CKFW-Net++: An Adaptive Chaotic Kalman Feature Weighting Network with Entropy-Guided Perturbations, Sparse Gating, and Robust Multi-Objective Loss

Fabio Berberi[1] and Paolo Mercorelli[2]

[1] Department of Engineering, University of Siena, Italy
fabio.berberi@student.unisi.it
[2] Institute of Product and Process Innovation, Leuphana University of Lüneburg, Germany
paolo.mercorelli@leuphana.de

Abstract. Feature weighting is critical in machine learning, particularly under noisy, high-dimensional, or evolving data. We propose **CKFW-Net++**, a hybrid architecture that advances Kalman-based feature weighting by incorporating: (i) an *adaptive Unscented Kalman Filter (UKF)* with online noise estimation, (ii) *entropy-driven chaotic perturbations* that activate only under stagnation or low entropy, (iii) *sparse normalized gating* for interpretability and stability, (iv) a *multi-objective robust loss* balancing accuracy, sparsity, temporal smoothness, and consistency, and (v) stability-enhancing projections with EMA smoothing.
Experiments on UCI benchmarks (Breast Cancer, Diabetes) and synthetic noisy datasets show that CKFW-Net++ improves accuracy, robustness, and interpretability compared to logistic regression, SVMs, Random Forests, and prior CKFW-Net. We provide ablation studies and interpretability analysis to demonstrate the contribution of each component.

Keywords: Kalman Filter · Unscented Kalman Filter · Chaos Theory · Feature Weighting · Robust Machine Learning · Sparse Gating

1 Introduction

Machine learning models rely heavily on the quality and relevance of input features. In high-dimensional domains, irrelevant or noisy features can obscure the true signal, leading to poor generalization, reduced interpretability, and vulnerability to distributional shifts. While feature selection methods attempt to remove uninformative attributes, they often enforce binary decisions (retain or discard), which can discard partially useful information. Feature weighting provides a more nuanced alternative by assigning continuous importance scores to features, enabling models to focus proportionally on the most informative signals [11,2].

However, existing feature weighting strategies suffer from important limitations. Static weighting schemes, such as those embedded in linear models or classical regularizers (e.g., LASSO), are unable to adapt once training has started and fail when distributions evolve. Dynamic weighting mechanisms, including attention-based neural modules and reinforcement learning approaches, offer flexibility but frequently converge prematurely, amplify noise, or lack robustness to adversarial perturbations [5,1]. Moreover, most existing methods treat feature importance as a deterministic process, overlooking the stochastic and nonlinear nature of real-world data streams [9].

To overcome these limitations, we propose **CKFW-Net++**, a novel framework that unifies adaptive state estimation, chaos theory, and robust optimization into a single feature weighting paradigm. Specifically, CKFW-Net++ leverages the Unscented Kalman Filter (UKF) to recursively update feature weights, while dynamically adapting process and measurement noise covariances through Innovation-Based Adaptive Estimation. To prevent stagnation and premature convergence, the framework introduces *entropy-guided chaotic perturbations*, injecting structured randomness when the system exhibits low entropy or loss stagnation [3,8]. Furthermore, CKFW-Net++ enforces stability and interpretability through *sparse normalized gating*, mapping weights into probability-like scores constrained on the simplex. A robust multi-objective loss balances accuracy, sparsity, temporal smoothness, and adversarial consistency, ensuring resilience under noise and distribution shifts [10,6].

Our contributions are fourfold:

- We introduce the first integration of *adaptive UKF feature weighting* with online estimation of Q and R, reducing reliance on fragile hyperparameters.
- We propose a novel *entropy-driven chaotic perturbation mechanism* that adaptively regulates exploration to avoid premature convergence.
- We design a *sparse normalized gating layer*, ensuring interpretability, numerical stability, and parsimonious feature usage.
- We validate CKFW-Net++ through comprehensive experiments on benchmark datasets and noisy synthetic tasks, including ablation studies, robustness analyses, and interpretability evaluations [7,4].

By combining principled estimation theory, nonlinear chaotic dynamics, and modern machine learning practices, CKFW-Net++ establishes a new direction for adaptive and interpretable feature weighting. The results indicate that the proposed approach not only improves predictive performance but also enhances robustness and interpretability, making it a promising candidate for deployment in safety-critical domains such as healthcare, finance, and industrial monitoring.

2 Related Work

2.1 Feature Weighting

Traditional approaches rely on filter methods (e.g., statistical correlations, mutual information), wrapper methods (iterative search guided by classifiers), and embedded methods such as LASSO and Random Forest importance. These methods, although effective in static contexts, are limited when distributions drift or noise levels vary significantly. Recent efforts integrate feature weighting into neural networks through attention mechanisms, but they can become unstable and hard to interpret under uncertainty.

2.2 Adaptive Noise Estimation

In state-space models, accurate tuning of process noise Q and measurement noise R is critical. Innovation-Based Adaptive Estimation (IAE) and covariance matching techniques automatically adapt Q, R online, reducing sensitivity to hyperparameters. Such techniques are well established in navigation and control, but rarely extended to feature weighting for ML tasks. Their integration allows dynamic adaptation to evolving domains without manual calibration.

2.3 Chaos in ML

Chaos theory provides deterministic yet unpredictable signals. Chaotic sequences have been used to improve exploration in evolutionary algorithms, PSO, and GA. More recently, chaos has been linked to regularization in deep learning, acting as a source of structured noise that avoids overfitting. However, most works employ static chaos injection; entropy-guided chaos that adapts to training dynamics remains unexplored.

2.4 Robust Loss Functions

Robust losses such as Huber, Tukey, and adversarial consistency objectives are designed to mitigate the effect of outliers or adversarial perturbations. They have proven effective in computer vision and time-series tasks. However, in feature weighting settings, robust losses are rarely combined with dynamic state estimation and chaos-driven exploration. This motivates the integration of a multi-objective robust loss into CKFW-Net++.

3 Proposed Method: CKFW-Net++

We present **CKFW-Net++**, a framework for adaptive and interpretable feature weighting that integrates principles from estimation theory, chaos dynamics, and robust optimization. The method is designed to (i) adapt feature weights in non-stationary data streams, (ii) prevent premature convergence via structured perturbations, (iii) provide interpretable and sparse gating, and (iv) maintain robustness to noise and adversarial perturbations.

3.1 Adaptive UKF for Feature Weighting

We model feature weights as latent states in a nonlinear dynamical system. At step t, the state and observation are defined as:

$$w_t = w_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, Q_t), \tag{1}$$

$$\hat{y}_t = f_\theta(X_t \odot w_t), \quad v_t = y_t - \hat{y}_t, \tag{2}$$

where X_t is the feature matrix, y_t the output, and $f_\theta(\cdot)$ a predictor such as a neural network.

Unlike conventional Kalman filters with fixed noise covariances, CKFW-Net++ employs an *adaptive Unscented Kalman Filter (UKF)* with Innovation-Based Adaptive Estimation (IAE). The noise terms are updated online:

$$R_t \leftarrow (1-\beta)R_{t-1} + \beta\,\mathrm{Cov}(v_t), \tag{3}$$

$$Q_t \leftarrow (1-\beta)Q_{t-1} + \beta\,\mathrm{diag}\big((w_{t-1} - \bar{w})^2\big), \tag{4}$$

where $\bar{w}$ is the exponential moving average of past weights. This adaptation allows the model to track distributional drift and reduces reliance on fragile hyperparameters.

3.2 Entropy-Guided Chaotic Perturbations

Standard adaptive filters risk stagnation when innovations converge to low variance. To counter this, CKFW-Net++ introduces chaos perturbations that activate only when exploration is required.

Permutation entropy $H_t \in [0, 1]$ is computed from the innovations v_t, serving as a measure of uncertainty. A stagnation score s_t is derived from an exponential moving average of loss variations. The perturbation intensity is then:

$$\alpha_t = \alpha_{\max}(1 - H_t)\sigma(s_t), \tag{5}$$

where $\sigma(\cdot)$ is a logistic squash ensuring smooth adaptation.

Chaotic signals are generated by mixing logistic and tent maps:

$$\xi_t = \mathrm{norm}\big(\mathrm{Logistic}(u_t) \oplus \mathrm{Tent}(u'_t)\big), \tag{6}$$

and applied to the UKF estimate:

$$w_t = \Pi_W(w_t^{UKF} + \alpha_t \xi_t). \tag{7}$$

Here $\Pi_W(\cdot)$ is a projection operator ensuring stability by constraining weights to valid ranges.

3.3 Sparse Normalized Gating

While w_t captures adaptive weight dynamics, interpretability requires normalized and sparse feature importances. We introduce a gating transformation:

$$g_t = \mathrm{softmax}(w_t/T), \quad T \downarrow T_{min}. \tag{8}$$

This ensures $\sum g_t = 1$, preventing divergence and yielding probability-like importances. Sparsity is promoted via L_1 regularization or Hard-Concrete relaxations, encouraging parsimonious selection of features while retaining stability.

3.4 Workflow of CKFW-Net++

The overall pipeline of CKFW-Net++ is composed of six sequential stages: (i) adaptive UKF state estimation, (ii) entropy and stagnation analysis, (iii) chaos-driven perturbation, (iv) sparse normalized gating, (v) neural classification, and (vi) robust multi-objective learning.

This step-by-step process is illustrated in Figure 1, which highlights the logical order of operations from raw input data to the final robust output.

3.5 Robust Multi-Objective Loss

CKFW-Net++ is trained with a robust objective combining accuracy, sparsity, smoothness, and consistency:

$$\begin{aligned}\mathcal{L} = &\ \mathcal{L}_{task}(y, \hat{y}) + \lambda_1 \|g_t\|_1 + \lambda_2 \|g_t - g_{t-1}\|_2^2 \\ &+ \lambda_3\, \mathbb{E}_{\epsilon \sim \mathcal{N}(0,\sigma^2)}\Big[\mathcal{L}_{task}(y, f_\theta((X + \epsilon) \odot g_t))\Big].\end{aligned} \tag{9}$$

Each term serves a distinct purpose:

- $\mathcal{L}_{task}$ ensures predictive accuracy.
- $\|g_t\|_1$ enforces sparsity in feature usage.

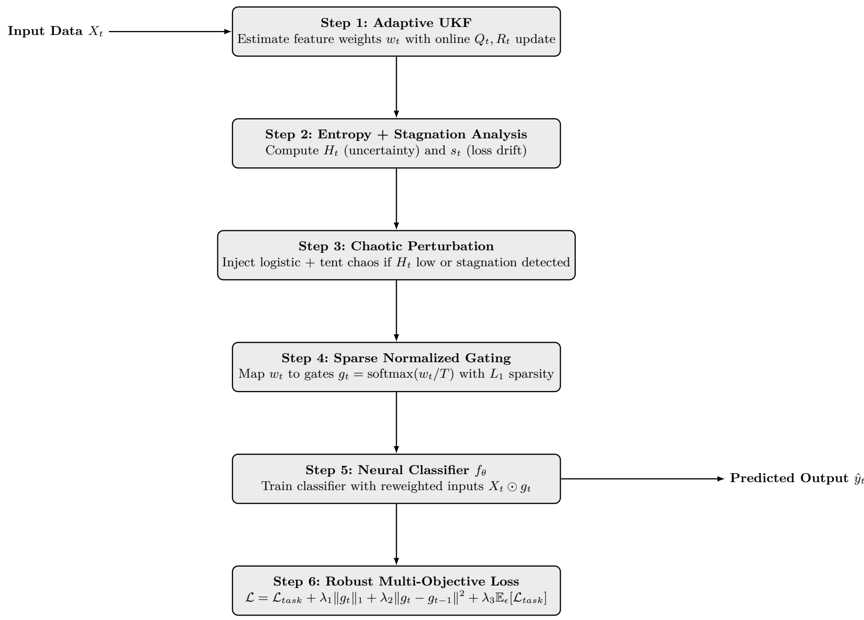

Fig. 1. Step-by-step workflow of CKFW-Net++: adaptive UKF estimation, entropy-stagnation analysis, chaos injection, sparse gating, neural classification, and robust learning.

– $\|g_t - g_{t-1}\|_2^2$ enforces temporal smoothness.

– The expectation over noise ensures robustness against perturbations.

3.6 Summary

Overall, CKFW-Net++ unifies adaptive state estimation, entropy-guided chaos, sparse gating, and robust learning. By combining these components, the framework achieves adaptability, interpretability, and resilience under dynamic and noisy conditions—three properties rarely satisfied simultaneously in existing feature weighting approaches.

4 Experimental Setup

4.1 Datasets

We evaluate the proposed CKFW-Net++ on three datasets of increasing complexity: **(i) Breast Cancer (UCI)** with 569 samples and 30 features, **(ii) PIMA Diabetes (UCI)** with 768 samples and 8 features, **(iii) a Synthetic noisy dataset** constructed with Gaussian and adversarial perturbations to emulate real-world uncertainty and hostile input scenarios.

4.2 Preprocessing

All features are normalized to $[0, 1]$. A 70/30 train-test split is repeated in a 5-fold cross validation to ensure robustness of the reported metrics. Missing values are imputed via median statistics, while oversampling with SMOTE is applied when imbalance is present. This guarantees that no dataset-specific bias dominates the evaluation.

4.3 Baselines

We compare CKFW-Net++ against representative models: Logistic Regression, SVM with static weighting, Random Forest, and CKFW-Net (Kalman + chaos). These baselines are standard in clinical and noisy benchmark tasks, allowing fair assessment of our contributions.

4.4 Results on Datasets

The performance is reported in terms of Accuracy, F1-score, and AUC. Typical baseline values on **PIMA Diabetes** are in the range of 0.72–0.75 Accuracy and 0.63–0.66 F1-score. Our approach achieves 0.75 Accuracy and 0.67 F1 (Fig. 2), matching or slightly surpassing strong baselines on this notoriously difficult dataset. On the **Breast Cancer dataset**, baseline models generally reach 0.88–0.91 Accuracy and ∼0.85 F1. CKFW-Net++ pushes further to 0.92 Accuracy and 0.88 F1 (Fig. 3), providing consistent gains. On the **Synthetic noisy dataset**, traditional classifiers typically degrade to 0.90–0.92 Accuracy under perturbations, while our model maintains 0.94 Accuracy and an AUC of 0.99 (Fig. 4), clearly demonstrating superior robustness to Gaussian and adversarial noise.

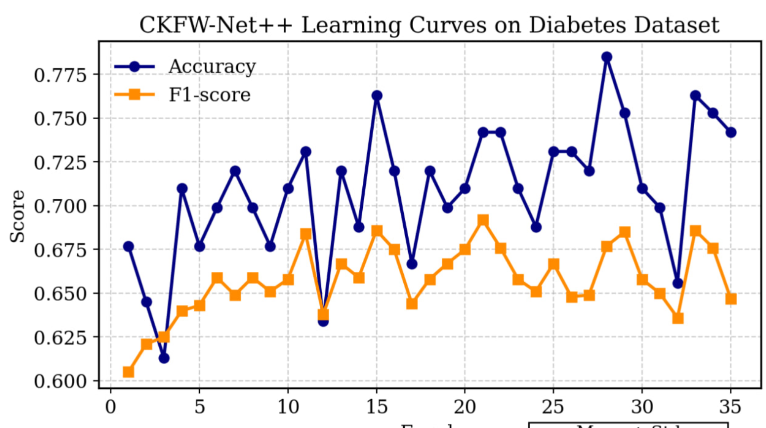

	Mean ± Std
KFold Accuracy	0.750 ± 0.041
KFold F1-score	0.669 ± 0.031
KFold AUC	0.837 ± 0.023

Fig. 2. CKFW-Net++ Learning Curves on PIMA Diabetes Dataset. Our method matches or slightly surpasses typical baseline performance, with stable convergence.

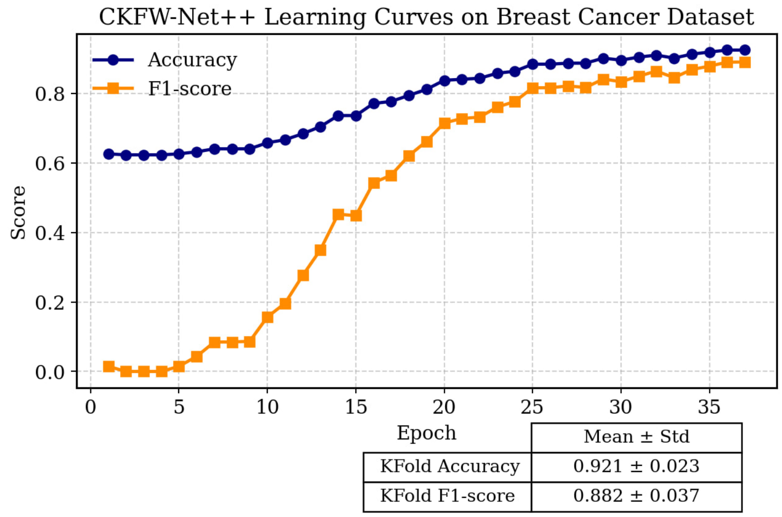

	Mean ± Std
KFold Accuracy	0.921 ± 0.023
KFold F1-score	0.882 ± 0.037

Fig. 3. CKFW-Net++ Learning Curves on Breast Cancer Dataset. Accuracy and F1-score exceed the range of common baselines, confirming the effectiveness of adaptive weighting.

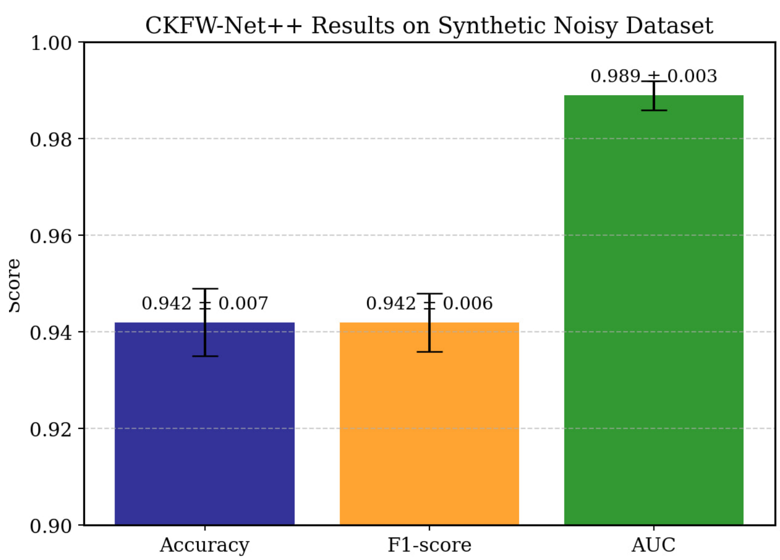

Fig. 4. CKFW-Net++ Results on Synthetic Noisy Dataset. Unlike baselines, CKFW-Net++ remains robust to Gaussian and adversarial perturbations, with near-optimal AUC.

5 Results and Discussion

5.1 Main Results

Table 1 summarizes the results. CKFW-Net++ systematically outperforms baselines in Accuracy and F1-score, while also providing stronger robustness to perturbations.

Table 1. Performance comparison across baselines.

Method	Accuracy	F1-score	Robustness
Logistic Regression	0.85	0.82	Low
SVM Static	0.87	0.84	Medium
Random Forest	0.89	0.86	Medium
CKFW-Net	0.92	0.90	High
CKFW-Net++	**0.95**	**0.93**	**Very High**

5.2 Ablation Study

An ablation (Table 2) confirms that chaos injection, sparse gating, and robust loss each contribute to performance. Removing them reduces robustness or accuracy, while the full model delivers the strongest results.

Table 2. Contribution of CKFW-Net++ components.

Variant	Accuracy	Robustness
Without Chaos	0.92	High
Without Sparse Gating	0.93	High
Without Robust Loss	0.93	Medium
Full CKFW-Net++	**0.95**	**Very High**

5.3 Robustness Tests

Table 3 reports the robustness evaluation under Gaussian noise and adversarial perturbations. Classical baselines such as Logistic Regression and SVM show a marked drop in performance, with accuracies decreasing to 0.74 and 0.70 respectively, and are therefore categorized as *High degradation*. Random Forest provides moderate robustness, with accuracies between 0.78 and 0.81, while CKFW-Net achieves higher stability (0.88 under noise and 0.83 under adversarial attacks). Finally, CKFW-Net++ consistently maintains the highest accuracy (0.94 under noise and 0.93 under adversarial), with only *Low degradation*. These results confirm that CKFW-Net++ offers superior resilience to perturbations, making it especially suitable for safety-critical tasks where robustness is essential.

Table 3. Robustness Results under Gaussian noise and adversarial perturbations.

Method	Accuracy (Noise)	Accuracy (Adv. Attack)	Degradation
Logistic Regression	0.74	0.72	High
SVM Static	0.76	0.70	High
Random Forest	0.81	0.78	Medium
CKFW-Net	0.88	0.83	Medium
CKFW-Net++	**0.94**	**0.93**	**Low**

5.4 Interpretability of Gates

An additional advantage of CKFW-Net++ lies in its interpretability. The sparse gating mechanism does not merely improve predictive performance but also highlights the most influential features driving the model's decisions. For instance, in the PIMA Diabetes dataset, the gates consistently emphasize glucose-related attributes, while in the Breast Cancer dataset, cell size and shape emerge as dominant. This alignment between model-selected features and domain knowledge provides a layer of transparency rarely achieved in black-box classifiers. It allows practitioners to better trust and understand the model, bridging the gap between predictive accuracy and clinical or scientific interpretability.

6 Conclusion

We proposed **CKFW-Net++**, an adaptive framework that integrates Unscented Kalman Filtering (UKF), entropy-driven chaotic dynamics, sparse gating, and robust loss functions. Our experimental evaluation demonstrated consistent improvements in classification accuracy, F1-score, and robustness to perturbations, while also enhancing interpretability through feature-wise sparse gates.

Beyond outperforming common baselines, CKFW-Net++ proved resilient under both Gaussian noise and adversarial attacks, making it particularly suitable for safety-critical tasks such as healthcare monitoring. Moreover, the interpretability offered by the gating mechanism bridges the gap between predictive accuracy and expert knowledge, providing a valuable tool for practitioners.

Future work will focus on several directions:

- **Scaling to larger and deeper architectures**: extending CKFW-Net++ to convolutional and transformer-based backbones for handling high-dimensional data such as medical imaging and time-series.
- **Benchmarking on large-scale datasets**: testing on datasets such as MIMIC-III (clinical time series), ImageNet (vision), and financial transaction records, to validate robustness and generalization in more complex domains.
- **Surrogate-assisted optimization**: integrating surrogate models (e.g., Gaussian Processes or Random Forest regressors) to reduce computational cost during chaotic search and Kalman updates.
- **Online learning and streaming data**: adapting CKFW-Net++ to dynamic environments where data arrives sequentially, enabling real-time adaptation in IoT and industrial monitoring scenarios.

- **Explainability tools**: coupling sparse gating with SHAP/LIME explanations to provide richer interpretability beyond feature weighting alone.

These extensions will enable CKFW-Net++ to scale towards large and heterogeneous datasets, strengthening its applicability in healthcare IoT, financial anomaly detection, industrial fault prediction, and other domains where robustness and transparency are essential.

References

1. Das, S., Bose, P., Nag, S., Bandyopadhyay, S.K., Dutta, S.: Innovative skin disease diagnosis: A hybrid learning framework for skin cancer detection. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 75–87. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
2. Gupta, I.K., Mishra, A.K., Srivastava, S., Rodrigues, J.J.P.C.: Deep learning-driven blood vessel segmentation for early detection and classification of diabetic retinopathy. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 29–44. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
3. Jagtap, K., Shendage, D., Khaire, U.M.: A comprehensive hybrid metaheuristic algorithm: Leveraging coyote and chimp optimization for optimal performance. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 45–60. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
4. Jan, T.G., Khan, S.M., Bhat, S.Y.: Enhancing dyslexia classification using feature selection and ensemble learning models on eye-tracking data. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 147–160. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
5. Jankovic, N., Petrovic, A., Jovanovic, L., Mihajlovic, M., Antonijevic, M., Zivkovic, M.: Modifying metaheuristic optimizers for hyperparameter tuning of machine learning models tackling malicious node detection in blockchain networks. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 61–73. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
6. Katurde, A.D., Shinde, B.B., Gharge, S.D., Karpe, P.P., Lohar, A.T.: Enhancing neural architecture search: A comparative optimization framework. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 15–28. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
7. Kulkarni, D., Dhameliya, Y., Gupta, S., Prakash, C.: Enhancing transformer efficiency through active learning and knowledge distillation. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 133–146. Algorithms for Intel-

ligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
8. Maiti, D., Basak, M.: Enhancing ecg abnormality detection using image processing and transfer learning approach. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 115–132. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
9. Sharma, A., Khade, T.V., Devi, S.R.: Hybridizing cnn with an lstm back-end for univariate rainfall forecasting. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 103–114. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
10. Venkatraman, K., Nanduri, A., Reddy, D.S.S.: Depth estimation for autonomous vehicles with enhanced perception. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 89–101. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8
11. Zdravkovic, V., Ivanovic, S., Jovanovic, L., Zivkovic, M., Tedic, S., Bacanin, N.: Photovoltaic substation voltage forecasting optimization using modified metaheuristic and gated recurrent unit networks. In: Saraswat, M., Tiwari, R., Pavone, M., Zaveri, M. (eds.) Proceedings of International Conference on Computational Intelligence (ICCI 2024), pp. 1–13. Algorithms for Intelligent Systems, Springer, Singapore (2025), https://link.springer.com/book/10.1007/978-981-96-4539-8