

# Learning Interpretable Rules from Neural Networks: Neurosymbolic AI for Radar-Based Hand Gesture Recognition

Sarah Seifi<sup>1,2,\*</sup>, Tobias Sukianto<sup>2,3</sup>, Cecilia Carbonelli<sup>2</sup>, Lorenzo Servadei<sup>1</sup> and Robert Wille<sup>1,4</sup>

<sup>1</sup>Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany

<sup>2</sup>Infineon Technologies AG, Am Campeon 1-15, 85579 Neubiberg, Germany

<sup>3</sup>Johannes Kepler University Linz, Altenbergerstraße 69, 4040 Linz, Austria

<sup>4</sup>Software Competence Center Hagenberg GmbH (SCCH), Softwarepark 32a, 4232 Hagenberg, Austria

## Abstract

Rule-based models offer interpretability but struggle with complex data, while deep neural networks excel in performance yet lack transparency. This work investigates RL-Net, a neuro-symbolic architecture that learns interpretable rule lists through neural optimization, applied for the first time to radar-based hand gesture recognition (HGR). We benchmark RL-Net against a fully transparent rule-based system (MIRA) and an explainable black-box model (XentricAI), evaluating accuracy, interpretability, and user adaptability via transfer learning. Our results show that RL-Net achieves a favorable trade-off, maintaining strong performance (93% F1) while significantly reducing rule complexity. We identify optimization challenges specific to rule pruning and hierarchy bias and propose stability-enhancing modifications. Compared to MIRA and XentricAI, RL-Net emerges as a practical middle ground between transparency and performance. This study highlights the real-world feasibility of neuro-symbolic models for interpretable HGR and offers insights for extending explainable AI to edge-deployable sensing systems.

## Keywords

Interpretable Classification, Rule-Based Learning, Neuro-Symbolic AI, FMCW Radar, Hand Gesture Recognition

## 1. Introduction

Hand gestures offer a natural modality for human-computer interaction, with applications in automotive safety, healthcare, and augmented reality [1, 2, 3]. Radar-based hand gesture recognition (HGR) is particularly attractive due to the radar's robustness, privacy-preserving properties, and compact form factor [4].

Despite recent success of deep learning for radar-based HGR [5], these models are often opaque and lack interpretability, posing challenges in safety-critical and regulated environments, particularly under emerging legislation like the European Union's Artificial Intelligence Act<sup>1</sup>.

Fully transparent, rule-based models provide interpretability but fail to generalize in complex high-dimensional settings. Neuro-symbolic artificial intelligence (AI) offers a promising middle ground by integrating symbolic logic with neural architectures [6], enabling models that learn structured, interpretable rules through gradient-based optimization.

In this work, we explore the RL-Net architecture [7], a neuro-symbolic model that learns ordered rule lists for classification. We apply it to radar-based HGR and evaluate its trade-offs between accuracy and interpretability. To our knowledge, this is the first practical application of neuro-symbolic learning to

---

*Late-breaking work, Demos and Doctoral Consortium, colocated with The 3rd World Conference on eXplainable Artificial Intelligence: July 09–11, 2025, Istanbul, Turkey*

\*Corresponding author.

✉ sarah.seifi@tum.de (S. Seifi); tobias.sukianto@infineon.com (T. Sukianto); cecilia.carbonelli@infineon.com (C. Carbonelli); lorenzo.servadei@tum.de (L. Servadei); robert.wille@tum.de (R. Wille)

🆔 0009-0003-0448-599X (S. Seifi); 0000-0003-4136-8453 (T. Sukianto); 0000-0001-7832-7597 (C. Carbonelli); 0000-0003-4322-834X (L. Servadei); 0000-0002-4993-7860 (R. Wille)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

<sup>1</sup><https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

real-world Frequency-Modulated Continuous Wave (FMCW) radar data, with a focus on interpretability and user adaptation.

Our main contributions are:

1. **First application of RL-Net to radar-based HGR:** We demonstrate the first real-world deployment of a neuro-symbolic model on FMCW radar gesture data.
2. **Training stabilization and architectural improvements:** We enhance RL-Net with batch normalization and validation-time regularization to improve robustness and reduce rule complexity.
3. **Personalized transfer learning:** We apply user-specific fine-tuning to RL-Net, improving accuracy while simplifying learned rule sets.
4. **Comparative evaluation:** We benchmark RL-Net against MIRA (white-box) and XentricAI (black-box), assessing accuracy, interpretability, and adaptability.
5. **Limitations and future work:** We identify structural training issues in RL-Net's fixed rule hierarchy and outline lightweight improvements to support scalable, interpretable deployment.

## 2. Background and Related Work

### 2.1. Interpretable White-box Models

Transparent models like decision trees, logistic regression, and rule-based systems are essential for high-stakes applications due to their interpretability [8]. Rule-based models, structured as "if-then" statements, align with human reasoning and are typically organized as rule sets or hierarchical rule lists [9].

A rule comprises multiple feature-based conditions leading to a prediction, e.g., "IF ( $x_1 > 0.5$  AND  $x_2 \leq 3.0$ ) THEN class = 1". This clarity supports traceability and user trust.

MIRA [10] exemplifies a transparent rule-based classifier designed for multi-class HGR. It incorporates foundational and user-personalized rules but suffers from limited scalability and overfitting as dimensionality increases.

### 2.2. Explainable Black-box Models

Deep Neural Networks (NN) (convolutional NNs [11], recurrent NNs [4], Transformers [12]) yield strong radar HGR performance but lack interpretability. XAI techniques like SHAP [13] offer post-hoc feature attributions but are faithful to the model, not the underlying data [14]. They fail to reveal intermediate reasoning steps.

Recent work such as XentricAI [15] applies SHAP to recurrent NNs for gesture explanation and anomaly feedback, combining gesture detection and classification per frame. While insightful, it remains a black-box at its core.

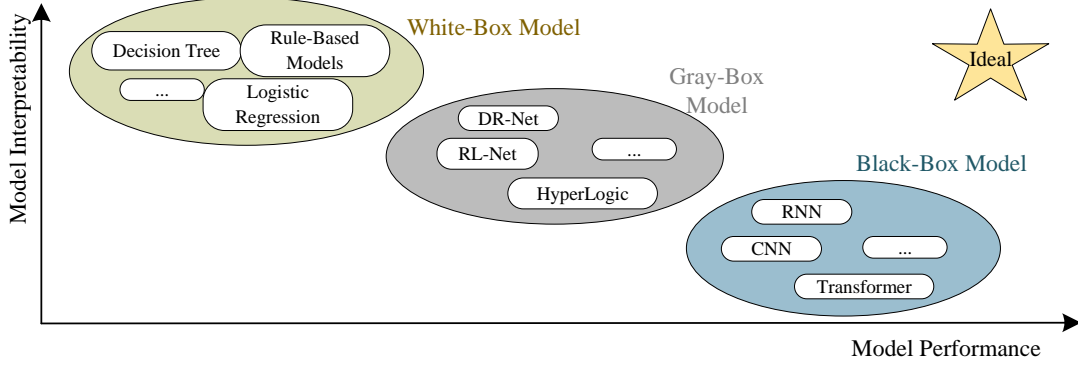
### 2.3. Neuro-Symbolic AI (Gray-box Models)

Neuro-symbolic AI bridges rule-based interpretability and neural flexibility [6]. These models generate readable rules through trainable NNs, producing interpretable outputs from opaque training, often termed "gray-box" systems [16].

DR-Net [17] models rules using a two-layer NN: an AND-based rules layer followed by an OR layer. Rules emerge from binary neuron activations with sparsity-based regularization to control complexity.

RL-Net [7] extends DR-Net with fixed-order rule hierarchy, producing interpretable rule lists suitable for edge deployment. More advanced systems like HyperLogic [18] improve scalability via hypernetworks but at the cost of interpretability and complexity, making them less ideal for transparent and on-the-edge applications.

Figure 1 summarizes model trade-offs between interpretability and performance, positioning RL-Net between black-box deep models and fully symbolic approaches.



**Figure 1:** Comparison of model categories (white-box, gray-box, black-box) based on model complexity and predictive performance. Neuro-symbolic approaches like RL-Net occupy an intermediate "gray-box" space, balancing interpretability and performance. An ideal model has high model interpretability as well as performance.


## 2.4. Motivation for Our Work

To date, neuro-symbolic models have not been applied to real-world FMCW radar gesture data. Their adaptability via Transfer learning (TL) also remains underexplored. This work fills that gap by evaluating RL-Net in comparison to MIRA and XentricAI, analyzing the performance–interpretability trade-off and exploring its potential for user-specific gesture adaptation.

## 3. Methodology

### 3.1. RL-Net Overview

We apply RL-Net [7] to FMCW radar gesture data. The pipeline involves signal preprocessing, feature extraction, model training, and interpretable rule generation, as illustrated in Figure 2.



**Figure 2:** Overview of the proposed interpretable gesture recognition pipeline.

**Input & Rule Layers.** The input layer receives binarized features  $x_i \in \{0, 1\}$ . Each rule neuron functions as a logical AND, with ternary weights  $\mathbf{W}_T$  encoding positive, negative, or excluded feature use. To enable differentiable training, these weights are reparameterized as  $\mathbf{W}_T = \mathbf{W}_S \circ \mathbf{W}_H$ , where  $\mathbf{W}_H \in [0, 1]$  are approximated binary masks following [19].

A neuron activates only when all conditions are satisfied:

$$y = \sum_i w_i x_i - \sum_{w_i > 0} w_i + 1$$

where  $w_i \in \mathbf{W}_T$  and the step function  $b(y) = 1$  if  $y = 1$ , else 0, binarizes the output.

To improve training stability, we add batch normalization.

**Hierarchy & Output Layers.** A fixed hierarchy enforces ordered rule evaluation: a rule  $R_k$  activates only if all previous rules are inactive. The final neuron acts as a default rule. Only one rule fires per sample, and a softmax assigns its associated class label.

**Loss Function.** The regularization loss is defined as  $\mathcal{L}_{\text{sparse}} = \sum_{i,j} \sigma\left(\frac{\text{loc}_{ij} - \gamma}{\zeta - \gamma}\right)$ , where  $\text{loc}_{ij}$  are the logits controlling the binary mask sampling, and  $\gamma, \zeta$  are stretch parameters controlling the thresholding behavior. This penalty encourages many weights to converge to zero and thus shortens the learned rules. The training optimizes the loss based on the cross-entropy loss, the sparsity and L2 regularization terms:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda_1 \mathcal{L}_{\text{sparse}} + \lambda_2 \|\mathbf{W}_{\text{out}}\|_2^2$$

### 3.2. Comparison to Baselines

We compare RL-Net to MIRA [10], a white-box rule-based system using foundational and user-specific rules, and XentricAI [15], a black-box GRU model explained post-hoc via SHAP.

A summary of capabilities is shown in Table 1.

**Table 1**

Comparison of explainable and interpretable HGR approaches.

Method	Explainability Type	Adaptability	Scalability
MIRA [10]	White-box (fully interpretable)	High (personalized rules)	Low
XentricAI [15]	Black-box (post-hoc explanations)	High (TL)	High
RL-Net [7]	Gray-box (semi-interpretable)	Medium-high (TL)	Medium-high

## 4. Dataset and Preprocessing

### 4.1. FMCW-Radar Gesture Dataset

The gesture dataset used in this study was collected using Infineon Technologies’ XENSIV™ BGT60LTR13C 60 GHz FMCW radar operating in indoor environments. Twelve users performed five predefined hand gestures (*SwipeLeft*, *SwipeRight*, *SwipeUp*, *SwipeDown*, and *Push*) across six different room types. A *Background* class was added to represent the absence of gestures. Each participant completed 1,000 samples, totaling 31,000 gesture recordings.

Recordings were captured over 100 frames per gesture using three receive antennas, with dimensions  $[100 \times 3 \times 32 \times 64]$  representing frames, antennas, chirps, and fast-time samples, respectively. The dataset is publicly available via IEEE Dataport [20].

### 4.2. Frame-Level Gesture Labeling

Following the method in [4], we assign gesture labels based on the frame with minimum radial distance (gesture anchor point). A fixed 10-frame window centered around this point defines the gesture duration, while all other frames are labeled as background. This results in a labeled dataset with dimensions  $[M \times F \times D]$ , where  $M$  is the number of samples,  $F$  the number of frames per sample, and  $D = 5$  the number of features.

### 4.3. Signal Processing and Feature Extraction

We applied a lightweight, real-time-capable preprocessing pipeline based on [4] to extract gesture-relevant features from the raw radar data. The main steps include: (i) **Range FFT**, applied after DC removal to generate range profiles; (ii) **Moving Target Indication (MTI)**, which removes static target reflections; (iii) **Target Localization**, using peak detection in the range profile to find the closest moving object (i.e., the hand); (iv) **Doppler FFT**, applied to the identified hand bin to extract radial velocity; and (v) **Angle Estimation**, computing azimuth and elevation from antenna phase differences.

From this processing, we extract five features per frame: radial distance (range), radial velocity (Doppler), azimuth angle, elevation angle, and signal magnitude. These were averaged over the ten gesture frames to construct the input to the models.

## 5. Experimental Setup

To evaluate both the generalization capabilities and adaptability of RL-Net, we conducted two main experiments: (i) optimized training of a user-agnostic pretrained model, and (ii) user-specific TL. Gesture data from twelve users (12,000 samples) was used in total. Data from six users (6,000 samples) was used to train and validate the base model, while data from the remaining six users was reserved for TL and evaluation.

**Dataset Split.** Six users (6,000 gestures) were used to pretrain all models; six different users (6,000 gestures) were held out for TL and testing.

**Training Protocol.** An extensive hyperparameter search was performed using Grid Search, selecting values based on prior work and empirical stability. We used the Adam optimizer with a learning rate of 0.01 and batch size of 40. Training was run for 200 epochs with early stopping on validation loss. A batch normalization layer was added after the rule layer for improved stability. Sparsity was controlled via a regularization weight ( $\lambda_{1,\text{train}} = 0.025$ ), while a higher weight ( $\lambda_{1,\text{val}} = 0.3$ ) was applied during validation to promote simpler models. L2 regularization was disabled. Hard concrete parameters were  $\gamma = -0.1$ ,  $\zeta = 1.1$ , and  $\beta = \frac{2}{3}$ .

**Custom Early Stopping Strategy.** Standard early stopping based on validation cross-entropy often favored overly complex models with minimal gains in accuracy. On the other hand, strong regularization led to premature convergence and vanishing gradients. To balance interpretability and performance, we introduced a custom early stopping criterion using an increased validation-time regularization weight ( $\lambda_{1,\text{val}} = 0.3$ ). This steered model selection toward sparser rule sets without compromising predictive accuracy.

**Transfer Learning.** For user adaptation, pretrained weights were partially frozen: rule masks and output weights were updated while base rule weights were fixed. Each user had 1,000 samples (20% test, 64% train, 16% val).

**Evaluation Metrics.** We report accuracy and F1 score, along with rule complexity measured by the number of active rules and total rule conditions. These metrics are tracked before and after fine-tuning.

**Baselines.** RL-Net was compared to MIRA [10], and XentricAI [15]. Note that XentricAI performs both gesture detection and classification by explicitly modeling a background class, so comparisons should consider this broader output scope. For both baselines, we adopt the hyperparameters and training protocols reported in their original publications.

## 6. Experimental Results and Discussion

### 6.1. General Training and Model Robustness

While we initially experimented with DR-Net by training independent rule sets for each gesture class, the resulting rules were not mutually exclusive, making them unsuitable for robust multi-class classification. As a result, we focused exclusively on RL-Net, which provides a hierarchical rule list enabling consistent multi-class predictions.

We evaluate the model’s robustness and interpretability across three architectural variants: the original RL-Net, RL-Net with batch normalization, and RL-Net with both batch normalization and a modified validation loss incorporating regularization weight  $\lambda_{1,\text{val}}$ . Results, averaged over four runs per configuration with consistent seeding, are summarized in Table 2.

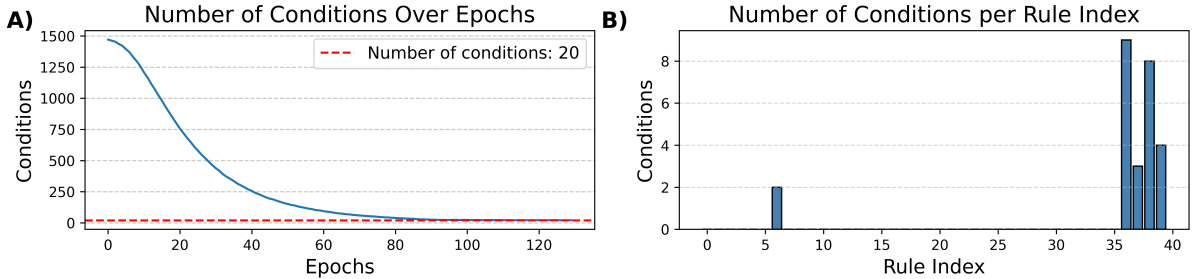
Introducing batch normalization significantly improved both performance stability and model simplicity. Further incorporating validation-time regularization reduced model complexity even more,

**Table 2**Performance and complexity across RL-Net architectural variants (mean  $\pm$  std).

Architecture	F1 Score [%]	Conditions
Original	89.25 $\pm$ 3.80	270 $\pm$ 64.03
+ BatchNorm	<b>90.74 <math>\pm</math> 1.00</b>	98 $\pm$ 59.20
+ BatchNorm + $\lambda_{1,\text{val}}$	90.03 $\pm$ 1.50	<b>69 <math>\pm</math> 22.00</b>

with only a slight trade-off in F1 score. The reduction in standard deviation across runs also indicates improved training consistency.

**Optimization Bottleneck from Rule Ordering and Vanishing Gradients.** While training improves overall sparsity (Figure 3, Panel A), we observed a persistent structural issue: RL-Net’s fixed hierarchy layer causes early rules to being suppressed and pruned. As shown in Panel B, this results in only high-index rules remaining active, often with long, complex conditions, compromising rule diversity and interpretability. This bottleneck stems from the fixed top-down rule evaluation, which prioritizes later rules during optimization. Addressing this may require adaptive or learnable rule ordering, ideally without increasing model complexity to maintain suitability for edge deployment.



**Figure 3:** Training dynamics of RL-Net. **(A)** Total number of active rule conditions over training epochs, showing progressive sparsity and improved interpretability. **(B)** Histogram of surviving rule indices and their number of conditions in a representative run. Due to the fixed rule evaluation hierarchy, early rules are pruned, while later rules dominate, highlighting a structural imbalance.

**Comparison with Baselines.** As a reference, the white-box model MIRA achieves an F1 score of 79.7%, while the GRU-based classification backbone of XentricAI reaches 85%. RL-Net demonstrates a compelling trade-off: It achieves strong performance ( $\sim 90\%$ ) while maintaining interpretable, compact rule lists. This positions RL-Net effectively between the extremes of full transparency and pure black-box modeling.

## 6.2. Transfer Learning

For these experiments, we initialized the model using a randomly selected baseline model consisting of seven rules and 36 total conditions.

Table 3 summarizes the classification performance and rule complexity before and after TL. All results are reported for the best-performing epoch on the validation set, using accuracy and F1-score as primary metrics, and the number of rules and total rule conditions as proxies for interpretability.

Across all users, TL consistently improved model performance, with F1-score increases ranging from modest (+0.71% for user<sub>3</sub>) to substantial (+9.08% for user<sub>1</sub>). Importantly, these gains were accompanied by a consistent reduction in rule complexity. All adapted models converged to five rules, and in several cases, the total number of rule conditions decreased significantly, highlighting the dual benefit of enhanced personalization and improved interpretability.



**Table 3**

TL results per user: Accuracy (Acc), F1-score (F1), and rule complexity.

User	Initial Acc [%]	Initial F1 [%]	TL Acc [%]	TL F1 [%]	#Rules	#Conditions
1	84.5	84.52	95.5	95.46	5	29
2	91.0	91.11	93.0	92.97	5	11
3	89.5	89.50	90.5	90.21	5	30
4	95.3	95.15	95.8	95.81	5	36
5	94.0	94.01	95.0	95.02	5	18
6	85.0	84.18	91.0	90.91	5	36
<b>Avg.</b>	<b>89.9</b>	<b>88.75</b>	<b>93.05</b>	<b>93.06</b>	<b>5</b>	<b>26.67</b>

**Comparison with Baselines.** RL-Net achieved an average user-specific performance of 93.05%, surpassing the fine-tuned XentricAI model, which reached 90.2%. However, it is important to note that XentricAI includes a *Background* class and performs both gesture detection and classification at the frame level, an extended functionality that goes beyond RL-Net’s current scope. MIRA achieved the highest average accuracy of 94.9% after user-specific calibration. While deterministic, MIRA relies on handcrafted rule tuning and lacks the learning flexibility of neural approaches.

Overall, RL-Net strikes a favorable balance between interpretability and performance, offering greater modeling flexibility than MIRA and better classification accuracy than XentricAI. That said, one limitation may lie in the fixed thresholds used for binarizing continuous features, which could constrain the model’s adaptability to subtle user-specific variations. Future research should explore adaptive thresholding or learnable binarization strategies to further improve generalization and personalization.

### 6.3. Limitations and Future Work

RL-Net inherits optimization challenges from DR-Net, notably vanishing gradients caused by multiplicative logical activations. This leads to early-rule pruning and dominance of later, often longer rules, limiting generalization. Fixed rule ordering further amplifies this by prioritizing high-indexed neurons and suppressing earlier ones.

While batch normalization improved stability, training remains inconsistent across runs. Attempts like L2 regularization worsened gradient flow and performance. Though more advanced techniques (e.g., HyperLogic [18]) exist, they add model complexity, compromising interpretability and edge-suitability.

Future work should explore gradient-stable activations, adaptive thresholding for input binarization, and learnable rule hierarchies. Physics-informed constraints tailored to gesture kinematics may further improve robustness without added complexity. These insights extend beyond HGR, pointing to broader improvements in neuro-symbolic learning.

## 7. Conclusion

This work presents the first real-world application of RL-Net, a neuro-symbolic model, to radar-based HGR. RL-Net achieves a strong balance between interpretability and performance, outperforming fully transparent models and approaching the accuracy of black-box methods.

Through optimized training and TL, we demonstrate RL-Net’s potential for user-adaptive, explainable gesture sensing. However, challenges in optimization and rule hierarchy remain. Addressing these without increasing model complexity is key to advancing interpretable AI for edge deployment and beyond.

## Declaration on Generative AI

The author has not employed any Generative AI tools.

## References

- [1] E. Ohn-Bar, M. M. Trivedi, Hand gesture recognition in real time for automotive interfaces: A multimodal vision-based approach and evaluations, *IEEE transactions on intelligent transportation systems* 15 (2014) 2368–2377.
- [2] N. Al Mudawi, H. Ansar, A. Alazeb, H. Aljuaid, Y. AlQahtani, A. Algarni, A. Jalal, H. Liu, Innovative healthcare solutions: robust hand gesture recognition of daily life routines using 1d cnn, *Frontiers in Bioengineering and Biotechnology* 12 (2024) 1401803.
- [3] L. Yang, J. Huang, T. Feng, W. Hong-An, D. Guo-Zhong, Gesture interaction in virtual reality, *Virtual Reality & Intelligent Hardware* 1 (2019) 84–112.
- [4] M. Strobel, S. Schoenfeldt, J. Daugalas, Gesture recognition for fmcw radar on the edge, in: 2024 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNeT), 2024, pp. 45–48. doi:10.1109/WiSNeT59910.2024.10438579.
- [5] B. Yan, P. Wang, L. Du, X. Chen, Z. Fang, Y. Wu, mmgesture: Semi-supervised gesture recognition system using mmwave radar, *Expert Systems with Applications* 213 (2023) 119042.
- [6] Y. Ji, Y. Sun, Y. Zhang, Z. Wang, Y. Zhuang, Z. Gong, D. Shen, C. Qin, H. Zhu, H. Xiong, A comprehensive survey on self-interpretable neural networks, *arXiv preprint arXiv:2501.15638* (2025).
- [7] L. Dierckx, R. Veroneze, S. Nijssen, Rl-net: Interpretable rule learning with neural networks, in: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, Springer, 2023, pp. 95–107.
- [8] M. Frasca, D. La Torre, G. Pravettoni, I. Cutica, Explainable and interpretable artificial intelligence in medicine: a systematic bibliometric review, *Discover Artificial Intelligence* 4 (2024) 15.
- [9] F. Gardin, R. Gautier, N. Goix, B. Ndiaye, J.-M. Schertzer, Skoperules, <https://github.com/scikit-learn-contrib/skope-rules>, 2017.
- [10] S. Seifi, T. Sukianto, C. Carbonelli, L. Servadei, R. Wille, Interpretable rule-based system for radar-based gesture sensing: Enhancing transparency and personalization in ai, in: 2024 21st European Radar Conference (EuRAD), 2024, pp. 156–159. doi:10.23919/EuRAD61604.2024.10734943.
- [11] S. Franceschini, M. Ambrosanio, S. Vitale, F. Baselice, A. Gifuni, G. Grassini, V. Pascasio, Hand gesture recognition via radar sensors and convolutional neural networks, in: 2020 IEEE Radar Conference (RadarConf20), IEEE, 2020, pp. 1–5.
- [12] B. Jin, X. Ma, Z. Zhang, Z. Lian, B. Wang, Interference-robust millimeter-wave radar-based dynamic hand gesture recognition using 2d cnn-transformer networks, *IEEE Internet of Things Journal* (2023).
- [13] S. M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, *Advances in neural information processing systems* 30 (2017).
- [14] H. Chen, J. D. Janizek, S. Lundberg, S.-I. Lee, True to the model or true to the data?, *arXiv preprint arXiv:2006.16234* (2020).
- [15] S. Seifi, T. Sukianto, M. Strobel, C. Carbonelli, L. Servadei, R. Wille, Xentricai: A gesture sensing calibration approach through explainable and user-centric ai, in: *World Conference on Explainable Artificial Intelligence*, Springer, 2024, pp. 232–246.
- [16] Y. Li, Z. O'Neill, L. Zhang, J. Chen, P. Im, J. DeGraw, Grey-box modeling and application for building energy simulations-a critical review, *Renewable and Sustainable Energy Reviews* 146 (2021) 111174.
- [17] L. Qiao, W. Wang, B. Lin, Learning accurate and interpretable decision rule sets from neural networks, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2021, pp. 4303–4311.
- [18] Y. Yang, W. Ren, S. Li, Hyperlogic: Enhancing diversity and accuracy in rule learning with hypernets, *Advances in Neural Information Processing Systems* 37 (2024) 3564–3587.
- [19] C. Louizos, M. Welling, D. P. Kingma, Learning sparse neural networks through  $l_0$  regularization, *arXiv preprint arXiv:1712.01312* (2017).
- [20] S. Seifi, T. Sukianto, C. Carbonelli, 60 ghz fmcw radar gesture dataset, 2024. URL: <https://dx.doi.org/10.21227/s12w-cc46>. doi:10.21227/s12w-cc46.