriculture;
Image Enhancement;
Cost-Effectiveness;
Weed Detection;
Technology Adoption



This is an open access article under the
[CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.

Copyright © 2026 by Author. Published
by Politeknik Siber Cerdika
International

Abstract

Background: Unmanned aerial vehicle (UAV) imagery is increasingly used for weed detection in rice fields, but limited spatial resolution can constrain segmentation performance. Image enhancement may improve detection; however, it remains uncertain whether computationally complex, deep learning-based enhancement provides task-level gains beyond those achieved through classical interpolation.

Objective: This study evaluates whether increasing enhancement complexity yields additional weed-detection benefits across four image conditions.

Methods: Using 148 UAV images from the WeedyRice-RGBMS-DB dataset, intersection over union (IoU) was measured under four image conditions: no enhancement, nearest-neighbor interpolation, bicubic interpolation, and an image-restoration enhancement (IRE) method. The U-Net model with an EfficientNet-B5 backbone was held constant across all conditions. Paired differences were evaluated using the Wilcoxon signed-rank test with Holm-Bonferroni correction, and effect sizes were quantified using Cohen's *d*.

Results: All enhancement conditions outperformed the no-enhancement condition (adjusted $p < 0.001$). Nearest-neighbor interpolation performed worse than both bicubic interpolation and IRE (adjusted $p < 0.02$), whereas bicubic interpolation and IRE did not differ significantly (adjusted $p = 0.937$; $d = 0.019$), despite IRE achieving higher peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) values.

Conclusion: For this dataset and detector, the task-level benefits of image enhancement plateaued at bicubic interpolation. This finding supports a cautious technology-selection principle: more complex enhancement methods should be justified by measurable downstream gains, while claims of economic superiority require direct measurement of computing, hardware, licensing, and implementation costs.

To cite this article: Huda, M. S., & Setiyono, B. (2026). Diminishing returns of algorithmic complexity: A cost-benefit evaluation of UAV image enhancement for precision weed detection in rice farming. *Journal of Business, Social and Technology*, 7(4), 2372–2383. <https://doi.org/10.59261/bustechno.v7i4.814>

INTRODUCTION

Rice is one of Indonesia's most important crops, both as a staple food and as a component of national food security policy. Weeds pose a major threat to rice yields, particularly when they are not detected early enough Utami and Purdyaningrum (2012) Imaizumi (2018) and Hsu et al. (2025). Farmers typically inspect their fields for weeds manually, but this process is time-consuming and difficult to carry out across the large, geographically dispersed rice-growing areas common in Indonesia. Unmanned aerial vehicles (UAVs), commonly known as drones, offer a

2372 | Journal of Business, Social and Technology

faster way to monitor extensive agricultural areas from the air (Colomina & Molina, 2014; Tahir et al., 2023). This approach aligns with Indonesia's efforts to promote digital agriculture (Kementerian Komunikasi dan Informatika, 2024).

However, UAV-based monitoring depends heavily on image quality. Cameras that capture higher-resolution images can facilitate the detection of small or early-stage weed patches, but such equipment may be prohibitively expensive for smallholder farmers and farmer groups with limited budgets (Ahmed et al., 2026). To address this limitation, researchers have increasingly explored image enhancement methods, ranging from conventional image-processing techniques to more advanced artificial intelligence (AI)-based approaches, to improve the visual quality of images captured by lower-cost cameras (Jonak et al., 2024). These AI-based methods often perform well on standard image-quality benchmarks, contributing to growing interest in their agricultural applications. However, their adoption may be driven by the assumption that more sophisticated enhancement methods consistently deliver greater practical benefits.

This assumption, however, has rarely been evaluated from a task-specific economic perspective. In particular, it remains unclear whether the additional computational complexity of deep learning-based image enhancement yields measurable improvements in weed-detection accuracy sufficient to justify its higher computational and resource requirements. Existing studies frequently emphasize image reconstruction quality metrics, whereas the relationship between reconstruction quality and downstream segmentation performance remains less clearly established.

This study addresses this question directly by comparing several image enhancement methods, ranging from conventional techniques to AI-based approaches, and evaluating their effects on UAV-based weed-detection accuracy in rice fields using the same detection model across all experiments. Rather than focusing solely on which method produces the sharpest images, we investigate whether improvements in image enhancement consistently translate into better weed detection or whether a point of diminishing returns exists beyond which additional algorithmic complexity yields negligible practical benefits. The goal is to provide farmer groups, technology providers, and policymakers with a straightforward, reproducible framework for evaluating the trade-offs between costs and benefits when selecting digital agriculture tools, particularly in settings where financial constraints, rather than technological availability, constitute the primary barrier to adoption (Mhlanga & Ndhlovu, 2023).

METHOD

Economic interpretation in this study is deliberately comparative rather than a full cost-accounting exercise. The relevant parameters include computational demand, processing time, hardware requirements, software or licensing requirements, and implementation complexity. Because the manuscript does not provide direct monetary measurements for each enhancement method, claims regarding cost-effectiveness are treated as conditional inferences based on comparable Intersection over Union (IoU) performance and lower computational complexity, rather than as evidence of measured net cost savings.

The analysis used 148 UAV test images from WeedyRice-RGBMS-DB. The dataset was selected because it provides aligned field imagery suitable for controlled comparisons of image enhancement conditions using an identical downstream detector. Therefore, the findings are most directly generalizable to rice-field imagery with characteristics represented in this dataset, rather than to all crops, growth stages, illumination conditions, or weed distributions.

Figure 1 outlines the workflow used in this study. Each of the four image conditions is processed by the same U-Net model with an EfficientNet-B5 encoder before the results are compared statistically.




Figure 1. Overview of the study steps, from the raw drone image, through the four enhancement conditions, the shared detection model, to the statistical comparison

This study used a quantitative experimental design to investigate how different levels of image enhancement affect UAV-based weed detection in rice fields. The detection model is held constant across all conditions so that observed differences in detection performance can be attributed to the image-enhancement methods rather than to variations in the model architecture.

The data were obtained from the WeedyRice-RGBMS-DB dataset Nguyen et al. (2025), which contains UAV-acquired images of rice paddies at a resolution of 1280×960 pixels and is divided into 438 training images, 148 validation images, and 148 test images. After enhancement, all three methods—nearest-neighbor interpolation, bicubic interpolation, and IRE—produce images at a resolution of 1707×1280 pixels. This resolution is also used as the high-resolution reference for the image-quality metrics reported later in the study. All results presented here are based on the 148 test images, which were held out from model training to enable an unbiased evaluation. Four versions of these test images were prepared: (1) no enhancement, (2) nearest-neighbor interpolation, (3) bicubic interpolation, and (4) a deep learning-based image-enhancement method called IRE Zhu et al. (2023), which is built on the Real-ESRGAN architecture Wang et al. (2021) and Nandal et al. (2024) serves here as an example of an AI-based approach. The internal development and implementation of IRE are beyond the scope of this study.

Each of the four image sets was processed using the same weed-detection model, which consists of a U-Net architecture Ronneberger et al. (2015) and Zhang et al. (2023) with an EfficientNet-B5 encoder (Tan & Le, 2019). The model was trained to segment weed regions in rice-field images. Using the same detection model across all four conditions ensures that differences in detection performance can be evaluated in relation to the image-enhancement condition while minimizing confounding effects from model variation. For each condition, weed-detection performance was evaluated using the Intersection over Union (IoU) score for each of the 148 test images, yielding four sets of 148 paired scores. IoU measures the ratio of the intersection to the union of the predicted and ground-truth weed masks, as expressed in Equation (1):

$$\text{IoU} = \frac{|P \cap G|}{|P \cup G|} \dots \dots \dots (1)$$

Here, P denotes the predicted weed mask, and G denotes the ground-truth weed mask. The term $|P \cap G|$ represents the number of overlapping pixels, whereas $|P \cup G|$ represents the total number of pixels belonging to either mask. IoU ranges from 0, indicating no overlap, to 1, indicating a perfect match.

To determine whether differences between the four conditions were statistically significant, all six pairwise comparisons were performed using the Wilcoxon signed-rank test, which is appropriate for paired observations when the assumptions of parametric alternatives

are not satisfied. Holm–Bonferroni correction was applied to account for multiple comparisons. Cohen’s (d) was also calculated for each comparison to assess effect magnitude and provide an indication of practical significance alongside statistical significance.

This study did not involve human or animal participants; therefore, formal ethical approval was not required. Nevertheless, all data were collected and handled in accordance with applicable institutional ethical practices.

RESULTS AND DISCUSSION

Result

Before examining how well the four conditions support weed detection, it is useful to assess the performance of each enhancement method using standard image-reconstruction quality metrics. Table 1 reports the Peak Signal-to-Noise Ratio (PSNR) Hore and Ziou (2010) and Dohmen et al. (2025) Structural Similarity Index Measure (SSIM) Wang et al. (2021) and Nandal et al. (2024) for the three enhancement methods. These metrics measure how closely each reconstructed image matches a high-resolution reference. Each enhancement method upsamples the original 1280×960 -pixel input to the same 1707×1280 -pixel output dimensions as the reference image, allowing reconstruction quality to be compared at a consistent image size. IRE achieved the highest scores on both metrics (PSNR = 27.54 dB, SSIM = 0.92), followed by Bicubic interpolation (PSNR = 25.68 dB, SSIM = 0.87) and Nearest Neighbor interpolation (PSNR = 23.08 dB, SSIM = 0.79). According to these standard image-quality metrics, IRE produces the reconstruction most similar to the reference image. The ranking is consistent with the expectation that more sophisticated enhancement methods may achieve higher reconstruction quality. However, whether this ranking also holds for the downstream task of weed detection is examined in the following sections.

Table 1

Image Reconstruction Quality by Enhancement Method (PSNR and SSIM)

Model	PSNR (dB)	SSIM
Nearest Neighbor	23.08	0.79
Bicubic	25.68	0.87
IRE	27.54	0.92

Expressed as absolute differences, IRE’s advantage over Bicubic is 1.86 dB in PSNR and 0.05 in SSIM, while its advantage over Nearest Neighbor is considerably larger, at 4.46 dB and 0.13, respectively. PSNR quantifies pixel-level reconstruction error on a logarithmic decibel scale, with higher values generally indicating lower distortion relative to the reference image. SSIM measures structural similarity on a scale from 0 to 1, with values closer to 1 indicating greater similarity in luminance, contrast, and local structure. Both metrics consistently rank the methods in the same order: IRE performs best, followed by Bicubic and Nearest Neighbor, with no ambiguity in their numerical rankings.

In terms of PSNR, IRE outperforms Bicubic by 1.86 dB and Nearest Neighbor by 4.46 dB. Because PSNR is expressed on a logarithmic scale, these differences are more appropriately reported in decibels than as percentages of improvement. The same pattern holds for SSIM: IRE’s 0.05-point advantage over Bicubic is less than half its 0.13-point advantage over Nearest Neighbor. Bicubic’s advantage over Nearest Neighbor is 2.60 dB in PSNR and 0.08 in SSIM, both smaller than IRE’s advantage over Nearest Neighbor but larger than the corresponding SSIM difference between IRE and Bicubic.

This numerical pattern is also visible in the images themselves. Figure 2 shows a close-up comparison of two sample regions under the original input, Nearest Neighbor, Bicubic, and IRE conditions. Nearest Neighbor produces a visibly blocky, pixelated appearance; Bicubic produces smoother edges and a more continuous appearance; and IRE produces sharper-looking fine details, consistent with its higher PSNR and SSIM scores in Table 1.




Figure 2. Close-up comparison of image detail for the input and each enhancement method (Nearest Neighbor, Bicubic, IRE), showing the visual differences behind the numbers in Table 1

A closer examination of the two enlarged regions in Figure 2 reveals the specific visual differences underlying these scores. In the Nearest Neighbor panels, individual pixel blocks are clearly visible as sharp-edged squares, particularly along the diagonal furrows separating rice rows, producing a stair-step effect along curved or angled boundaries. Bicubic interpolation smooths these blocky edges by calculating weighted averages of neighboring pixel values, producing more continuous boundaries but potentially softening fine leaf-level textures. IRE produces additional high-frequency detail relative to the Bicubic output, visible as sharper-looking leaf outlines and finer variations in canopy texture, consistent with the learning-based nature of the method (Fang & Bao, 2024; Goodfellow et al., 2014). However, such apparent detail should not necessarily be interpreted as the recovery of genuine scene information that was absent from the original image.

These differences are particularly visible along the narrow furrows separating adjacent rice rows. Nearest Neighbor renders these furrows as sharp but jagged steps, whereas Bicubic renders them as more continuous but visibly softened lines. IRE produces comparably continuous edges without the prominent blocky artifacts associated with Nearest Neighbor, while retaining more apparent fine-scale leaf texture than Bicubic.

Having established how the enhancement methods differ in image-reconstruction quality, we now turn to the downstream weed-detection task. Before comparing detection accuracy, it is useful to assess whether each condition-specific segmentation run converged to a stable solution. Figure 3 presents the training and validation losses alongside the weed-class intersection over union (IoU) trajectories for all four scenarios. All four runs exhibit a broadly consistent convergence pattern, with validation IoU approaching a plateau near the respective best values (0.7344–0.7440) by approximately epochs 30–45. This pattern suggests that the differences reported below are unlikely to be attributable solely to incomplete or visibly unstable training.

U-Net Training Curves Across the Four Segmentation Scenarios




Figure 3. Training and validation loss, and weed-detection IoU, for the four training scenarios (No Enhancement, Nearest Neighbor, Bicubic, and IRE)

Across all four training runs, the loss curves follow a similar trajectory: a steep decline during the first 5–10 epochs as the model learns to distinguish weeds from rice, followed by a more gradual decline and eventual plateau. The gap between training and validation losses remains relatively small throughout training under all four conditions, providing no obvious indication of substantial overfitting based on these curves alone. The No Enhancement and Nearest Neighbor conditions reach their best validation IoU slightly earlier, at approximately epochs 31 and 39, respectively, while Bicubic and IRE continue to improve marginally until approximately epoch 45. Nevertheless, the differences in final converged IoU across conditions are small relative to the overall improvement observed during training.

Within this plateau range, the ordering of the four conditions' converged validation IoU values broadly mirrors the ordering subsequently observed on the held-out test set. The enhancement methods that achieve the highest test-set accuracy also attain the highest IoU values during this training-time range, whereas No Enhancement consistently occupies the lower end. This suggests that the observed pattern emerges during training rather than appearing exclusively during final evaluation.

Table 2

Weed Detection IoU by Enhancement Condition

Condition	Mean IoU	SD	Median
No enhancement	0.7264	0.1523	0.7594
Nearest Neighbor interpolation	0.7396	0.1418	0.7732
Bicubic interpolation	0.7451	0.1388	0.7827
Deep learning-based enhancement (IRE)	0.7446	0.1419	0.7813

To illustrate how Equation (1) is calculated, suppose that the ground-truth mask in a test image covers 850 weed pixels and the model's predicted mask covers 870 pixels, with 720 pixels of overlap between the two masks. The union of the masks is then $850 + 870 - 720 = 1,000$ pixels, yielding an IoU of $720/1,000 = 0.72$. This value is close to the mean IoU obtained under the No Enhancement condition, as reported in Table 2.

The deep learning-based method (IRE) and Bicubic interpolation achieved the highest and nearly identical mean weed-detection IoU values (0.7446 and 0.7451, respectively). These represent absolute improvements of 0.0182 and 0.0187 over the No Enhancement baseline (0.7264), equivalent to approximately 1.82 and 1.87 percentage points, respectively. Nearest Neighbor interpolation (0.7396) fell between the two methods, improving on the baseline by 0.0132, or 1.32 percentage points. Enhancement also reduced the variability in scores across images: the standard deviation decreased from 0.1523 under No Enhancement to between 0.1388 and 0.1419 under the three enhancement methods. Thus, enhancement was associated with both higher mean IoU and less variability in performance across the test images. Under all four

conditions, the median score exceeded the mean, suggesting that the distributions were negatively skewed, although the extent and causes of this skewness would require further examination. The reduction in standard deviation indicates that the changes in performance were not simply a uniform upward shift in every image's score. However, establishing whether enhancement disproportionately benefited the most challenging images would require a direct analysis of per-image score changes.

Figure 4 illustrates this pattern using a sample test image, with the predicted weed regions displayed for each of the four conditions alongside the corresponding IoU scores. Consistent with the aggregate results, the IoU scores increase from No Enhancement (0.6244) to Nearest Neighbor (0.6557) and Bicubic (0.6684). The IRE result (0.6808) is also relatively close to the Bicubic result, despite the sharper-looking reconstruction produced by IRE in Figure 2. Visually, the predicted weed boundaries under Bicubic and IRE are difficult to distinguish.




Figure 4. Predicted weed areas for each enhancement condition on a sample test image, with the IoU score for that image

Beyond the aggregate IoU scores, the overlays in Figure 4 also illustrate the spatial differences among the four conditions. Under No Enhancement, the predicted weed regions exhibit visibly rougher, less continuous boundaries, with small gaps along the edges of larger weed patches. Nearest Neighbor produces somewhat smoother boundaries but retains some of these gaps. Both Bicubic and IRE produce more continuous, filled-in weed regions with fewer apparent boundary gaps, consistent with their higher IoU scores. The two conditions are visually similar enough that distinguishing between their overlays would be difficult without the accompanying labels.

A comparable pattern appears in the smaller, more scattered weed patches near the edges of the sample image. These patches are only partially captured under No Enhancement, are more completely delineated under Nearest Neighbor and Bicubic, and exhibit a similar degree of completeness under IRE. This visual pattern is consistent with the relatively close aggregate IoU scores of the two highest-performing enhancement conditions reported above.

Taken together, these results establish two distinct patterns. For standard image-quality metrics (Table 1), the three enhancement methods are clearly differentiated, with IRE consistently outperforming Bicubic and Bicubic outperforming Nearest Neighbor. For the downstream weed-detection task (Table 2), however, the same ordering does not fully hold: Nearest Neighbor underperforms Bicubic and IRE, whereas Bicubic and IRE achieve nearly identical mean IoU scores. The following section examines this divergence between reconstruction quality and downstream task performance in greater detail.

Discussion

Cost-effectiveness should therefore be interpreted cautiously. The experiment establishes that the difference in downstream detection performance between Bicubic and IRE was not statistically significant under the tested conditions. However, the economic comparison is based on differences in algorithmic complexity and general implementation requirements rather than a direct accounting of electricity consumption, hardware amortization, licensing, processing latency, or maintenance. A definitive economic conclusion would require these cost components to be measured prospectively.

The divergence between PSNR/SSIM and IoU is consistent with the distinction between reconstruction fidelity and task utility. PSNR and SSIM quantify pixel-level or structural similarity

to a reference image, whereas semantic segmentation depends on whether the features used by the detector—such as boundaries, shape, color contrast, and spatial context—become more discriminative. Once Bicubic interpolation restores sufficient image structure for U-Net with an EfficientNet-B5 encoder to identify these features, additional high-frequency detail from IRE may improve reconstruction metrics without producing a corresponding improvement in the representations used by the segmentation network.

To determine whether the differences in weed-detection accuracy reported above are statistically significant, Table 3 presents pairwise comparisons among all four conditions using the Wilcoxon signed-rank test, with Holm–Bonferroni correction for multiple comparisons.

Table 3

Pairwise Statistical Comparisons (Wilcoxon Signed-Rank Test, Holm-Bonferroni Corrected, $n = 148$)

Comparison	Raw p	Adjusted p	Cohen's d	Sig.
No enhancement vs Nearest	<0.001	<0.001	-0.320	Yes
No enhancement vs Bicubic	<0.001	<0.001	-0.450	Yes
No enhancement vs IRE	<0.001	<0.001	-0.452	Yes
Nearest vs Bicubic	0.00195	0.00586	-0.230	Yes
Nearest vs IRE	0.00839	0.01678	-0.171	Yes
Bicubic vs IRE	0.93668	0.93668	0.019	No

The statistical tests indicated that all three enhancement methods achieved significantly better weed-detection results than No Enhancement (adjusted $p < 0.001$ in all three comparisons), with reported effect sizes ranging from 0.320 to 0.452 in absolute magnitude. Among the enhancement methods, Nearest Neighbor performed significantly worse than both Bicubic (adjusted $p = 0.00586$) and IRE (adjusted $p = 0.01678$). These results indicate that the choice of enhancement method affects detection performance. However, no statistically significant difference was detected between Bicubic and IRE (adjusted $p = 0.93668$, $d = 0.019$), despite the greater algorithmic complexity of the deep learning-based method. This nonsignificant result does not, by itself, establish statistical equivalence between the two methods.

Using Cohen (1988) and Uehata and Rein (2025) conventional guidelines for interpreting effect sizes, the absolute effect sizes for enhancement versus No Enhancement (0.320–0.452) fall within the small-to-medium range, depending on the specific comparison. Thus, although enhancement produced statistically significant improvements, the magnitudes of these improvements were not large. This distinction is important because statistical significance alone does not indicate practical importance. The very small reported effect size between Bicubic and IRE ($d = 0.019$) is also noteworthy: it is substantially smaller than even the smallest effect size reported for enhancement versus the baseline ($d = 0.320$ for Nearest Neighbor). This suggests that the observed difference between Bicubic and IRE is negligible in magnitude in this sample, although the absence of a statistically significant difference should not be interpreted as proof that the methods are equivalent.

These results are based on 148 paired test images and use a nonparametric test appropriate for paired observations, together with a correction for multiple comparisons. These methodological choices strengthen the evidence supporting the observed pattern, although the conclusions remain subject to the assumptions and limitations of the analysis. Three findings point in the same direction: the nearly identical mean IoU scores for Bicubic and IRE in Table 2, the absence of a statistically significant difference between them in Table 3, and the very small reported effect size. Together, these findings suggest that the additional reconstruction quality provided by IRE did not translate into a meaningful improvement in weed-detection performance under the tested conditions.

Taken together, the statistical results reinforce the pattern observed in the Results section: enhancement improves detection relative to the unenhanced baseline, but the additional benefit appears to diminish before the most complex method tested is reached. The next question is why this occurs.

One possible explanation is that both Bicubic and IRE start with the same original 1280×960 -pixel image and produce outputs with the same dimensions of 1707×1280 pixels. Neither method can directly recover scene information that was never captured by the camera. A learning-

based method can, however, synthesize plausible fine-scale details based on patterns learned during training. Such details may improve perceptual quality or benefit tasks that depend on fine textures, particularly when the model can exploit those details reliably. In contrast, distinguishing weeds from rice may depend primarily on relatively coarse visual cues, such as leaf shape, color, and plant boundaries. Bicubic interpolation may already preserve these cues sufficiently for the segmentation network to perform the task effectively. If the weed-detection model does not use the additional fine-scale details introduced by IRE, those details may improve conventional image-quality metrics without improving weed detection. This explanation is consistent with the observed results, although it would require targeted experiments to establish the precise mechanism.

This pattern has a clear potential economic implication: it does not mean that all enhancement methods perform equally well. A sufficient level of image quality appears necessary to obtain a detection benefit, which may explain why the computationally inexpensive Nearest Neighbor method does not achieve the same performance as Bicubic and IRE. However, once a sufficient level of image quality is reached, a more complex method may provide little or no additional return for this particular task. In this study, Bicubic interpolation, a classical and widely available method, achieved detection results that were not statistically distinguishable from those of a deep learning-based method with greater computational and implementation requirements. This finding has practical implications for different stakeholders. Farmer cooperatives could consider Bicubic a reasonable default for comparable applications and allocate limited resources to other priorities, such as additional drone flight time or field-staff training. Technology providers marketing AI-based enhancement methods, meanwhile, should demonstrate measurable improvements in downstream task performance rather than relying solely on general image-quality metrics to justify higher costs.

The potential cost difference is substantial. Recent reviews of drone systems for smallholder farming report that a basic drone setup costs approximately USD 10,000–20,000, with an additional USD 2,000–4,000 per year for maintenance, data processing, and staff training ([Ahmed et al., 2026](#)). These figures do not necessarily include the additional computing resources and specialized personnel that may be required to develop or deploy a deep learning-based method such as IRE. Bicubic interpolation, by contrast, is a relatively simple, fixed mathematical operation that can run on many ordinary computers without model training or specialized inference hardware. Although implementation and maintenance requirements are generally lower, their actual costs depend on the deployment environment and should not be assumed to be zero. This distinction matters for smallholder farmers and cooperatives deciding which technologies to adopt. If a deep learning-based method does not provide a measurable downstream benefit over a simpler alternative once a sufficient level of image quality has been reached, the additional expense may be difficult to justify on cost-effectiveness grounds alone. This consideration is particularly relevant for farmers with limited capital, who constitute a substantial share of Indonesia's rice-farming population ([Badan Pusat Statistik, 2023](#)).

These findings do not imply that advanced image enhancement has no value in precision agriculture. Rather, its value should be evaluated for each task and deployment context instead of being assumed based on general image-quality metrics alone. For applications requiring finer visual detail than basic weed detection, such as identifying specific weed species or detecting early-stage plant diseases, the additional details produced by a learning-based enhancement method may provide a meaningful benefit. Determining whether such benefits occur is an important direction for future research and does not contradict the findings of the present study.

The findings may also depend on the specific detection architecture used. U-Net with an EfficientNet-B5 encoder is an established segmentation architecture, but its performance does not necessarily represent that of all contemporary detection and segmentation models. Whether a more powerful architecture, such as a transformer-based model, would exhibit the same diminishing-returns pattern or better exploit the additional details produced by a learning-based enhancement method remains an open question. Consequently, the cost-effectiveness implications of this study should be restricted to the tested combination of task, enhancement methods, dataset, and segmentation architecture rather than generalized to all image-enhancement applications.

Within the broader context of precision-agriculture research, these findings are relevant

to concerns about the rapid adoption of AI-based farming technologies without sufficient evidence of their practical effectiveness, particularly in developing countries, where cost can be a major barrier to adoption (Rejeb et al., 2022). This concern is consistent with broader bibliometric reviews indicating that technical and remote-sensing themes have received more attention than economic considerations and technology adoption. Research conducted in wealthier agricultural economies may also emphasize technical performance more heavily than the cost constraints that influence adoption among smallholder farmers, who account for a substantial share of the world's agricultural workforce (Kovács & Szöllősi, 2026). By comparing enhancement methods with different levels of complexity in a downstream agricultural task, this study provides evidence that a simpler method can achieve detection performance comparable to that of a more complex alternative under specific conditions. Whether this finding extends to other crops, environments, and agricultural tasks remains to be established.

More broadly, these findings support an important principle in applied computer vision: general image-quality metrics such as PSNR and SSIM are insufficient, on their own, to establish the practical utility of a method for a specific task. Evaluation protocols for agricultural computer vision would therefore benefit from reporting downstream task performance alongside conventional image-quality measurements. This principle is consistent with the broader understanding in machine learning that strong performance on a proxy metric does not necessarily translate into improved performance on the actual objective a system is intended to serve. Future studies proposing image-enhancement methods for agricultural applications should, where feasible, include downstream task evaluation as a standard component of their experimental design.

This study has several limitations. First, it examines only one crop (rice) and one task (weed detection), using images from a specific set of drone-flight conditions in the WeedyRice-RGBMS-DB dataset. Whether the same diminishing-returns pattern holds for other crops, weed species, environmental conditions, or tasks requiring finer visual detail remains unclear. Second, the deep learning-based enhancement method evaluated here, IRE, represents only one possible approach among the many learning-based methods available, including architectures developed specifically for aerial imagery. Other methods may produce different cost-effectiveness outcomes. Third, the detection architecture was fixed to U-Net with an EfficientNet-B5 encoder. Because the study did not evaluate more powerful or structurally different segmentation models, it remains unclear whether the point at which additional enhancement ceases to provide a measurable benefit would change with the detection architecture. Fourth, the economic argument relies on the observed similarity in detection performance and general cost estimates from the broader drone literature rather than on a detailed economic analysis of the specific methods compared. Direct measurements of computational time, energy consumption, hardware requirements, licensing, and deployment costs would strengthen these conclusions and provide a valuable direction for future research.

CONCLUSION

Image enhancement significantly improved weed-detection IoU relative to No Enhancement, but increasing algorithmic complexity beyond Bicubic interpolation did not yield a statistically significant additional benefit at the task level under the tested conditions. IRE achieved higher image-reconstruction quality, yet its weed-detection IoU was not statistically distinguishable from that of Bicubic (adjusted $p = 0.93668$; reported $d = 0.019$). The principal academic contribution is therefore the demonstration that improvements in reconstruction quality do not necessarily translate into corresponding improvements in downstream segmentation performance, supporting task-relevant evaluation rather than reliance on PSNR or SSIM alone.

Practically, these findings suggest that a classical image-enhancement method may be sufficient for comparable rice-field weed-detection tasks. However, the study does not establish comprehensive monetary cost-effectiveness because implementation costs were not measured directly. Future research should validate the observed diminishing-returns pattern across other crop species, growth stages, field conditions, transformer-based segmentation architectures, and tasks requiring finer visual detail. It should also conduct explicit economic analyses of computational time, hardware, energy consumption, licensing, maintenance, and deployment

costs to determine whether the observed performance differences justify the additional expenditure associated with more complex enhancement methods.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the creators and maintainers of the WeedyRice-RGBMS-DB dataset for providing the UAV imagery used in this study. The availability of this dataset made the experimental evaluation and comparative analysis of image enhancement methods for precision weed detection possible. The authors also acknowledge Institut Teknologi Sepuluh Nopember for its academic and institutional support throughout the research process. This research received no external funding from public, commercial, or not-for-profit funding agencies.

AUTHOR CONTRIBUTION STATEMENT

Muhamad Syaiful Huda contributed to the methodology, software development, formal analysis, investigation, data curation, visualization, and preparation of the original manuscript draft. Budi Setiyono contributed to the conceptualization, methodology, validation, supervision, and project administration. Both authors contributed to resource provision, interpretation of the findings, and critical review and editing of the manuscript. Both authors read and approved the final version of the manuscript and agreed to be accountable for the integrity and accuracy of the work.

REFERENCES

- Ahmed, W. A., Ojerinde, J. S., Olatoyinbo, S. F., & Ogaleye, F. J. (2026). UAV technologies for precision agriculture: Capabilities, constraints, and deployment models for smallholder systems in Sub-Saharan Africa. *Drones*, 10(2), 115. <https://doi.org/10.3390/drones10020115>
- Badan Pusat Statistik. (2023). *Hasil Pencacahan Lengkap Sensus Pertanian 2023 - Tahap I*. Badan Pusat Statistik. <https://www.bps.go.id/id/publication/2023/12/15/def0edfb13a6b16411ec8c80/hasil-pencacahan-lengkap-sensus-pertanian-2023---tahap-i.html>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Lawrence Erlbaum Associates.
- Colomina, I., & Molina, P. (2014). Unmanned aerial systems for photogrammetry and remote sensing: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 92, 79–97. <https://doi.org/10.1016/j.isprsjprs.2014.02.013>
- Dohmen, M., Klemens, M. A., Baltruschat, I. M., Truong, T., & Lenga, M. (2025). Similarity and quality metrics for MR image-to-image translation. *Scientific Reports*, 15(1). <https://doi.org/10.1038/s41598-025-87358-0>
- Fang, F., & Bao, S. (2024). FragmGAN: generative adversarial nets for fragmentary data imputation and prediction. *Statistical Theory and Related Fields*, 8(1). <https://doi.org/10.1080/24754269.2023.2272554>
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27. <https://papers.nips.cc/paper/5423-generative-adversarial-nets>
- Hore, A., & Ziou, D. (2010). Image quality metrics: PSNR vs. SSIM. *2010 20th International Conference on Pattern Recognition*, 2366–2369. <https://doi.org/10.1109/ICPR.2010.579>
- Hsu, Y.-T., Ting, C.-Y., Du, P.-R., Li, C.-P., & Wu, D.-H. (2025). Implications of Weedy Rice in Various Smallholder Transplanting Systems. *Agronomy*, 15(12), 2754.
- Imaizumi, T. (2018). Weedy rice represents an emerging threat to transplanted rice production systems in Japan. *Weed Biology and Management*, 18(2), 99–102. <https://doi.org/10.1111/wbm.12146>
- Jonak, M., Mucha, J., Jezek, S., Kovac, D., & Cziria, K. (2024). SPAGRI-AI: Smart precision agriculture dataset of aerial images at different heights for crop and weed detection using super-resolution. *Agricultural Systems*, 216, 103876. <https://doi.org/10.1016/j.agry.2024.103876>
- Kementerian Komunikasi dan Informatika. (2024). Dorong Penggunaan Smart Farming, Kominfo Salurkan Bantuan Layanan IoT untuk Petani.

- Kovács, E., & Szöllősi, L. (2026). Economic aspects of precision crop production: A systematic literature review. *Agriculture*, 16(7), 820. <https://doi.org/10.3390/agriculture16070820>
- Mhlanga, D., & Ndhlovu, E. (2023). Digital technology adoption in the agriculture sector: Challenges and complexities in Africa. *Human Behavior and Emerging Technologies*, 2023, 6951879. <https://doi.org/10.1155/2023/6951879>
- Nandal, P., Pahal, S., Khanna, A., & Rogerio Pinheiro, P. (2024). Super-Resolution of Medical Images Using Real ESRGAN. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3497002>
- Nguyen, V.-H., Le, C.-D., Truong, M.-T., Bui, M.-P. T., & Le, T.-P. (2025). A dataset of aligned RGB and multispectral UAV imagery for semantic segmentation of weedy rice. *Data in Brief*, 63, 112237. <https://doi.org/10.1016/j.dib.2025.112237>
- Rejeb, A., Abdollahi, A., Rejeb, K., & Treiblmaier, H. (2022). Drones in agriculture: A review and bibliometric analysis. *Computers and Electronics in Agriculture*, 198, 107017. <https://doi.org/10.1016/j.compag.2022.107017>
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28
- Tahir, M. N., Lan, Y., Zhang, Y., Huang, W., Wang, Y., & Naqvi, S. M. Z. A. (2023). Application of unmanned aerial vehicles in precision agriculture. In Q. Zaman (Ed.), *Precision Agriculture: Evolution, Insights and Emerging Trends* (pp. 55–70). Academic Press. <https://doi.org/10.1016/B978-0-443-18953-1.00001-5>
- Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning*, 97, 6105–6114. <https://proceedings.mlr.press/v97/tan19a.html>
- Uehata, A., & Rein, R. (2025). Statistical power analysis in exercise science. *Journal of Sports Sciences*, 43(24). <https://doi.org/10.1080/02640414.2025.2571841>
- Utami, S., & Purdyaningrum, L. R. (2012). Struktur Komunitas Gulma Padi (*Oryza sativa* L.) Sawah Organik dan Sawah Anorganik di Desa Ketapang, Kecamatan Susukan, Kabupaten Semarang. *Bioma: Berkala Ilmiah Biologi*, 14(2), 91–95. <https://doi.org/10.14710/bioma.14.2.91-95>
- Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., & Loy, C. C. (2021). Real-ESRGAN: Training real-world blind super-resolution with pure synthetic data. *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 1905–1914. <https://doi.org/10.1109/ICCVW54120.2021.00217>
- Zhang, J., Zhang, Y., Jin, Y., Xu, J., & Xu, X. (2023). MDU-Net: multi-scale densely connected U-Net for biomedical image segmentation. *Health Information Science and Systems*, 11(1). <https://doi.org/10.1007/s13755-022-00204-9>
- Zhu, Z., Lei, Y., Qin, Y., Zhu, C., & Zhu, Y. (2023). IRE: Improved image super-resolution based on Real-ESRGAN. *IEEE Access*, 11, 45334–45348. <https://doi.org/10.1109/ACCESS.2023.3256086>