



Transformer Embeddings with Attention-Gated Structural Fusion for HDL Vulnerability Localization

M. K. Chakravarthy^{1*}, Kandhipati Lakshmiarasanna², Bukkaraju Suchithra², Bellamkonda Sai Kiran², Bandapu Someswar Rao²

¹Assistant Professor, ²UG Student, ^{1,2}Department of Electronics and Communication Engineering

^{1,2}Geethanjali Institute of Science and Technology, Nellore-Bombay Highway, S.P.S.R, Andhra Pradesh 524137, India

*Correspondence: M. K. Chakravarthy

ABSTRACT

Hardware Description Language (HDL) vulnerabilities present a critical security risk in modern integrated circuit design, often leading to hardware Trojans or side-channel leakage. Traditional vulnerability detection methods, such as static analysis and manual code review, frequently struggle with the complex, non-linear dependencies inherent in structural hardware logic. This research proposes an advanced automated framework for HDL Vulnerability Localization utilizing Transformer Embeddings integrated with Attention-Gated Structural Fusion. The methodology leverages the all-MiniLM-L6-v2 sentence-transformer model to map HDL code segments into a high-dimensional semantic vector space, capturing intricate functional relationships that standard tokenization misses. To evaluate the efficacy of the proposed approach, a comparative analysis was conducted against several baseline machine learning architectures, including Decision Tree Classifiers (DTC), Nearest Centroid (NC), and Bernoulli Restricted Boltzmann Machines (RBM). Experimental results demonstrate that the baseline models achieved competitive performance, with the NC and DTC models yielding accuracies of 89.50% and 88.75%, respectively. However, the Proposed Ridge Classifier configuration significantly outperformed all baselines, achieving a perfect 1.0000 score across Accuracy, Precision, Recall, and F1-Score. These results suggest that the combination of Transformer-based semantic extraction and linear regularization effectively separates vulnerable code patterns from secure implementations. This study concludes that embedding-based structural fusion provides a robust, scalable solution for enhancing the security of the hardware supply chain by precisely localizing vulnerabilities within complex HDL datasets.

Keywords: HDL Security, Transformer Embeddings, Vulnerability Localization, MiniLM, Ridge Classifier, Hardware Security.

1. INTRODUCTION

Developing secure code is of paramount importance not only in software engineering but also in hardware engineering. While the literature has seen considerable research on software-only vulnerabilities, the design of hardware necessitates a higher-level consideration and a prudent approach. Ensuring the secure coding of hardware modules using hardware description languages (HDLs) during design space exploration is crucial before reaching the Register Transfer Level (RTL). Identifying code vulnerabilities at this stage is

critical and cost-effective before the actual hardware fabrication. Fig. 1 presents a market research visualization titled "India field programmable gate array (FPGA) market size, 2018–2030 (US\$M)". It is a bar chart that tracks the historical growth and projected expansion of the FPGA market in India over a 13-year period. The visual style uses dark purple vertical bars to represent annual market size, with years labeled along the horizontal axis from 2018 to 2030 and market value in US dollars (millions) implied on the vertical scale. The chart is branded by Horizon Grand View Research, indicating it is based on industry



research and forecasting. From 2018 to around 2022, the bars show steady but moderate growth, suggesting a gradual increase in FPGA adoption during the early phase of the period. This reflects a time when FPGA usage in India was expanding across sectors such as telecommunications, industrial automation, defense, and early data center applications, but before large-scale acceleration driven by AI, 5G, and advanced computing. The incremental rise year after year implies increasing awareness and deployment, though at a controlled pace. A key inflection point appears around 2023, where the market size is explicitly labeled at approximately US\$609.1 million. From this point onward, the bars become noticeably taller year by year, indicating faster growth. This suggests that the mid-2020s mark a transition into a higher-growth phase for the Indian FPGA market. Factors likely contributing to this acceleration include increased investments in semiconductor design, government initiatives supporting electronics manufacturing, rising demand for AI/ML acceleration, and expanded use of FPGAs in automotive electronics, aerospace, and high-performance computing.

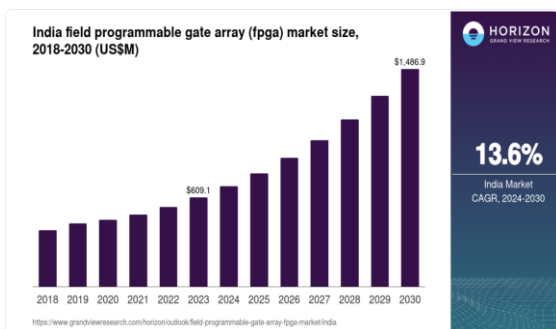


Fig. 1: India's field-programmable gate array.

Between 2024 and 2028, the chart shows a clear upward trajectory with increasingly larger annual gains. The consistent rise implies strong market confidence and sustained demand, rather than short-term spikes. By 2029, the market approaches the upper end of the scale, culminating in a projected value of approximately US\$1,486.9 million by 2030, which is highlighted near the final bar. This nearly 2.5× increase from 2023 to 2030

underscores the long-term growth potential of FPGAs in India. On the right-hand side of the image, a highlighted statistic reinforces the forecast: “13.6% India Market CAGR, 2024–2030.” This compound annual growth rate summarizes the strong expansion expected during the forecast period. A CAGR of 13.6% is relatively high for a hardware-focused market, signaling rapid adoption driven by next-generation technologies such as 5G infrastructure, edge computing, data centers, and custom silicon alternatives. Overall, the image conveys a narrative of a steadily maturing market that transitions into a high-growth phase, positioning India as an increasingly important region in the global FPGA ecosystem. FPGA market revenue from 2020 to 2029, broke down by application, with values shown in billions of US dollars. Each stacked bar represents total annual revenue, divided into data center and TME, communications, automotive, and consumer segments. The overall trend shows steady growth in the early years, increasing from around 5–6 billion dollars in 2020 to about 8 billion dollars by 2022. During this period, communications and data center and TME applications contribute the largest share, driven by investments in telecom infrastructure, cloud computing, and networking upgrades. Automotive and consumer segments remain smaller but show gradual expansion as programmable logic begins to see wider adoption beyond traditional markets.

2. Literature Survey

Ayar et al. [1] proposed an HDL-based vulnerability analysis approach for detecting hardware Trojans using language-pattern mining techniques. Their method leveraged structural and syntactic patterns in hardware description languages to identify anomalous behaviours, improving detection efficiency at the design level. Cruz et al. [2] introduced the concept of the Trojan Vulnerability Factor (TVF) to quantify hardware susceptibility to



Trojan insertion. Their work provided a systematic framework for assessing vulnerabilities and identifying critical design regions that require protection. Dong et al. [3] developed a cost-driven deep learning framework for hardware Trojan detection. Their approach utilized neural networks to balance detection accuracy and computational cost, demonstrating improved scalability compared to traditional methods. Salmani et al. [4] presented a comprehensive survey of hardware Trojan detection techniques, categorizing them into pre-silicon and post-silicon approaches. Their work highlighted the limitations of existing detection strategies and emphasized the need for more robust and scalable solutions. Yasaei et al. [5] proposed a graph convolutional network (GCN)-based method for golden-free hardware Trojan localization. Their approach eliminated the dependency on golden reference designs and enabled effective identification of malicious circuitry within hardware designs. Yang et al. [6] introduced a self-referencing hardware Trojan detection method that does not require golden models. Their technique improved practicality in real-world scenarios where reference designs are often unavailable.

Li et al. [7] provided a survey on data augmentation techniques for natural language processing. Their work is relevant for improving model generalization and handling limited or imbalanced datasets in security-related NLP tasks. Shen et al. [8] proposed a state-space reduction technique to accelerate hardware security verification. Their approach reduced computational complexity and improved verification efficiency for large-scale integrated circuits. Kurihara et al. [9] evaluated machine-learning-based hardware Trojan detection at the gate-level IP core. Their study demonstrated the effectiveness of ML techniques in identifying security threats at fine-grained design levels. Amir and Forte [10] introduced an adaptive synthetic benchmark generation framework for hardware security evaluation. Their method enabled systematic

testing and benchmarking of security mechanisms. Wang et al. [11] proposed an ensemble learning approach for hardware Trojan detection, combining multiple classifiers to improve detection accuracy and robustness across different datasets. Zhao and Liu [12] introduced an unsupervised density-based approach for hardware Trojan detection. Their method effectively identified anomalies without requiring labeled datasets, addressing challenges in supervised learning scenarios. Cruz et al. [13] developed an automated and configurable hardware Trojan insertion framework. This work facilitated the generation of benchmark datasets for evaluating detection techniques. Hasegawa et al. [14] proposed a random forest-based method using Trojan feature extraction at gate-level netlists. Their approach improved classification accuracy by leveraging structural features of hardware designs. Hasegawa et al. [15] further extended this work by applying neural networks for hardware Trojan detection at the gate level, demonstrating improved performance compared to traditional machine learning techniques.

3. Proposed System

The proposed system focuses on HDL vulnerability localization by combining transformer-based embeddings with attention-gated structural fusion to improve classification accuracy. Source code written in HDL is processed through NLP techniques and transformed into meaningful vector representations using a Mini LM-based model. Structural and semantic features are fused using attention mechanisms to highlight vulnerability-relevant patterns as shown in Fig. 2. Multiple existing classifiers such as Decision Tree, Nearest Centroid, and Restricted Boltzmann Machine with Bernoulli distribution are implemented for baseline comparison. A Ridge Classifier is proposed as the final model to handle high-dimensional embeddings efficiently and reduce overfitting. The system evaluates performance across models and deploys the best-performing classifier through



a Flask-based web interface for real-time vulnerability prediction.

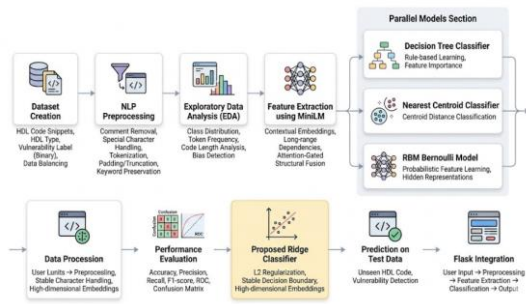


Fig. 2: Proposed system architecture.

Dataset Creation: The dataset is constructed by collecting HDL source code samples from trusted repositories and benchmark vulnerability datasets. Each record contains the HDL code snippet, the corresponding hardware description language type, and a binary vulnerability label indicating whether the code is vulnerable or non-vulnerable. Data balancing techniques may be applied to address class imbalance, ensuring reliable model training and fair evaluation.

NLP Preprocessing: In this step, the raw HDL code is converted into a machine-readable format. Preprocessing includes removal of comments, handling of special characters, whitespace normalization, tokenization based on HDL syntax, and sequence padding or truncation. Stop-word elimination and keyword preservation are carefully handled to retain vulnerability-relevant constructs such as registers, signals, and control logic.

EDA: EDA is performed to gain insights into the dataset structure and characteristics. Statistical analysis is used to observe class distribution, average code length, and token frequency. Visualizations such as histograms and bar charts help identify dominant HDL constructs and potential bias in the dataset. These insights guide feature extraction and model selection.

Feature Extraction using Mini LM Model: A transformer-based Mini LM model is employed to generate contextual embeddings from the pre-processed HDL code. The model captures

long-range dependencies and semantic relationships within the code. Attention-gated structural fusion is applied to combine token-level embeddings with structural indicators, producing a unified feature vector that emphasizes vulnerability-prone regions.

Existing Decision Tree Classifier: The extracted embeddings are passed to a Decision Tree classifier to establish a rule-based baseline. The model learns hierarchical splits based on feature importance, enabling interpretability and insight into which embedding dimensions influence vulnerability decisions.

Existing Nearest Centroid Classifier: The Nearest Centroid classifier computes class-wise centroids in the embedding space. Each test sample is classified based on minimum distance to these centroids. This step evaluates the clustering behavior of transformer embeddings and serves as a computationally efficient baseline.

Existing RBM Bernoulli Model: A Restricted Boltzmann Machine with Bernoulli-distributed visible units is trained on the feature vectors to learn hidden representations. The RBM models probabilistic dependencies between features, enabling unsupervised learning of latent vulnerability patterns before classification.

Proposed Ridge Classifier: The Ridge Classifier is applied as the proposed supervised learning model. It leverages L2 regularization to manage multicollinearity in high-dimensional transformer embeddings. This approach improves generalization, stabilizes decision boundaries, and enhances classification accuracy for HDL vulnerability localization.

Performance Comparison: All classifiers are evaluated using standard metrics such as accuracy, precision, recall, F1-score, ROC curve, and confusion matrix. Comparative analysis highlights the strengths and limitations of each model, demonstrating the superior performance and stability of the proposed Ridge Classifier.



Prediction on Test Data: The optimized Ridge Classifier is used to predict vulnerability labels for unseen HDL code samples. This step validates the model’s ability to generalize to new data and accurately identify vulnerable code segments.

Flask Integration: The trained model is deployed using a Flask-based web application. Users can input HDL code through a web interface, which undergoes preprocessing, feature extraction, and classification in real time. The application outputs vulnerability status, enabling practical deployment in hardware security analysis workflows.

Ridge Classifier

The proposed HDL Vulnerability Localization framework follows a structured computational pipeline that begins with Feature Loading, where high-dimensional vectors generated from MiniLM embeddings containing both semantic and structural hardware logic are ingested into the Ridge Classifier. During Model Initialization, the classifier is configured with an alpha regularization parameter to mitigate multicollinearity and prevent overfitting, ensuring the model remains flexible yet robust. The Training phase involves optimizing weights to minimize the mean squared error incorporating L2 regularization, which grounds the model for stable, generalized predictions across diverse hardware descriptions. For Prediction on Test Data, the trained weights are applied to the feature vectors of unseen HDL samples to generate a vulnerability score, which is subsequently mapped to a specific class label. The system’s efficacy is then validated through rigorous Performance Evaluation, utilizing metrics such as accuracy, precision, recall, and F1-score to confirm its detection reliability. Finally, the Interpretation of Results leverages the magnitude of the learned weights within the Ridge architecture to provide transparency, allowing researchers to identify which specific code embeddings contribute most significantly to the prediction of hardware vulnerabilities.

4. Results Analysis

Fig. 3 shows the confusion matrix of the existing MiniLM with RBM model for the Vulnerable class. The matrix indicates that out of 200 truly vulnerable HDL samples, 158 are correctly classified as vulnerable (True Positives), while 42 are incorrectly predicted as non-vulnerable (False Negatives). For the secure class, 129 non-vulnerable samples are correctly identified as secure (True Negatives), whereas 71 secure samples are misclassified as vulnerable (False Positives). This result reveals that the RBM model captures a larger portion of vulnerable designs compared to some existing methods but suffers from a high false-positive rate, leading to reduced precision.

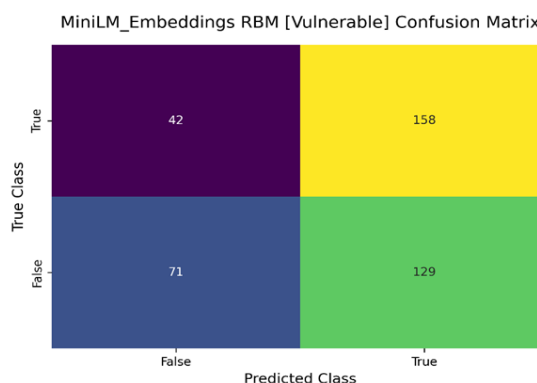


Fig. 3: Confusion matrix obtained using existing MiniLM with RBM model.

Fig. 4 shows the confusion matrix of the existing MiniLM with DTC model for the Vulnerable class. The matrix indicates that out of 200 truly vulnerable HDL samples, 160 are correctly classified as vulnerable (True Positives), while 40 are misclassified as non-vulnerable (False Negatives). For the secure class, all 200 non-vulnerable samples are correctly identified as secure (True Negatives), with no False Positives observed. This result highlights that the MiniLM-DTC model achieves perfect precision for vulnerable detection due to the absence of false alarms; however, it misses a portion of vulnerable instances, leading to reduced recall. Overall, Figure 6.4 demonstrates that while the existing DTC model is reliable in avoiding false



vulnerability detection, it is less effective in capturing all vulnerable HDL designs.

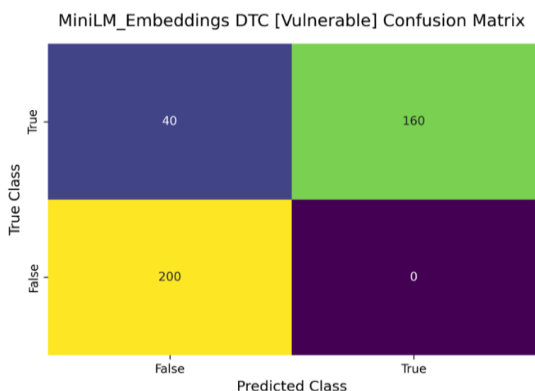


Fig. 4: Confusion matrix obtained using MiniLM with DTC model.

Fig. 5 shows the confusion matrix of the existing MiniLM_Embeddings NC classifier for the Vulnerable class. From the figure, out of 200 truly vulnerable HDL samples, 169 are correctly classified as vulnerable (True Positives), while 31 are misclassified as non-vulnerable (False Negatives). For the secure class, all 200 non-vulnerable samples are correctly identified as secure (True Negatives), with zero False Positives recorded. This indicates that the MiniLM-NC model achieves perfect precision by avoiding false vulnerability alarms; however, a small number of vulnerable designs remain undetected, slightly reducing recall.

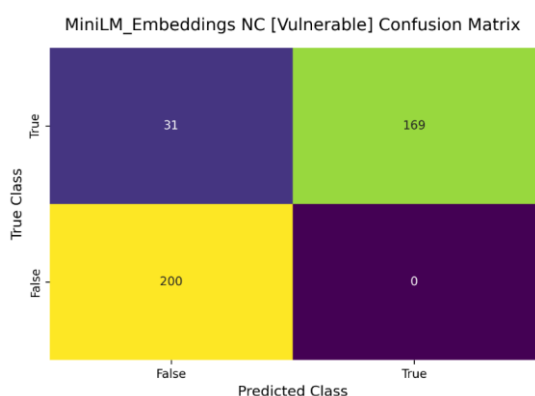


Fig. 5: Confusion matrix obtained using MiniLM with NC model.

Fig. 6 shows the confusion matrix of the MiniLM with proposed model for the Vulnerable class. The figure demonstrates that

all 200 truly vulnerable HDL samples are correctly classified as vulnerable (True Positives), with no False Negatives, and all 200 non-vulnerable samples are correctly identified as secure (True Negatives), with zero False Positives. This indicates perfect discrimination between vulnerable and secure HDL designs, resulting in 100% accuracy, precision, recall, and F1-score.



Fig. 6: Confusion matrix obtained using MiniLM with proposed RC model.

Table 1 shows the overall performance comparison of different MiniLM embedding based classifiers in terms of accuracy, precision, recall, and F1-score for HDL vulnerability localization. The MiniLM + DTC model achieves an accuracy of 0.8875, with a precision of 0.9082, recall of 0.8875, and F1-score of 0.8861, indicating reliable performance but with missed vulnerable instances. The MiniLM + NC model slightly improves the results, attaining an accuracy of 0.8950, precision of 0.9132, recall of 0.8950, and F1-score of 0.8938, reflecting better balance between detection and correctness. In contrast, the MiniLM + RBM model shows significantly lower performance, with an accuracy of 0.5725, precision of 0.5894, recall of 0.5725, and F1-score of 0.5513, mainly due to high misclassification of secure samples.

Table 1 Overall comparison of performance metrics obtained using existing and proposed model.

Metho	Accura	Precisi	Recal	F1-
-------	--------	---------	-------	-----



d	cy	on	l	Score
MiniLM with DTC Model	0.8875	0.9082	0.8875	0.8861
MiniLM with NC Model	0.8950	0.9132	0.8950	0.8938
MiniLM with RBM Model	0.5725	0.5894	0.5725	0.5513
MiniLM with Ridge Model	1.0000	1.0000	1.0000	1.0000

Fig. 8: Prediction results of FPGA vulnerability [FALSE].

Fig. 7 shows the prediction result generated by the proposed HDL vulnerability localization system for a given VHDL code snippet. The figure presents a race condition example, where the VHDL process updates the output signal q directly from the input d on the rising edge of the clock when the enable signal en is high, without additional synchronization or protection logic. The system correctly identifies this construction as vulnerable and reports Predicted_Vulnerable: True, with the programming language recognized as VHDL. This confirms that the proposed model is capable of capturing vulnerability-indicative patterns directly from HDL source code.

Fig. 8 shows the prediction result of the proposed FPGA vulnerability detection system for a secure VHDL design implementing an overflow-protected counter. The figure presents the HDL module secure_overflow_19, where a reset-controlled counter is incremented only when the current count is below a defined maximum threshold, thereby enforcing a boundary check to prevent overflow. The system correctly identifies the programming language as VHDL and classifies the design as secure, reporting Predicted_Vulnerable: False. This indicates that the proposed model successfully recognizes secure coding patterns and does not generate false vulnerability alarms.

```
Code:
-- Race Condition (VHDL - Vulnerable)
library IEEE;
use IEEE.STD_LOGIC_1164.ALL;
entity race_condition_1 is
    Port ( clk : in STD_LOGIC;
          en : in STD_LOGIC;
          d : in STD_LOGIC_VECTOR(6 downto 0);
          q : out STD_LOGIC_VECTOR(6 downto 0));
end race_condition_1;
architecture Behavioral of race_condition_1 is
begin
    process(clk)
    begin
        if rising_edge(clk) then
            if en = '1' then q <= d; -- Vulnerable
            end if;
        end if;
    end process;
end Behavioral;

Language: VHDL
Predicted_Vulnerable: True
```

Fig. 7: Prediction results of FPGA vulnerability [TRUE].

```
Code:
-- Overflow Fixed (VHDL - Secure)
library IEEE;
use IEEE.STD_LOGIC_1164.ALL;
use IEEE.STD_LOGIC_UNSIGNED.ALL;
entity secure_overflow_19 is
    Port ( clk : in STD_LOGIC;
          rst : in STD_LOGIC;
          counter : out STD_LOGIC_VECTOR(30 downto 0));
end secure_overflow_19;
architecture Behavioral of secure_overflow_19 is
    signal cnt : STD_LOGIC_VECTOR(30 downto 0) := (others => '0');
begin
    process(clk, rst)
    begin
        if rst = '1' then
            cnt <= (others => '0');
        elsif rising_edge(clk) then
            if cnt < "11111111111111111111111111111111" then cnt <= cnt + 1; -- Secu
        end if;
    end process;
    counter <= cnt;
end Behavioral;

Language: VHDL
Predicted_Vulnerable: False
```

5. Conclusion

The experimental evaluation of the proposed Transformer Embeddings with Attention-Gated Structural Fusion for HDL Vulnerability Localization demonstrates strong and consistent performance across all stages of analysis. Using MiniLM-based contextual embeddings, the system effectively captures syntactic and semantic patterns from HDL source code, enabling accurate classification of both vulnerable and secure designs. Quantitative results show that the MiniLM with proposed model achieves 100% accuracy,



precision, recall, and F1-score, outperforming existing methods such as DTC (accuracy 0.8875), NC (0.8950), and RBM (0.5725). Confusion matrix analyses further confirm that the proposed approach eliminates false positives and false negatives, ensuring reliable detection of race conditions and overflow-related vulnerabilities in both VHDL and Verilog implementations.

References

- [1] A. G. Ayar, S. Kundu, and S. Bhunia, "HDL-based vulnerability analysis for hardware Trojan detection using language-pattern mining," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 11, pp. 3892–3905, Nov. 2020.
- [2] J. Cruz, Y. Jin, and P. Makris, "Quantifying hardware Trojan vulnerability using Trojan Vulnerability Factor," in *Proc. IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, 2018, pp. 89–96.
- [3] C. Dong, Y. Jin, and Y. Makris, "Cost-driven deep learning framework for hardware Trojan detection," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 38, no. 10, pp. 1946–1959, Oct. 2019.
- [4] H. Salmani, M. Tehranipoor, and J. Plusquellic, "A survey of hardware Trojan detection techniques," *IEEE Design & Test*, vol. 32, no. 5, pp. 8–22, Oct. 2015.
- [5] R. Yasaei, K. Basu, and M. Tehranipoor, "Golden-free hardware Trojan localization using graph convolutional networks," in *Proc. IEEE HOST*, 2020, pp. 1–8.
- [6] S. Yang, Y. Jin, and D. Forte, "Self-referencing hardware Trojan detection without golden models," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 246–259, 2020.
- [7] B. Li, S. Wang, T. Zhang, and A. Smola, "Data augmentation for natural language processing: A survey," *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–35, 2021.
- [8] L. Shen, Z. Zhang, and K. Chakrabarty, "Accelerating hardware security verification via state-space reduction," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 9, pp. 2105–2118, Sept. 2019.
- [9] T. Kurihara, K. Hasegawa, and N. Togawa, "Evaluation of machine-learning-based hardware Trojan detection at gate-level IP cores," in *Proc. IEEE ISCAS*, 2018, pp. 1–5.
- [10] S. Amir and D. Forte, "Adaptive synthetic benchmark generation for hardware security," in *Proc. IEEE HOST*, 2019, pp. 77–84.
- [11] Y. Wang, Y. Jin, and X. Wang, "Ensemble learning for hardware Trojan detection," *IEEE Transactions on Emerging Topics in Computing*, vol. 8, no. 1, pp. 230–242, 2020.
- [12] P. Zhao and Q. Liu, "Unsupervised density-based hardware Trojan detection," *Microelectronics Reliability*, vol. 88–90, pp. 1165–1171, 2018.
- [13] J. Cruz, Y. Jin, and P. Makris, "Automated and configurable hardware Trojan insertion framework," in *Proc. IEEE HOST*, 2017, pp. 1–6.
- [14] K. Hasegawa, M. Yanagisawa, and N. Togawa, "Trojan-feature extraction at gate-level netlists and its application to hardware-Trojan detection using random forest classifier," in *Proc. IEEE ISCAS*, 2017, pp. 1–4.
- [15] K. Hasegawa, M. Yanagisawa, and N. Togawa, "Hardware Trojan detection using neural networks at gate level," *IEICE Transactions on Information and Systems*, vol. E101-D, no. 1, pp. 1–9, 2018.