

# Causal Shortcut Control for Reliable Traffic Sign Recognition across Weather and Road Scenes

## Authors

Lukas Meier<sup>1\*</sup>, Camille Rochat<sup>2</sup>, Matteo Bernasconi<sup>3</sup>

## Affiliations

Department of Information Technology and Electrical Engineering, ETH Zurich, Gloriastrasse 35, 8092 Zurich, Switzerland

\*Corresponding author: [lukas.meier@ethz.ch](mailto:lukas.meier@ethz.ch)

## Abstract

Traffic sign recognition systems may learn shortcut features from background scenery, camera exposure, road texture, weather conditions, or country-specific sign placement instead of sign shape and symbol content. This limits model reliability in unseen driving environments. This study develops a Causal Shortcut Control Network (CSCN) for robust traffic sign recognition under weather and scene shift. The model identifies shortcut variables by measuring prediction instability across weather, illumination, camera angle, and road-background environments. It then applies sign-region causal attention, background counterfactual replacement, and environment-invariant loss to encourage decisions from sign-relevant visual evidence. The dataset contained 148,300 traffic sign images from 11 cities, covering 43 sign categories, six weather conditions, daytime and nighttime scenes, and three camera types. External testing used 26,400 images from two unseen cities and unseen rainy-night conditions. Compared with EfficientNet-B3, CSCN improved external accuracy from 82.7% to 90.4% and reduced misclassification under heavy rain from 21.8% to 11.6%. Attention evaluation showed that the proportion of model focus located within the sign region increased from 68.2% to 84.9%. These results indicate that causal shortcut control can strengthen traffic sign recognition under real-world driving variation.

**Keywords:** Traffic sign recognition; shortcut control; causal attention; autonomous driving; weather shift; visual robustness

## 1. Introduction

Traffic sign recognition is an essential perception task in advanced driver-assistance systems and automated driving because traffic signs provide explicit information about road regulations, warnings, speed limits, and driving restrictions. Reliable recognition requires not only accurate classification but also stable detection under changes in sign scale, viewpoint, illumination, occlusion, and surrounding traffic conditions. Convolutional neural networks have substantially improved feature extraction for traffic sign recognition, particularly in structured road environments [1]. YOLO-based detectors and related real-time architectures have further improved the balance between detection accuracy and computational efficiency, making traffic

sign recognition more suitable for on-board applications [2]. In China, the CCTSDB 2021 benchmark extended earlier traffic sign datasets by introducing more challenging road scenes and evaluation subsets organized according to sign size, category, and environmental conditions [3]. Recent research on robust machine learning has also shown that predictive models may simultaneously exploit multiple shortcut features and that identifying and suppressing these shortcuts can improve robustness when data distributions change [4-7]. For traffic sign detection specifically, high-resolution attention mechanisms have been introduced to preserve discriminative information from small signs that occupy only a limited number of pixels in an image [5,6]. Multi-frame temporal information has also been used to recover useful visual cues when signs are partly occluded, blurred, or difficult to distinguish in a single frame [7]. Beyond conventional supervised learning, synthetic traffic sign images have been generated to increase training diversity and reduce dependence on large manually labelled datasets [8-12]. Template-based and synthetic-data strategies have also been explored to improve transferability across countries, where sign appearance, road layout, and imaging conditions may differ substantially [10].

Environmental degradation remains one of the major causes of recognition failure in real-world traffic scenes. Rain can introduce streaks, reflections, and local contrast changes, while fog reduces visibility and weakens object boundaries. Snow may partially cover sign surfaces, and low-light conditions can alter colour distributions and suppress fine details. Strong sunlight and glare may produce overexposed regions, whereas dirt or water on the camera lens can generate spatially irregular image corruption. Image restoration has therefore been combined with traffic sign detection to recover degraded visual features before recognition [11]. Related approaches have incorporated enhancement or restoration modules into detection pipelines so that adverse-weather images can be processed more effectively. Rain-removal techniques have been used to recover sign contours and local texture under wet-weather conditions. Multiscale feature extraction has also been adopted to preserve information from small or blurred signs across different feature resolutions [12]. Transformer-based modules can further model long-range relationships and improve feature aggregation in visually complex road scenes [13-17]. Domain adaptation provides another solution by reducing the representation gap between clear-weather training images and adverse-weather target images [14]. These methods have improved recognition performance, but most of them formulate adverse weather mainly as an image degradation or domain discrepancy problem[18-23]. Such a formulation does not directly determine whether a prediction is based on intrinsic sign information or on contextual signals that happen to correlate with a sign class. A model may rely on road texture, surrounding buildings, lane structure, vegetation, camera exposure, typical sign position[24-28], or other background objects when these features provide an easier predictive signal than the sign itself. These correlations may remain stable within the training data while becoming unreliable when the vehicle enters a different city, road system, or weather environment. This limitation is closely

related to distribution shift, spurious correlation, and shortcut learning. Large-scale robustness benchmarks have demonstrated that models with strong in-distribution accuracy can experience substantial performance degradation when the geographic location, acquisition environment, or statistical composition of the test data changes [29-33]. Domain generalization attempts to address this problem by learning representations that remain useful in domains that are unavailable during training [16]. Flat-minimum optimization has been explored as a way to improve generalization by reducing sensitivity to domain-specific parameter configurations [17,18]. Gradient-based matching strategies have similarly been proposed to encourage learning signals that remain consistent across different training environments [19,20]. Although these approaches can reduce the overall performance gap between domains, they do not explicitly distinguish causal sign features from background features that are only statistically associated with particular sign categories. Contrastive representation learning can weaken some spurious associations by encouraging semantically related samples to remain close while separating irrelevant representations [21,22]. Classifier retraining with carefully selected features has also been shown to reduce reliance on misleading background correlations after representation learning [34-37]. However, traffic sign recognition presents a particularly demanding shortcut-learning problem because the target object is often small relative to the full image, while the surrounding scene occupies most of the available pixels. Background information can therefore dominate the learned representation even when the detector correctly localizes the sign. Adverse weather further complicates this problem because degradation may affect the sign and the surrounding environment differently [38-40], creating new correlations that were absent under clear conditions. Robust traffic sign recognition consequently requires a mechanism that not only improves visual quality or feature invariance but also actively controls the information used to make the final prediction [41-43].

This study develops a shortcut-controlled causal network (CSCN) for robust traffic sign recognition across changing road and weather environments. The central objective is to encourage the model to base its decisions on stable information contained within the traffic sign while reducing dependence on background cues that may vary across cities or environmental conditions. CSCN integrates sign-region causal attention to strengthen features that are directly associated with the target sign, background replacement to disrupt statistical relationships between sign categories and surrounding scenery, and an environment-invariant loss to encourage consistent predictive representations across different acquisition conditions. The framework is evaluated on 148,300 images collected from 11 cities, providing substantial variation in road structure, background appearance, weather, illumination, and traffic context. Its generalization capability is further examined using an external test set containing 26,400 images from two previously unseen cities, allowing robustness to be assessed under geographic and environmental distribution shifts rather than only through random train-test partitioning. The study therefore investigates whether explicit shortcut control can improve traffic sign recognition when both weather conditions and

road environments change. Beyond improving recognition accuracy, the work aims to clarify how background dependence contributes to model failure and whether interventions that separate sign-related evidence from contextual correlations can provide more stable predictions. This perspective extends traffic sign recognition from conventional image enhancement and domain adaptation toward causal robustness, offering a practical framework for developing perception systems that remain reliable when automated vehicles encounter road environments that differ from those represented in their training data.

## **2. Materials and Methods**

### **2.1 Samples and Study Areas**

The development dataset included 148,300 traffic sign images from 11 cities. Images were collected on urban streets, suburban roads, and highways. The dataset covered 43 sign classes, six weather conditions, day and night scenes, and three vehicle-mounted RGB camera types. The weather conditions included clear, cloudy, light rain, heavy rain, fog, and snow. Images were recorded at different driving speeds, viewing distances, and camera angles. Each image retained part of the road background. This allowed us to measure the effect of the surrounding scene. Images with missing labels, damaged files, or signs smaller than  $12 \times 12$  pixels were removed. The external test set contained 26,400 images from two other cities. Data from these cities were not used for training or model selection. The external set also included rainy-night scenes. This combined condition was not included in the development data.

### **2.2 Experimental Design and Control Models**

CSCN was the study model, and EfficientNet-B3 was the main control. ResNet-50, ConvNeXt-T, and ViT-B/16 were added for further comparison. All models used the same training data, image size, optimizer, and stopping rule. Three reduced versions of CSCN were also tested. Each version removed one model part: sign-region attention, background replacement, or the environment-invariant loss. During background replacement, the traffic sign remained unchanged, while the road scene was replaced with a scene from another city or weather group. The class label was not changed. This test measured the model's reliance on the road background. The development data were divided into training, validation, and internal test sets at a ratio of 70:15:15. Consecutive frames and images of the same sign were placed in the same set. This prevented the same sign from appearing in both training and testing.

### **2.3 Image Measurement and Quality Control**

Each image was labelled with the sign class, bounding box, sign mask, city, weather, light condition, camera type, and road type. Two trained reviewers checked all labels separately. A third reviewer resolved any differences. Another 10% of the images were checked after labelling. Duplicate images were found by comparing timestamps, sign tracks, and image hashes. Camera

clocks were matched before each recording trip. Image time, location, and camera information were checked after collection. The sign masks were used to measure attention within the sign area. Background-replaced images were checked for sign size, edge quality, colour, and position. Images with clear cutting marks or unnatural sign placement were removed. Model attention was measured with normalized Grad-CAM maps from the last feature layer. Model performance was calculated for each weather, light, camera, city, road, and sign-size group.

## 2.4 Data Processing and Model Formulation

All images were resized to  $300 \times 300$  pixels. Pixel values were normalized with the ImageNet channel means and standard deviations. Training changes included small rotations, position shifts, scale changes, brightness changes, and mild blur. Changes that could alter the sign meaning were not used. CSCN compared predictions across weather, light, camera, and background groups. Features with large prediction changes were marked as possible shortcuts. The training loss was defined as follows:

$$\mathbf{L}_{\text{total}} = \mathbf{L}_c \mathbf{E} + \lambda_a \text{tt}[\mathbf{1} - (\sum_u \mathbf{A}_i(\mathbf{u}) \mathbf{M}_i(\mathbf{u})) / (\sum_u \mathbf{A}_i(\mathbf{u}) + \epsilon)] + \lambda_c \mathbf{f}(\mathbf{1}/N) \sum_i \mathbf{D}_K \mathbf{L}(\mathbf{p}_i || \mathbf{p}_i^{\text{cf}}) + \lambda_i \mathbf{nv}(\mathbf{1}/E) \sum_e (\mathbf{R}_e - \mathbf{R}_{\text{mean}})^2$$

Here,  $\mathbf{L}_c \mathbf{E}$  is the class loss,  $\mathbf{A}_i$  is the attention map,  $\mathbf{M}_i$  is the sign mask,  $\mathbf{p}_i$  and  $\mathbf{p}_i^{\text{cf}}$  are the predictions for the original and background-replaced images, and  $\mathbf{R}_e$  is the prediction risk in environment  $e$ . Attention within the sign region was calculated as follows:

$$\mathbf{ARF} = (\mathbf{1}/N) \sum_i [(\sum_u \mathbf{A}_i(\mathbf{u}) \mathbf{M}_i(\mathbf{u})) / (\sum_u \mathbf{A}_i(\mathbf{u}) + \epsilon)].$$

A higher ARF means that more model attention falls within the sign area

## 2.5 Statistical Analysis

Top-1 accuracy was the main measure. Macro-F1, class recall, heavy-rain error rate, and attention-region fraction were also calculated. Each model was trained five times with different random seeds. Results were reported as the mean and standard deviation. The 95% confidence intervals were calculated from 1,000 bootstrap samples. Sampling was performed at the driving-sequence level because nearby frames were related. Accuracy differences between CSCN and each control model were tested with McNemar's test. Differences in attention-region fraction were tested with a paired permutation test. Separate results were reported for each weather, light, camera, road, city, and sign-size group. The Holm method was used when several tests were performed. All tests were two-sided, and  $p < 0.05$  was considered statistically significant.

# 3. Results and Discussion

## 3.1 Overall Classification Results

CSCN achieved an external accuracy of 90.4%, while EfficientNet-B3 achieved 82.7% (Fig. 1). The gain was 7.7 percentage points. The overall error rate fell from 17.3% to 9.6%, which

represents a relative reduction of 44.5%. The external set came from two unseen cities and included new rainy-night scenes. Thus, the improvement was not caused by familiar road backgrounds or camera settings. Dalveren et al. (2026) improved small-sign detection by changing the feature extraction layers of YOLOv5. Their method achieved 95.0% precision and 91.6% recall on benchmark data. However, their tests mainly examined detection within known datasets. The present study focused on classification under changes in city, weather, camera, and road scene [25,26].




**Fig. 1.** Comparison of external accuracy and heavy-rain error between EfficientNet-B3 and CSCN.

### 3.2 Response to Weather and Scene Changes

Under heavy rain, the error rate of EfficientNet-B3 was 21.8%. CSCN reduced this rate to 11.6%. The absolute decrease was 10.2 percentage points, and the relative decrease was 46.8%. Rain can reduce contrast, weaken sign edges, and hide small symbols. It also changes the appearance of roads and nearby objects. Models that use these scene features may therefore produce unstable predictions. In CSCN, background replacement changed the road scene while keeping the sign unchanged. The invariant loss also reduced differences among weather groups. Wang et al. (2024) combined low-light image enhancement with small-target detection and obtained an mAP gain of about 15% over RetinaNet on the GLARE dataset. Their method was designed mainly for low-light images. CSCN covered six weather types and tested an unseen rainy-night setting. It therefore addressed a wider range of scene changes [27,28].

### 3.3 Attention within the Sign Region

The share of attention within the sign region increased from 68.2% to 84.9% after shortcut control (Fig. 2). This was an increase of 16.7 percentage points. Attention outside the sign fell from 31.8% to 15.1%, equal to a relative decrease of 52.5%. EfficientNet-B3 often focused on lane markings, sign poles, buildings, and bright road areas. CSCN placed more attention on the sign border, shape, and symbol. This change was consistent with the lower error rates in new weather and road scenes [29,30]. Attention maps alone cannot confirm that a model uses causal features. However, CSCN also maintained more stable predictions after background replacement. These results provide stronger evidence that the model relied less on scene-related cues.




**Fig. 2.** Grad-CAM maps of the traffic sign and background regions for EfficientNet-B3 and CSCN.

### 3.4 Practical Value and Limitations

The findings show that shortcut control can improve traffic sign classification in new driving settings. CSCN increased external accuracy without using a larger feature extractor or a separate image enhancement model. This property may support its use with new cameras and in different cities. Several limitations remain. Sign masks were required during training, which increased annotation time. The generated backgrounds may not represent every real road scene. The dataset also covered only 43 sign classes and a limited number of cities. Damaged signs, full occlusion, snow, glare, and country-specific sign designs require further study. Grad-CAM produces coarse attention maps and may miss small symbol details. Future tests should include more countries, seasons, cameras, and rare sign types. Processing speed and safety should also be tested on vehicle hardware before road use.

## 4. Conclusion

This study examined shortcut control in traffic sign recognition across different weather and road scenes. CSCN uses sign-region attention, background replacement, and an invariant loss to reduce the effect of scene features. External accuracy increased from 82.7% to 90.4%, while the heavy-rain error rate decreased from 21.8% to 11.6%. Attention within the sign region increased from 68.2% to 84.9%. These results show that limiting background and weather cues can improve recognition in new cities. CSCN may be used in driver-assistance systems and autonomous vehicles that operate under changing road conditions. However, sign masks require extra

annotation, and the dataset covered only 43 sign classes and a limited number of cities. Future work should test more countries, cameras, sign types, and severe weather conditions. Tests on vehicle hardware are also needed to assess processing speed and road safety.

## References

1. Wu, J., Zhang, J., Wu, D., & Peng, Y. (2026). Visual Transformation Mechanisms for Integrating Digital Cultural Heritage Resources into Junior High School Art Curricula.
2. Huang, J., Xu, T., Yang, J., & Yin, J. (2026). Blockage Early Warning and Financial Impact Quantification in the Monthly Closing Process of Fintech Companies. Available at SSRN 7179198.
3. Que, H., Gao, H., Shan, W., Liu, M., An, J., Deng, F., ... & Mu, L. (2026). FM-SAM: individual tree crown delineation and classification based on Segmentation Anything Model (SAM) and YOLOv10 in UAV imagery for forest monitoring. *Computers and Electronics in Agriculture*, 240, 111162.
4. Zheng, J., & Makar, M. (2022). Causally motivated multi-shortcut identification and removal. *Advances in Neural Information Processing Systems*, 35, 12800-12812.
5. Fan, Y., Miyazaki, T., Tang, Z., Wang, J., Huang, Y., & Omachi, S. (2025, November). Scene Text Reconstructor: A Contextual-Aware Masking Framework for Pre-training Text Detectors. In *International Conference on Neural Information Processing* (pp. 181-195). Singapore: Springer Nature Singapore.
6. 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.
7. 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.
8. You, S. (2026). Verifiable Audit Mechanisms in AI Compliance Automation: Scalability. Available at SSRN 6547458.
9. Mini Han Wang , " AI-Powered Innovations in Ophthalmic Diagnosis and Treatment ", Bentham Science Publishers (2025). <https://doi.org/10.2174/97988988122871250101>.
10. Wu, Y., & Su, D. (2026). An Evaluation of Parental Question Characteristics and the Quality of AI Recommendations in Family Counseling for Children with Autism.
11. Jiao, Y., Shi, T., Zhao, B., & Wang, A. (2026). Retrieval-Guided Structured Reasoning and Interpretable Representation Learning for Large-Scale Video–Language Models.
12. Liang, S., Du, Y., Chen, W., & Liu, Z. (2026). Order Allocation and Emergency Transportation Decisions Amid Multi-Source Procurement Disruptions.
13. Zhao, R., Liu, W., Sha, Z., Su, N., Zhang, Y., & Long, Y. (2026). Culturally uneven urban perception in large language models [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2604.20048v3>.

14. Bao, Y., & Wang, H. (2026). Orion: A High-Throughput, Fault-Tolerant Event Routing Architecture for Hyperscale Data Streams. *Fault-Tolerant Event Routing Architecture for Hyperscale Data Streams* (January 01, 2026).
15. 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.
16. Han, Z., Chen, W., Han, Y., Mao, R., & Qin, J. (2026). Fast Diversified Top-k Rule Discovery via User-Guided Embeddings. *IEEE Transactions on Knowledge and Data Engineering*.
17. Yang, Q., & Du, Y. (2026). Predicting Rollback Risks and Controlling Gradual Rollout Traffic for Online Ranking Feature Deployments.
18. Zhao, Z., & Welsch, R. E. (2024). Hierarchical reinforced trader (hrt): A bi-level approach for optimizing stock selection and execution. *arXiv preprint arXiv:2410.14927*.
19. Yang, Z., Alexandrova, A. N., & Sautet, P. (2026). Modeling CO<sub>2</sub> Hydrogenation to Methanol on an Ensemble of Inverse ZrO<sub>2</sub> on Cu Catalytic Sites: Mechanism, Reactivity, and Deactivation. *Angewandte Chemie*, e5448247.
20. 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.
21. Qi, C., & Qiao, X. (2026). Efficient Data Sampling and Feature Selection Algorithms for Scalable Machine Learning Pipelines.
22. Xiong, W., Zeng, Y., & Guo, Y. (2026). A Study on the Characterization of Building MEP System States and the Learning of Behavioral Patterns Driven by Large-Scale Operational Data. Available at SSRN 7121883.
23. Wang, Z., Yu, J., Liu, H., Zheng, Z., Jin, Y., Li, S., ... & Zhang, K. (2025, July). Generative Music Models' Alignment with Professional and Amateur Users' Expectations. In *Findings of the Association for Computational Linguistics: ACL 2025* (pp. 6909-6920).
24. Hong, Z., Liang, S., Xu, T., & Chen, H. (2026). A Study on Failover Verification and Recovery Objective Prediction for Cross-Region Cloud Services.
25. Huang, R. (2026). Edge-Centric Generative AI: A Survey on Efficient Inference for Large Language Models in Resource-Constrained Environments. *Journal of Computer and Communications*, 14(4), 238-253.
26. Xu, Y., Sun, Y., Zhai, B., Li, M., Liang, W., Li, Y., & Du, S. (2025, April). Zero-shot video moment retrieval via off-the-shelf multimodal large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 39, No. 9, pp. 8978-8986).
27. Liang, S., Hsu, K. H., & Liu, Z. (2026). Rolling Scheduling and Capacity Balancing for Server Assembly Under Engineering Change Disturbances.

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. 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.
30. Bao, Y., & Wang, H. (2026). Large-scale Metrics Migration: A Systematic Approach to Rebuilding Production Measurement Infrastructure. Available at SSRN 7185660.
31. Su, J., Cai, C., Zhu, F., He, C., Xu, X., Guan, D., & Si, C. (2024, September). Momentum auxiliary network for supervised local learning. In *European Conference on Computer Vision* (pp. 276-292). Cham: Springer Nature Switzerland.
32. 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.
33. 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).
34. 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*.
35. 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*.
36. Jia, Z., Huang, C., Wang, Z., Fei, H., Wu, S., & Feng, J. (2024). Finger recovery transformer: toward better incomplete fingerprint identification. *IEEE Transactions on Information Forensics and Security*, 19, 8860-8874.
37. Huang, C., Jia, Z., Fei, H., Zhu, Y., Yuan, Z., Zhang, J., & Zhou, J. (2025, October). ArtFRD: A Fisher-Rao Mixture Metric for Generative Model Aesthetic Evaluation. In *Proceedings of the 33rd ACM International Conference on Multimedia* (pp. 6654-6662).
38. Cui, J., Zhang, G., Chen, Z., & Yu, N. (2022). Multi-homed abnormal behavior detection algorithm based on fuzzy particle swarm cluster in user and entity behavior analytics. *Scientific Reports*, 12(1), 22349.
39. 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.

40. Wang, T., & Xia, Z. (2025). Stability of In-Context Learning: A Spectral Coverage Perspective. arXiv preprint arXiv:2509.20677.
41. 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
42. Yuyan Li, Lin Chen, Mengyu Yang, Qiaomei Li and Taifu Li, Polarization Converters Composed of Metasurfaces with Independent Manipulation of Polarization and Phase, International Journal of Comprehensive Engineering, 2023, 12(1), 7-14
43. Jiaxun Du, Zhenbao Fu, Jiachen Liu, Huaqing Wang, Zhentao Ge and Liuyang Song, Bearing Fault Diagnosis Method Based on Attention Residual Network Under Small Sample Condition, International Journal of Comprehensive Engineering, 2023, 12(1), 15-24