

PropertyGPT: LLM-driven Formal Verification of Smart Contracts through Retrieval-Augmented Property Generation

Ye Liu

Nanyang Technological University
Singapore, Singapore
li0003ye@ntu.edu.sg

Yue Xue*

MetaTrust Labs
Singapore, Singapore
nerbonic@gmail.com

Daoyuan Wu[†]

The Hong Kong University of Science
and Technology
Hong Kong SAR, China
daoyuan@cse.ust.hk

Yuqiang Sun

Nanyang Technological University
Singapore, Singapore
suny0056@e.ntu.edu.sg

Yi Li

Nanyang Technological University
Singapore, Singapore
yi_li@ntu.edu.sg

Miaolei Shi

MetaTrust Labs
Singapore, Singapore
stan@metatrust.io

Yang Liu

Nanyang Technological University
Singapore, Singapore
yangliu@ntu.edu.sg

ABSTRACT

Formal verification is a technique that can prove the correctness of a system with respect to a certain specification or property. It is especially valuable for security-sensitive smart contracts that manage billions in cryptocurrency assets. Although existing research has developed various static provers for smart contracts, a key missing component is the *automated* generation of *comprehensive* properties, including invariants, pre-/post-conditions, and rules. Hence, industry-leading players like Certora have to rely on their own or crowdsourced experts to manually write properties case by case.

With recent advances in large language models (LLMs), this paper explores the potential of leveraging state-of-the-art LLMs, such as GPT-4, to transfer existing human-written properties (e.g., those from Certora auditing reports) and automatically generate customized properties for unknown code. To this end, we embed existing properties into a vector database and retrieve a reference property for LLM-based in-context learning to generate a new property for a given code. While this basic process is relatively straightforward, ensuring that the generated properties are (i) *compilable*, (ii) *appropriate*, and (iii) *runtime-verifiable* presents challenges. To address (i), we use the compilation and static analysis feedback as an external oracle to guide LLMs in iteratively revising the generated properties. For (ii), we consider multiple dimensions of similarity to rank the properties and employ a weighted algorithm to identify the top-K properties as the final result. For (iii), we design a dedicated prover to formally verify the correctness of the generated properties. We have implemented these strategies into a novel system called PropertyGPT, with 623 human-written properties collected from 23 Certora projects. Our experiments show that PropertyGPT can generate comprehensive and high-quality properties, achieving an 80% recall compared to the ground truth. It successfully detected

26 CVEs/attack incidents out of 37 tested and also uncovered 12 zero-day vulnerabilities, resulting in \$8,256 bug bounty rewards.

1 INTRODUCTION

Smart contracts are transaction-driven programs deployed and executed on blockchain platforms, automating the execution of digital agreements among users. Most smart contracts are written in Turing-complete programming languages, such as Solidity [45], and have been widely adopted on popular blockchain platforms like Ethereum [58] and BSC [1]. Smart contracts are extensively used in decentralized applications such as DeFi [24] and NFTs [5]. However, they are susceptible to various types of attacks, including integer overflow [49], re-entrancy [41], front-running [48], and access control vulnerabilities [15, 32]. These vulnerabilities primarily arise from loopholes in smart contracts due to programming errors, incorrect implementations, and logical bugs [62].

Formal verification is one of the most advanced approaches to identify contract loopholes by performing a comprehensive examination with different kinds of specifications. To perform formal verification, it is necessary to generate customized formal specifications for different smart contracts. Formal specifications for smart contracts usually include temporal logic properties and Hoare logic properties, as surveyed in [51]. *Invariants* are the most common contract specification, stating a property that holds for any contract execution, followed by *function pre-/post-conditions* for particular functional usage, as well as *rules* that cover cross-function properties. In most cases, temporal logic properties can be transformed into Hoare properties that could be instrumented into smart contract code [40]. Hence, existing works typically use only Hoare-style specifications for vulnerability detection [53, 55], inconsistency detection [11], and correctness validation [40, 57].

*Ye Liu and Yue Xue are the co-first authors.

[†]Daoyuan Wu is the corresponding author.

Despite the promise of formal verification in enhancing the security and reliability of smart contracts, one notable challenge remains: the community still lacks the automated generation of comprehensive properties for effective formal verification of smart contracts. While several works have attempted this, they have not yet achieved the ultimate goal of automatically generating necessary properties, including invariants, pre-/post-conditions, and rules, for an unknown contract code. For example, InvCon [31, 34] can dynamically infer *likely* contract invariants and function pre-/post-conditions, but it requires historical transaction information. Likewise, Cider [30] and SmartInv [56] employed a machine-learning-based approach to generate specifications through the training-and-inference paradigm, but only for the *invariant* properties. As a result, industry-leading players like Certora [10] have to rely on their own or crowdsourced experts [9] to manually write properties case by case, which hinders the effective formal verification of smart contracts on a large scale.

In this paper, we aim to explore how recent advances in large language models (LLMs), such as the state-of-the-art GPT-4 [4], could enable a new approach to the automated generation of smart contract properties. Given LLMs' strong capability for in-context learning (see background in §2), we try to achieve effective transfer learning from existing human-written properties to customized properties for unknown code. More specifically, we embed existing properties into a vector database and retrieve a reference property for LLM-based in-context learning to generate a new property for a given code. In this way, we can generate diverse types of properties as long as there are existing samples for each type in the collected vector database. Moreover, compared to the training-and-inference paradigm mentioned above [30, 56], our approach does not require the error-prone labeling process (we can directly use existing raw property results, such as those from Certora auditing reports), nor the "re-training" efforts when there is updated data.

While the basic property generation process is relatively straightforward, it is challenging to ensure that the retrieval-augmented properties are (i) *compilable*, (ii) *appropriate*, and (iii) *runtime-verifiable*. To address these challenges, we employ three novel designs and implement them in an LLM-driven system called PropertyGPT. *First*, we use compilation and static analysis feedback as an external oracle to guide LLMs in iteratively revising the generated properties. *Second*, we consider multiple dimensions of similarity to rank the properties and find a balanced metric for all these dimensions. The resulting weighted algorithm thus identifies the top-K properties as the final result. *Third*, we design a dedicated prover to formally verify the correctness of the generated top-K properties.

To evaluate PropertyGPT, we collected 623 human-written properties from 23 Certora projects. We first split 90 of them as a ground-truth testing set and used the rest as reference properties. We found that PropertyGPT can cover 80% equivalent properties in the ground truth as judged by human experts, with a reasonable precision of 64%. Note that the additional properties (FPs) produced by PropertyGPT generally also hold, complementing the human-written ones. We further used all 623 properties as a knowledge base to supply PropertyGPT for detecting real-world CVEs and past attack incidents. Our results demonstrated that PropertyGPT successfully detected 9 out of 13 CVEs and 17 out of 24 attack incidents. Furthermore, we ran PropertyGPT on four real-world bounty projects to

demonstrate its ability to find zero-day bugs. PropertyGPT successfully generated 22 bug findings, out of which 12 have been both confirmed and fixed, earning us a total of \$8,256 in bounty rewards.

We summarize the following main contributions in this paper:

- We proposed an end-to-end LLM-driven formal verification pipeline, PropertyGPT, for smart contracts, with the major retrieval-augmented property generation described in §5.
- To facilitate PropertyGPT, we also designed a property specification language (PSL) for smart contracts (§4) and a dedicated prover for property verification (§6).
- We conducted extensive experiments and ablation studies to evaluate PropertyGPT in real-world settings; see §7 and §8.

2 PRELIMINARY

Large language models (LLMs), such as GPT-3.5 [38] and CodeLlama [42], have been widely used in many natural language processing tasks, such as text generation, translation, and summarization. GPT series models are trained on a large corpus of text data and have the potential to generate human-like text, while CodeLlama is a fine-tuned version of Llama 2 [52] on open-source code. The LLMs are pre-trained on a large corpus of text data and then fine-tuned on specific tasks and these datasets usually contain code from different programming languages. Additionally, the pre-trained LLMs have exercised its potential to revolutionize the traditional software tasks, e.g., code generation [8], repairing [39, 60], vulnerability detection [48].

In-context learning (ICL). Based on the pre-trained knowledge, LLMs could leverage existing human-written properties written with various specification languages. Yet, due to the limitations of the pre-training data and the efforts needed for training, LLMs may not be able to include the real-time information. To address this problem, in-context learning (ICL) mechanism have been proposed by offering LLMs with the ability to learn from the latest conversation or task context [44, 63]. In essence, in-context learning is a specialized kind of few-shot learning [7], basing itself on a few examples or a small amount of data to learn a new task.

Instead of using fine-tuning [56], in this work, we employ the in-context learning ability from the state-of-the-art GPT-4 model for retrieval-augmented property generation.

3 PROPERTYGPT OVERVIEW

In this section, we present the overall design of PropertyGPT, which leverages LLMs' ICL capability to transfer existing human-written properties and generate customized properties for formally verifying unknown code. At a high level, PropertyGPT takes a piece of subject smart contract code as input and ultimately produces its corresponding properties along with the verification results.

As illustrated in Figure 1, PropertyGPT consists of eight major steps: ① PropertyGPT first creates a vector database for reference properties by embedding their corresponding critical code. Note that the reference properties themselves will not be embedded because they are not the search key. ② Given a piece of subject code under testing (typically one function), PropertyGPT queries the vector database to ③ retrieve all similar code within the threshold and map each code to their original reference properties. ④ All the reference properties are then tested with the subject code one by one

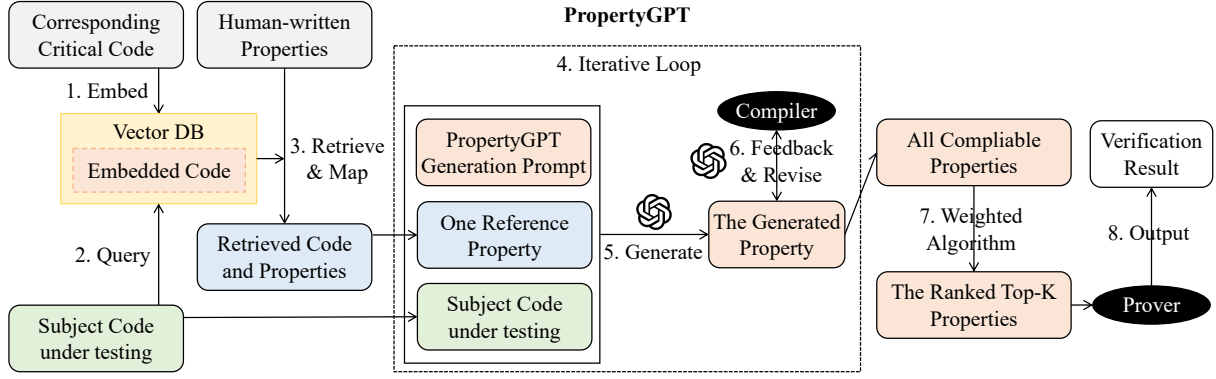


Figure 1: A high-level workflow of PropertyGPT.

$$\begin{aligned}
 v &\in \text{StateVar} \quad tmp \in \text{TemporalVar} \quad C \in \text{Constant} \\
 SC &= v*; func* \\
 func &\in \text{Function} = param*; stmt* \\
 expr &\in \text{Expression} = tmp \mid v \mid \text{old}(v) \mid param \mid C \mid expr \bowtie expr \\
 Spec(SC) &= inv*; \{p*\} func \{q*\}; rule* \\
 inv &\in \text{Invariant} = bool_expr \mid_{(v*, C*)} \\
 p &\in \text{Precondition} = bool_expr \mid_{(param, \text{old}(v), C)} \\
 q &\in \text{Postcondition} = bool_expr \mid_{(param, v, C)} \\
 rule &\in \text{Rule} = \text{assume}(expr) * \mid func(expr*) \mid \text{assert}(expr)*
 \end{aligned}$$

Figure 2: Property Specification Language (PSL).

in an iterative loop. ⑤ For each reference property, PropertyGPT employs a generation prompt to generate a candidate property for the subject code. ⑥ This candidate property is then checked by the compiler for grammar, and if it is not grammatically correct, it will be further revised according to the compiler’s feedback using a revising prompt. ⑦ Eventually, we obtain a list of compilable properties and rank them according to our weighted algorithm for the top-K appropriate properties. ⑧ These properties are finally formally verified by our prover, aiming to discover smart contract vulnerabilities.

To explore and understand the details of PropertyGPT, we first introduce its property specification language in §4. Following this, we describe the main process of LLM-based property generation and refinement in §5. Finally, we connect the generated properties with our dedicated formal prover in §6 for property verification.

4 PROPERTY SPECIFICATION LANGUAGE

To bridge the gap between property generation in §5 and formal verification in §6, we propose an intermediate language in this section to specify the properties of smart contracts.

Figure 2 illustrates our property specification language (PSL), which extends the popular smart contract programming language Solidity. In Solidity, a smart contract (SC) consists of a group of state variables recording persistent program state and a list of public functions allowing user interactions. The symbol \bowtie represents a set

of arithmetic, comparison, or logical operators, namely $\{+, -, /, >, <, =, ! =, >=, <=, \&, \mid\}$. PSL includes three kinds of properties with respect to different purposes. Invariants are properties that always hold true during contract execution and are defined over state variables; function pre-/post-conditions are properties that can be expressed in Hoare triples $\{p*\} func \{q*\}$, checking whether parameters and the modification of state variables satisfy functionality requirements. Scenario-based properties, defined on restricted environments, can be implemented as different rules by enforcing varied assumptions and customized assertions.

The proposed PSL for smart contracts offers several advantages over existing works [10, 57]. Firstly, it is easy to write contract specifications expressed in PSL since it shares most logic structures with contract code, without any tricky logic bindings from specifications to code. For example, Cetora verification language (CVL)¹ introduces many low-level primitives to map storage read/write, such as *hook*. Additionally, each function in CVL has to explicitly insert an additional context variable called *env*, and each function must also be declared state modification free or not. This non-trivial process may jeopardize specification integrity, thus being error-prone. Secondly, PSL specifications can be easily reused and upgraded. This is because PSL is a customized variant of Solidity, able to extend itself by including other contract specifications through inheritance or modularization management. More importantly, PSL specifications are easier to automate directly from contract code, along with efforts to ensure specification quality.

5 PROPERTY GENERATION & REFINEMENT

With the targeted PSL introduced in §4, our objective is to automatically generate properties written in PSL for the given code. The generated properties are the result of LLM-based transfer learning from existing human-written properties, which can be written in any specification language, not limited to PSL, such as CVL.

PropertyGPT can achieve such powerful transfer learning fundamentally rooted in LLMs’ capability for in-context learning (see §2). Nevertheless, we need to design a novel pipeline to facilitate this.

¹<https://docs.certora.com/en/latest/docs/cvl/index.html>

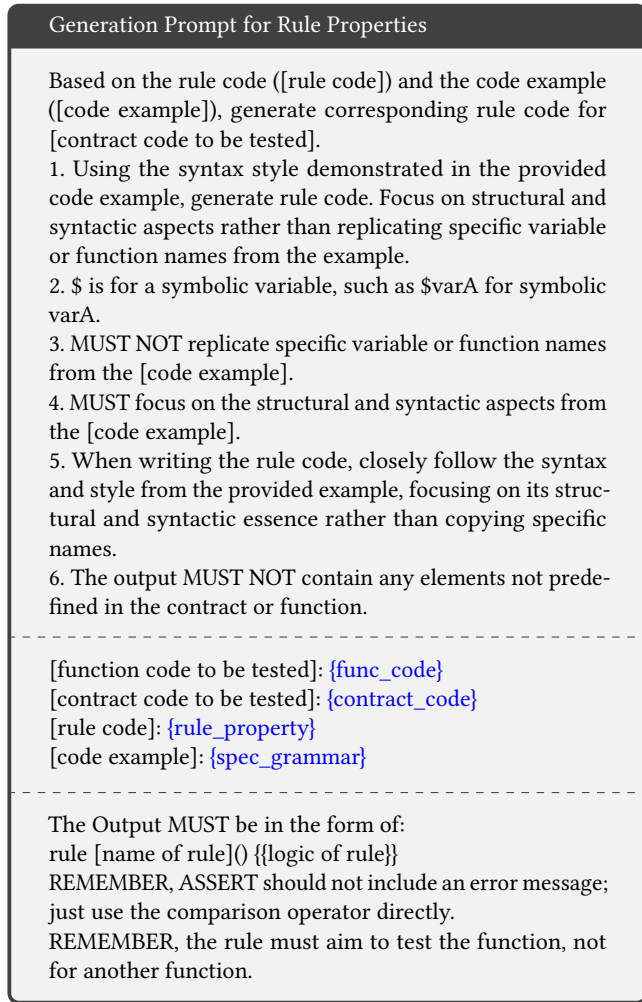


Figure 3: The prompt for generating rule properties.

Our idea is to mimic the RAG (retrieval-augmented generation) process in the NLP [27] or code [47] domain, using the reference properties retrieved to augment the generation of new properties. As previously illustrated in Figure 1, we first detail how such retrieval-augmented property generation is conducted in PropertyGPT in §5.1. After that, PropertyGPT iteratively revises the LLM-generated properties to fix their compilation errors in §5.2. Furthermore, we design a weighted algorithm to help PropertyGPT rank all compilable properties and obtain only the top-K appropriate properties for the prover’s verification in §6.

5.1 Retrieval-Augmented Property Generation

Here we focus on the basic retrieval-augmented property generation that occurs from step ② to step ⑤ in Figure 1. But the entire property generation process also includes property ⑥ revising and ⑦ ranking, which will be introduced in §5.2 and §5.3, respectively.

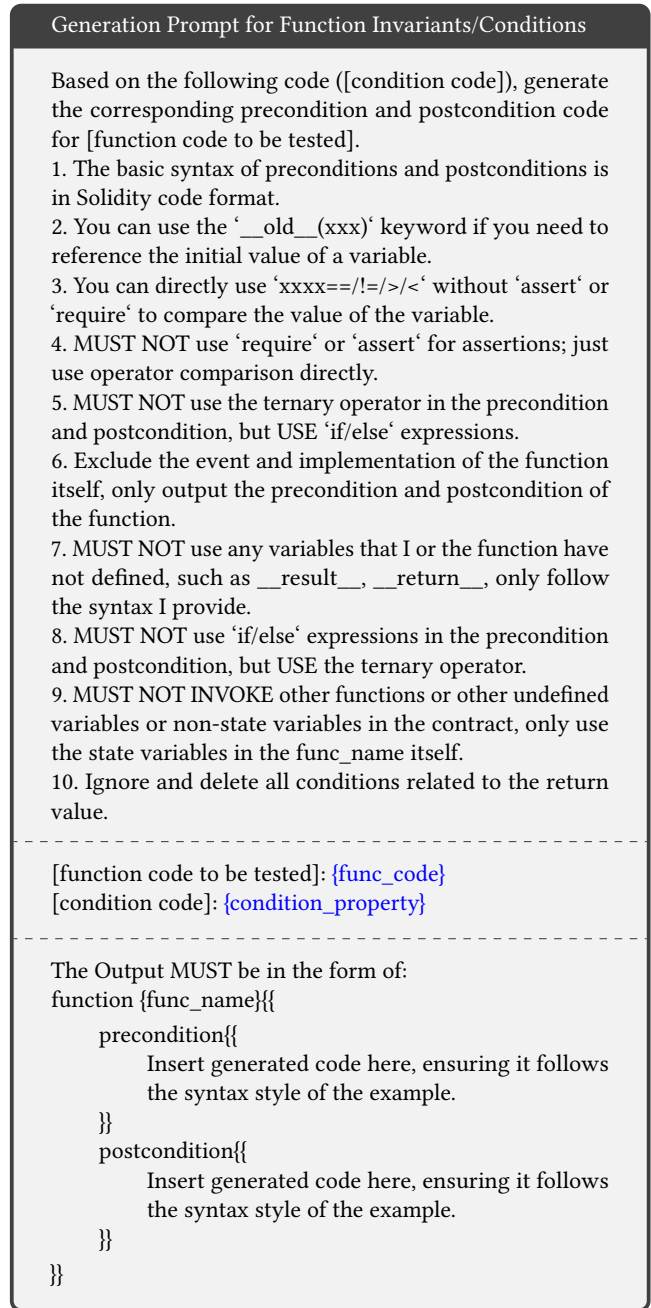


Figure 4: The prompt for generating invariants/conditions.

① **Knowledge Preprocessing.** One critical step in RAG-based systems is to first build a knowledge base, typically a vector database [47]. In PropertyGPT’s scenario, we do not aim to extract “knowledge” from the existing human-written properties; instead, we use them as reference properties for LLMs’ in-context learning. Therefore, we directly use the raw information from human-written reference properties to construct our vector database. As shown

in Figure 1, we embed the corresponding critical code of existing properties to build the search key used by RAG. Note that the reference properties themselves will not be embedded because they cannot be queried against by the subject test code.

②-③ **Similar Example Retrieval.** With the vector database, we can retrieve similar reference properties given the subject code to enable one-shot LLM learning in subsequent steps. To do this, the subject code is also embedded in step ② and the dot product is calculated [23] with all the vectors in the database. The top similar code with the highest dot products are then retrieved, and their corresponding properties are returned as the result in step ③. Here, we use a conservative code similarity threshold (e.g., 0.8) to limit the number of retrieved reference properties (typically 10 to 20 properties), which is acceptable because PropertyGPT eventually uses a weighted algorithm in §5.3 to rank only the top-K generated properties as the final result.

④-⑤ **In-context Learning.** All the reference properties are then tested with the subject code one by one in an iterative loop. For each reference property, PropertyGPT employs a generation prompt to generate a candidate property for the subject code, using the reference property as a one-shot example. Specifically, there are two types of generation prompts. One is used to generate global, cross-function rule properties, as shown in Figure 3, and the other is used to generate function-level pre-/post-conditions, as illustrated in Figure 4. Note that here we omit the generation prompt template for contract-level invariants because they are usually in a simpler form and equivalent to the one-for-all pre-/post-conditions for every public contract function. Both prompt templates consist of three parts: the first part details the generation instructions, the second part lists the code and reference property, and the third part defines the output property format. In particular, we supply a rule example in Figure 3 to help LLMs understand the grammar of our rule properties, while the grammar of invariants/conditions is directly specified using natural language instructions in Figure 4 as it is relatively simple. Additionally, since rule properties are cross-function, we provide not only the function code but also the entire contract code in the prompt template shown in Figure 3.

To determine which generation prompt should be used, PropertyGPT leverages the type of the retrieved reference property. If the reference property is a rule, PropertyGPT uses the first type of prompt template to generate the property. Otherwise, if the reference property is classified as pre-/post-conditions, PropertyGPT uses the second type of prompt template.

5.2 Revising Property to Fix Compilation Errors

While the basic property generation process is relatively straightforward, one particular challenge is how to guarantee that the generated property is *compilable*. To address this challenge, we leverage the feedback from the compiler and static checking to iteratively revise the property until it is compilable or until it reaches the maximum number of attempts, as shown in step ⑥.

Leveraging Compiler Feedback. Our PSL compiler can provide compilation error information, including detailed error locations and reasons, if the property cannot be successfully compiled. PropertyGPT thus leverages this feedback to instruct LLMs to iteratively revise the property. Specifically, we design and employ a common

Common Prompt for Revising (Rule) Properties

Here is the rule I provided: {spec_res}.
When this code is compiled with a solc-like program, an error occurs: {error_info}.

Your task is to understand the rule I provided, fix the rule code, and correct the error within the rule. Refer to the contract code provided above.

Note, only modify the rule code; do not add other code. If the error is due to a non-existent variable, find feasible methods to reimplement it, or if it is not implementable, delete this line.

Here is the function code to be tested: {func_code}

Here is the contract code to be tested: {contract_code}

Provide me with the repaired rule code. The revised rule code must not be the same as the old rule code.

1. Using the syntax style demonstrated in the provided code example, generate a rule code. Focus on the structural and syntactic aspects rather than replicating specific variable or function names from the example.
2. \$ is for a symbolic variable, such as \$varA which symbolizes varA.

Rule Code Output MUST be in the form of:

```
rule [name of rule](){{logic of rule}}
```

REMEMBER, ASSERT does not include an error message, just use the operator comparison directly.

REMEMBER, the rule must aim to test the function [{function_name}], not for other functions.

Figure 5: The common prompt for revising (rule) properties.

prompt template as shown in Figure 5 for revising rule properties; note that the prompt for revising function pre-/post-conditions is similar and is therefore not shown here. In this prompt, we ask LLMs to first understand the generated property code, identify and fix errors, maintain stylistic consistency throughout the process, and finally ensure that the revised rule code meets specific formatting requirements. We set a threshold for the maximum number of attempts to avoid an endless loop. In our experiment, as shown in §8.3, we found that 74% of properties could be successfully compiled with no revisions (63%) or with only one attempt, and 84% of all properties could be successfully revised within five attempts. This makes the iterative process manageable.

Employing Static Checking. However, we found that even when the compiler does not report any errors, it does not mean that the generated property is fully correct. One notable issue is that LLM-generated properties could fail to include the target subject function, which renders the property meaningless. To address this, we perform additional static checking for all compilable properties passed in the above step. If PropertyGPT identifies that the property is missing the target subject function, it employs a special prompt template as shown in Figure 6 (listing only the scenario for

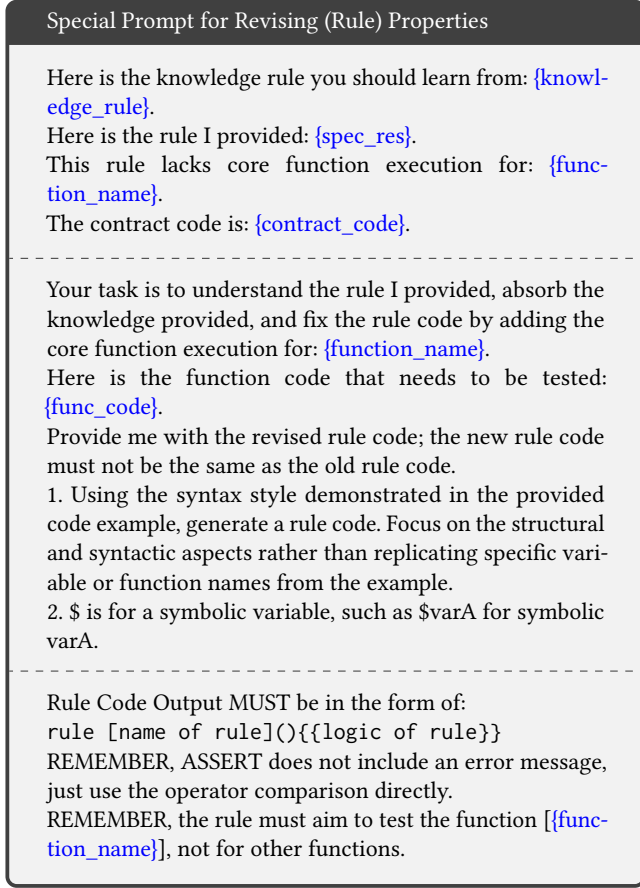


Figure 6: The special prompt for revising (rule) properties.

rule properties, similar to Figure 5). It is generally similar to the common revising prompt but explicitly addresses the rule's lack of testing for the core function execution. In this way, we not only guarantee that the generated property is grammatically correct but also functionally meaningful.

5.3 Ranking the Top-K Appropriate Properties

Another challenge is how to select the appropriate properties from all compilable properties as the final generation result. To do so, we propose a weighted algorithm to rank all the resulting properties, as shown in step ⑦. Specifically, we rank the properties based on the following four embedding-based metrics:

$X_{raw}(f, g)$: Similarity between contract code f and g .

$X_{summary}(f, g)$: Similarity between high-level functionality summaries of code f and g .

$Y_{raw}(\phi_1, \phi_2)$: Similarity between raw properties ϕ_1 and ϕ_2 .

$Y_{summary}(\phi_1, \phi_2)$: Similarity between high-level property summaries for ϕ_1 and ϕ_2 .

Note that we introduce $X_{summary}$ and $Y_{summary}$ to cope with that variety could exist for same-functionality code or same-semantic




Figure 7: Workflow of Property Verification.

properties, where high-level natural language summaries are made by large language model for given code or properties.

Given an unknown code f , let ϕ_1 be its property generated corresponding to reference code g having property ϕ_2 , we score ϕ_1 using a weighted algorithm as below.

$$\begin{aligned} \text{Score}(f, \phi_1) = & \alpha * X_{raw}(f, g) + \beta * X_{summary}(f, g) \\ & + \gamma * Y_{raw}(\phi_1, \phi_2) + \eta * Y_{summary}(\phi_1, \phi_2) \end{aligned} \quad (1)$$

where $\alpha, \beta, \gamma, \eta$ are hyperparameters and $\alpha + \beta + \gamma + \eta = 1$.

To tune these hyperparameters, we train a linear regression model by approximating actual property score $\text{Score}(f, \phi_1)$. In this work, for simplicity, we consider $\text{Score}(f, \phi_1) = Y_{summary}(\phi_1, \hat{\phi}_1)$, where $\hat{\phi}_1$ is the corresponding ground truth property of ϕ_1 . We have conducted a primitive experiment on 3622 properties generated by PropertyGPT, and the results show α : 0.134, β : 0.556, γ : 0.141, and η : 0.168 are the optimal weights.

Consequently, properties with different score are ranked in descending order, where we believe, properties with higher rank are likely to be the important security properties for prover to verify.

6 PROPERTY VERIFICATION

The properties generated by PropertyGPT are not only *compilable* and *appropriate* but also *verifiable*. Figure 7 illustrates the workflow of our property verification process. Our prover accepts smart contracts written in Solidity, along with their corresponding PSL specifications. We employ forward symbolic execution to conduct strongest postcondition analysis for each contract statement. Subsequently, we perform modular verification to ascertain whether these formal specifications have been accurately implemented in the smart contract and can produce a proof if the properties hold. In cases where the properties are violated, we employ bounded model checking to validate whether the violated properties genuinely remain unfulfilled during contract execution. Upon encountering counter-examples, we can confidently conclude that the properties indeed fail to hold, suggesting the presence of vulnerabilities in the smart contracts, which necessitate further manual confirmation.

We utilize source-code level symbolic execution to conduct strongest postcondition analysis for Solidity smart contracts. Distinguishing ourselves from existing research [29], our novel symbolic execution approach implements comprehensive small-step semantics, enabling automated analysis for real-world, complex smart contracts. Although previous work [29] has made strides in symbolic execution for Solidity smart contracts, it lacks support for certain critical features commonly used in smart contracts, such as the

aggregated effect of intricate expressions and polymorphic handling during complex inheritance relationships. We address these limitations by meticulously adhering to the practices of the Solidity compiler, ensuring precise semantics of complicated expressions. For instance, expressions are evaluated from left to right as per the compiler's specification. Furthermore, our approach accurately resolves polymorphism during both the compilation and execution stages of smart contract. To mimic actual contract execution, our symbolic execution approach maintains a comprehensive list of function signatures and meticulously revisits contract inheritance chains to determine the exact function implementation for ambiguous calls, such as "super().call()", where "super" refers to an unknown parent contract.

We have also implemented abstraction techniques to handle unknown or non-linear operation semantics. Firstly, the behaviors of function calls to on-chain smart contracts remain unknown, and we assume that all such calls succeed but make their return data symbolic to accommodate any possible outcome. This approach is necessary because on-chain contracts may not be open source, and their inter-contract interaction could be overly complex, falling beyond the scope of our current research. Secondly, non-linear native functions, such as `sha3`, which computes the hash value of a string, are challenging to model precisely. To address this issue, we utilize uninterpreted functions [19] to capture their primary features, such as an injective function for `sha3`.

We integrate modular verification and bounded model checking to establish the correctness of contract specifications. Smart contract execution relies on the initial contract state, current transaction environment, and the called functions. During modular verification, we lift all state constraints by making all state variables symbolic, even if it includes unreachable contract states. Then, correctness can be safely ensured when specification properties hold accordingly. Otherwise, we employ bounded model checking to systematically explore all feasible states to find counterexamples regarding the property to be verified within a given bound. The properties violated and their counterexamples will be investigated manually to confirm the existence of vulnerabilities. Similar to the practice in existing works [29, 57], in our prover, we set the bound by default to be three, and we also cap the maximum loop iterations to be five in case of non-terminated execution.

7 IMPLEMENTATION AND SETUP

We implemented PropertyGPT in around 3K lines of Python code for LLM-based property generation, and around 38K lines of C++ code for grammar support and verification of PSL property specifications. Additionally, for applying symbolic execution to smart contracts written in various versions of Solidity, we developed a converter to map smart contracts written in Solidity versions 0.6.x and 0.7.x into abstract syntax trees compatible with the latest Solidity version 0.8.x, where we have systematically investigated their syntax and semantic differences. We use the Z3 solver, version 4.11.2, to discharge symbolic constraints for path feasibility checking and property satisfiability checking.

7.1 Property Knowledge Collection

To obtain high-quality human-written properties as the knowledge base for in-context learning, we systematically analyzed 61 audit reports from the Certora platform, for which experts have written property specifications to facilitate formal analysis of smart contracts. These audit reports were published from 2019 to 2023. Through further investigation, we removed 38 projects whose contract code and raw properties were not available, and eventually, we collected 23 Certora projects, including 623 human-written properties, which will be detailed in Table 7 in Appendix A.

To study the characteristics of these properties, we employed the affinity propagation clustering algorithm [18] from the `sklearn`² library to discern property categories, based on pairwise embedding similarity across the properties. Specifically, we performed some preliminary experiments and found that this setting `AffinityPropagation(damping=0.5, preference=-75, random_state=5)` could establish a good result for property clustering. Figure 8 illustrates the distribution where six clusters are labeled with different colors. However, it is clear that overlapping exists among clusters, especially for clusters #3 and #5.

Furthermore, we investigated all the clusters and have the following classifications of human-written properties as shown in Table 1. There are six property categories as follows:

- *DeFi*, involving the management of its essential protocol components including reserves, collateral, and liquidity pools;
- *Token*, which is the cornerstone of entire DeFi ecosystems, specifying standard behaviors such as token balance and critical operations like transfer and minting;
- *Arithmetic*, focusing on the correctness of numerical conversions and the consistency of asset splitting;
- *Usability*, examining the validity of operations containing timestamp-based constraints (Temporal Use) and operations with contract state-based constraints (State-dependent Use);
- *Governance*, which plays an important role in the management of decentralized applications, usually through a voting mechanism, concerning issues such as the transfer of voting power to a delegator and the double-voting problem;
- *Security*, checking for the presence of common vulnerability types including front running and overflow.

7.2 Experimental Setup

We use the large language model GPT-4-turbo provided by OpenAI through its API `gpt-4-0125-preview`. Regarding the model configuration, we adhere to the default settings where the temperature is 0.8, top-p is 1, frequency penalty and presence penalty are both 0, and the maximum response length is 2000. Moreover, we calculate all embedding similarities using the pre-trained model `text-embedding-ada-002` from OpenAI. For property generation, we cap revising attempts at nine to limit LLM usage for better economics. All experiments were conducted on a Docker with Ubuntu 20.04 OS, an Intel Core Xeon 2.2 GHz processor, and 2GB RAM.

Data Availability. The property dataset and raw experimental data are available at <https://sites.google.com/view/propertygpt>, while the prototype will be released once the paper is published.

²<https://scikit-learn.org/>




Figure 8: The property cluster distribution after two-dimension PCA [20] reduction.

Table 1: Characteristics of the collected Certora properties.

Category	Classification	Property Examples
DeFi	Reserve	setReserveFactorIntegrity
	Collateral	integrityOfisUsingAsCollateralAny
	Liquidity	checkBurnAllLiquidity
Token	Balance	total_supply_is_sum_of_balances
	Transfer/TransferFrom	transferBalanceIncreaseEffect
	Mint/Burn	integrityMint, additiveBurn
Arithmetic	Approve	approvedTokensAreTemporary
	Numerical Conversion	toElasticAndToBaseAreInverse1down
	Asset Splitting	moneyNotLostOrCreatedDuringSplit
Usability	Temporal Use	timestamp_constrains_fromBlock
	State-dependent Use	unsetPendingTransitionMethods
Governance	Delegation	integrityDelegationWithSig
	Voting	totalNonVotingGEAccountNonVoting
Security	Front run	cannotFrontRunSplitTwoSameUsers
	Overflow	integrityOfMulDivNoOverflow

8 EVALUATION

In this work, we aim to answer the following research questions:

- RQ1: (**Property Generation**) How accurately does PropertyGPT generate properties for smart contracts?
- RQ2: (**Vulnerability Detection**) How effectively does PropertyGPT discover smart contract vulnerabilities? Can PropertyGPT achieve state-of-the-art results?
- RQ3: (**Ablation Study**) What factors influence the performance of PropertyGPT?
- RQ4: (**Impact**) How well does PropertyGPT find zero-day vulnerabilities in real-world smart contract projects?

Methodology. To answer RQ1, we divide Certora properties into a testing dataset and a “training” dataset as the knowledge base. We instruct PropertyGPT to generate properties for smart contracts in the testing dataset using smart contracts and their properties from the knowledge base. We compare the resulting properties by PropertyGPT with ground-truth human-written ones to investigate its effectiveness. Specifically, we randomly selected nine (40%) Certora projects as our testing dataset and then picked 10 properties for

Table 2: The evaluation benchmarks.

Benchmark	Research questions
23 Certora audited projects (623 properties)	RQ1, RQ3
13 CVEs, curated SmartInv benchmark (24 projects)	RQ2

each project. Consequently, our testing dataset includes 90 ground truth properties from nine projects. During this experiment, PropertyGPT first extracts the subject function code where the ground truth properties are specified, and then queries the knowledge base to enable ICL to automate property generation.

To answer RQ2, we compare PropertyGPT with SmartInv [56], a concurrent work with ours to be published in May 2024. We contacted the authors to obtain a copy of their source code and benchmark³, which includes 60 attack incident projects that have suffered significant losses. Upon reviewing their benchmark, we identified several issues. Among the listed cases, 2 are repeated, 9 lack public exploit transactions (e.g., *sherlockYields*), 2 are not open-sourced, and 2 have incomplete code. Of the remaining cases, 11 are reentrancy attacks that could be easily remedied by adding the widely-used *nonReentrancy* modifier. Furthermore, 8 cases involve price manipulation attacks, which may be impractical to identify using the simple invariant properties that SmartInv generated. For example, Tolmach et al.[50] proposes a semi-automated formal composite analysis for DeFi protocols that detects such problems with fairness properties, while others use either runtime monitoring[59] to identify attack behavior or static analysis [25] to flag vulnerable code with predefined patterns. Through this deep analysis of their benchmark, we curated 24 attack incidents from the SmartInv benchmark for our evaluation. Additionally, we compare PropertyGPT with state-of-the-art tools [2, 16, 36, 48] on well-known smart contract CVEs. As of April 21, 2024, there are 577 smart contract CVEs, predominantly integer overflows. To avoid bias, we selected 13 representative CVEs: three are integer overflow cases, three involve access control vulnerabilities, four are other logic bugs, etc., details of which are provided in Table 4.

In summary, as shown in Table 2, we evaluate the property generation process using Certora audited projects and test the applicability of PropertyGPT in vulnerability detection using well-known CVEs and the attack incident projects studied by SmartInv.

8.1 RQ1: Property Generation

We evaluated PropertyGPT on 90 ground-truth properties from nine Certora projects to investigate the effectiveness of property generation. Table 3 shows PropertyGPT’s property generation results using the rest of Certora properties as the knowledge base.

The first two columns show the project name and the number of properties written by Certora experts, i.e., #Property (Certora). The middle five columns list the number of properties generated by PropertyGPT, i.e., #Property (ours), true positives that are equivalent to the ground-truth properties (TP), the number of ground-truth properties hit by the properties generated (#Hit), the number of missed ground-truth properties (FN), as well as false positives (FP). The last three columns are recall, precision,

³<https://github.com/sallywang147/attackDB>

Table 3: The property generation results for 90 ground-truth properties from nine Certora projects.

Project	#Property (Certora)	#Property (ours)	TP	#Hit	FN	FP	Recall	Precision	F1-score
aave_proof_of_reserve	3*	38	25	3	0	13	1.00	0.66	0.79
aave_v3	17	61	32	15	2	29	0.88	0.52	0.66
celo_governance	10	39	29	10	0	10	1.00	0.74	0.85
furucombo	10	23	11	7	3	12	0.70	0.48	0.57
openzeppelin	10	2	2	1	9	0	0.10	1.00	0.18
opyn_gamma_protocol	10	30	14	8	2	16	0.80	0.47	0.59
ousd	10	100	67	10	0	33	1.00	0.67	0.80
radicle_drips	10	17	9	7	3	8	0.70	0.53	0.60
sushi_benttoibox	10	70	49	10	0	21	1.00	0.70	0.82
Average	10	42	26	8	2	16	0.80	0.64	0.71

* This project contains only three human-written properties, so we picked seven more from aave_v3, both of which are developed by the same institution.

and F1-score metrics where $\text{recall} = \frac{\#Hit}{\#Hit + FN}$, $\text{precision} = \frac{TP}{TP + FP}$, and $\text{F1-score} = \frac{2 \times \text{recall} \times \text{precision}}{\text{recall} + \text{precision}}$. Because Certora properties are written in the proprietary Certora Verification Language (CVL) that supports formal verification of smart contracts at the EVM bytecode level, while PropertyGPT uses PSL to facilitate property formulation and verification at the source code level, there is currently no automated analysis tool available for equivalence checking between properties from these two distinct specification systems. Therefore, two authors with 5 years of research and auditing experience independently examined the equivalence between the ground-truth properties by Certora and the properties generated by PropertyGPT, with a third author breaking ties in case of disagreement. We welcome other researchers to conduct replication and verification using our released data, available at <https://sites.google.com/view/propertygpt>.

Table 3 shows that PropertyGPT can generate comprehensive properties with relatively high recall and reasonable precision. Most properties generated (26/42) are true positives, and most ground truth properties (8/10) can be successfully reproduced, achieving a satisfactory recall (0.80), reasonable precision (0.64), and fairly good F1-score (0.71). Delving into project-specific results, PropertyGPT was able to reproduce all the ground-truth properties for four projects including aave_proof_of_reserve, celo_governance, ousd, and sushi_benttoibox. In contrast, PropertyGPT reproduced only two ground-truth properties for openzeppelin, suffering the lowest recall and F1-score, although with the highest precision. We investigated the results and found this is largely because OpenZeppelin [3] is a foundational contract library that has been directly imported by nearly all real-world applications, and client code is unlikely to re-implement similar functionality, thus leading to the scarcity of reliable reference properties. In terms of precision, opyn_gamma_protocol achieves the lowest, reaching only 0.47. We investigated all 16 false positives about it and later recognized that 11 of these false positives are properties that hold for smart contracts but are not documented in the ground truth by Certora.

Answer to RQ1: PropertyGPT can generate comprehensive and high-quality properties, covering 80% equivalent properties in the ground truth as judged by human experts. Moreover, the additional properties (FPs) produced by PropertyGPT generally also hold, complementing the human-written ones.

8.2 RQ2: Vulnerability Detection

We investigate the applicability of PropertyGPT in the vulnerability detection task on well-known smart contract CVEs and the attack incident projects studied by SmartInv. Note that to mimic the situation of real-world deployment, we set the top-K to top-2, as measured by §8.3, as the best configuration starting from this section.

CVEs. Table 4 demonstrates PropertyGPT’s effectiveness in detecting 13 smart contract CVEs. We compared PropertyGPT with GPTScan [48], which employs the variable recognition ability of LLMs to instantiate high-level detection patterns for logic bugs, and Slither [16], a popular static analysis tool used to detect a wide range of common vulnerability types. In particular, since the original GPTScan covers only ten types of logic bugs, we have enhanced it with the recent unsupervised paradigm [47] and refer to the enhanced version as GPTScan+. Additionally, Manticore [36] and Mythril [2] are two bytecode-level symbolic execution tools that automate comprehensive program state exploration and exploit generation of smart contract vulnerabilities. In Table 4, the first two columns list CVE names and their vulnerability types, while the remaining columns show the detection results by each tool.

The detection results presented in Table 4 illustrate that PropertyGPT outperforms all the comparison tools by detecting 9 out of 13 CVEs, followed by GPTScan detecting five CVEs, Slither detecting only one delegatecall-related CVE, Mythril detecting three overflow-related CVEs, and Manticore detecting zero CVEs. We also investigated the remaining four CVEs that PropertyGPT failed to detect. It is unknown what valid properties can express the expectation of proper randomness and delegatecall use. CVE-2018-17111 is caused by the misuse of access control rather than the lack of access control, which is quite challenging for PropertyGPT to recognize this subtle difference during property generation.

Table 4: Vulnerability detection results for 13 CVEs.

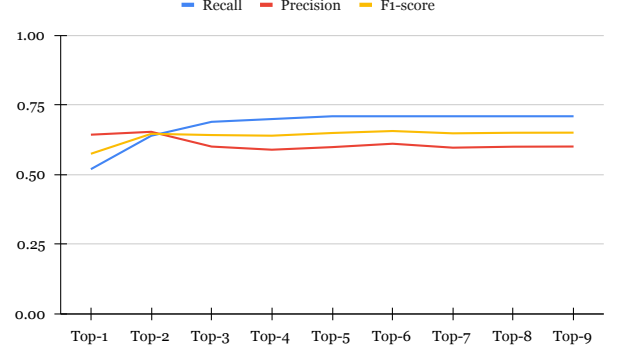
CVE	Description	PropertyGPT	GPTScan+	Slither	Manticore	Mythril
CVE-2021-34273	access control	✓	✓	×	×	×
CVE-2021-33403	overflow	✓	×	×	×	✓
CVE-2018-18425	logic error	✓	×	×	×	×
CVE-2021-3004	logic error	×	×	×	×	×
CVE-2018-14085	delegatecall	×	×	✓	×	×
CVE-2018-14089	logic error	✓	✓	×	×	×
CVE-2018-17111	access control	×	×	×	×	×
CVE-2018-17987	bad randomness	×	✓	×	×	×
CVE-2019-15079	access control	✓	×	×	×	×
CVE-2023-26488	logic error	✓	×	×	×	×
CVE-2021-34272	access control	✓	✓	×	×	×
CVE-2021-34270	overflow	✓	✓	×	×	✓
CVE-2018-14087	overflow	✓	×	×	×	✓

Table 5: Evaluation results for 24 attack incident projects from the curated SmartInv benchmark.

Contracts	Loss (million \$)	Detection	#Property	Generation (seconds)	Verification (seconds)
dfxFinance	4.00	✓	8	235	7
AnySwap	1.40	×	11	518	7
Dodo	0.70	✓	17	1,182	19
Bancor	0.55	✓	19	1,948	9
BeautyChain	900.00	✓	5	104	9
Melo	0.09	✓	9	252	8
BGLD	0.02	×	9	229	39
GYMNetwork	2.00	✓	21	274	71
elasticSwap	0.85	✓	37	1,136	120
EulerFinance	200.00	×	23	376	43
monoSwap	31.00	✓	5	69	12
nimBus	0.01	✓	32	4,288	30
VTF	0.05	×	8	358	21
Nomad	152.00	✓	14	590	70
Umbrella	0.70	✓	14	404	25
Fortress Loan	3.00	✓	2	71	5
ShadowFinance	0.30	✓	25	551	80
Revest	-	✓	4	75	10
Cartel	-	✓	11	401	20
sushiSwap	3.30	×	10	419	20
ChainSwap	8.00	×	9	307	25
Ragnarok	0.04	✓	42	1,890	88
templeDao	2.30	✓	13	302	30
BabySwap	-	×	33	1,842	50
Overall	1,310.30	17	16	743	34

The ability of PropertyGPT can be enhanced by the introduction of newly confirmed vulnerable code and properties into our knowledge database. As shown in Table 1, the studied properties written by Certora experts seem to lack support for access control, which could limit the effectiveness of PropertyGPT in detecting other wild access control vulnerabilities, even though we realized that PropertyGPT has demonstrated a certain level of generalization capability in the aforementioned CVE detection results.

Attack Incidents. Table 5 shows the evaluation results on 24 attack incident projects from the curated SmartInv benchmark. Because the authors of SmartInv did not share their instrumented buggy contract code, and the ground truth and raw experimental results

**Figure 9: The impact of Top-K settings on property accuracy.**

about their generated invariant properties are also missing, we perform a qualitative rather than a quantitative comparison with SmartInv and will discuss this in Section 9. In Table 5, the first two columns list project names and the amount of attack loss. The remaining columns show the detection results, the number of properties generated, the time used for property generation, and formal verification, respectively.

PropertyGPT successfully identified vulnerabilities in 17 out of 24 real-world attack incidents, on average generating 16 properties per project, spending around 12 minutes for property generation and only 34 seconds for formal verification. For the remaining seven projects that PropertyGPT failed to detect, we studied the root causes of their reported vulnerabilities. We recognized that PropertyGPT does not support the runtime context of smart contracts, which may be essential for generating properties in particular use scenarios, for example, deflationary token abuse for the BGLD project, which we leave as future work.

Answer to RQ2: PropertyGPT can effectively detect vulnerabilities in both simple and complex smart contracts. Specifically, PropertyGPT has distinguished itself from the current state-of-the-art by detecting 9 out of 13 CVEs. Additionally, PropertyGPT achieved relatively good results in identifying logic bugs in 17 out of 24 attack incidents.

8.3 RQ3: Ablation Study

In our ablation study, we first systematically explored the impact of varying Top-K settings on the property selection process. Conducting different trials on the same Certora projects in RQ1, in Figure 9, we plotted recall, precision, and F1-score for the resulting properties accordingly. It is clear that all the metrics are above 0.5. More importantly, when moving from top-1 to top-2, all the metrics increase, although precision has only a very slight increase from top-1 (0.64) to top-2 (0.65), and afterward, precision goes down and finally fluctuates around 0.60. Therefore, for the sake of efficiency, we suggest selecting top-2 properties generated for use by external community experts or other compatible program analysis tools.

In addition, we delved into our property generation process with a focus on the success rate of property generation and property repair times using compiler feedback information. Table 6 shows that GPT-4.0-turbo without revising or repair achieves a 63% success rate, which is quite lower compared to PropertyGPT (87%).

Table 6: The success rate of property generation.

Method	#Compilable	#Failed	Success Rate
GPT-4.0-turbo w/o fix	234	136	0.63
PropertyGPT w/ §5.2	321	49	0.87


**Figure 10: The distribution of property fix times.**

Figure 10 visualizes the distribution of property repair times, where we capped repair time at nine. We can see that most compilable properties (84%) can be generated with no more than five fix attempts. We also investigated the remaining 49 properties that could not be fixed by PropertyGPT and discovered the main compiler error message to be the use of undeclared variables, which may be addressed with a pattern-based approach [26].

Answer to RQ3: PropertyGPT can effectively generate compilable properties, with 84% of all properties being successfully revised within five attempts, and the highest success rate reaching 87%. Among them, the top-2 properties achieve the best balance between precision and recall.

8.4 RQ4: Real-world Impact

To demonstrate PropertyGPT’s ability to identify zero-day vulnerabilities, we ran PropertyGPT on real-world bounty projects hosted by popular platforms such as Secure3 [43] and Code4Rena [12]. PropertyGPT successfully generated 22 bug findings for 4 projects, 12 of which have been both confirmed and fixed. In return, we received \$8,256 in bug bounties from vendors. In this section, for case studies, we list two zero-day bugs that have been fixed for responsible disclosure, and we do not mention their project source to respect the anonymity policy.

Figure 11 shows that the `withdrawForwardFee` function contains a critical vulnerability allowing validators to withdraw more than their allocated share of forwarding fees, potentially leading to unfair distributions and loss of funds. The vulnerability arises because the function fails to track and limit individual validators’ withdrawals according to their proportionate share. It calculates the new total withdrawn fee by simply adding the requested withdrawal amount `_amount` to `totalValidatorForwardFeeWithdrawn` (Lines 10-13), without considering the requesting validator’s entitled share. The only check performed is against the total collected forwarding fees, ensuring that the new total withdrawn does not

```

1  pragma solidity ^0.8.0;
2  contract SimplifiedStandaloneZkLink {
3      address private _owner;
4      mapping(address => bool) private _validators;
5      uint256 public totalValidatorForwardFee;
6      uint256 public totalValidatorForwardFeeWithdrawn;
7
8      function withdrawForwardFee(uint256 _amount) external
9          ↪ nonReentrant onlyValidator {
10         require(_amount > 0, "Invalid amount");
11         uint256 newWithdrawnFee =
12             ↪ totalValidatorForwardFeeWithdrawn + _amount;
13         require(totalValidatorForwardFee >= newWithdrawnFee,
14             ↪ "Withdraw exceed");
15
16         totalValidatorForwardFeeWithdrawn = newWithdrawnFee;
17         (bool success, ) = msg.sender.call{value: _amount}("");
18         require(success, "Withdraw failed");
19         emit WithdrawForwardFee(_amount);
20     }
21 }

```

Figure 11: The vulnerable `withdrawForwardFee` function.

```

1  function withdrawForwardFee(uint256 _amount)
2  precondition {
3      _validators[msg.sender] == true;
4      _amount > 0;
5      old(totalValidatorForwardFee) >=
6          ↪ old(totalValidatorForwardFeeWithdrawn) + _amount;
7  }
8  postcondition {
9      totalValidatorForwardFeeWithdrawn ==
10         ↪ old(totalValidatorForwardFeeWithdrawn) + _amount;
11      totalValidatorForwardFee -
12         ↪ totalValidatorForwardFeeWithdrawn ==
13         ↪ old(totalValidatorForwardFee) -
14         ↪ old(totalValidatorForwardFeeWithdrawn) - amount;
15  }

```

Figure 12: The property generated for the case in Figure 11.

exceed this amount (Line 11). However, this does not prevent individual validators from withdrawing more than their share.

PropertyGPT detected this vulnerability through the generation and verification of function pre-/post-conditions listed in Figure 12. The pre-conditions (Lines 3-5) hold for this contract as they precisely capture the constraints `onlyModifier` and the other two `require` statements. The post-conditions describe the expected functionality. However, one of the post-conditions (Line 9), `totalValidatorForwardFee - totalValidatorForwardFeeWithdrawn == old(totalValidatorForwardFee) - old(totalValidatorForwardFeeWithdrawn) - amount`, does not hold, thus identifying the contract vulnerability in Figure 11.

Figure 13 shows a vulnerable `addEnvelope` function where it does not enforce uniqueness of envelope (Line 7), where existing storage can be overwritten arbitrarily. Figure 14 presents the property generated by PropertyGPT to detect such issue. Interestingly,

```

1 function addEnvelope(
2   string calldata envelopeID, bytes32 hashedMerkleRoot,
3   uint32 bitarraySize, address ERC721ContractAddress,
4   uint256[] calldata tokenIDs
5 ) public {
6   require(tokenIDs.length > 0, "Trying to create an empty
   ↪ envelope!");
7   MerkleEnvelopeERC721 storage envelope =
   ↪ idToEnvelopes[envelopeID]; // bug: overwrite storage.
8   envelope.creator = msg.sender;
9   envelope.unclaimedPasswords = hashedMerkleRoot;
10  envelope.isPasswordClaimed = new uint8[](bitarraySize / 8 +
   ↪ 1);
11  envelope.tokenAddress = ERC721ContractAddress;
12  envelope.tokenIDs = tokenIDs;
13  ...
14 }

```

Figure 13: The vulnerable addEnvelope function.

```

1 rule checkAddEnvelopeCorrectSenderAndCreator() {
2   assume(msg.sender ==
   ↪ 0x0000000000000000000000000000000000000000000000000000000000000001);
3   string memory envelopeID = "uniqueID";
4   bytes32 hashedMerkleRoot = 0x1234567890abcdef1234567890abc;
   ↪ def1234567890abcdef1234567890abcdef;
5   uint32 bitarraySize = 128;
6   address ERC721ContractAddress =
   ↪ 0x0000000000000000000000000000000000000000000000000000000000000002;
7   uint256[] memory tokenIDs = new uint256[](1);
8   tokenIDs[0] = 12345;
9
10  MerkleEnvelopeERC721 storage $envelopeBefore =
   ↪ idToEnvelopes[envelopeID];
11  bool $existsBefore = ($envelopeBefore.creator !=
   ↪ address(0));
12
13  addEnvelope(envelopeID, hashedMerkleRoot, bitarraySize,
   ↪ ERC721ContractAddress, tokenIDs);
14
15  MerkleEnvelopeERC721 storage $envelopeAfter =
   ↪ idToEnvelopes[envelopeID];
16  bool $correctlyAdded = ($envelopeAfter.creator ==
   ↪ msg.sender);
17  bool $notExistsBefore = ! $existsBefore;
18
19  assert($correctlyAdded && $notExistsBefore);
20 }

```

Figure 14: The property generated for the case in Figure 13.

PropertyGPT can skillfully construct varied input data (Line 3-8). When the function call (Line 13) succeeds, we check the condition (Line 19) that same-id envelope does not exist before and the envelope creator equals to the current caller. In other words, this condition disallows calling addEnvelope function with same envelope id and ensures the effect of addEnvelope will set envelope

creator to be the function caller. Due to the overwrite bug in Figure 13, this assertion does not hold for this function.

8.5 Threats to Validity

Internal Validity. We evaluated the effectiveness of PropertyGPT on an established Certora property dataset. Nevertheless, there is lack of equivalence checking tool between Certora-style properties and our proposed PSL-style properties generated by PropertyGPT. To mitigate this issue, three authors independently reviewed these properties to determine equivalence and we release all the properties generated and the according labeling results for public use.

External Validity. Our findings in vulnerability detection may not apply to other kinds of smart contracts and other types of contract vulnerabilities. In this work, we evaluated PropertyGPT on 13 representative smart contract CVEs covering many kinds of vulnerabilities and 24 real-world victim projects of different application domains. Moreover, we generated 24 bug findings for high-profile projects to audit and 12 have been confirmed and fixed, with \$8,256 bounty reward. Therefore, we believe PropertyGPT develops a practical formal verification technique for the detection of a wider range of smart contract vulnerabilities.

9 RELATED WORK

Vulnerability Detection. Numerous automated and semi-automated analysis tools have been proposed to detect smart contract vulnerabilities. On the one hand, static analysis tools analyze code sequences along the abstract syntax tree of contracts [16] or use a fact-based transformation and query system [6, 54] to flag weaknesses and vulnerabilities against expert-written patterns. In contrast, with test oracles, fuzzers examine runtime behaviors including operation traces [22, 37] and execution effects [33, 55] for exploit generation, usually leading to higher precision but lower recall compared with static analyses. On the other hand, formal verification has been widely employed in techniques to ensure smart contract correctness. Automated tools like Manticore [36] and Mythril [2] use symbolic execution to explore as many program states as possible to identify vulnerable behavior with a set of predefined detection rules. Semi-automated tools require users to provide specification properties, including invariants [53], function pre-/post-conditions [57], temporal properties [40, 46], and other customized rules [10, 21].

PropertyGPT distinguishes itself by automating property generation using a large language model and proposing a powerful prover based on source code-level symbolic execution of smart contracts, supporting the detection of a wide range of contract vulnerabilities.

Property Generation. Static inference [49, 57] and dynamic inference [31, 34] have been applied in property generation for smart contracts, and recently, machine-learning-based models have also been used for invariant property generation [30, 56]. VeriSol [57] applies the Houdini algorithm [17] to reason about correct invariant properties from a set of hypothesized candidates. SolType [49] discovers type invariants for Solidity smart contracts, requiring developers to add refinement type annotations to the contracts. However, their properties are limited to arithmetic operations to secure smart contracts from integer overflow and underflow. InvCon [31] and its subsequent work [34] apply dynamic invariant detection and static inference to produce contract invariants and

function pre-/post-conditions, but they necessitate contracts having sufficient transaction histories.

Our work aligns with previous efforts in machine-learning-based approaches, i.e., Cider [30] and SmartInv [56]. Cider uses a deep reinforcement learning approach but only generates *likely* invariant properties, while PropertyGPT can verify all the properties generated with a prover. Both SmartInv and PropertyGPT are powered by large language models. PropertyGPT differs from SmartInv in that we use in-context learning rather than fine-tuning, and our properties generated extend beyond function pre-/post-conditions. **LLM-based Security Systems.** By combining LLMs, various security tasks have been addressed more effectively. Sun *et al.* [48] proposed GPTScan, and Li *et al.* [28] introduced methods that combine LLMs with static program analysis for vulnerability detection, covering more types of vulnerabilities than traditional tools. Beyond vulnerability detection, LLMs have been used for other security tasks. Deng *et al.* [13] proposed TitanFuzz, which utilizes LLMs to guide the fuzzing of deep learning libraries such as PyTorch and TensorFlow. They also introduced FuzzGPT [14] to synthesize unusual programs for fuzzing vulnerabilities. ChatAFL [35] utilizes LLMs to guide the fuzzing of protocols by interpreting the protocol documents. Additionally, LLMs have been applied to program repairing tasks, such as ACFix [61] and ChatRepair [60].

10 CONCLUSION

In this paper, we proposed retrieval-augmented property generation for smart contracts by utilizing LLMs' in-context learning capabilities. We implemented this approach in a tool called PropertyGPT and addressed challenges to ensure the generated properties are compilable, appropriate, and runtime-verifiable. Our evaluation results indicate that PropertyGPT can detect many real-world contract vulnerabilities, especially in high-profile projects, collectively receiving \$8,256 in bounty rewards from vendors. For future work, we plan to include more comprehensive contract context information, such as documentation, in our approach and enhance PropertyGPT with richer property knowledge from various sources.

REFERENCES

- [1] 2020. Binance Smart Chain. <https://docs.binance.org/smart-chain/guides/bsc-intro.html>. Introduction of Binance Smart Chain.
- [2] 2024. Mythril. <https://github.com/Consensys/mythril>.
- [3] 2024. OpenZeppelin—A library for secure smart contract development. <https://github.com/OpenZeppelin/openzeppelin-contracts>.
- [4] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [5] Akash Arora, Kanisk, and Shailender Kumar. 2022. Smart contracts and NFTs: non-fungible tokens as a core component of blockchain to be used as collectibles. In *Cyber Security and Digital Forensics: Proceedings of ICCSDF 2021*. Springer, 401–422.
- [6] Lexi Brent, Neville Grech, Sifis Lagouvardos, Bernhard Scholz, and Yannis Smaragdakis. 2020. Ethainter: a smart contract security analyzer for composite vulnerabilities. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 454–469.
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, and Others. 2020. Language Models are Few-Shot Learners. *arXiv:2005.14165* (July 2020). <https://doi.org/10.48550/arXiv.2005.14165> arXiv:2005.14165 [cs].
- [8] Cristian Cadar, Daniel Dunbar, and Dawson Engler. 2008. KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation* (San Diego, California) (OSDI'08). USENIX Association, USA, 209–224.
- [9] Certora. [n.d.]. A community of hackers putting their formal verification skills to the test and earning rewards from leading protocols. <https://www.certora.com/contests>.
- [10] Certora. [n.d.]. Securing Web3 with Decentralized Intelligence. <https://www.certora.com/>.
- [11] Ting Chen, Yufei Zhang, Zihao Li, Xiapu Luo, Ting Wang, Rong Cao, Xiuzhuo Xiao, and Xiaosong Zhang. 2019. Tokenscope: Automatically detecting inconsistent behaviors of cryptocurrency tokens in Ethereum. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*. 1503–1520.
- [12] Code4rena. [n.d.]. Keeping high severity bugs out of production. <https://code4rena.com/>.
- [13] Yinlin Deng, Chunqiu Steven Xia, Haoran Peng, Chenyuan Yang, and Lingming Zhang. 2023. Large Language Models Are Zero-Shot Fuzzers: Fuzzing Deep-Learning Libraries via Large Language Models. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. ACM, Seattle WA USA, 423–435. <https://doi.org/10.1145/3597926.3598067>
- [14] Yinlin Deng, Chunqiu Steven Xia, Chenyuan Yang, Shizhuo Dylan Zhang, Shujing Yang, and Lingming Zhang. 2023. Large Language Models are Edge-Case Generators: Crafting Unusual Programs for Fuzzing Deep Learning Libraries. *IEEE Computer Society*, 830–842.
- [15] Yuzhou Fang, Daoyuan Wu, Xiao Yi, Shuai Wang, Yufan Chen, Mengjie Chen, Yang Liu, and Lingxiao Jiang. 2023. Beyond "Protected" and "Private": An Empirical Security Analysis of Custom Function Modifiers in Smart Contracts. In *Proc. ACM ISSTA*.
- [16] Josselin Feist, Gustavo Grieco, and Alex Groce. 2019. Slither: a static analysis framework for smart contracts. In *2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*. IEEE, 8–15.
- [17] Cormac Flanagan and K Rustan M Leino. 2001. Houdini, an annotation assistant for ESC/Java. In *International Symposium of Formal Methods Europe*. Springer, 500–517.
- [18] Brendan J Frey and Delbert Dueck. 2007. Clustering by passing messages between data points. *science* 315, 5814 (2007), 972–976.
- [19] Graeme Gange, Jorge A Navas, Peter Schachte, Harald Søndergaard, and Peter J Stuckey. 2016. An abstract domain of uninterpreted functions. In *Verification, Model Checking, and Abstract Interpretation: 17th International Conference, VMCAI 2016, St. Petersburg, FL, USA, January 17–19, 2016. Proceedings 17*. Springer, 85–103.
- [20] Basma Mohammed Salih Hasan and Adnan Mohsin Abdulazeez. 2021. A review of principal component analysis algorithm for dimensionality reduction. *Journal of Soft Computing and Data Mining* 2, 1 (2021), 20–30.
- [21] Everett Hildenbrandt, Manasvi Saxena, Nishant Rodrigues, Xiaoran Zhu, Philip Daian, Dwight Guth, Brandon Moore, Daejun Park, Yi Zhang, Andrei Stefanescu, et al. 2018. KEVM: A complete formal semantics of the Ethereum virtual machine. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. IEEE, 204–217.
- [22] Bo Jiang, Ye Liu, and WK Chan. 2018. ContractFuzzer: Fuzzing Smart Contracts for Vulnerability Detection. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. ACM, 259–269.
- [23] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* 7, 3 (2019), 535–547.
- [24] Junsang Kim and Seyong Kim. 2021. A survey of decentralized finance (defi) based on blockchain. *Journal of the Korea Society of Computer and Information* 26, 3 (2021), 59–67.
- [25] Queping Kong, Jiachi Chen, Yanlin Wang, Zigui Jiang, and Zibin Zheng. 2023. DeFiTainter: Detecting Price Manipulation Vulnerabilities in DeFi Protocols. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 1144–1156.
- [26] Anil Koyuncu, Kui Liu, Tegawendé F Bissyandé, Dongsun Kim, Jacques Klein, Martin Monperrus, and Yves Le Traon. 2020. Fixminer: Mining relevant fix patterns for automated program repair. *Empirical Software Engineering* 25 (2020), 1980–2024.
- [27] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *arXiv:2005.11401* (April 2021). <https://doi.org/10.48550/arXiv.2005.11401> arXiv:2005.11401 [cs].
- [28] Haonan Li, Yu Hao, Yizhuo Zhai, and Zhiyun Qian. 2023. Enhancing Static Analysis For Practical Bug Detection: An LLM-Integrated Approach. In *Proc. ACM Program. Lang.*
- [29] Shang-Wei Lin, Palina Tolmach, Ye Liu, and Yi Li. 2022. Solsee: a source-level symbolic execution engine for solidity. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1687–1691.
- [30] Junrui Liu, Yanju Chen, Bryan Tan, Isil Dillig, and Yu Feng. 2022. Learning Contract Invariants Using Reinforcement Learning. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*. 1–11.
- [31] Ye Liu and Yi Li. 2022. InvCon: A Dynamic Invariant Detector for Ethereum Smart Contracts. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*. 1–4.

- [32] Ye Liu, Yi Li, Shang-Wei Lin, and Cyrille Artho. 2022. Finding permission bugs in smart contracts with role mining. In *Proc. ACM ISSTA*.
- [33] Ye Liu, Yi Li, Shang-Wei Lin, and Cyrille Artho. 2022. Finding Permission Bugs in Smart Contracts with Role Mining. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA)*. ACM, New York, NY, USA, 716–727.
- [34] Ye Liu, Chengxuan Zhang, and Yi Li. 2024. Automated Invariant Generation for Solidity Smart Contracts. *arXiv preprint arXiv:2401.00650* (2024).
- [35] Ruijie Meng, Martin Mirchev, Marcel Bohme, and Abhik Roychoudhury. [n.d.]. Large Language Model guided Protocol Fuzzing. In *Proceedings of the Symposium on Network and Distributed System Security 2024*. <https://doi.org/10.14722/ndss.2024.24556>
- [36] Mark Mossberg, Felipe Manzano, Eric Hennenfent, Alex Groce, Gustavo Grieco, Josselin Feist, Trent Brunson, and Artem Dinaburg. 2019. Manticore: A user-friendly symbolic execution framework for binaries and smart contracts. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 1186–1189.
- [37] Tai D Nguyen, Long H Pham, Jun Sun, Yun Lin, and Quang Tran Minh. 2020. sfuzz: An efficient adaptive fuzzer for solidity smart contracts. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 778–788.
- [38] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *arXiv:2203.02155*.
- [39] Hammond Pearce, Benjamin Tan, Baleegh Ahmad, Ramesh Karri, and Brendan Dolan-Gavitt. 2023. Examining Zero-Shot Vulnerability Repair with Large Language Models. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 2339–2356. <https://doi.org/10.1109/SP46215.2023.10179324>
- [40] Anton Permenev, Dimitar Dimitrov, Petar Tsankov, Dana Drachler-Cohen, and Martin Vechev. 2020. Verx: Safety verification of smart contracts. In *2020 IEEE symposium on security and privacy (SP)*. IEEE, 1661–1677.
- [41] Michael Rodler, Wenting Li, Ghassan O Karam, and Lucas Davi. 2018. Sereum: Protecting existing smart contracts against re-entrancy attacks. *arXiv preprint arXiv:1812.05934* (2018).
- [42] Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. Code Llama: Open Foundation Models for Code. *arXiv:2308.12950*.
- [43] Secure3. [n.d.]. Securing Web3 with Decentralized Intelligence. <https://www.secure3.io/>.
- [44] Seongjin Shin, Sang-Woo Lee, Hwijee Ahn, Sungdong Kim, Hyoungseok Kim, Boseop Kim, Kyunghyun Cho, Gichang Lee, Woomyoung Park, Jung-Woo Ha, and Nako Sung. 2022. On the Effect of Pretraining Corpora on In-context Learning by a Large-scale Language Model. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, 5168–5186. <https://doi.org/10.18653/v1/2022.naacl-main.380>
- [45] Solidity 2022. Solidity. <https://solidity.readthedocs.io/en/v0.5.1/>.
- [46] Jon Stephens, Kostas Ferles, Benjamin Mariano, Shuvendu Lahiri, and Isil Dillig. 2021. Smartpulse: Automated checking of temporal properties in smart contracts. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 555–571.
- [47] Yuqiang Sun, Daoyuan Wu, Yue Xue, Han Liu, Wei Ma, Lyuye Zhang, Miaolei Shi, and Yang Liu. 2024. LLM4Vuln: A Unified Evaluation Framework for Decoupling and Enhancing LLMs' Vulnerability Reasoning. *arXiv:2401.16185* (Jan. 2024). <https://doi.org/10.48550/arXiv.2401.16185> *arXiv:2401.16185* [cs].
- [48] Yuqiang Sun, Daoyuan Wu, Yue Xue, Han Liu, Haijun Wang, Zhengzi Xu, Xiaofei Xie, and Yang Liu. 2024. Gptscan: Detecting logic vulnerabilities in smart contracts by combining gpt with program analysis. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.
- [49] Bryan Tan, Benjamin Mariano, Shuvendu K Lahiri, Isil Dillig, and Yu Feng. 2022. SolType: refinement types for arithmetic overflow in solidity. *Proceedings of the ACM on Programming Languages* 6, POPL (2022), 1–29.
- [50] Palina Tolmach, Yi Li, Shang-Wei Lin, and Yang Liu. 2021. Formal analysis of composable DeFi protocols. In *Financial Cryptography and Data Security: FC 2021 International Workshops: CoDecFin, DeFi, VOTING, and WTSC, Virtual Event, March 5, 2021, Revised Selected Papers 25*. Springer, 149–161.
- [51] Palina Tolmach, Yi Li, Shang-Wei Lin, Yang Liu, and Zengxiang Li. 2021. A survey of smart contract formal specification and verification. *ACM Computing Surveys (CSUR)* 54, 7 (2021), 1–38.
- [52] Hugo Touvron, Louis Martin, Kevin Stone, and et al. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. (2023). *arXiv:2307.09288*.
- [53] Trail of Bits 2019. *Echidna*. Trail of Bits. <https://github.com/trailofbits/echidna>
- [54] Petar Tsankov, Andrei Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Buenzli, and Martin Vechev. 2018. Securify: Practical security analysis of smart contracts. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*. 67–82.
- [55] Haijun Wang, Ye Liu, Yi Li, Shang-Wei Lin, Cyrille Artho, Lei Ma, and Yang Liu. 2020. Oracle-Supported Dynamic Exploit Generation for Smart Contracts. *IEEE Transactions on Dependable and Secure Computing* (2020).
- [56] Sally Junsong Wang, Kexin Pei, and Junfeng Yang. 2024. SMARTINV: Multimodal Learning for Smart Contract Invariant Inference. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 126–126.
- [57] Yuepeng Wang, Shuvendu K Lahiri, Shuo Chen, Rong Pan, Isil Dillig, Cody Born, and Immad Naseer. 2018. Formal specification and verification of smart contracts for azure blockchain. *arXiv preprint arXiv:1812.08829* (2018).
- [58] Gavin Wood. 2014. Ethereum: A Secure Decentralised Generalised Transaction Ledger. *Ethereum project yellow paper* 151 (2014), 1–32.
- [59] Siwei Wu, Dabao Wang, Jianting He, Yajin Zhou, Lei Wu, Xingliang Yuan, Qiming He, and Kui Ren. 2021. Defiranger: Detecting price manipulation attacks on defi applications. *arXiv preprint arXiv:2104.15068* (2021).
- [60] Chunqiu Steven Xia and Lingming Zhang. 2023. Keep the Conversation Going: Fixing 162 out of 337 bugs for \$0.42 each using ChatGPT. *arXiv:2304.00385* (April 2023). <https://doi.org/10.48550/arXiv.2304.00385> *arXiv:2304.00385* [cs].
- [61] Lyuye Zhang, Kaixuan Li, Kairan Sun, Daoyuan Wu, Ye Liu, Haoye Tian, and Yang Liu. 2024. ACFIX: Guiding LLMs with Mined Common RBAC Practices for Context-Aware Repair of Access Control Vulnerabilities in Smart Contracts. (2024). *arXiv:2403.06838* [cs.SE]
- [62] Zhuo Zhang, Brian Zhang, Wen Xu, and Zhiqiang Lin. 2023. Demystifying exploitable bugs in smart contracts. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 615–627.
- [63] Yilun Zhu, Joel Ruben Antony Moniz, Shruti Bhargava, Jiarui Lu, Dhivya Piraviperumal, Site Li, Yuan Zhang, Hong Yu, and Bo-Hsiang Tseng. 2024. Can Large Language Models Understand Context? *arXiv:2402.00858* (Feb. 2024). <https://doi.org/10.48550/arXiv.2402.00858> *arXiv:2402.00858* [cs].

APPENDIX

A SUPPLEMENTARY MATERIAL

Table 7 lists the raw information for all the 61 Certora projects, of which we collected 23 projects with available code and properties.

Table 7: The raw information for all the 61 Certora projects.

Report Name	Year	Month	Included	#Property
Aave CLSynchronicity Price Adapter	2022	December	×	35
Aave GHO Stablecoin	2023	March	✓	
Aave Governance V2 Update	2022	September	×	42
Aave L2 Bridge	2022	July	✓	
Aave Proof of Reserve	2022	November	✓	3
Aave Protocol V2	2020	December	✓	17
Aave Rescue Mission Phase 1	2023	January	✓	1
Aave Staked Token v1.5	2023	February	✓	11
Aave Static aToken	2023	April	✓	24
AAVE Token V3	2022	September	×	59
Aave V2 ASiETH	2022	August	×	
Aave V3	2022	January	✓	16
Aave V3 BTC.b Listing Steward	2022	September	×	
Aave V3 MAI & FRAX Listing Stewards	2022	August	×	10
Aave V3 PR #820	2023	March	×	
Aave V3 sAVAX Listing Steward	2022	July	×	35
Aave V3 sUSD Listing Steward	2022	August	×	
Aave V3.0.1	2022	December	×	41
Aave Vault	2023	June	✓	
Aave-StarkNet L1-L2 Bridge	2022	October	✓	20
Balancer	2022	September	×	
Balancer V2	2021	April	×	13
Balancer V2 (Issues only)	2021	April	×	
Balancer's Timelock Authorizer Verification Report	2023	May	×	1
Benqi's Liquid Staking Contracts	2022	April	×	
Celo Core Contracts Release 4	2021	May	×	30
Celo Governance Protocol	2020	May	✓	
Compound V1 Price Oracle	2018	September	×	80
Compound V3 Comet	2022	July	×	
Compound's MoneyMarket v2 formal verification report	2019	August	✓	33
Compound's Open-Oracle with Uniswap Anchor	2020	August	×	
Daoism	2022	October	×	16
dcSpark	2022	December	×	
dForce Lending Protocol	2021	February	×	21
Euler	2021	November	×	
Furucombo	2021	May	✓	36
Kashi Lending Protocol	2021	March	×	
Keep's Fully-backed bonding contract	2020	November	✓	22
Lido V2	2023	April	✓	
Lyra	2021	May	×	98
Master Chef V2	2021	April	×	
Notional Finance V2	2021	November	✓	2
OOPSLA'2020	2020	-	×	
Open Zeppelin	2022	April	✓	1
Open Zeppelin	2022	June	×	
OpenZeppelin Governance contracts	2021	December	×	1
Opyn Gamma Protocol	2020	December	✓	
Orchid's Smart Contracts	2019	December	×	1
Origin OUSD Token	2021	February	✓	
Popsicle V3 Optimizer	2021	November	✓	1
Radicale Drips	2023	January	✓	
Rolla Finance	2021	August	×	1
SaaS Verification Report by Blockswap Labs	2022	July	×	
SaaS Verification Report by Silo	2022	July	×	1
Sushi BentoBox	2021	February	✓	
Sushi Compound Strategy	2021	April	×	1
SushiSwap ConstantProductPool	2021	November	×	
SushiSwap TridentRouter	2021	November	×	1
Synthetix Multi-Collateral Loans	2020	December	×	
Trader Joe	2022	March	✓	1
Zesty	2021	July	×	