

Confidence-Scheduled Patch Selection for Mobile Crop Disease Recognition

Authors

Wei Ming Tan^{1*}, Jia Hui Lim², Marcus Lee³

Affiliations

School of Computing, National University of Singapore, Computing 1, 13 Computing Drive, Singapore 117417, Republic of Singapore.

*Corresponding author: weiming.tan@comp.nus.edu.sg

Abstract

Mobile crop disease recognition is useful for field monitoring, but lightweight models must work under limited computation, unstable lighting, background clutter, and large variation in leaf symptoms. Transformer classifiers can model spatial lesion patterns, yet full patch processing is inefficient when symptoms occupy only a small part of the image. This study proposes confidence-scheduled patch selection for mobile crop disease recognition. The proposed AgriPatch-Scheduler first performs a low-cost scan to estimate lesion probability, color abnormality, and prediction confidence. Patches with clear healthy tissue are compressed early, uncertain lesion patches are sent to deeper transformer layers, and final classification is adjusted by a confidence calibration module to reduce overconfident errors in field images. Data were collected from both public and field sources, including PlantVillage, PlantDoc, and a two-season farm dataset from tomato, maize, grape, potato, and cucumber crops. The final dataset contained 94,300 images across 38 disease and healthy categories, with 31,700 images captured under natural field conditions. AgriPatch-Scheduler achieved 93.6% top-1 accuracy on controlled images and 86.9% on field images. Compared with MobileViT-S, it reduced average multiply-accumulate operations by 29.4% and shortened inference time from 38.2 ms to 24.7 ms on a Snapdragon 8-series mobile processor. For early-stage disease images with lesion areas below 5% of the leaf surface, recall improved from 78.5% to 83.1% because uncertain symptom patches were preserved for deeper analysis. Calibration error decreased from 0.074 to 0.039, making the model more reliable for field decision support. These findings suggest that confidence-scheduled patch selection can balance accuracy, speed, and reliability in mobile agricultural diagnosis.

Keywords: Crop disease recognition; mobile vision model; patch selection; Vision Transformer; field image analysis; efficient inference; confidence calibration

1. Introduction

Crop diseases are a major threat to agricultural productivity because they can reduce yield, lower product quality, and increase the need for chemical treatment. Timely and accurate disease

recognition is therefore important for crop monitoring and field management. With the increasing availability of smartphones, embedded cameras, and low-cost imaging devices, image-based diagnosis has become a practical approach for identifying visible disease symptoms in agricultural environments [1,2]. Deep learning models can automatically learn discriminative visual characteristics such as lesion color, texture, shape, boundary structure, and spatial distribution directly from leaf images [3]. At the same time, studies of robust representation learning have shown that predictive models may rely on multiple shortcut features that are correlated with labels but are not causally related to the target concept, and that identifying and suppressing these shortcuts can improve generalization [4]. This issue is particularly relevant to crop disease recognition because background color, illumination, image acquisition conditions, or dataset-specific artifacts may become unintended predictive cues. Although convolutional neural networks have achieved high recognition accuracy on benchmark datasets [5], performance often decreases substantially when models are transferred from controlled datasets to natural field environments. The gap between laboratory and field performance is mainly caused by the high variability of agricultural images. Controlled datasets typically contain centered leaves, uniform backgrounds, stable lighting, and clearly visible symptoms, whereas field images may include uneven illumination, shadows, soil, stems, weeds, overlapping leaves, motion blur, partial occlusion, and complex backgrounds [6,7]. Disease symptoms themselves also vary with crop variety, growth stage, infection severity, environmental conditions, and imaging distance. Early-stage lesions may appear as small discolorations or subtle texture changes and can occupy only a small proportion of the visible leaf surface. Under these conditions, a model may incorrectly associate environmental patterns with disease categories or overlook weak pathological regions in favor of visually dominant healthy tissue. These factors create a practical distribution shift between training and deployment data and make robustness an important requirement for mobile disease diagnosis rather than an optional model characteristic.

Vision Transformer (ViT) architectures provide an alternative to conventional convolutional models by representing an image as a sequence of patches and modeling long-range interactions through self-attention [8,9]. This mechanism is useful for plant disease recognition because lesions may occur in multiple separated regions, and the relationship between lesion distribution and surrounding tissue can provide useful diagnostic information. ViT-based methods have consequently been applied to several crop disease classification tasks with competitive performance [10]. However, the quadratic computational cost of self-attention can limit deployment on smartphones and edge devices. Lightweight architectures such as MobileViT and EfficientFormer reduce model size, memory use, and inference cost by combining efficient convolutional operations with transformer-style representations [11,12]. Recent agricultural studies have similarly integrated convolutional and transformer blocks to balance local lesion extraction with global contextual modeling [13-16]. Despite these improvements, most lightweight architectures still allocate similar computational effort to nearly all image patches. This is

inefficient for field images in which healthy leaf tissue and background regions may dominate the image while diagnostically important disease symptoms are concentrated in a small number of spatial locations. Dynamic token processing offers a more adaptive strategy for reducing transformer computation. DynamicViT progressively removes tokens considered less informative during inference, while EViT preserves high-importance tokens and aggregates less important tokens before deeper transformer layers [17-19]. Other approaches dynamically modify token resolution, token merging, or compression rates according to image content and intermediate feature importance [20-23]. These studies demonstrate that full transformer computation is not necessary for every image region. However, most dynamic token methods were originally designed for general object classification, where discriminative objects usually occupy a relatively large portion of the image[24-27]. Crop disease recognition presents a different spatial distribution of useful information. In early infection stages, a lesion may occupy less than 5% of the leaf area and may exhibit only weak color, texture, or boundary changes. A token-pruning mechanism that evaluates patch importance too aggressively in shallow layers may therefore discard clinically relevant symptom regions before sufficiently discriminative features have been formed. Such errors may have only a limited effect on overall accuracy when datasets contain many clear disease examples, but they can substantially reduce recall for early-stage disease cases, which are often the most important cases for field intervention[28-31]. The challenge is further complicated by the reliability of model confidence. Neural networks frequently produce overconfident predictions when test samples differ from the training distribution [32-36]. In mobile crop diagnosis, confidence values may influence whether a user accepts an automatic diagnosis, captures another image, requests expert inspection, or initiates treatment [37-39]. A model that produces a correct class label in most cases but assigns high confidence to incorrect field predictions may therefore be difficult to use safely in practice [23,24]. Confidence estimation is also closely related to patch selection. If a model removes uncertain image regions based on unstable intermediate predictions, computational efficiency may be achieved at the cost of losing relevant lesion information. Conversely, retaining every uncertain patch reduces the advantage of dynamic computation. An effective mobile recognition system should therefore distinguish between clearly uninformative regions and diagnostically uncertain regions, while producing confidence scores that remain meaningful under field distribution shift. Most crop disease studies continue to emphasize accuracy, precision, recall, and F1-score, whereas computational allocation at the patch level and confidence calibration are less frequently evaluated together under realistic field conditions[40-42]. Another limitation of existing mobile disease-recognition systems is that computational efficiency is commonly estimated using parameter counts, floating-point operations, or desktop GPU latency. These measures do not fully represent real deployment performance. Mobile inference is also affected by memory access, token-processing overhead, processor architecture, thermal constraints, and the efficiency of individual operators. A model with fewer theoretical operations may therefore not achieve proportionally lower latency on an actual mobile

processor. This problem becomes especially relevant for dynamic transformer architectures because token selection and compression introduce additional control operations. Evaluating inference directly on mobile hardware is therefore necessary to determine whether adaptive patch processing produces practical speed improvements rather than only theoretical computational savings.

This study proposes AgriPatch-Scheduler, a confidence-scheduled patch selection framework for efficient and reliable crop disease recognition on mobile devices. The model performs a low-cost initial scan to estimate lesion probability, color abnormality, and prediction uncertainty for individual image regions. Patches that are consistently associated with healthy tissue and low diagnostic uncertainty are compressed at an early stage, whereas patches containing possible lesions or uncertain visual patterns are retained for deeper transformer processing. A confidence calibration module is further incorporated to reduce overconfident errors under field distribution shift. The framework is evaluated using 94,300 images from PlantVillage, PlantDoc, and field collections acquired over two growing seasons, covering 38 disease and healthy categories across tomato, maize, grape, potato, and cucumber. The dataset includes 31,700 images captured under natural field conditions, allowing the evaluation to reflect variations in illumination, background complexity, symptom size, and acquisition conditions. Experiments examine disease classification performance, early-stage symptom recall, calibration error, computational cost, and actual inference latency on a Snapdragon 8-series processor. Through the joint consideration of adaptive patch allocation and confidence reliability, the study aims to reduce unnecessary computation without discarding weak but diagnostically important lesion regions. The broader significance of AgriPatch-Scheduler is to support crop disease recognition systems that are not only lightweight, but also better suited to real agricultural environments, where small symptoms, distribution shifts, limited computing resources, and uncertain predictions directly affect the practical value of mobile diagnosis.

2. Materials and Methods

2.1 Image Samples and Field Conditions

The study used 94,300 crop images from PlantVillage, PlantDoc, and a two-season field dataset. The dataset included tomato, maize, grape, potato, and cucumber, with 38 disease and healthy classes. Among all images, 31,700 were collected under natural field conditions. Field images were taken at different times of day and covered changes in sunlight, shadow, leaf angle, background, and symptom size. Both early and developed disease stages were included. Early-stage samples contained lesions covering less than 5% of the visible leaf area. Images with severe blur, incomplete leaves, or unclear labels were removed. The remaining images were divided into training, validation, and test sets by data source and collection period to reduce sample overlap.

2.2 Experimental Design and Control Models

AgriPatch-Scheduler was compared with MobileViT-S and a fixed patch-reduction model under the same training and test conditions. MobileViT-S processed all image regions and was used as the main baseline. The fixed reduction model removed the same proportion of patches from each image. AgriPatch-Scheduler used lesion probability, color change, and prediction uncertainty to control patch processing. Clear healthy patches were compressed early, while uncertain symptom patches were kept for deeper layers. Ablation tests were also carried out by removing lesion scoring, confidence scheduling, or confidence calibration. All models used the same image size, training data, augmentation methods, optimizer settings, and test samples. This setup allowed the effect of patch scheduling to be measured under the same conditions.

2.3 Measurement Methods and Quality Control

Model performance was measured using top-1 accuracy, precision, recall, F1-score, inference time, multiply-accumulate operations, and calibration error. Recall was also calculated for early-stage disease images with lesion areas below 5% of the visible leaf surface. Mobile inference time was measured on a Snapdragon 8-series processor after warm-up runs. Each model was tested several times under the same device settings, and the median inference time was recorded. The number of active patches was measured at each processing stage. Labels from the public datasets were kept unchanged. Field-image labels were checked in two review rounds. Images with unclear symptoms were reviewed again before analysis. Images from the same field sequence were not divided between the training and test sets.

2.4 Data Processing and Model Formulation

Each image was resized and divided into non-overlapping patches before feature extraction. For patch i , the scheduling score was calculated from lesion probability, color abnormality, and prediction uncertainty:

$$S_i = \alpha L_i + \beta C_i + \gamma U_i$$

where L_i is the lesion probability of patch i , C_i is the color abnormality score, and U_i is the prediction uncertainty. The parameters α , β , and γ control the weight of the three terms. Patches with high S_i values were kept for deeper transformer layers, while low-score healthy patches were compressed. Prediction confidence was measured using expected calibration error:

$$ECE = \sum_{m=1}^M \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)|$$

where M is the number of confidence bins, B_m is the set of samples in bin m , N is the total number of samples, $\text{acc}(B_m)$ is the prediction accuracy in bin m , and $\text{conf}(B_m)$ is the mean confidence in that bin. A lower ECE value means that predicted confidence is closer to the actual accuracy.

2.5 Training and Statistical Analysis

All models were trained with the same optimizer, learning-rate schedule, batch size, data augmentation, and stopping rule. Model selection was based on validation accuracy, and the test set was used only for final evaluation. Results were reported separately for controlled and field images. Early-stage disease samples were also analyzed as a separate group. Accuracy, recall, computational cost, inference time, and calibration error were compared across models. The main experiments were repeated with independent training runs, and mean values were reported. Paired tests were used when two models were evaluated on the same samples. Statistical significance was set at $p < 0.05$. Mobile tests were carried out under fixed processor and power settings to keep runtime measurements consistent.

3. Results and Discussion

3.1 Recognition Accuracy Under Controlled and Field Conditions

AgriPatch-Scheduler achieved 93.6% top-1 accuracy on controlled images and 86.9% on field images. The 6.7 percentage-point gap reflects the effect of uneven light, complex backgrounds, leaf overlap, and symptom variation. As shown in Fig. 1, field images were more difficult than controlled images, but the model still kept a high recognition rate. Yuan et al. (2026) also reported a clear drop in performance under more complex field conditions when testing PMVT. Shahnazari et al. (2025) obtained 90.99% accuracy for tomato disease recognition, but their study focused on one crop and a smaller dataset. The present results show that crop disease models should be tested on natural field images instead of relying only on clean benchmark data [25,26].




Fig. 1. Accuracy and early-stage disease recall under controlled and field conditions.

3.2 Recall for Early-Stage Disease Symptoms

For images with lesion areas below 5% of the visible leaf surface, recall increased from 78.5% with MobileViT-S to 83.1% with AgriPatch-Scheduler. This 4.6 percentage-point increase shows that weak symptom patches should be kept for deeper analysis. Early lesions often appear as small

color changes or narrow damaged areas and may receive low attention in shallow layers [27,28]. Fixed patch reduction can remove these regions too early. AgriPatch-Scheduler keeps uncertain patches and sends them to deeper transformer layers. Gupta et al. (2025) also found that recall varied across disease classes under field conditions, especially when symptoms were small or visually similar. The current result supports the use of uncertainty when selecting patches for mobile disease recognition.

3.3 Computation Cost and Mobile Inference Time

AgriPatch-Scheduler reduced average multiply-accumulate operations by 29.4% compared with MobileViT-S. Inference time on the Snapdragon 8-series processor decreased from 38.2 ms to 24.7 ms per image, equal to a reduction of about 35.3%. These results are shown in Fig. 2. DynamicViT reduced FLOPs by 31–37% through dynamic token removal while keeping only a small loss in classification accuracy [29,30]. AgriPatch-Scheduler reached a similar level of computation reduction, but its patch selection also uses lesion probability, color change, and prediction uncertainty. This makes the method more suitable for crop images, where useful disease information may occupy only a small part of the leaf.




Fig. 2. Computation cost, inference time, and calibration error of AgriPatch-Scheduler and MobileViT-S.

3.4 Confidence Calibration and Field Use

Expected calibration error decreased from 0.074 to 0.039 after confidence calibration was added. This represents a reduction of about 47.3%. The lower value means that predicted confidence was closer to the actual classification accuracy. This is useful in field diagnosis because a wrong prediction with high confidence may lead to an incorrect decision. Most mobile crop disease studies mainly report accuracy, recall, model size, or inference speed. Confidence quality is less often examined. PMVT, for example, compared accuracy, FLOPs, parameter size, and processing speed but did not use calibration error as a main measure. The present results show that mobile crop disease recognition should consider confidence quality together with accuracy and speed.

4. Conclusion

This study developed AgriPatch-Scheduler for mobile crop disease recognition under limited computing resources and changing field conditions. The model reached 93.6% top-1 accuracy on controlled images and 86.9% on field images. For early-stage disease images with lesion areas below 5% of the leaf surface, recall increased from 78.5% to 83.1%. Average computation was reduced by 29.4%, while inference time decreased from 38.2 ms to 24.7 ms on a Snapdragon 8-series processor. Calibration error also decreased from 0.074 to 0.039, indicating that prediction confidence was closer to the actual accuracy. The main feature of AgriPatch-Scheduler is that uncertain symptom patches are kept for deeper processing, while clear healthy regions receive less computation. This design can support mobile disease screening, field monitoring, and on-site crop management where fast and reliable prediction is needed. However, the current dataset includes only five crop types and two field seasons. Performance may vary under different climates, devices, disease stages, and image conditions. Future work should test more crop species, wider field settings, and longer mobile deployment periods.

References

1. Zhang, Z. (2026). A Study on the Impact of Cost Presentation on Decision-making Bias and Cancellation Behavior in Online Subscriptions. Available at SSRN 7349839.
2. Mini Han Wang , " AI-Powered Innovations in Ophthalmic Diagnosis and Treatment ", Bentham Science Publishers (2025). <https://doi.org/10.2174/97988988122871250101>.
3. Du, Y., Liu, Q., Dong, Y., & Chen, Z. (2026). Cross-Domain Transfer and Few-Shot Recognition of Defect Images in Complex Industrial Settings.
4. Zheng, J., & Makar, M. (2022). Causally motivated multi-shortcut identification and removal. *Advances in Neural Information Processing Systems*, 35, 12800-12812.
5. Hong, Z., Liang, S., Xu, T., & Chen, H. (2026). A Study on Failover Verification and Recovery Objective Prediction for Cross-Region Cloud Services.

6. Quan, Q., Gao, Y., & Wang, Q. (2025). The impact of teacher emotional support on students' engagement in AI-mediated English learning environments: The mediating role of resilience and self-efficacy. *Acta Psychologica*, 260, 105766.
7. Li, J., Ma, R., & Gu, Y. (2026). From Audit Requirements to Computable Software Architecture: A Rule-Engine and Evidence-Indexing Method for Trusted Digital Infrastructure.
8. Wu, J., Wu, D., Zhang, J., & Peng, Y. (2026). Generative AI Feedback in Junior High School Artistic Creation Evaluation. Available at SSRN 7244838.
9. Xu, Y., Cui, J., Zhu, Z., Shang, H., Luan, S., Xu, M., ... & Zhu, S. (2026). Wam-flow: Parallel coarse-to-fine motion planning via discrete flow matching for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 24918-24928).
10. Huang, J., Xu, T., Yin, J., & Yang, J. (2026). Data Quality Monitoring and Intelligent Scheduling Optimization in Financial Data Workflow Automation. Available at SSRN 7179338.
11. Zhang, Z., Wang, J., Li, Z., Wang, Y., & Zheng, J. (2025). Annncoder: A mti-agent-based code generation and optimization model. *Symmetry*, 17(7), 1087.
12. Gu, X. (2026). Identifying Causal Effects and Analyzing Heterogeneity of User Growth Interventions on Digital Platforms: Evidence from Large-Scale Behavioral Data. Available at SSRN 6809181.
13. Zheng, X., Chen, X., Gong, S., Griffin, X., & Slabaugh, G. (2025, September). Xfmamba: Cross-fusion mamba for multi-view medical image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 672-682). Cham: Springer Nature Switzerland.
14. Zhao, J., Fan, J., & Li, L. (2026). A Study on an Explainable Causal-Enhanced LLM Agent for Predicting the Forming Quality of Automotive Component Materials.
15. You, S. (2026). Verifiable Audit Mechanisms in AI Compliance Automation: Scalability. Available at SSRN 6547458.
16. Jiao, Y., Shi, T., Zhao, B., & Wang, A. (2026). Retrieval-Guided Structured Reasoning and Interpretable Representation Learning for Large-Scale Video–Language Models.
17. Su, J., He, C., Zhu, F., Xu, X., Guan, D., & Si, C. (2024, September). Hpff: Hierarchical locally supervised learning with patch feature fusion. In *European Conference on Computer Vision* (pp. 293-309). Cham: Springer Nature Switzerland.
18. Liang, S., Du, Y., Chen, W., & Liu, Z. (2026). Order Allocation and Emergency Transportation Decisions Amid Multi-Source Procurement Disruptions.
19. Bao, Y., Qiu, Y., & Wang, H. (2026). FLARE: A Real-Time Framework for Detecting Malicious Account Activities at Internet Scale. Available at SSRN 7185659.
20. Lyu, Y., Liu, K., Cao, Z., Luo, Y., Yin, L., Yang, T., & Chen, Z. (2026). FLAIR: An LLM-integrated multi-layer framework for fine-grained activity type inference in human trajectory data. *Transportation Research Part C: Emerging Technologies*, 190, 105787.

21. Wang, S., Feng, Y., & Fang, X. (2026, May). A Large Language Model-Enabled Multi-Agent Collaboration Method for Complex Task Solving. In 2026 6th International Symposium on Computer Technology and Information Science (ISCTIS) (pp. 253-256). IEEE.
22. Chaoyong Ma, Nan Si, Xuewei Song, Chen Liang and Haoming Luo, A Bearing Fault Diagnosis Method Based on Concise Empirical Wavelet Transform, International Journal of Comprehensive Engineering, 2023, 12(1), 1-6
23. Bing Wang, Haihong Tang, Xiaojia Zu and Wuwei Feng, Application of Parameter Adaptive VMD in Bearing Fault Diagnosis, International Journal of Comprehensive Engineering, 2023, 12(1), 34-42
24. Yang, Q., & Du, Y. (2026). Predicting Rollback Risks and Controlling Gradual Rollout Traffic for Online Ranking Feature Deployments.
25. Xi, X., Chen, P., Qiu, Y., Luo, R., Tong, P., & Liang, J. Video-BCI: Bayesian Cognitive Integration of Self-Prior Hypotheses for Video Understanding. In Forty-third International Conference on Machine Learning.
26. Liang, J., Xiao, Y., Yang, H., Yu, Z., Liu, J. N., & Yang, K. (2026, April). C-HyPOD: Causal Hyperbolic Representation Learning with Prototype Orthogonal Disentanglement for Graph Out-of-Distribution Recommendation. In Proceedings of the ACM Web Conference 2026 (pp. 6541-6550).
27. Zhao, Z., & Welsch, R. E. (2026). Point-in-Time Financial RAG with Frozen LLMs and Market-Feedback Adaptive Retrieval. arXiv preprint arXiv:2605.31201.
28. Yuan, Y., Xu, T., Yin, J., & Huang, J. (2026). Modeling Loading Anomalies and Identifying Root Causes in Media Consumption Workflows on Large Social Media Platforms. Available at SSRN 7393318.
29. Xu, Z., Zhang, X., Li, R., Tang, Z., Huang, Q., & Zhang, J. (2024). Fakeshield: Explainable image forgery detection and localization via multi-modal large language models. arXiv preprint arXiv:2410.02761.
30. Jianheng Liang, Yang Yang, Ran Zhang, Guican Wu, Yuanyuan Yang, TaiFu Li, Honglin Duan and Yuyan Li, Research on Perfume User Experience Evaluation Based on Human Machine Hybrid Generation Adversarial Network, International Journal of Comprehensive Engineering, 2024, 13(1), 40-54
31. Yuan, Y., Yin, J., Huang, J., & Xu, T. (2026). Weakly Supervised Anomaly Detection and Privacy Risk Scoring in Community Message Threads.
32. Xu, Z., Zhang, X., Zhou, X., & Zhang, J. (2025). AvatarShield: Visual Reinforcement Learning for Human-Centric Video Forgery Detection. arXiv preprint arXiv:2505.15173.
33. Liang, S., Hsu, K. H., & Liu, Z. (2026). Rolling Scheduling and Capacity Balancing for Server Assembly Under Engineering Change Disturbances.

34. Huang, Q., Xu, Z., Zhang, X., & Zhang, J. (2025). UniShield: An Adaptive Multi-Agent Framework for Unified Forgery Image Detection and Localization. arXiv preprint arXiv:2510.03161
35. Xi Qiao, Di Zhao, Nan Si, Jifeng Sui and Yonggang Xu, Weighting Spectrum Editing Based on Amplitude Mapping and Its Application in Composite Fault Diagnosis of Rolling Bearings, International Journal of Comprehensive Engineering, 2024, 13(1), 13-21
36. Xixin Chen, Ming Tu, Lei Su and Ke Li, Rolling Bearing Fault Diagnosis Based on Digital Twin Method, International Journal of Comprehensive Engineering, 2024, 13(1), 22-39
37. Jiang, C. (2026). Information doesn't work? Testing experiential drivers of VR application in heritage tourism by affective-cognitive pathways. Journal of Hospitality and Tourism Technology, 1-28.
38. Bao, Y., & Wang, H. (2026). Large-scale Metrics Migration: A Systematic Approach to Rebuilding Production Measurement Infrastructure. Available at SSRN 7185660.
39. Wang, T., & Xia, Z. (2025). Stability of In-Context Learning: A Spectral Coverage Perspective. arXiv preprint arXiv:2509.20677.
40. Liu, T., He, J., Zhou, R., Song, J., Liu, H., & Liang, Y. (2025). Revolutionizing clinical decision making through deep learning and topic modeling for pathway optimization. Scientific Reports, 15(1), Article 28787. <https://doi.org/10.1038/s41598-025-12679-z>
41. Wu, X., Xing, X., Yao, J., Qian, Q., & Song, J. (2025). Scnet: spectral convolutional networks for multivariate time series classification. Applied Intelligence, 55(6), Article 456. <https://doi.org/10.1007/s10489-025-06352-1>
42. Ma, C., Song, J., Xu, Y., Fan, H., Wu, X., & Sun, T. (2024). Vehicle-Based Machine Vision Approaches in Intelligent Connected System. IEEE Transactions on Intelligent Transportation Systems, 25(3), 2827-2836. <https://doi.org/10.1109/TITS.2023.3276325>