

Enhancing the resilience of CNN-KAN hybrid architectures to environmental distortions in traffic sign classification via focused data augmentation

Urszula Libal, Pawel Biernacki, and Szymon Sołtysiak

Abstract—This paper investigates the efficacy of integrating Kolmogorov-Arnold Networks (KANs) into deep learning architectures to enhance traffic sign recognition robustness against environmental distortions. We evaluate three architectures—a baseline CNN, a hybrid CNN-KAN, and a fully integrated CKAN—using a stochastic augmentation pipeline comprising rotation, simulated rain, blur, and brightness shifts. Experimental results on the GTSRB dataset demonstrate that the hybrid CNN-KAN achieves superior accuracy on clean dataset. KAN-based models excel at handling geometric transformations when trained with augmentation but exhibit higher sensitivity to signal degradation (heavy blur, low light) compared to standard convolutions. We conclude that while KANs offer sharper decision boundaries, their learnable activation functions require extensive augmentation to match the inherent stability of traditional CNNs in low-signal-to-noise environments.

Keywords—Kolmogorov-Arnold networks (KAN), convolutional neural networks (CNN), deep learning, robustness, data augmentation, hybrid architectures

I. INTRODUCTION

REAL-WORLD traffic sign recognition (TSR) systems operate in uncontrolled and highly dynamic environments where ideal visibility is the exception rather than the norm. This presents significant challenges for autonomous driving perception. Physical instability caused by wind or accidents often leads to traffic sign tilt, introducing geometric distortions. Adverse weather conditions introduce rain streaks or fog that occlude critical features. Furthermore, hardware limitations or vehicle motion can result in camera focus problems (blur), and the variable times of day create drastic lighting conditions, ranging from overexposed glare to low-light underexposure.

To address these inherent difficulties and serve as a critical testbed for autonomous driving perception systems, the German Traffic Sign Recognition Benchmark (GTSRB) dataset was introduced [1]. Characterized by severe class asymmetry, varying object scales, and complex environmental surroundings, GTSRB has driven the advancement of state-of-the-art machine learning techniques.

U. Libal and P. Biernacki are with the Wrocław University of Science and Technology, Faculty of Electronics, Photonics and Microsystems, Poland (e-mail: urszula.libal@pwr.edu.pl, ORCID: 0000-0002-3348-510X; pawel.biernacki@pwr.edu.pl, ORCID: 0000-0002-0818-5981).

S. Sołtysiak is with the Wrocław University of Science and Technology, Faculty of Mathematics, Poland (e-mail: szymon.soltysiak.research@gmail.com, ORCID: 0009-0001-2054-2857).

Recent literature has focused on two primary, often conflicting, objectives in TSR: achieving high classification accuracy and ensuring architectural efficiency and robustness against the aforementioned environmental degradations.

A key research direction focuses on designing neural network architectures that are suitable for real-time and resource-constrained embedded systems. These lightweight and efficient models aim to minimize the number of parameters and computational requirements. For instance, enhanced LeNet-5 variants [2] have achieved high accuracy (up to 99.84%) with minimal parameter counts (only 0.38 million). Similarly, ENet [3] has demonstrated performance (up to 98.6% accuracy) with approximately 0.9 million parameters.

To further improve reliability, complex models often incorporate Stacked and Ensemble Approaches. Stacking ensemble models, such as the Hybrid Image Improving and CNN (HI-ICNN) [4], combine the predictions of multiple preprocessed and trained CNNs. This approach has yielded state-of-the-art accuracy (up to 99.75%) and improved robustness under adverse conditions. Alternatively, multi-expert Switching Systems with Gating Networks [5] utilize a combination of several specialized CNNs, with a gating network performing decision fusion, demonstrating superior performance (up to 99.10% accuracy) particularly when facing challenging real-world scenarios.

The comparison of the compact, lightweight models and the high-performing ensemble architectures (performed on the GTSRB dataset) is presented in Table I.

TABLE I: Comparison of recent CNN architectures and their performance on clean GTSRB dataset with only standard augmentation applied, including image resizing, zooming, squeezing, transform to grayscale (for Enhanced LeNet-5), histogram equalization and rotation up to 10° (for HIICNN).

CNN Model / Approach	Accuracy [%] for Clean GTSRB Dataset
Enhanced LeNet-5 [2]	99.84
ENet [3]	98.60
HIICNN Stacking Ensemble [4]	99.75
Multi-expert CNN + Gating [5]	99.10

The models, presented in Table I, were trained using standard augmentation protocols, focusing on balancing the classes and increasing the number of samples in the training set [2]–[4]. The authors of [5] have not mentioned applying any augmentation of images.

Building upon prior research [6], [7], which utilized a single-source augmentation strategy to enhance model robustness, this study addresses the complexity of multi-modal environmental degradation. Specifically, we propose a **generalized data augmentation** approach that simultaneously incorporates various realistic corruptions (e.g., rotation, rain, blur, and brightness shifts). While previous experiments confirmed that models adapt well to isolated disturbances, real-world driving scenarios frequently present concurrent challenges, such as a tilted sign viewed through rain under poor lighting. Consequently, this work advances the training methodology by employing a universal augmentation pipeline where rotation, rain simulation, blur, and brightness shifts are applied simultaneously to every training sample. This approach aims to synthesize a rigorous feature learning environment that better approximates the compound entropy of physical operating conditions.

The paper is structured as follows. Section II provides the theoretical background of Kolmogorov-Arnold Networks and introduces the proposed hybrid models. Section III details the dataset preparation and the stochastic augmentation pipeline used to model environmental disturbances. Section IV reports the experimental results, comparing the architectures in terms of parameter efficiency and resilience to image degradation. Finally, Section V offers concluding remarks and a summary of the contributions.

II. KOLMOGOROV-ARNOLD CONVOLUTIONAL NETWORKS

This study integrates standard CNN components with Kolmogorov-Arnold (KA) layers to create hybrid architectures. We briefly define the three fundamental building blocks used in our models: standard convolutions, KA convolutional layers and KA dense layers.

1) Standard Convolutional Layers:

The standard CNN layer extracts features via discrete linear convolution using kernels of fixed scalar weights. For an input I and a learnable kernel K , the output S is computed as a dot product followed by a static activation function (e.g., ReLU):

$$S(i, j) = (I * K)(i, j) = \sum_m \sum_n I(i + m, j + n) \cdot K(m, n). \quad (1)$$

Here, the complexity is limited by the static nature of the non-linearity applied after the linear transformation.

2) Kolmogorov-Arnold Convolutional Layers: The Convolutional KAN (CKAN) adapts this principle to spatial data [8]. While maintaining the sliding window structure of CNNs, the kernel elements are not scalars but learnable non-linear functions ϕ_{mn} . The output at a specific location is the sum of these functions applied to the input pixels:

$$S_{CKAN}(i, j) = \sum_{m,n} \phi_{mn}(I(i + m, j + n)). \quad (2)$$

This replaces the linear dot product with a non-linear transformation at every kernel position, significantly increasing the feature extraction capacity per layer.

3) Kolmogorov-Arnold Dense Layers: Based on the Kolmogorov-Arnold theorem [9], KA layers replace the fixed activation functions of standard Multilayer Perceptrons (MLP) with learnable univariate functions on every edge. Unlike a standard connection defined by a scalar weight w , a KA connection learns the activation profile itself, typically parametrized as a B-spline:

$$f(x) = \sum_i w_i B_i(x), \quad (3)$$

where $B_i(x)$ are basis functions and w_i are trainable coefficients. This allows the network to dynamically optimize its non-linearities during training.

III. EXPERIMENTS

Experiments were conducted using the German Traffic Sign Recognition Benchmark (GTSRB) [1], selected for its comprehensive representation of real-world environments. This archive contains 51,839 images divided into 43 unevenly distributed classes (39,209 for training and 12,630 for testing). To standardize the input pipeline, we resized all samples to 32×32 pixels and applied tensor normalization.

A. Convolutional Network Architectures

This study evaluates three distinct architectures: a baseline CNN, the proposed CNN-KAN, and a CKAN variant. As illustrated in Fig. 1, all three models share a common topological framework consisting of a feature extractor and a classification head. The distinctions are as follows:

1) CNN: The reference model utilizes a traditional feature extractor composed of three convolutional blocks. Blocks 1 and 2 utilize 3×3 kernels (stride 1, padding 1) to produce 32 and 64 feature maps respectively. Block 3 increases the depth to 128 channels. Each convolution is followed by batch normalization, ReLU activation, and 2×2 max-pooling. The classification head is a standard Multi-Layer Perceptron (MLP) consisting of an input flattening layer and three fully connected layers decreasing in size ($2048 \rightarrow 512 \rightarrow 128 \rightarrow 43$).

2) CNN-KAN: This hybrid architecture retains the identical feature extraction backbone (Blocks 1–3) as the baseline CNN to ensure comparable feature representation. However, the traditional MLP classification head is replaced by a Kolmogorov-Arnold Network (KAN). The flattened feature vector (2048) is processed by a KANet with layers of size 512, 128, and 43. The KAN layers utilize B-splines with a grid size of 5 knots and a spline order of 3.

3) CKAN: This model integrates KAN mechanisms into both feature extraction and classification. The first two blocks remain standard convolutions (32 and 64 channels). The third block replaces the standard convolution with a KAN convolutional layer, expanding the receptive field to 192 channels using a 3×3 kernel. Following the final max-pooling, the 3072-dimensional flattened vector passes through a standard linear projection (reducing dimensionality to 512) before




Fig. 1: The three network models share a common general topology but differ in layer implementation: (a) **CNN**: Standard convolutional backbone (Blocks 1–3) with a dense MLP classifier. (b) **CNN-KAN**: Standard convolutional backbone (Blocks 1–3) with a KAN classifier. (c) **CKAN**: Hybrid backbone introducing a KA-convolution (Block 3), followed by a hybrid MLP/KAN classifier.

entering a final KAN-based classifier (KANet) with layers 512→128→43 (5 knots, spline order 3).

B. Data Augmentation

To rigorously evaluate the impact of training dynamics on model robustness, we compare two distinct learning strategies:

- baseline models trained on the original **clean dataset** (CNN, CNN-KAN, CKAN) versus
- robust variants trained on an **augmented dataset** (Augm-CNN, Augm-CNN-KAN, Augm-CKAN).

To enhance the generalization capability of the Augm- model variants, we employed a focused data augmentation pipeline where four distinct environmental disturbances were simultaneously applied to every training sample, as visualized in Fig. 2. The intensity of each disturbance was sampled stochastically from probability distributions designed to balance frequent mild conditions with rare, extreme scenarios.

Specifically, the **rotation angle** ϕ was drawn from a zero-mean Normal distribution $\phi \sim \mathcal{N}(0, 60^\circ)$, effectively covering the full range of $[-180^\circ, 180^\circ]$ while centering density around the canonical upright orientation. **Brightness changes** were modeled using a multiplicative factor v sampled from a Gamma distribution $v \sim \Gamma(2, 1)$; this distribution peaks at $v = 1.0$ (original brightness) but is right-skewed to simulate common overexposure events.

For image degradation, we utilized Exponential distributions to model strictly non-negative parameters where mild disturbances are most probable. The **rain coverage factor** α followed $\alpha \sim \text{Exp}(1/3)$, resulting in a mean occlusion of



Fig. 2: Examples of the **generalized data augmentation** (right) applied to GTSRB traffic sign images (left). Each sample was simultaneously subjected to four distinct **environmental distortions**: random rotation, simulated rain, Gaussian blur, and brightness level shifts.

33%, while the **blur radius** r was sampled from $r \sim \text{Exp}(8)$ to simulate varying degrees of focus loss. By compositing these random transformations onto each image in the training set, the Augm- models were forced to learn features invariant to geometric, lighting, and textural corruptions, as opposed to the standard models which were trained solely on pristine GTSRB data.

C. Analysis of Learning Dynamics

We employed a unified training configuration: 10 epochs, a batch size of 64, and a 90/10 training-validation split. Model parameters were updated via the Adam optimizer based on categorical cross-entropy loss. To track convergence, we recorded the loss at every step and the validation accuracy




Fig. 3: Learning curves (training vs validation loss) for all models. Top row: standard models. Bottom row: models with augmentation.




Fig. 4: Validation accuracy over epochs.

after every full pass through the dataset, which are shown for each network model in Fig. 3.

The analysis of categorical cross-entropy loss reveals a consistent convergence pattern across all configurations, characterized by a sharp loss reduction at epoch 6 driven by a step-based learning rate decay. For models trained on clean data, all architectures achieved rapid convergence with near-zero training losses by epoch 10. While the baseline CNN followed a stable trajectory, the KAN-based variants displayed minor volatility in early epochs—attributed to the sensitivity of learnable spline activations to high gradients—yet the CKAN ultimately secured the lowest final validation loss (0.0067), confirming its superior precision on pristine data.

Under the augmented training regime, a characteristic inversion emerged where training exceeded validation loss, reflecting the heavy regularization imposed by stochastic disturbances. The Augm-CNN-KAN exhibited the highest opti-

mization difficulty (final training loss of 0.5288), suggesting that fitting B-splines to highly distorted features is computationally challenging. However, the Augm-CKAN demonstrated superior generalization; despite elevated training loss, it achieved a lower validation loss (0.0292) compared to the Augm-CNN (0.0365), indicating that KAN-Convolution layers, while harder to optimize under noise, offer more robust generalization once convergence is achieved.

The validation accuracy progression, presented in Fig. 4, highlights the distinct impact of optimization dynamics on each architecture’s capacity. For the standard models, the CKAN ultimately secures the top performance (99.74%), yet it displays instability during the initial high-learning-rate phase (notably dipping to 95.80% at epoch 5), whereas the baseline CNN follows a smoother, more monotonic trajectory. The application of the learning rate scheduler at epoch 6 triggers a decisive performance jump across all configurations, effectively stabilizing the KAN-based architectures. Under the augmented regime, the Augm-CKAN demonstrates superior long-term convergence (99.40%) compared to the Augm-CNN-KAN, which lags significantly (97.64%). This suggests that fitting learnable activation functions in the classification head is more challenging than optimizing a KAN-based feature extractor when the input distribution is highly destabilized by augmentation.

IV. DISCUSSION

The augmented training set incorporated four random disturbances applied simultaneously to every image: rotation, rain simulation, blur, and brightness intensity changes. We evaluate the six models (CNN, CNN-KAN, CKAN, Augm-CNN, Augm-CNN-KAN, Augm-CKAN) across four perturbation factors: (a) rotation angle ϕ , (b) rain coverage factor α , (c)




Fig. 5: Test accuracy for models (CNN, CNN-KAN, CKAN) trained on clean and augmented data. The accuracy depends on: (a) the rotation angle ϕ , (b) the rain coverage factor α , (c) the radius of blur r , (d) the brightness factor v .

blur radius r , and (d) brightness factor v . These results are illustrated in Fig. 5(a–d).

A. Traffic Sign Tilt

The comparative analysis of models' accuracy under rotational disturbance highlights that integrating Kolmogorov-Arnold Networks (KAN) into the architectures offers distinct advantages over the traditional CNN baseline. The CNN-KAN model, which retains the standard convolutional backbone but employs a KAN-based classifier, achieved the highest performance on clean data (97.70%), surpassing the baseline CNN (95.11%) – see Table II. This suggests that KAN layers are more effective than standard Multi-Layer Perceptrons (MLP) at classifying complex feature sets extracted by convolutional layers. Conversely, the CKAN model, which introduces KAN mechanisms into the feature extraction stage itself, showed a slightly lower baseline accuracy (93.12%), indicating that KAN convolutional layers may be more difficult to optimize or require different training dynamics than standard convolutions on clean datasets.

Regarding robustness, all models demonstrated high sensitivity to rotational disturbances when trained on standard

data, with accuracy dropping precipitously (below 20% at 90°) regardless of the architecture. However, under the training on augmented data, the KAN-based variants exhibited superior generalization capabilities. While the augmentation significantly recovered performance for all models, the Augm-CKAN and Augm-CNN-KAN consistently outperformed the baseline Augm-CNN in challenging scenarios (e.g., 30° and 180° rotations). Notably, the CKAN model achieved the highest stability under extreme rotation (59.02% at 180°), suggesting that replacing spatial convolutions with KAN operations may contribute to better feature invariance when supported by appropriate data augmentation.

B. Rain Occlusion

As presented in Fig. 5b, synthetic rain introduces severe performance degradation across all architectures when trained solely on clean data. The baseline models exhibit a catastrophic drop in accuracy, losing nearly half their performance at moderate rain levels ($\alpha = 0.5$), with the CNN-KAN maintaining only a marginal lead (50.89%) over the standard CNN (48.97%), as shown in Table III. The CKAN model proves the least robust in this regime, suggesting that its

TABLE II: Accuracy [%] of models under image rotations.

Network Model	Rotation Angle ϕ			
	0°	30°	90°	180°
CNN	95.11	45.53	17.07	25.80
CNNKAN	97.70	51.03	14.15	29.40
CKAN	93.12	43.37	18.02	28.38
Augm-CNN	91.52	71.65	49.87	57.05
Augm-CNN-KAN	95.08	74.97	48.69	57.99
Augm-CKAN	91.32	75.55	49.68	59.02

TABLE III: Accuracy [%] of models under rain simulation.

Network Model	Rain Coverage Ratio		
	$\alpha = 0.0$	$\alpha = 0.5$	$\alpha = 1.0$
CNN	95.11	48.97	38.77
CNN-KAN	97.70	50.89	41.08
CKAN	93.12	42.79	31.27
Augm-CNN	91.52	69.75	62.22
Augm-CNN-KAN	95.08	69.92	61.05
Augm-CKAN	91.32	66.41	57.65

TABLE IV: Accuracy [%] of models under blur.

Network Model	Blur Radius			
	$r = 0$	$r = 1$	$r = 4$	$r = 12$
CNN	95.11	91.20	39.05	10.47
CNN-KAN	97.70	92.45	40.73	12.39
CKAN	93.12	87.15	38.25	9.44
Augm-CNN	91.52	88.04	57.70	25.32
Augm-CNN-KAN	95.08	90.11	52.64	21.59
Augm-CKAN	84.84	81.88	52.28	22.68

TABLE V: Accuracy [%] of models under brightness changes.

Network Model	Brightness Coefficient			
	$v = 0.1$	$v = 1.0$	$v = 2.0$	$v = 5.0$
CNN	57.66	95.11	89.77	58.69
CNN-KAN	52.16	97.70	88.07	50.52
CKAN	49.57	93.12	84.55	50.23
Augm-CNN	69.36	91.52	86.73	64.73
Augm-CNN-KAN	59.32	95.08	90.23	65.15
Augm-CKAN	59.48	91.32	86.33	62.82

KAN-based feature extractor struggles to preserve semantic information when high-frequency noise and occlusions, such as rain streaks, disrupt the input image.

Data augmentation significantly mitigates this vulnerability, yet it reveals a divergence in model behavior compared to rotational disturbances. While the Augm-CNN-KAN retains the best performance on clean data, the standard Augm-CNN actually surpasses both KAN-based variants under extreme rain conditions ($\alpha = 1.0$), achieving an accuracy of 62.22% compared to 61.05% for Augm-CNN-KAN and 57.65% for

Augm-CKAN. This indicates that while KAN layers offer superior classification boundaries in stable environments, the traditional MLP classification head may be slightly more robust—or easier to optimize—against the specific textural distortions and signal loss caused by heavy rain overlay.

C. Camera Defocus and Blur

Based on the results, Gaussian blur proves to be a highly destructive perturbation, causing a catastrophic performance collapse in all baseline models (trained on clean data) once the radius exceeds $r = 1$ – compare Table IV. While the CNN-KAN achieves the highest fidelity on slightly blurred images (92.45% at $r = 1$), all standard models fail to generalize to moderate blurring, dropping to $\approx 40\%$ accuracy at $r = 4$. The augmented training strategy, which incorporates randomly blurred images, significantly improves robustness, lifting performance at $r = 4$ by over 12% across the board. However, this comes at the cost of clean-data accuracy for the fully KAN-based Augm-CKAN, which drops to 84.84%, significantly lower than the other architectures.

Interestingly, a trade-off emerges between peak precision and degradation resistance under heavy blurring. The Augm-CNN-KAN retains the highest accuracy on clean and low-blur inputs (90.11% at $r = 1$), confirming that KAN classifiers excel when feature details are preserved. However, as the blur radius increases to $r = 4$ and $r = 12$, the traditional Augm-CNN overtakes both KAN variants, achieving 57.70% and 25.32% respectively. This suggests that while KANs provide superior decision boundaries for high-fidelity signals, their spline-based activations may be more sensitive to the loss of high-frequency edge information than the standard MLP head of the baseline CNN, which appears to generalize better to the low-frequency structural shapes remaining in heavily blurred images.

D. Variable Lighting Conditions

All models exhibit an asymmetric sensitivity to lighting conditions, tolerating overexposure ($v > 1.0$) significantly better than underexposure ($v < 1.0$), what is shown in Fig. 5d and Table V. In the baseline scenario (clean training), the standard CNN demonstrates superior natural robustness to extreme luminosity shifts. While the CNN-KAN achieves the highest peak accuracy at standard brightness (97.70%), it proves more brittle than the conventional CNN when subjected to severe darkening ($v = 0.1$) or blinding brightness ($v = 5.0$), suggesting that the learnable activation functions in KANs may overfit to the specific intensity distributions of the training set more than standard ReLU activations.

Data augmentation successfully bridges this gap in high-brightness scenarios, where the Augm-CNN-KAN recovers to outperform the baseline CNN at $v = 5.0$ (65.15% vs. 64.73%). However, low-light conditions remain a distinct weakness for Kolmogorov-Arnold networks. Even with aggressive augmentation, the Augm-CNN outperforms both KAN-based variants by a wide margin ($\approx 10\%$) at $v = 0.1$. This indicates that standard convolutional layers with fixed activations are more efficient at extracting features from low-contrast, low-intensity

inputs, whereas KANs struggle to fit splines effectively when signal variance is compressed by darkness.

E. Specialized versus Generalized Augmentation

Comparing specific versus universal data augmentation strategies reveals a fundamental trade-off between peak robustness and broad generalization. In the work [7], models trained on single-source disturbances (i.e., **specialized data augmentation**) achieved exceptional resilience within their specific domains. For example, CNN-KAN_{rotation} maintained 82% accuracy at a 30° rotation, while CNN-KAN_{rain} achieved 80% under a 0.5 rain coverage ratio. In contrast, the Augm-CNN-KAN model presented here focuses on simultaneous multi-modal distortions to ensure generalized robustness. Consequently, it achieves slightly lower but stable scores, such as 75% at 30° rotation and 70% for the 0.5 rain coverage ratio.

Admittedly, models trained using the multi-modal augmentation approach (i.e., **generalized data augmentation**) do not achieve the specialized peak performance of models trained on a single disturbance type, typically resulting in a performance decrease in specific corruption tests. However, crucially, they prevent the catastrophic failure observed in baseline models trained on clean data. These baselines (CNN, CNN-KAN, CKAN) suffer performance deficits of 20–30 percentage points compared to the augmented variants Augm-CNN, Augm-CNN-KAN, and Augm-CKAN across all tested categories, including rotation, rain, blur, and brightness changes. This suggests that while specialized training maximizes performance for known constraints, the generalized augmentation pipeline provides a significantly safer baseline. It is therefore better suited for autonomous systems operating in unpredictable environments where the specific degradation type cannot be anticipated.

V. CONCLUSIONS

In this work, we investigated the efficacy of Kolmogorov-Arnold Networks (KANs) when integrated into deep learning architectures for robust image classification, specifically focusing on handling real-world environmental degradations in traffic sign recognition.

Our primary contribution was demonstrating the resilience of the proposed **CNN-KAN hybrid model** when trained using a **generalized multi-modal augmentation protocol**. This protocol involved the simultaneous application of four severe environmental distortions: sign rotation, simulated rain, blur, and brightness shifts. The results conclusively showed that the augmented variant, Augm-CNN-KAN, maintained a competitive classification accuracy of 95.08% even under this comprehensive degradation regimen. For comparison, the standard CNN-KAN trained on clean data achieved 97.70% on the standard benchmark. This robust performance affirms that the hybrid architecture successfully leverages the powerful feature extraction capabilities of Convolutional Neural

Networks (CNNs) with the highly adaptive and expressive classification strength of KANs. Specifically, the KAN-based models showed distinct advantages in mitigating the effects of geometric distortions (e.g., rotation) and environmental occlusions (e.g., simulated rain) when compared to traditional CNNs, offering superior general-purpose resilience.

Despite these gains, the study highlights a vulnerability of KANs to textural and intensity-based disturbances. The standard CNN maintained superior performance under severe blurring and underexposure, suggesting that the fixed activation functions of traditional architectures offer better stability against feature corruption than the dynamic splines of KANs. Therefore, while KANs offer a higher performance ceiling, their deployment in safety-critical systems requires robust training pipelines to mitigate their sensitivity to environmental noise.

While models trained with this generalized approach may not always achieve the specialized peak accuracy of models trained exclusively on a single type of corruption, the substantial overall improvement in robustness under multi-modal degradation makes the proposed strategy highly valuable for real-world autonomous driving applications.

REFERENCES

- [1] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel, "The german traffic sign recognition benchmark: A multi-class classification competition," in *The 2011 International Joint Conference on Neural Networks*, 2011, pp. 1453–1460. [Online]. Available: <https://doi.org/10.1109/IJCNN.2011.6033395>
- [2] A. Zaibi, A. Ladgham, and A. Sakly, "A lightweight model for traffic sign classification based on enhanced lenet-5 network," *Journal of Sensors*, vol. 2021, no. 1, p. 8870529, 2021. [Online]. Available: <https://doi.org/10.1155/2021/8870529>
- [3] X. Bangquan and W. Xiao Xiong, "Real-time embedded traffic sign recognition using efficient convolutional neural network," *IEEE Access*, vol. 7, pp. 53 330–53 346, 2019. [Online]. Available: <https://doi.org/10.1109/ACCESS.2019.2912311>
- [4] G. Yildiz, A. Ulu, B. Dızdaroğlu, and D. Yildiz, "Hybrid image improving and cnn (hiicnn) stacking ensemble method for traffic sign recognition," *IEEE Access*, vol. 11, pp. 69 536–69 552, 2023. [Online]. Available: <https://doi.org/10.1109/ACCESS.2023.3292955>
- [5] A. Ahangi and R. Möckel, "Switching network for mixing experts with application to traffic sign recognition," *Multimedia Tools and Applications*, vol. 82, pp. 43 841–43 864, 2023. [Online]. Available: <https://doi.org/10.1007/s11042-023-14959-0>
- [6] S. Sołtysiak, P. Biernacki, and U. Libal, "Traffic sign recognition with convolutional kolmogorov-arnold networks," in *2025 Signal Processing Symposium (SPSymposium)*, 2025, pp. 147–152. [Online]. Available: <https://doi.org/10.23919/SPSymposium63739.2025.11123989>
- [7] U. Libal, P. Biernacki, and S. Sołtysiak, "Robustness analysis of cnn and convolutional kan architectures against weather and geometric distortions in traffic sign recognition," *International Journal of Electronics and Telecommunications*, vol. 72, no. 2, pp. 1–8, 2026. [Online]. Available: <https://doi.org/10.24425/ijet.2026.157931>
- [8] A. D. Bodner, A. S. Tepsich, J. N. Spolski, and S. Pourteau, "Convolutional kolmogorov-arnold networks," 2024. [Online]. Available: <https://arxiv.org/abs/2406.13155>
- [9] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M. Tegmark, "Kan: Kolmogorov-arnold networks," 2024. [Online]. Available: <https://arxiv.org/abs/2404.19756>