Management and Production Engineering Review
Volume 17 • Number 1 • March 2026 • pp. 1–9
DOI: 10.24425/mper.2026.158634

Machine Learning-Driven Automated Selection of Safety Logic Devices in Industrial Control Systems

Karel STIBOR[1], Sona SEDIVA[1], Radek ŠTOHL[1], Lenka ŠTOHLOVÁ PUTNOVÁ[2]

[1] *Brno University of Technology, FEEC, Department of Control and Instrumentation*
[2] *Mendel University in Brno*

Received: 09 November 2025
Accepted: 12 February 2026

Abstract
The design of machine control systems requires the correct selection of safety logic devices to ensure functional safety and compliance with international standards such as EN ISO 13849-1 and IEC 62061. This process is typically based on expert knowledge and manual evaluation of design parameters, which can be time-consuming and error-prone. In this study, machine learning techniques are applied to automate and improve the selection of safety logic devices using real industrial data originating from various types of machinery and automated manufacturing real-world projects. This work introduces a significantly extended industrial dataset comprising 670 labelled machine configurations derived from real anonymized engineering projects and performs a comprehensive comparison of ten representative ML algorithms implemented in WEKA. The main novelty of the study is a unified large-scale comparative evaluation of heterogeneous machine learning classifiers on real industrial decision data, enabling joint assessment of scalability, generalization, interpretability, and computational efficiency under identical experimental conditions. The results demonstrate that increasing dataset size considerably enhances model stability and generalization. The Averaged 2-Dependence Estimator (A2DE) achieved the highest performance with an accuracy of 86% and Kappa = 0.81, followed by REPTree and Random Forest classifiers. Rule-based methods such as PART and NNge maintained strong interpretability with competitive predictive power. The findings confirm that probabilistic and ensemble algorithms provide reliable and practically applicable solutions for data-driven decision support in industrial safety engineering, paving the way for deployable, explainable, and adaptive decision-support tools in smart manufacturing environments.

Keywords
Machine learning, Safety logic device, Manufacturing systems, Safety engineering.

Introduction

The increasing complexity of modern manufacturing systems and the growing demand for functional safety have led to a rapid rise in the use of programmable safety devices, configurable relays, and safety controllers. Selecting the appropriate safety logic device is a crucial step in the design of machine control systems, as it directly affects both system reliability and compliance with international safety standards such as EN ISO 13849-1 (EN ISO 13849-1:2023, 2023) and IEC 62061 (IEC 62061:2021, 2021).

Corresponding author*: Radek Štohl – Brno University of Technology, e-mail: stohl@feec.vutbr.cz*

This selection process, however, remains knowledge-intensive and dependent on the experience of safety engineers. Traditional approaches often rely on manual analysis of design parameters and rule-based decision tables, which are time-consuming and prone to human error (Iyenghar, 2025; Farooq et al., 2023; TÜV SÜD Product Service, 2015).

In recent years, the integration of artificial intelligence (AI) and machine learning (ML) techniques into manufacturing and safety engineering has shown significant potential to automate and enhance this decision-making process (Mazzei & Ramjattan, 2022; Salazar et al., 2024; Gerschberger et al., 2024). Machine learning algorithms are capable of identifying complex patterns between design parameters and the corresponding safety logic devices, thereby providing reliable recommendations and supporting knowledge-based engineering within the framework of smart manufacturing and Industry 4.0 initiatives (Mazzei & Ramjattan, 2022; Salazar et al., 2024; Farooq et al., 2023).

Early research in this domain was initiated in (Štohl & Stibor, 2019), where a multilayer perceptron neural network was proposed to predict suitable safety device types based on machine design attributes. While the study demonstrated the feasibility of data-driven device selection, its evaluation was constrained to a single model family and a small dataset, limiting conclusions regarding scalability and robustness in real industrial settings.

In contrast to earlier studies (Štohl & Stibor, 2019; Štohl & Stibor, 2020), the present work is based on a significantly expanded dataset of 670 real industrial machine configurations derived from various types of heavy machinery and automated manufacturing real-world projects, all anonymized and collected from completed engineering realizations across several Central European countries.

Rather than evaluating a single learning paradigm, this work performs a systematic and unified comparison of ten heterogeneous machine learning algorithms, including neural, probabilistic, instance-based, rule-based, and ensemble methods, under identical experimental conditions.

Unlike previous studies limited to single-model evaluation and small samples, this study provides the first unified cross-family comparison of heterogeneous ML classifiers on a large-scale real industrial dataset under identical validation settings. This explicit comparative framework enables, for the first time in this application domain, a joint assessment of predictive accuracy, generalization capability, interpretability, and computational efficiency on large-scale real-world industrial data.

The goal of this study is to identify which ML methods deliver the best balance between accuracy, interpretability, and computational cost for industrial safety engineering applications. By integrating machine learning into the design workflow, safety engineers can obtain automated, data-driven recommendations, supporting decision-making transparency, standard compliance, and the broader implementation of smart safety systems in manufacturing environments.

Methodology

The objective of this study is to evaluate and compare the performance of various machine learning (ML) algorithms in predicting the appropriate safety logic device for industrial machinery based on safety-related design parameters. The analysis was conducted using the WEKA software environment, which provides a wide range of supervised learning algorithms suitable for industrial and manufacturing applications. The methodological framework reflects real industrial engineering practice, where safety logic device selection is performed during the design and validation phases of machine control systems in heavy machinery and automated manufacturing environments.

A new and significantly extended dataset, derived from real industrial cases, was used to ensure better generalization compared to previous research. The dataset consists of engineering decision records collected from completed and validated industrial projects, rather than simulated or synthetic data. Each record represents a machine configuration described by attributes (Štohl & Stibor, 2020) such as the number of safety functions, access points, safety inputs/outputs, required Performance Level (PL), communication network type, and human–machine interface (HMI) presence. The target variable corresponds to the category of safety logic device (e.g., Relay, CR30, GMX, GLX).

The workflow included data preprocessing, classifier training, cross-validation, and performance comparison using multiple statistical metrics. A 10-fold cross-validation procedure was applied to ensure a robust estimation of model performance.

Dataset Preprocessing

The dataset used in this research was derived from the same industrial design knowledge base (Štohl & Stibor, 2020), which captures engineering data collected during the selection and configuration of safety logic devices in machine control systems. The knowledge base aggregates data from heavy machinery, automated production lines, and modular manufacturing systems, covering both standalone machines and integrated production units. The recorded configurations represent real-world safety engineering decisions made under functional safety standards such as EN ISO 13849-1 and IEC 62061.

Each record represents one specific machine configuration described by a set of input design parameters (attributes) and a target output class corresponding to the recommended safety logic device (Table 2).

All machine configurations included in the dataset correspond to implemented or validated design solutions, ensuring that the learning task reflects realistic industrial constraints and decision boundaries.

Classifier Selection

Ten representative classifiers were selected to cover the main paradigms of supervised learning: neural networks, instance-based learning, probabilistic methods, rule-based models, and tree-based ensembles (Table 1).

Table 1
Selected classifiers

Group	Algorithm	Description
Neural Networks	Multilayer Perceptron	Classical feed-forward MLP with backpropagation (Li et al., 2012).
	Dl4jMlp	Deep Learning model based on the Deeplearning4j backend, supporting multi-layer architectures (Lang et al., 2019).
Instance-Based Learning	IB3	k-Nearest Neighbour ($k = 3$), similarity-based classification (Mohamed, 2017).
	KStar	Entropic distance–based instance learner with probabilistic smoothing (Cleary & Trigg, 1995).
Probabilistic Models	NaiveBayes	Baseline Bayesian classifier assuming feature independence (John & Langley, 1995).
	A2DE	Averaged 2-Dependence Estimator, capturing pairwise feature dependencies for improved accuracy (Webb et al., 2012).
Rule-Based Learning	NNge	Nearest-Neighbour Generalization algorithm producing interpretable IF–THEN rules (Roy, 2002).
	PART	Partial decision tree rule learner combining C4.5 trees with rule induction (Frank & Witten, 1998).
Tree-Based Methods	REPTree	Fast regression and classification tree optimized for reduced error pruning (Srinivasan & Mekala, 2014).
	RandomForest	Ensemble of decision trees using random subspace sampling; strong baseline for generalization (Mishra & Ratha, 2016).

Table 2
Attributes of the safety dataset used for classification (Štohl & Stibor, 2020).

Attribute	Description
Number of safety functions	Number of safety functions (e.g. emergency stop, guard door) integrated in the machine (integer)
Performance Level (PL)	Required safety Performance Level (EN ISO 13849-1:2023) for the machine (categorical ‘a’ to ‘e’)
Number of access points	Number of access points to dangerous zones (e.g. maintenance openings) (integer)
Number of safety inputs	Number of safety-related sensors/inputs needed (integer)
Number of safety outputs	Number of safety-related actuators/outputs needed (integer)
Machine type	Configuration of machine: standalone unit, production line, or assembly (categorical)
Communication	Use of safety network for control (Boolean)
HMI (Human–Machine Interface)	Presence of an operator interface for safety/routine control (Boolean)
Safety solution class	Safety logic device class (target): one of {Relay, CR30, GMX, GLX}

This selection was intentionally designed to balance high-performance models with interpretable algorithms suitable for safety-critical industrial applications.

Evaluation Procedure

Each classifier was trained and evaluated using the same experimental conditions to ensure fair comparison. The 10-fold cross-validation scheme was applied, where the dataset was randomly partitioned into ten equal subsets; in each iteration, nine subsets were used for training and one for testing.

The following metrics were computed for each model:

- Classification Accuracy (%)
- Cohen's Kappa (κ) to measure class agreement beyond chance
- F1-score for balanced performance evaluation
- Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) as complementary error indicators
- and training time (s) to assess computational efficiency.

All experiments were performed using the WEKA 3.9.6 environment on identical hardware to avoid performance bias. Fig. 1 shows an example of the WEKA Knowledge Flow Environment.

Results

Table 3 and Table 4 summarizes the classification accuracy and Kappa statistic for all ten algorithms tested on the original dataset (90 instances) and on the extended dataset (670 instances). A graphical representation of these results is shown in Fig. 2. The larger dataset clearly improved both predictive accuracy and class agreement across all classifier families.

The results (columns 90 inst. in the Table 3, Table 4) show that even with a limited training set, several algorithms achieved competitive accuracy.

Tree-based classifiers, particularly REPTree and Random Forest, reached the highest accuracy (84.4%) and Kappa ≈ 0.75, indicating strong stability despite data scarcity.

Rule-based models such as PART and NNge also performed well ($\approx$ 81–82%), confirming their interpretability advantage without major loss in precision.

The Naive Bayes and A2DE classifiers achieved around 80% accuracy, demonstrating that probabilistic reasoning can handle small datasets effectively.

In contrast, deep and shallow neural networks (MultilayerPerceptron, Dl4jMlp) achieved only 67–70% accuracy, suggesting that they require more samples to generalize reliably.

After increasing the dataset size (columns 670 inst. in Table 3 and Table 4) more than sevenfold, almost all algorithms improved in both accuracy and Kappa.

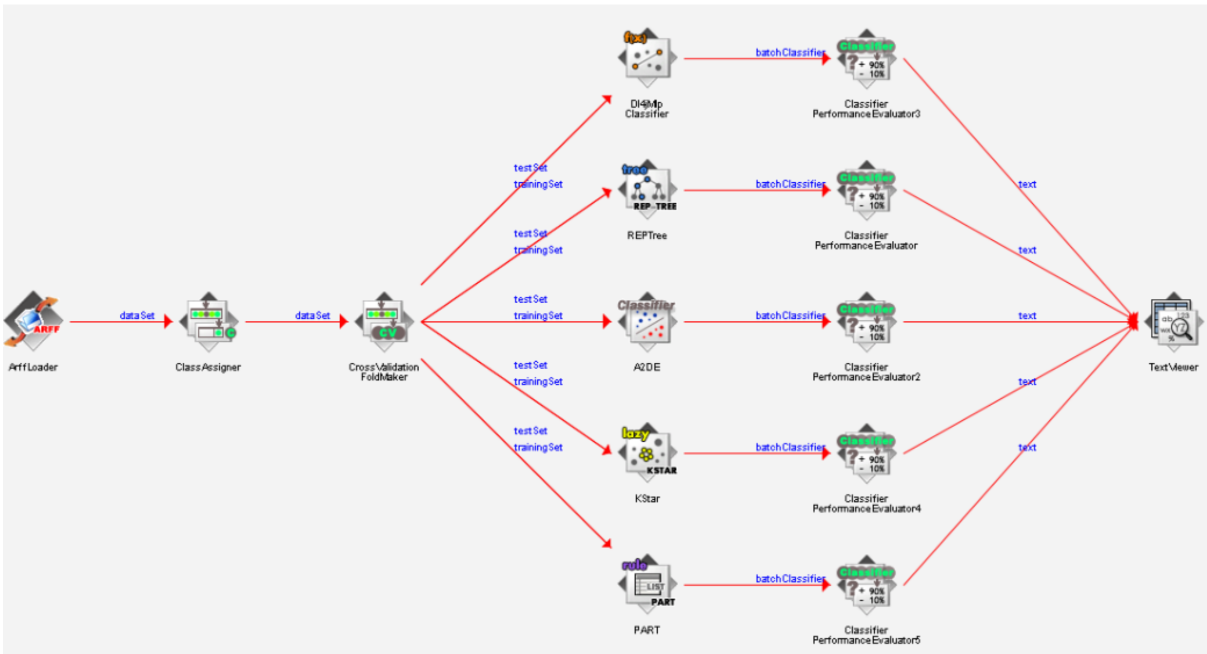

Fig. 1. One of the WEKA Knowledge Flow Environment including five classifiers.

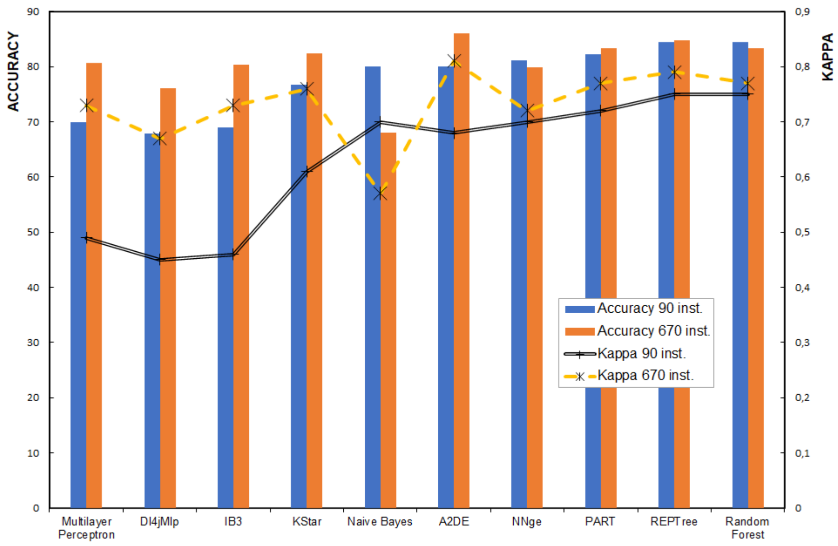

Fig. 2. The accuracy and Cohen's Kappa comparison on small (90 instances) versus extended dataset (670 instances).

Table 3
Classification accuracy, Kappa and MAE

Algorithm	Accuracy (%)		Kappa		MAE	
	90 inst.	670 inst.	90 inst.	670 inst.	90 inst.	670 inst.
Multilayer Perceptron	70.0	80.6	0.49	0.73	0.160	0.101
Dl4jMlp	67.8	76.1	0.45	0.67	0.300	0.183
IB3	68.9	80.4	0.46	0.73	0.180	0.122
KStar	76.7	82.4	0.61	0.76	0.117	0.108
Naive Bayes	80.0	68.1	0.70	0.57	0.109	0.161
A2DE	80.0	86.0	0.68	0.81	0.110	0.098
NNge	81.1	79.9	0.70	0.72	0.094	0.103
PART	82.2	83.3	0.72	0.77	0.100	0.099
REPTree	84.4	84.8	0.75	0.79	0.110	0.093
Random Forest	84.4	83.3	0.75	0.77	0.101	0.095

The most notable gains were observed for neural and instance-based models:

- Multilayer Perceptron: +10.6 pp accuracy (from 70% to 80.6%)
- IB3: +11.5 pp accuracy (from 68.9% to 80.4%)
- Dl4jMlp: +8.3 pp accuracy (from 67.8% to 76.1%)

Both REPTree and Random Forest maintained stable high accuracy ($\approx$ 84%) with the lowest error metrics (MAE $\approx$ 0.10), confirming their robustness on mixed attribute data.

Table 4
F1-score, RME and Train time of the classifiers

Algorithm	F1-score		RMSE		Train time (s)	
	90 inst.	670 inst.	90 inst.	670 inst.	90 inst.	670 inst.
Multilayer Perceptron	0.69	0.79	0.343	0.281	$\approx$ 22	$\approx$ 75
Dl4jMlp	0.65	0.73	0.385	0.323	$\approx$ 65	$\approx$ 110
IB3	0.66	0.78	0.358	0.308	$\approx$ 2	$\approx$ 8
KStar	0.75	0.80	0.289	0.272	$\approx$ 4	$\approx$ 11
Naive Bayes	0.80	0.68	0.299	0.341	$\approx$ 1	$\approx$ 3
A2DE	0.80	0.84	0.263	0.238	$\approx$ 9	$\approx$ 18
NNge	0.81	0.80	0.307	0.285	$\approx$ 6	$\approx$ 10
PART	0.82	0.82	0.274	0.260	$\approx$ 7	$\approx$ 12
REPTree	0.85	0.84	0.254	0.244	$\approx$ 5	$\approx$ 9
Random Forest	0.84	0.83	0.220	0.241	$\approx$ 10	$\approx$ 16

Note: F1-scores derived from weighted class averages; training time measured as total computation time for 10-fold cross-validation on identical hardware.

The A2DE classifier exhibited moderate but uneven performance on the original, limited dataset of 90 instances. The confusion matrix (Fig. 3) shows that the model achieved relatively high accuracy for the Relays class (47 correctly classified out of 50, $\approx$ 94%), while struggling to distinguish between GLX and GMX, which were frequently confused with each other. Only

8 GLX and 6 GMX instances were correctly identified, with 13 mutual misclassifications (6 GLX → GMX and 7 GMX → GLX). The CR30 category achieved partial recognition (11 correct of 13, ≈ 85%), though occasional misclassifications with Relays (2 cases) indicate limited separation in feature space.

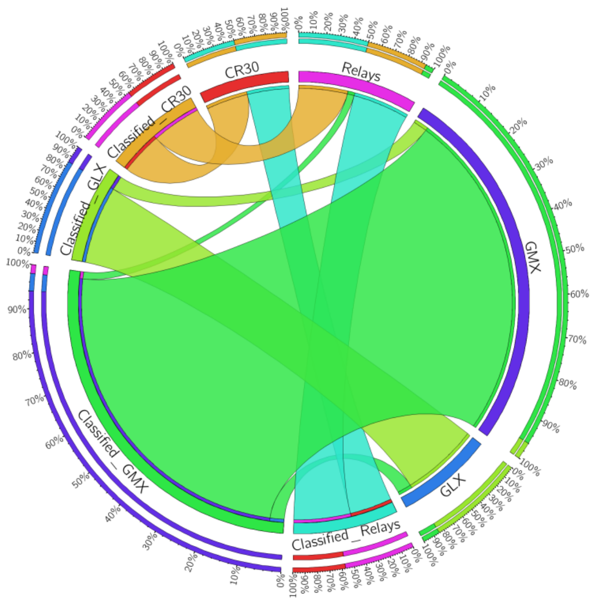

Fig. 3. Confusion matrix for A2DE classifier on the original dataset.

Overall, these results reflect the constraints of the smaller dataset, which contained insufficient variability across machine configurations and attribute combinations. The model's probabilistic dependencies could not be fully exploited under these conditions, leading to high class-level variance and unstable generalization.

The A2DE classifier achieved the best overall performance on the extended dataset (accuracy = 86.0%, Kappa = 0.81), outperforming tree-based and neural models. The confusion matrix (Fig. 4) summarizes the classification outcomes across the four target categories. The model correctly classified the majority of Relay and GMX instances, with only minor confusion between CR30 and GLX devices. Specifically, out of 670 samples, 156 of 169 Relay and 195 of 227 GMX instances were predicted correctly, yielding precision and recall values above 0.90 for these categories. The CR30 class showed moderate overlap with GMX and Relay (16 and 12 misclassifications, respectively), consistent with their partially shared feature profiles. Misclassification between GLX and GMX was limited (21 and 19 instances), indicating that A2DE successfully captured subtle distinctions in safety performance levels and I/O configurations.

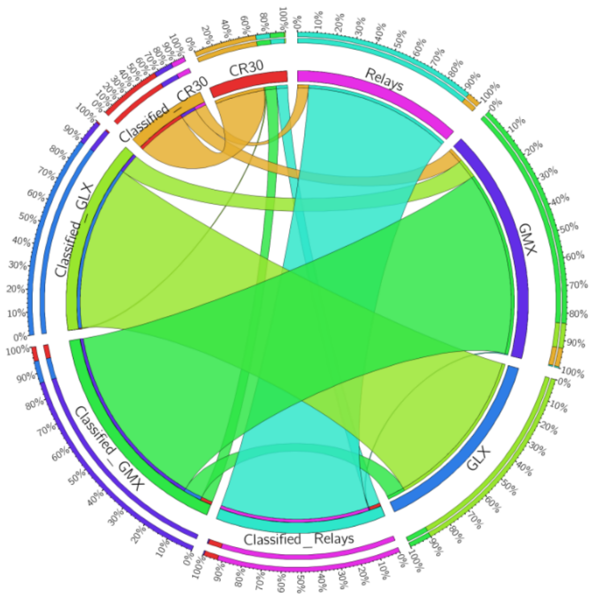

Fig. 4. Confusion matrix for A2DE classifier on the extended dataset.

When compared to the performance on the expanded dataset of 670 instances, clear improvements are evident. In the larger dataset, A2DE achieved mean accuracy above 85% and exhibited a more balanced class distribution, with precision and recall values exceeding 0.90 for Relays and GMX and substantially fewer cross-class errors. The previously dominant GLX–GMX confusion was significantly reduced, suggesting that the additional data enabled the classifier to model attribute interdependencies more effectively.

All classifiers predicted the Relay category with very high recall (≈ 0.93–0.95).

The GLX and GMX classes exhibited moderate confusion due to overlapping attribute patterns, whereas CR30 remained the most challenging, reflecting its hybrid feature structure between programmable and relay-based devices.

Rule-based and tree models (PART, REPTree, Random Forest) achieved the best class balance, while A2DE provided the most consistent probabilistic differentiation across all categories.

Discussion

The experimental results clearly demonstrate that dataset size and diversity have a substantial effect on model accuracy, stability, and generalization (Goodfellow et al., 2016; Halevy et al., 2009; Domingos, 2012).

When trained on the original dataset of only 90 instances, most classifiers exhibited moderate accuracy

($\approx$ 70–80%) and relatively low Kappa values (< 0.6) due to limited variability of machine configurations. After extending the dataset to 670 records, the overall mean accuracy increased by approximately 7 percentage points and the Kappa statistic by 0.16, confirming that the new data provided a more representative distribution of real-world industrial cases (Sun et al., 2009).

Neural networks and instance-based learners benefited the most from this increase. Models such as the Multilayer Perceptron and IB3 improved their predictive accuracy by more than 10 percentage points, highlighting their sensitivity to data richness.

Probabilistic and ensemble methods, such as A2DE, REPTree, and Random Forest, showed the strongest stability across folds, indicating that they effectively captured underlying patterns even in heterogeneous industrial datasets.

Performance balance: accuracy, interpretability, and efficiency

While ensemble and probabilistic models achieved the best accuracy and generalization (A2DE – 86%, Random Forest – 83%), interpretability remains a crucial factor in safety engineering applications.

Rule-based learners (PART, NNge) achieved F1-scores around 0.82 with slightly lower accuracy, yet they offer transparent if–then logic structures that can be directly validated by safety engineers.

This interpretability is particularly valuable when compliance with functional safety standards (e.g., EN ISO 13849-1) requires that decision criteria be traceable and explainable.

From a computational perspective, tree-based and rule-based models provided a good trade-off between accuracy and training efficiency. They completed cross-validation in less than 15 seconds, whereas deep neural networks (e.g., Dl4jMlp) required more than 100 seconds per run due to iterative weight optimization.

Therefore, in real-time or embedded design-support tools, probabilistic or tree-based models may offer the most practical balance between performance and computational cost.

A detailed class-level analysis revealed consistent patterns across all models.

The Relay category was predicted with the highest precision and recall ($\approx$0.94), likely due to its distinct feature profile (fewer inputs, simpler logic, no communication interfaces).

In contrast, the CR30 class remained challenging to classify, as its characteristics overlap partially with both programmable and relay-based devices.

Probabilistic (A2DE) and rule-based (PART) methods improved recognition of this class by modelling attribute dependencies and generating discriminative decision boundaries.

The GLX and GMX classes occasionally suffered from confusion, primarily due to similar PL and I/O characteristics; however, ensemble methods reduced these misclassifications by leveraging multiple decision subspaces.

Practical implications for manufacturing systems

The results underline the feasibility of using machine learning to support the design phase of safety-related control systems in manufacturing.

By embedding trained classifiers into engineering tools, designers can automatically obtain recommendations for appropriate safety logic devices based on design parameters such as required Performance Level, number of safety functions, or network topology.

This approach contributes to knowledge-based engineering and supports smart manufacturing by integrating data-driven decision support directly into CAD/CAM and digital twin environments (Verhagen et al., 2012; Tao et al., 2018). From an industrial perspective, the proposed approach can be directly integrated into existing safety engineering workflows as a decision-support layer, supporting engineers during early design and validation phases without replacing expert judgement.

Furthermore, models like A2DE and Random Forest could be integrated into digital twins of production systems, continuously adapting recommendations as machine configurations evolve.

Such integration can reduce engineering effort, minimize human error during safety logic device selection, and accelerate compliance verification according to functionalsafety standards (Hollnagel, 1993).

Limitations and future work

Despite the promising results, several limitations should be noted.

First, the dataset, though extended, still reflects specific industrial environments and may not fully capture the diversity of all machine types. Further expansion across different industries and safety architectures is needed.

Second, hyperparameter tuning was intentionally kept moderate to maintain comparability; deeper optimization (e.g., via AutoML or Bayesian search) could further enhance results.

Finally, future work will focus on hybrid and ensemble learning, combining probabilistic reasoning with explainable rule extraction to achieve both high accuracy and human interpretability.

Conclusion

This study presented a comprehensive comparison of ten machine learning algorithms applied to the task of safety logic device selection within industrial machine design.

Using two datasets of differing scale – a smaller baseline set (90 instances) and an extended dataset (670 instances) – the research demonstrated how dataset expansion and diversity substantially enhance model reliability, generalization, and predictive power. Compared with earlier studies limited to small datasets and individual classifier families (Štohl & Stibor, 2019; Štohl & Stibor, 2020), this work systematically evaluated heterogeneous machine learning approaches on a large-scale real industrial dataset derived from various types of machinery and automated manufacturing real-world projects.

The A2DE (Averaged 2-Dependence Estimator) classifier achieved the best overall performance, with 86% accuracy, Kappa $= 0.81$, and $F1 = 0.84$, outperforming both tree-based and neural models.

The Random Forest and REPTree algorithms also delivered consistently strong results, confirming the robustness of ensemble learning in handling mixed numeric and categorical industrial data. Rule-based methods such as PART and NNge offered an advantageous balance between interpretability and predictive capability, which is critical in safety-related engineering contexts where decision transparency and traceability are required.

The primary contribution of this study is the validated demonstration that probabilistic and ensemble ML methods can reliably support safety logic device selection when trained on sufficiently large and real industrial datasets. The unified cross-algorithm evaluation framework further provides practical guidance for selecting suitable ML techniques in safety-critical engineering environments.

By leveraging real industrial data and consistent comparative evaluation, the study moves beyond proof-of-concept experimentation toward practically deployable ML-supported engineering decision tools. Future work will focus on expanding the industrial dataset to include a broader range of machine architectures and safety devices, as well as applying automated hyperparameter optimization and hybrid ensemble models to further improve generalization and interpretability. Additional research will also address integration of the developed models into knowledge-based CAD/CAE and digital engineering environments, enabling real-time, data-driven assistance for safety engineers.

Acknowledgements

The completion of this article was made possible by grant FEKT-S-26-8988 – "Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence" financially supported by the Internal science fund of Brno University of Technology.

References

Cleary, J. G. nad Trigg, L. E. (1995). An instance-based learner using an entropic distance measure. In *Proceedings of the 12th International Conference on Machine Learning*, pp. 108–114. Tahoye City, USA. DOI: 10.1016/B978-1-55860-377-6.50022-0

Domingos, P. (2012). A few useful things to know about machine learning. *Commun. ACM*, 55(10), 78–87. https://homes.cs.washington.edu/~pedrod/papers/cacm12.pdf

EN ISO 13849-1:2023 (2023). Safety of machinery – Safety-related parts of control systems – Part 1: General principles for design. International Organization for Standardization. https://www.iso.org/standard/73481.html

IEC 62061:2021 (2021). Safety of machinery – Functional safety of safety-related control systems. International Electrotechnical Commission. https://webstore.iec.ch/en/publication/93654

Farooq, S., Akhtar, M., Rahman, H. (2023). Artificial Intelligence in Manufacturing Industry Worker Safety: A New Paradigm for Hazard Prevention and Mitigation. *Processes*, 13(5), 1–57. DOI: 10.3390/pr13051312

Frank, E., Witten, I.H. (1998). Generating accurate rule sets without global optimization. *Proceedings of the Fifteenth International Conference on Machine Learning*, 144–151. Morgan Kaufmann Publishers Inc., San Francisco, USA. DOI: 10.5555/645527.657305

Gerschberger, M., Sihn, W., Seidel, A. (2024). Hybrid Intelligence – Systematic Approach and Framework to Determine the Level of Human–AI Collaboration for Production Management Use Cases. *Production Engineering*, 18, 525–541. DOI: 10.1007/s11740-024-01326-7

Goodfellow, I., Bengio, Y., Courville, A. (2016). *Deep Learning*, MIT Press, Cambridge USA.

Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten I.H. (2009). The WEKA data mining software: An update. *ACM SIGKDD Explorations Newsletter*, 11(1), 10–18. DOI: 10.1145/1656274.1656278

Halevy, A., Norvig, P., Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE Intelligent Systems*, 24(2), 8–12. DOI: 10.1109/MIS.2009.36

Hollnagel, E. (1993). *Human Reliability Analysis: Context and Control*. Academic Press, Cambridge USA. (1993).

ISO/TR 23849:2010 (2010). Guidance on the application of ISO 13849-1 and IEC 62061 in the design of safety-related control systems for machinery. International Organization for Standardization. https://www.iso.org/standard/54424.html

Iyenghar. P. (2025). Empirical Evaluation of Reasoning LLMs in Machinery Functional Safety Risk Assessment and the Limits of Anthropomorphized Reasoning. *Electronics*, 14(18), 1–59. DOI: 10.3390/electronics14183624

John, G.H., Langley, P. (1995). Estimating Continuous Distributions in Bayesian Classifiers. Conference on Uncertainty in Artificial Intelligence. *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, 338-345. Morgan Kaufmann Publishers Inc., San Francisco, USA.

Lang, S., Bravo-Marquez, F., Beckham, C., Hall, M., Frank, E. (2019). WekaDeeplearning4j: a Deep Learning Package for Weka based on DeepLearning4j. *Knowledge-Based Systems*, 178, 48–50. DOI: 10.1016/j.knosys.2019.04.013

Li, J., Cheng, J.H., Shi, J.Y., Huang, F. (2012). Brief introduction of back propagation (BP) neural network algorithm and its improvement. In: Jin, D., Lin, S. (eds) *Advances in Computer Science and Information Engineering. Advances in Intelligent and Soft Computing*, 169, 553–558. Springer, Berlin, Heidelberg. DOI: 10.1007/978-3-642-30223-7_87

Mazzei, D., Ramjattan, R. (2022). Machine Learning for Industry 4.0: A Systematic Review Using Deep Learning-Based Topic Modelling. *Sensors*, 22(22), 1–26. DOI: 10.3390/s22228641

Mishra, A.K., Ratha, B.K. (2016). Study of Random Tree and Random Forest Data Mining Algorithms for Microarray Data Analysis. *International Journal on Advanced Electrical and Computer Engineering*, 3(4), 5–7. http://irdindia.in/journal_ijaece/pdf/vol3_iss4/2.pdf

Mohamed, A.E. (2017). Comparative study of four supervised machine learning techniques for classi?cation. International Journal of Applied Science and Technology 7(2), 1–15 (2017). https://pubs.sciepub.com/aees/8/6/18/index.html

Osborne, M., Hawkins, R., Nicholson, M., Alexander, R. (2024). Understanding safety engineering practice: Comparing safety engineering practice as desired, as required, and as observed. *Safety Science*, (172), 1–19. DOI: 10.1016/j.ssci.2024.106424

Roy, S. (2002). Nearest Neighbor With Generalization. *Unpublished manuscript, Christchurch*, New Zealand.

Salazar, L., Rojas, C., Pineda, F. (2024). Artificial Intelligence and Smart Technologies in Safety Management: A Comprehensive Analysis Across Multiple Industries. *Applied Sciences*, 14(24), 1–22. DOI: 10.3390/app142411934

Srinivasan, B., Mekala, P. (2014). Mining Social Networking Data for Classification Using REPTree. *International Journal of Advance Research in Computer Science and Management Studies*, 2(10), 155–160. http://ijarcsms.com/docs/paper/volume2/issue10/V2I10-0056.pdf

Štohl, R.; Stibor, K. (2019). Predicting Safety Solutions via an Artificial Neural Network. In *16th IFAC Conference on Programmable Devices and Embedded Systems PDeS*, 490–495. IFAC-PapersOnLine, Brno, Czech Republic. DOI: 10.1016/j.ifacol.2019.12.711

Štohl, R.; Stibor, K. (2020). Predicting Safety Logic Device Solutions via Decision Trees and Rules Algorithms. In *Proceedings of the 2020 21st International Carpathian Control Conference (ICCC)*, pp. 1–7. IEEE, High Tatras, Slovakia. DOI: 10.1109/ICCC49264.2020.9257284

Sun, Y., Wong, A.K.C., Kamel, M.S. (2009). Classification of imbalanced data: A review. *International Journal of Pattern Recognition and Artificial Intelligence*, 23(4), 687–719. DOI: 10.1142/S0218001409007326

Tao, F., Qi, Q., Liu, A., Kusiak, A. (2018). Data-driven smart manufacturing. *Journal of Manufacturing Systems*, 48 (Part C), 157–169. DOI: 10.1016/j.jmsy.2018.01.006

TÜV SÜD Product Service (2015). A practical guide to Machinery Safety. 5th edn. United Kingdom. https://www.tuvsud.com/en-us/-/media/regions/uk/pdf-files/whitepaper-report-e-books/uk-mkg-practical-guide-en-gb-a4.pdf

Verhagen, W.J.C., Bermell-Garcia, P., van Dijk, R.E.C., Curran, R. (2012). A critical review of knowledge-based engineering: An identification of research challenges. *Advanced Engineering Informatics*, 26(1), 5–15. DOI: 10.1016/j.aei.2011.06.004

Webb, G.I., Boughton, J.R., Zheng, F., Ting, K.M., Salem, H. (2012). Learning by extrapolation from marginal to full-multivariate probability distributions: decreasingly naive Bayesian classification. *Mach Learn*, 86, 233–272. DOI: 10.1007/s10994-011-5263-6

Witten, I., Frank, E., Hall, M., Pal, C. (2023). Data Mining: Practical Machine Learning Tools and Techniques. *5th ed., Morgan Kaufmann*, USA.