# Evolutionary Large Language Models for Hardware Security: A Comparative Survey

Mohammad Akyash mohammad.akyash@ucf.edu ECE Department, University of Central Florida Orlando, Florida, USA Hadi M Kamali

hadi.mardanikamali@ucf.edu ECE Department, University of Central Florida Orlando, Florida, USA

### ABSTRACT

Automating hardware (HW) security vulnerability detection and mitigation during the design phase is imperative for two reasons: (i) It must be before chip fabrication, as post-fabrication fixes can be costly or even impractical; (ii) The size and complexity of modern HW raise concerns about unknown vulnerabilities compromising CIA triad. While Large Language Models (LLMs) can revolutionize both HW design and testing processes, within the semiconductor context, LLMs can be harnessed to automatically rectify securityrelevant vulnerabilities inherent in HW designs. This study explores the seeds of LLM integration in register transfer level (RTL) designs, focusing on their capacity for autonomously resolving security-related vulnerabilities. The analysis involves comparing methodologies, assessing scalability, interpretability, and identifying future research directions. Potential areas for exploration include developing specialized LLM architectures for HW security tasks and enhancing model performance with domain-specific knowledge, leading to reliable automated security measurement and risk mitigation associated with HW vulnerabilities.

#### **KEYWORDS**

Large Language Models, Hardware Security, RTL Debugging

# **1** INTRODUCTION

In today's semiconductor technology landscape, As system-on-chip (SoC) designs integrate more and more intellectual property (IP) cores, each with unique functionality and security challenges, each from various vendors, each with ever-increasing complexity, we witness a growing challenge in detecting and fixing vulnerabilities. Given the pivotal role of SoCs, while substantial efforts have been invested in software (SW) testing and debugging, SoC (HW-based) testing, validation, and verification remain less mature [30]. The problem worsens while bugs are detected at lower levels of abstraction, which makes respins extremely difficult (and even impossible, e.g., post-silicon) [34]. Moreover, existing solutions, from simulation to formal verification, usually require expertise. Such solutions also suffer from scalability issues, unable to cope with the growing size and complexity of SoCs [2]. Furthermore, these solutions cannot address the majority of SoCs' vulnerabilities due to rapidly evolving threats, such as zero-day attacks.

With the rapid evolution of LLMs, their capabilities have expanded into the domain of SW code generation with remarkable success, e.g., OpenAl's Codex [36]. Moreover, the scope of LLMs extends to SW code testing and verification while outperforming techniques like fuzzing [32]. While significant progress has been achieved in SW through LLMs, studies at the HW/SoC level, particularly at RTL, have been dispersed. Many studies have initiated the LLMs' applicability at the HW/SoC level by raising questions like whether "LLM can generate HDL" or "LLM can validate HW designs". Just like in SW, LLMs have the potential to be utilized for both HW design, testing and validation (see Fig. 1). These studies show harnessing LLMs' capability to analyze, comprehend, and generate/validate complex code structures, might make them a right target vs. existing formal tools to identify potential security vulnerabilities within RTL codes [3, 37]. However, ensuring the integrity and security of HW designs, coupled with the potential for unknown vulnerabilities, presents broader challenges.

This survey aims to offer a useful and comprehensive snapshot of rapidly growing use of LLMs in HW/SoC designs, particularly for security. We explore advancements, analyzing the pros and cons of each method. By examining current approaches, this work highlights the innovative application of LLMs to automate the detection and resolution of security vulnerabilities in HW designs. Also, we investigate future research directions, emphasizing the need for specialized LLM architectures and domain-specific knowledge integration. Our goal is to outline a roadmap for harnessing the full potential of LLMs in addressing HW security challenges, setting the stage for more robust and secure HW systems.

#### 2 LLMS FOR SW: ENGINEERING AND TESTING

Since the 1950s, many research efforts have been undertaken to develop highly efficient automated code generation tools [38]. These efforts have spanned from traditional program synthesizers [38]<sup>1</sup>, either deductive or inductive, to current neural-based models, notably codebase-reliant generative models [31].

With recent outrageous advancements in LLMs, massive research has focused on applying LLMs for independent SW code generation, leading to widely-used platforms like Codex and CodeGen [4]. The foundation of these models lies in autonomously predicting the subsequent token by considering the preceding context, typically comprising function signatures and docstrings that describe the intended functionality of the program, translating human-written instructions into precise code snippets or entire programs [4].

While this code generation relies on natural language processing (NLP), unlike natural language that is typically parsed as a sequential array of words or tokens, code generation is scrutinized based on its syntactic and semantic structure, often depicted using tree structures, e.g., abstract syntax trees (AST) [39]. Also, programming languages have a limited set of keywords, symbols, and rules, unlike the broad and nuanced vocabulary of natural languages.

<sup>&</sup>lt;sup>1</sup>Synthesizers aim to automatically generate programs (SW codes), based on a space search over a variety of constraints relevant to domains known as Domain Specific Languages (DSLs). These techniques are mostly limited to pre-defined DSLs and thus suffer scalability, being general-purpose, and adaptability issues [1].

GLSVLSI '24, June 12-14, 2024, Clearwater, FL, USA



Figure 1: The Usage of LLMs for HDL (RTL) Generation/Validation.

Given such differences, the primary concern for LLM-generated code is (i) correctness (testing and verification process), and (ii) codebase data hungriness [39]. In terms of correctness, testing and validation from the viewpoint of LLMs require well-defined metrics, where traditional metrics, e.g., *BLEU* that widely used in NLP assessments [39], fail due to their focus on linguistic similarity. For example, *CodeBLEU* that evaluates the quality of code produced by LLMs, or *Pass@k* that quantitatively measures the functional accuracy of code generation models, are example of such new metrics [36]. Regarding codebase data for code generation, substantial codebase data<sup>2</sup> is required for enhanced training and/or fine-tuning to improve the efficacy of LLMs for code ganeration [4, 36].

## 3 LLMS FOR HW: DESIGN AND TESTING

Similar to SW engineering and testing, leveraging LLMs can significantly optimize and enhance circuit design processes, particularly within Electronic Design Automation (EDA) frameworks. LLMs can be used at high level abstraction, e.g., RTLs, to (i) reduce manual efforts for implementation<sup>3</sup>, (ii) address the challenge of lacking HDL codebase<sup>4</sup>, (iii) expedite time-to-market (TTM) in the competitive chip design process, and (iv) enable a more efficient and reliable system (by reducing human-induced faults) [40].

The current LLM-based methodologies in HW can be classified into two primary categories: (1) Development of automated AI agents aimed at streamlining EDA workflows (e.g., ASIC flow); (2) Derivation of SW code generation for RTL implementation. Regarding the former category, LLMs assist in various tasks such as script generation, architecture specification, and interpretation of compilation reports, thereby minimizing the workload of the design team. Within the latter category, solutions predominantly utilize LLMs in two manners: (i) refinement of design prompts, which entails the creation (engineering) of more precise prompts to guide LLMs towards RTL generation with increased effectiveness, and (ii) RTL-based tuning, which involves directly tuning LLMs through training on RTL code examples. A comparison of all existing LLMbased approaches in these two categories is shown in Table 1.

#### 3.1 LLM Agent for EDA Automation

Several studies have explored the potential of LLM in automating the ASIC design/implementation process [8, 14, 27, 29]. ChatEDA and ChipNeMo are two examples of task planning and execution agents that interpret natural language commands from the design team. ChipNeMo [29] implements a series of domain-specific training strategies for chip design tasks. It involves the deployment of bespoke tokenizers, domain-adaptive continued pretraining, and supervised fine-tuning guided by domain-specific instructions. ChatEDA [27] aims to facilitate optimal interaction with the EDA tools by comprehending instructions in natural language for generating and delivering executable programs.

Using such techniques, LLM agents can offer automated ASIC flow, from RTL generation to GDSII creation, by invoking necessary SW tools and utilizing required scripts/files. However, while promising, these techniques necessitate thorough analysis to truly enhance automation in EDA tools for the following reasons:

(1) Expert-Oriented Training and Fine-Tuning: Constructing such frameworks heavily relies on expert efforts for training or finetuning them to accommodate specific ASIC flows. Given the variety of technologies with their respective documentation, syntaxes, flows, and scripting methods, the pre-trained LLM may not offer a universally applicable model for all environments.

(2) Failure in Handling Unforeseen Incidents: Despite extensive finetuning, the LLM-based agent may inaccurately extract information from reports/specs or generate incorrect scripts/configs when confronted with new incidents in the flows. Technology advancements, EDA tools updates, etc., may worsen this issue, as the LLM agent may fail to provide the desired output under evolving conditions.
(3) Dependence on Technology: To clarify this, we raise a question! How similar is the EDA flow (i) from one design to another design, (ii) from one technology to another technology, (iii) from one vendor to another vendor? Now, the question becomes how deep is LLM fine-tuned based on these designs, technologies, and vendors? While chatbots may offer basic assistance, the prospect of achieving comprehensive automation seems to remain elusive.

### 3.2 LLM for RTL Generation and Refinement

The main LLM-based RTL-oriented research focuses on the generation and refinement of RTL, primarily transitioning from specification to RTL design (+optimization). Initial efforts emphasize prompt engineering, crucial to successful RTL generation while relying on the existing LLMs [8, 10, 25]. Other methods, e.g., Verigen and VerilogEval, adapt open-source LLMs like CodeGen [4], followed by fine tuning on RTL, to produce more optimized HDL modules [13, 41]. Additionally, studies such as ChipGPT and AutoChip explore use of feedback mechanisms to enhance HDL quality, addressing aspects like compilation errors and design optimization (PPA optimization) [10, 20]. While these methods often rely on static analysis, DeLorenzo et al. Introduce optimization techniques like Monte Carlo tree search (MCTS) to fine-tune LLM tokens even further for more tuned optimization at the backend of LLMs [12].

<sup>&</sup>lt;sup>2</sup>The data must be not only vast but also diverse, relevant, and of high integrity as the superioir quality codebase data enhances model performance significantly [32].
<sup>3</sup>It can potentially serve as an alternative to high level synthesis (HLS), thereby enabling designers with limited HDL expertise to swiftly generate HW designs [40].

<sup>&</sup>lt;sup>4</sup>Lack of HDL codebase is always a substantial barrier for Al-driven HW solutions, consequently enhancing the efficiency of the training phase [33].

Study	Target	LLM Engine	Input	Output	Comment (-Shortcomings-)		
Chang et al. [10]	RTL Generation + Refine- ment	GPT-3.5	Design Specification Prompts + Human Feedback for Correc- tions	RTL Module	- Static PPA analysis is post-LLM with no LLM-based improvement. - Human feedback is needed for manual correction per design.		
Thakur et al. [20]	RTL Generation w/ guaranteed Compilation	GPT-4, Llama2, GPT-3.5T, Claude 2	Design prompt + Com- pile/Synthesis Report	Compiled and Tested RTL Design	<ul> <li>Feedback addresses compilation/simulation errors but may alter function priority, leading to unintended func- tions.</li> <li>No Feedback for PPA Efficiency Matter</li> </ul>		
He et al. [27]	Automatic EDA Flow Scripting and Execution Calls	Llama2-70B	Natural Language Instructions + RTL Design	EDA Tool Commands & Reports + Scripts + Synthesized Design + Lay- out (GDSII)	- It is either design- or technology-Dependent. - Cannot be easily design/tool-agnostic.		
Li et al. [14]	Architecture Specifications Generation + Review	GPT-4	Architecture specifications + RTL Design	Hierarchical Reviewed Architecture Specifications	- Specifications are limited to the existing technologies. - It is mostly processor-based instructions. Not for generic HW.		
Lu et al. [25]	RTL Generation	GPT-3.5, GPT-4, VeriGen, StarCoder	Natural language instructions	RTL Design	<ul> <li>With no feedback, success rate is low for functional correctness.</li> <li>The reference designs are very limited and relatively small.</li> </ul>		
Liu et al. [18]	RTL Generation	RTLCoder	Natural language instructions	RTL Design	<ul> <li>Diversity rate is low in the training dataset.</li> <li>The functional correctness of training dataset is not ensured, leading to lower functional coverage in the generated outputs.</li> </ul>		
Thakur et al. [41]	Completing Partial RTL De- sign	MegatronLM-355M, Code- Gen, code-davinci-002, and J1-Large-7B	Partial RTL Design + Custom problem set with testbenches	RTL Design	<ul> <li>Lack of Organized Dataset.</li> <li>RTLLM shows the performance does not surpass existing commercial models.</li> <li>Completion necessarily does not provide correct functionalities.</li> </ul>		
Cheng et al. [11]	RTL Generation + Repair + EDA Script Generation	Llama2-7B, Llama2-13B	Natural language descriptions + Verilog files + EDA scripts	Corrected Verilog code + Verilog code from descriptions + EDA scripts	- For refinement, it is for syntactic errors (compilation issues).		
DeLo et al. [12]	RTL Generation	VeriGen-2B	Natural language instruction + RTL modules description	Compiled, Tested, and PPA Improved RTL Design	<ul> <li>Tested on Small Toy Circuits, e.g., adders and MAC units.</li> <li>Stochastic behavior of MCTS. Less Improvement in More Iterations.</li> </ul>		
Li et al. [42]	RTL Synthesis (Mapping)	Circuit Transformer	Gate-Level Design (AIG)	Design Model (Truth Table) + Synthe- sized AIG	- Low Accuracy for Larger Circuits. - Low Performance with no MCTS (Low Scalability).		

More recent advancements have shifted the focus from fine tuning and prompt engineering in existing LLMs to the development of dedicated circuit transformers, e.g., Li et al. Introduce "Circuit Transformer" with 88M parameters and integrated MCTS for optimization, leading to a fully open-source independent LLMs for RTL [42]. Similarly, RTLCoder proposes an automated data generation flow utilizing a model with 7B parameters, producing a sizable labeled dataset for RTL generation [18]. These endeavors have led to the emergence of large circuit models (LCM), enhancing the expression of circuit data's semantics and structures, thus creating more robust, efficient, and innovative design approaches.

Despite its promise, more research is needed as follows: (1) Universality Issues: LLM-based RTL generation faces limitations due to scarce codebase knowledge available for model fine-tuning and training per application [18]. As an example, developing security enclaves or fully-debugged Verilog modules is incredibly challenging as there are not many training datasets available for it. (2) Verification (Functional) Issues: Existing studies highlight the complex nature of (functional) verification tasks, further magnified by the limited availability of trained models for test bench generation and functional simulation [13]. The complexity of circuit designs, which involve both functional and structural attributes, worsens the challenge, as even small changes to the structure (a code line) can have significant effects on functionality, underscoring the complexity of testbench generation and simulation of circuits. (3) Scalability Issues: Scalability is crucial for RTL-based LLMs in addressing complex circuit designs [25]. Efforts to enhance computational efficiency and model architecture sophistication are essential to accommodate larger designs and meet evolving electronic

device demands. Further research is necessary to overcome scalability challenges and maximize LLM potential in RTL generation.

# 4 LLM FOR HW: SECURITY (VERIFICATION)

Given the paramount significance of security of HW designs in modern SoCs, and in light of the earlier discussion emphasizing the importance of verification over LLMs, several studies have commenced employing LLM for SoC verification (moving towards bug-free designs, either functional or security-oriented). Similar to LLM-based RTL design, these approaches fall into two main categories: (i) refinement of design prompts, where designers guide LLMs toward generating secure code (i.e. prompt engineering), and (ii) <u>RTL-based tuning</u>, which is about altering the LLM's framework itself to generate output bug-free code. In advancing HW security, researchers have leveraged LLMs using either pure natural language prompts (i.e. description of the code) or a blend of natural language (i.e. comments designed by human experts) and code. The following describes these two categories in detail and how each category can enhance verification and security for HW designs.

#### 4.1 **Prompt Engineering**

Prompt engineering is the practice of designing inputs for LLMs, to obtain specific, desirable outputs. This technique optimizes the interaction with LLMs to improve its performance on various tasks, leveraging strategies like few-shot [21], and chain-of-thought [9] prompting to guide the model's responses effectively. A few recent studies in HW explore the applications of prompt engineering for enhancing vulnerability detection and repair, as well as design verification. For example, [3] employs a range of detailed instruction

prompts for various LLMs, aiming to evaluate the efficacy of each model in correcting HW vulnerabilities<sup>5</sup>. Fig. 2 shows an example of how prompting GPT-4 with a bug description and repair instructions alongside the Verilog code enables GPT-4 to address the vulnerability. Here are two important lessons to be learned:

(1) The example shows that being super specific is crucial in engineering the prompt to ensure the generated code is devoid of vulnerabilities. Thus, it is vital to have careful crafting by human experts to generate such prompts. This requirement for human input could become a tedious process, posing challenges in scaling and automating the approach for broader applications.

(2) The performance and efficacy of LLMs depends on the infrastructure of LLM used. While commercial LLMs like GPT-4 tend to outperform models trained on coding datasets, including Codegen and VeriGen, in terms of repair accuracy and efficacy, this advantage comes at the cost of increased number of parameters.

The importance of precision in prompt generation is also shown in [15], relying on ChatGPT, revealing the fact that the success rate can be degraded significantly while the model is more limited<sup>6</sup>. This study also demonstrates models misguiding the designers while the Verilog code of various CWE scenarios as part of instruction can lead to new form of vulnerabilities from prompts (may not fully represent the capture of potential vulnerabilities in SoC designs).

To enhance verification capability, some studies focus on the use of LLMs for verification assertion generation (e.g., SystemVerilog Assertions (SVAs)). For instance, [16] uses GPT-4 in an iterative mechanism to refine prompts for GPT-4, enabling it to generate more accurate and complete SVA properties from RTL code. This approach coupled with AutoSVA2, which automatically generates formal verification testbenches, enables LLM-guided formal verification towards more automation. However, the major obstacle to this automation is the reliance of this approach on iterative refinement by an expert, which requires a deep understanding of both HW verification and prompt engineering.

Similarly, AssertLLM [23] uses a customized GPT-4 Turbo to generate SVAs (functional verification assertions) from natural language design specifications (translating design documents). Although results show high success rate, this model is also heavily dependent to the quality and completeness of the design documents. This is while richness of documentation is always a critical issue in HW design, thus AssertLLM might struggle to generate assertions that fully capture the intended design behavior.

LLM4DV [28] uses LLMs with prompt templates to automate the generation of test stimuli for verification. LLM4DV integrates LLMs with a systematic method that includes a stimulus generation agent, prompt templates, and four LLM-based improvements, e.g., summarizing prompts, resetting, etc. Evaluated using three custom-designed large-scale DUTs, this framework demonstrated promising results and achieved high coverage rates in simple scenarios. However, this approach focuses more on coverage-related metrics, overlooking security-oriented vulnerabilities.

Similar to these formal-based mechanisms, [37] proposes designing an evaluation framework that includes generating natural

#### Mohammad Akyash and Hadi M Kamali



```
Based on the provided instruction, correct the security bug in this Verilog module.
BUG: Access Control Check Implemented After Asset is Accessed
Ensure that access is granted before data is accessed.
module user grant access(data out, usr id, data in, clk, rst n);
       user_grant_access(data_
output reg [7:0] data_out;
input wire [2:0] usr_id;
input wire [7:0] data_in;
input wire clk, rst_n;
reg grant_access;
               ays @(posedge clk or negedge rst_n) begin
if (!rst_n)
                      data_out = 0;
                else begi
                       grant_access = (usr_id == 3'h4) ? 1'b1 : 1'b0;
data_out = (grant_access) ? data_in : data_out;
               end
  en dm odu 1e
Response
module user_grant_access(data_out, usr_id, data_in, clk, rst_n);
output reg [7:0] data_out;
input wire [2:0] usr_id;
input wire [7:0] data_in;
input wire [7:0] data_in;
input wire clk, rst_n;
        reg grant_access;
        always @(posedge clk or negedge rst_n) begin
if (!rst_n) begin
data_out <= 0;
end else begin
                      else begin
grant_access = (usr_id == 3'h4) ? 1'b1 : 1'b0;
if (grant_access) begin
data_out <= data_in;</pre>
                       end
               end
endmodule
```

Figure 2: An Examplary Case in GPT-4 for Security Debugging.

language prompts that mimic code comments in assertion files, using these prompts to generate SVAs with LLMs, and then assessing the correctness of these assertions against a benchmark suite of real-world HW designs and corresponding golden reference assertions. The results demonstrate that LLMs, with varying levels of detail in the prompts, can generate valid HW security assertions.

More recent use of LLMs for RTL debugging aimed to enhance automation in the domain. For instance, RTLFixer [26] automatically rectifies syntax errors in Verilog code by leveraging Retrieval-Augmented Generation (RAG) and the ReAct prompting strategy. RTLFixer employs a retrieval database filled with expert knowledge of syntax errors. ReAct also introduces an iterative approach involving reasoning, action, and observation, mimicking experts' debugging techniques. This combination builds a more effective system for automating the debugging. However, it still heavily relies on the comprehensiveness and currentness of the external knowledge database, which is collected by human experts.

Some LLM-based studies focus on the use of such models at the SoC level. DIVAS [19] uses LLMs to analyze SoC specifications and crafts precise queries that encapsulate potential security vulnerabilities related to the SoC. These queries are submitted to LLMs, e.g., ChatGPT and Google's BARD, and the LLMs map these queries to relevant CWE vulnerabilities that could compromise the SoC. Once CWEs have been identified, DIVAS utilizes LLMs to construct SVAs for each. These SVAs are designed to act as security verification mechanisms, ensuring the SoC's design complies with security standards and is safeguarded against identified vulnerabilities.

Similarly, [5] explores how GPTs are utilized in SoC level for security vulnerability insertion, detection, assessment, and mitigation. This study, focusing on smaller models, e.g., ChatGPT-3.5, and relying on a sub-set of CWEs, evaluates the modification possibility over RTL using one- and few-shot learning. By comprehensive

<sup>&</sup>lt;sup>5</sup>These prompts must provide a thorough description of the bug, strategies for debugging, and illustrative examples that contrast insecure code with its secure counterpart.
<sup>6</sup>The number of parameters was restricted to a range of millions instead of billions.

#### Evolutionary LLMs for Hardware Security: A Comparative Survey

Study	Target	LLM Engine	# of Bugs	Success Rate	Source of Bench- marks	Expert Knowledge Needed?	Reference (for Eval)	Comment	
Nair et al. [15]	Prompt gen- eration for Debugging RTL	ChatGPT	10	100%*1	CWE (Descrip- tions)	For the Whole Process	Manual expert interven- tion per debugging	- Cannot be automated. - Limited evaluation on CWEs	
Kande et al. [37]	Detection (Gener- ate Assertion)	OpenAI Codex (code-davinci- 002)	10	~25%	Hack@DAC21, OpenTitan	For manually building detailed se- curity constraints	Golden Assertion	<ul> <li>High success rate only when bug and securi policy is known. Otherwise, it is below 10%.</li> <li>Only for single endmodule, No Hierarchical a Recursive SVA.</li> </ul>	
Ahmad et al. [3]	Repair (pre-detected bugs)	OpenAI Codex (code-davinci- 001, code- davinci-002, code-cushman- 001), CodeGen	15	~31%	CWE (Bench- mark), OpenTitan, Hack@DAC21	- For training (dataset generation for assisting repairs) - For CWEAT static analyze verifi- cation		<ul> <li>Only applicable on pre-observed cases with high similarity (to be detected by CWEAT)</li> </ul>	
Saha et al. [5]	Detection (Gener- ate Assertion), se- curity vulnerabil- ity insertion	GPT 3.5, GPT 4	N/R*2	N/R*2	CWE, Trust-Hub	For prompt engineering and evalu- ation	Manual expert interven- tion per debugging	- Limited evaluation on CWEs and smart toy circuits.	
Fu et al. [22]	Detection and/or Repair	StableLM, Falcon, LLama2	1 (different models)	~35%	Open-Source SoCs and Micro- processors	For fine-tuning (Open-source code classifications)	Repaired Code (Pre- and Post-correction of Git (CVA6, Opentitan,))	<ul> <li>Detailed enhancement for training is needed Per design, a new training might be required.</li> <li>Raw dataset is limited and not design-agnostic).</li> </ul>	
Meng et al. [24]	Detection (Gener- ate Assertion)	HS-BERT	8	326 Bugs from 1723 sentences	RISC-V, Open- RISC, MIPS, OpenSPARC, OpenTitan (docu- mentation)	For classifying security rules in doc- uments Manual expert labling for security property validation		- Limited by the quality of the input HW docu mentation. - Limited to the design/verification team knowl edge.	
Fang et al. [23]	Detection (Gener- ate Assertion)	GPT4 Turbo	N/A	89%	Open-source CPUs, SoCs, Xbars, arith- metic.	For extracting verification-required information from documents	Golden RTL Implemen- tation	- Limited by the quality of the input HW docu mentation. - Mostly syntactic and basic functional verifica tion.	
Paria <i>et al.</i> [19]	Detection (Gener- ate Assertion)	ChatGPT, BART	N/A	N/A	CEP SoC (MIT- LL)	For assumptions (CWE-based secu- rity rules)	N/R*2	- Expert review for Spec Generation is needed per design.	
Vera <i>et al.</i> [16]	Detection (Gener- ate Assertion)	GPT-4	N/R*2	N/R*2	RISC-V CVA6	For building rules related to asser- tions Previously development of the formal tools (Auto		- The success rate heavily depends on expert's input for prompt engineering.	
Zhang et al. [28]	Test Stimuli Gen- eration	GPT-3.5-turbo	N/A	small: ~98%, large: ~65%	Self-designed RTL Designs	For prompts generation	Coverage Monitoring	- Not for security purposes. Coverage-based test ing.	
Tsai et al. [26]	Syntax Errors Re- pair	GPT-3.5, GPT-4	212	98.5%	VerilogEval benchmarks, RTLLM bench- marks	For retrieval database (debugging reference)	VerilogEval, RTLLM	- Not for security purposes. Only for Syntax e rors.	

Table 2: A	Тор	Comparison of	LLM-based	нw	Security	Validation	Solutions
10010 2011	-•P	companioon or	LLIII buotu		occurry	, and a state of the	ooranono

\*1: It is 100% as all the debugging is done manually. Bug is known, the debugging instruction (flow) is known, and GPT is used for generation. N/R\*2: Not Reported.

exploration, the study suggests specific prompt guidelines for effectively using LLMs in SoC security-related tasks.

LLMs possess a dual-use nature; While advancing HW security initiatives, LLM can also present new threats simultaneously. [7] delves into the potential of general-purpose models like ChatGPT in the offensive HW security domain This study involves employing prompt engineering techniques to guide LLMs in filtering complex HW design databases, correlating system-level concepts with specific HW modules, identifying security-critical design modules, and modifying them to introduce HW Trojans. This study initiates the possibility of using LLMs for building more stealthy and undetectable HW Trojans, reshaping the characteristics of HW Trojan implementation, detection, and mitigation.

#### **Fine-Tuning** 4.2

As mentioned previously, some of these LLM-based HW verification solutions rely on fine-tuning, which involves adjusting a pre-trained language model by training it on Verilog/SVA data. However, LLMs require extensive datasets for effective training, posing a significant challenge in specialized domains, particularly in HW security due to the scarcity of targeted data. LLM4SecHW [22] is one example, which leverages a dataset compiled from defects and remediation steps in open-source HW designs, using version control data from GitHub. This dataset was created by selecting significant HW projects such as CVA6, CVA5, OpenTitan, etc., and extracting commits, issues, and pull requests (PRs) related to HW designs. This approach provides a rich source of domain-specific

data for training models, specifically tailored to identifying and fixing bugs in HW designs. Although innovative and promising, the quality of this data is dependent on the filtering process accuracy. The effectiveness of LLMs in debugging HW designs is thus directly tied to how precisely the data is curated and processed.

The NSPG framework [24] is another example of LLM solution for HW verification that offers a novel methodology for automating the generation of HW security properties utilizing fine-tuned LLMs. This approach is anchored by the development of a specialized language model for HW security, HS-BERT, which is trained on domain-specific data. Through deep evaluation on previously unseen design documents from OpenTitan, NSPG has proven its capability by extracting and validating security properties, showing security vulnerabilities within the OpenTitan design. However, a notable limitation of not only NSPG, but also all HW-oriented finetuned model for now lies in its dependency on the quality and scope of the HW documentation provided as input (which is almost super limited). As in the realm of HW/SoC design, this documentation often remains incomplete, inconsistent, or lack necessary detail, the precision and efficacy of the solution could be adversely affected.

#### 5 TAKEAWAYS AND FUTURE DIRECTIONS

In all facets of using LLMs for HW security, it becomes apparent that a significant hurdle, whether in HW design or in testing/verification, whether stemming from prompt engineering or fine-tuning, lies in the procurement and effective utilization of quality data [17]. Also, as depicted in Table 2, creating specialized LLMs (e.g., LCMs) or

employing pre-existing ones necessitates a deep expert knowledge to achieve a high success rate for generation, detection, and mitigation. Considering these two obstacles, despite being promising, the endeavor requires rigorous effort across multiple facets.

Creating a standard database reference is crucial for both training and evaluating the methods proposed in this domain. It facilitates a fair comparison among different techniques, ensuring that the pros/cons of each approach can be accurately assessed. Moreover, high-quality RTL data is indispensable for the optimal training of LLMs. It enables these models to learn the intricacies of RTL designs effectively, thereby enhancing their efficiency in security tasks.

Given the distinct characteristics of RTL codes as opposed to natural language texts, it becomes crucial to consider domain-specific models for handling HW codes. Incorporating concepts such as graphs and ASTs into LLMs can bridge the gap between the structural nuances of RTL codes and the inherently sequential processing of conventional language models. It is crucial to devise a novel metric specifically for evaluating the security coverage of RTL code examined by LLMs. This metric would serve as a critical feedback mechanism for LLMs, enabling them to assess and refine their output continually. By quantitatively measuring the security of RTL designs, the metric would allow LLMs to optimize their learning process towards generating code that is not only functionally correct but also adheres to high security standards.

Building on the foundational strategies mentioned above, further refinement can be achieved through the optimization of continuous prompts<sup>7</sup>. Such strategies also open the doors for mechanisms to enhance prompt automation for LLMs, e.g., auto-prompting<sup>8</sup>. These optimizations are open research directions potentially presenting a more feasible and efficient alternative to LLM fine-tuning.

### 6 CONCLUSION

This paper examined the use of LLMs in detecting/addressing security flaws in HW designs. We specifically analyzed their incorporation into RTL, revealing their independent problem-solving abilities in this domain. Our examination of existing approaches highlights both their benefits and drawbacks, notably scalability and accuracy issues. Also, we identified potential areas for future research. Our suggestion involves developing dedicated LLM architectures and datasets focused on HW security, indicating a path toward targeted improvements that could mitigate HW vulnerabilities.

#### REFERENCES

- A. Desai et al. 2016. Program synthesis using natural language. In International Conference on Software Engineering. 345–356.
- [2] A. Inamdar et al. 2021. Development of superconductor advanced integrated circuit design flow using synopsys tools. IEEE Transactions on Applied Superconductivity 31, 5 (2021), 1–7.
- [3] B. Ahmad et al. 2024. On Hardware Security Bug Code Fixes By Prompting Large Language Models. IEEE Transactions on Information Forensics and Security (2024).
- [4] E. Nijkamp *et al.* 2022. Codegen: An open large language model for code with multi-turn program synthesis. *arXiv preprint arXiv:2203.13474* (2022).
- [5] D. Saha et al. 2023. LLM for SoC Security: A Paradigm Shift. arXiv:2310.06046 [cs.CR]
- [6] D. Yin et al. 2023. Dynosaur: A Dynamic Growth Paradigm for Instruction-Tuning Data Curation. arXiv:2305.14327 [cs.CL]

- [7] G. Kokolakis et al. 2024. Harnessing the Power of General-Purpose LLMs in Hardware Trojan Design. In Proceedings of the 5th Workshop on Artificial Intelligence in Hardware Security, in conjunction with ACNS.
- [8] J. Blocklove et al. 2023. Chip-Chat: Challenges and Opportunities in Conversational Hardware Design. In 2023 ACM/IEEE 5th Workshop on Machine Learning for CAD (MLCAD). IEEE. https://doi.org/10.1109/mlcad58807.2023.10299874
- [9] J. Wei *et al.* 2022. Chain-of-thought prompting elicits reasoning in large language models. 35 (2022), 24824–24837.
- [10] K. Chang et al. 2023. ChipGPT: How far are we from natural language hardware design. arXiv:2305.14019 [cs.AI]
- [11] K. Chang *et al.* 2024. Data is all you need: Finetuning LLMs for Chip Design via an Automated design-data augmentation framework. arXiv:2403.11202 [cs.AR]
- [12] M. DeLorenzo *et al.* 2024. Make Every Move Count: LLM-based High-Quality RTL Code Generation Using MCTS. arXiv:2402.03289 [cs.LG]
- [13] M. Liu et al. 2023. VerilogEval: Evaluating Large Language Models for Verilog Code Generation. In 2023 IEEE/ACM International Conference on Computer-Aided Design (ICCAD).
- [14] M. Li et al. 2024. SpecLLM: Exploring Generation and Review of VLSI Design Specification with Large Language Model. arXiv:2401.13266 [cs.AR]
- [15] M. Nair et al. 2023. Generating Secure Hardware using ChatGPT Resistant to CWEs. Cryptology ePrint Archive, Paper 2023/212. https://eprint.iacr.org/2023/ 212 https://eprint.iacr.org/2023/212.
- [16] M. Orenes-Vera et al. 2023. Using LLMs to Facilitate Formal Verification of RTL. arXiv:2309.09437 [cs.AR]
- [17] Suriya Gunasekar et al. 2023. Textbooks Are All You Need. arXiv:2306.11644 [cs.CL]
- [18] S. Liu et al. 2024. RTLCoder: Outperforming GPT-3.5 in Design RTL Generation with Our Open-Source Dataset and Lightweight Solution. arXiv:2312.08617 [cs.PL]
- [19] S. Paria et al. 2023. DIVAS: An LLM-based End-to-End Framework for SoC Security Analysis and Policy-based Protection. arXiv:2308.06932 [cs.CR]
- [20] S. Thakur et al. 2023. AutoChip: Automating HDL Generation Using LLM Feedback. arXiv:2311.04887 [cs.PL]
- [21] Tom B. Brown et al. 2020. Language Models are Few-Shot Learners. CoRR abs/2005.14165 (2020). arXiv:2005.14165 https://arxiv.org/abs/2005.14165
- [22] W. Fu et al. 2023. LLM4SecHW: Leveraging Domain-Specific Large Language Model for Hardware Debugging. In AsianHOST.
- [23] W. Fang et al. 2024. AssertLLM: Generating and Evaluating Hardware Verification Assertions from Design Specifications via Multi-LLMs. arXiv:2402.00386 [cs.AR]
- [24] X. Meng et al. 2023. Unlocking Hardware Security Assurance: The Potential of LLMs. arXiv:2308.11042 [cs.CR]
- [25] Y. Lu et al. 2023. RTLLM: An Open-Source Benchmark for Design RTL Generation with Large Language Model. arXiv:2308.05345 [cs.LG]
- [26] Y. Tsai et al. 2024. RTLFixer: Automatically Fixing RTL Syntax Errors with Large Language Models. arXiv:2311.16543 [cs.AR]
- [27] Z. He et al. 2024. ChatEDA: A Large Language Model Powered Autonomous Agent for EDA. arXiv:2308.10204 [cs.AR]
- [28] Z. Zhang et al. 2023. LLM4DV: Using Large Language Models for Hardware Test Stimuli Generation. arXiv:2310.04535 [cs.LG]
- [29] M. Liu et al. 2023. ChipNeMo: Domain-Adapted LLMs for Chip Design. arXiv:2311.00176 [cs.CL]
- [30] H. Witharana et al. 2022. A survey on assertion-based hardware verification. ACM Computing Surveys (CSUR) 54, 11s (2022), 1–33.
- [31] J. Austin et al. 2021. Program synthesis with large language models. arXiv preprint arXiv:2108.07732 (2021).
- [32] J. Liu et al. 2024. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. Advances in Neural Information Processing Systems 36 (2024).
- [33] K. Z. Azar et al. 2020. NNgSAT: Neural network guided SAT attack on logic locked complex structures. In XInternational Conference on Computer-Aided Design. 1–9.
   [34] K. Z. Azar et al. 2022. Fuzz. penetration, and ai testing for soc security verification:
- [34] K. Z. Azar et al. 2022. Fuzz, penetration, and ai testing for soc security verification: Challenges and solutions. Cryptology ePrint Archive 2022, 394 (2022), 1–22.
- [35] Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. arXiv preprint arXiv:2101.00190 (2021).
- [36] M. Chen, et al. 2021. Evaluating large language models trained on code. arXiv preprint arXiv:2107.03374 (2021).
- [37] R. Kande et al. 2024. (Security) Assertions by Large Language Models. IEEE Transactions on Information Forensics and Security (2024).
- [38] S. Gulwani et al. 2017. Program synthesis. Foundations and Trends in Programming Languages 4, 1-2 (2017), 1–119.
- [39] S. Ren et al. 2020. Codebleu: a method for automatic evaluation of code synthesis. arXiv preprint arXiv:2009.10297 (2020).
- [40] S. Shi et al. 2023. Sechls: Enabling security awareness in high-level synthesis. In Asia and South Pacific Design Automation Conference. 585–590.
- [41] S. Thakur et al. 2023. Verigen: A large language model for verilog code generation. ACM Transactions on Design Automation of Electronic Systems (2023).
- [42] X. Li et al. 2024. Circuit Transformer: End-to-end Circuit Design by Predicting the Next Gate. arXiv preprint arXiv:2403.13838 (2024).

<sup>&</sup>lt;sup>7</sup>For instance, the Prefix-Tuning concept [35] involves the addition of trainable tokens to prompts, thus enabling more task-specific model responses.

<sup>&</sup>lt;sup>8</sup>Auto prompting could significantly mitigate the automation challenge and enhance the feasibility of (secure) code (RTL) generation [6].