



VERISTRUCT: AI-assisted Automated Verification of Data-Structure Modules in Verus

Chuyue Sun¹, Yican Sun²*, Daneshvar Amrollahi¹, Ethan Zhang¹,
Shuvendu Lahiri³, Shan Lu³, David Dill¹, and Clark Barrett¹

¹ Stanford University

² School of Computer Science, Peking University, Beijing, China

³ Microsoft Research

Abstract. We introduce VERISTRUCT, a novel framework that extends AI-assisted automated verification from single functions to more complex data structure modules in Verus. VERISTRUCT employs a planner module to orchestrate the systematic generation of abstractions, type invariants, specifications, and proof code. To address the challenge that LLMs often misunderstand Verus' annotation syntax and verification-specific semantics, VERISTRUCT embeds syntax guidance within prompts and includes a repair stage to automatically correct annotation errors. In an evaluation on eleven Rust data structure modules, VERISTRUCT succeeds on ten of the eleven, successfully verifying 128 out of 129 functions (99.2%) in total. These results represent an important step toward the goal of automatic AI-assisted formal verification.

1 Introduction

The power of generative AI introduces new risks to our critical digital infrastructure. AI systems have a remarkable ability to analyze code, making it much easier for a malicious actor to find and exploit vulnerabilities [3]. At the same time, more and more code is being written by or with the assistance of AI systems. While it is true that such AI systems can greatly increase the productivity of programmers, they can also introduce additional correctness errors and security vulnerabilities [25, 36]. Thus, productivity gains come at the risk of churning out mountains of buggy, insecure code.

Program verification can mitigate these risks by mathematically proving that code—whether human-written or AI-generated—is free of bugs and vulnerabilities. However, program verification relies on the addition of sophisticated *logical annotations* to code, including preconditions, postconditions, invariants, and auxiliary proof blocks. These annotations, which we also refer to as *specifications*, are used to generate verification conditions, which are then proved automatically by automated reasoning tools like satisfiability modulo theories (SMT) solvers [4]. The application of program verification in practice has been severely limited because of the enormous effort and expertise required to create such annotations.

* Work was accomplished while the author was visiting Stanford University.

Fortunately, the power of generative AI is also shifting the calculus on the practicality of program verification. Recently, there have been encouraging results on using large language models (LLMs) to *automatically* specify and verify software [20, 29, 30, 37, 39]. Our work aims to build on these efforts with the vision of one day making it possible to have AI-assisted program verification at scale and verified libraries of reusable code [34].

So far, AI-assisted verification efforts have focused on simple textbook algorithms consisting of a single function. In this paper, we extend this by developing a novel workflow for the verification of data structure modules. The ability to verify data structures is a crucial advance since: (1) data structures are ubiquitous in everyday programming, forming the foundation of many software systems; and (2) verifying library data structures lowers the burden for verifying more complex code, as client code can then rely on the correctness of these already-proven components.

We implement our workflow in VERISTRUCT, an AI-assisted automated verification framework for Verus [12, 13]. Verus is a state-of-the-art program verification extension for Rust. We target Verus because it has been successfully applied to the construction of verified systems [12], and its design allows developers to write verification annotations using a syntax closely aligned with Rust. This design not only lowers the barrier for Rust engineers to adopt formal verification but also positions Verus to have an increasing impact as Rust continues to gain traction in mainstream software development.

The input to VERISTRUCT consists of the Rust source code for the module to be verified and a unit test suite that illustrates the intended usage and expected behavior of the data structure. The test suite serves a dual purpose: it helps ensure that the generated formal specifications align with the developer’s intent, and it helps rule out trivial or vacuous specifications [11]. VERISTRUCT automatically augments the provided Rust implementation with annotations needed for program verification. This augmented code, together with the provided unit test suite, are then sent to the Verus verifier, which attempts to formally verify the assertions in the test suite. If the verification check succeeds, the formally verified code is returned to the user. Otherwise, VERISTRUCT attempts to extend or repair the annotations based on the feedback from the verifier.

We highlight two challenges for LLM-assisted verification of data-structure modules. The first challenge is that verifying data structures is inherently more complex than verifying a single function. Specifically, data-structure verification often requires two additional elements: (i) a suitable mathematical abstraction that allows the verifier to reason logically about the data structure [22, 24, 27]; and (ii) so-called *type invariants* that must be preserved by operations on the data structure [19, 29]. Moreover, data-structure modules usually expose multiple methods that must be verified jointly under a shared type invariant. This stands in contrast to other recent approaches, which focus on generating a single specification artifact (e.g., proof blocks for an individual function or automated synthesis of class invariants) [9, 20, 26, 29, 37–39]. In short, our verification task extends beyond the scope of these existing techniques.

The second challenge lies in the LLMs’ limited understanding of Verus’ specialized syntax and verification-specific semantics. For instance, annotations in Verus are only allowed to invoke “specification functions,” which cannot modify global state. However, the LLM sometimes suggests invoking a regular Rust function. Such deficiencies primarily stem from the scarcity of Verus code in existing training data.

To address the first challenge, we build on previous work [39] by adding two new modules: (1) a *View* module responsible for producing a mathematical abstraction of the data structure; and (2) a *Type Invariant* module, responsible for producing type invariants. For each component, we design dedicated system and user prompts that clearly articulate the task. These prompts include: (1) the objective of the component; (2) relevant background information to help the LLM better understand the task; (3) step-by-step instructions and the required output format; and (4) in-context learning examples [6]. Additionally, because not all verification components are necessary for every task, we introduce a *planner* module, which analyzes the verification task as an initial step and decides which components need to be generated.

To address the second challenge, we pursued two approaches. First, we developed a suite of automated repair modules to identify and correct errors in the generated annotations. Compared with the previous work [39] which focused solely on repairing proof blocks, we further introduce dedicated repair modules targeting views, type invariants, and specifications. Specifically, to facilitate the joint verification of multiple methods and tests, VERISTRUCT includes a repair module that refines specifications whenever the verifier reports a failed test verification. Second, we embedded structured syntax guidelines into the prompt to mitigate syntax errors in generated annotations. We sourced these guidelines from the Verus tutorial [32] and from standard Verus libraries [33].

We evaluate VERISTRUCT on eleven data-structure benchmarks drawn from Verus examples and third-party code. Our framework automatically produces verified implementations for ten of these benchmarks while substantially reducing the amount of handwritten annotation required from developers. When writing unit tests, we observe that achieving high coverage is both challenging and labor-intensive. To alleviate this burden, we apply LLM-assisted techniques to automatically generate unit test cases.

To summarize, this paper makes the following key contributions:

1. we introduce a new LLM-assisted workflow for generating program verification annotations for data-structure modules;
2. we provide an implementation of our workflow in the VERISTRUCT tool; and
3. we evaluate VERISTRUCT on eleven data structure benchmarks, demonstrating its effectiveness.

The prompt and source code of VERISTRUCT is available on GitHub [1].

```

1  verus! { // Indicates the Verus environment
2
3  pub struct RingBuffer<T: Copy> {
4      ring: Vec<T>,
5      head: usize,
6      tail: usize,
7  }
8  // The View trait, which tells the Verus verifier
9  // how to logically represent RingBuffer
10 impl<T: Copy> View for RingBuffer<T> {
11     type V = (Seq<T>, usize);
12     closed spec fn view(&self) -> Self::V {
13         let cap = self.ring.len();
14         let content =
15             if self.tail >= self.head {
16                 (self.ring)@.subrange(
17                     self.head as int, self.tail as int
18                 )
19             } else {
20                 (self.ring)@.subrange(
21                     self.head as int, cap as int
22                 ).add((self.ring)@.subrange(
23                     0, self.tail as int
24                 ))
25             };
26         (content, cap)
27     }
28 }
29
30 impl RingBuffer<T> {
31     // The type invariant, specifying the properties that
32     // RingBuffer object must satisfy
33     #[verifier::type_invariant]
34     closed spec fn inv(&self) -> bool {
35         &&& self.head < self.ring.len()
36         &&& self.tail < self.ring.len()
37         &&& self.ring.len() > 0
38     }
39     pub fn new(ring: Vec<T>) -> (ret: RingBuffer<T>)
40     requires
41         ring.len() >= 1
42     ensures
43         ret@.0.len() == 0,
44         ret@.1 == ring.len() as nat
45     {
46         // Create an empty ring buffer
47         RingBuffer { head: 0, tail: 0, ring }
48     }
49     pub fn is_full(&self) -> (ret: bool)
50     ensures
51         ret == (self@.0.len() == (self@.1 - 1) as nat)
52     {
53         proof {
54             use_type_invariant(&self);
55         }
56         self.head == ((self.tail + 1) % self.ring.len())
57     }
58     pub fn enqueue(&mut self, val: T) -> (succ: bool) { ... }
59     pub fn dequeue(&self) { ... }
60 }
61
62 fn test_ring_buffer() { /* Unit tests (omitted). */ }
63 fn main() { test_ring_buffer(); }
64 // End of verus!

```

Fig. 1. Verified Ring Buffer, Lines Highlighted in Green are Annotations

2 Preliminaries: Verifying a Data-Structure in Verus

In this section, we provide background on deductive verification and Verus. We start by introducing an example that will be used to illustrate key concepts throughout the paper.

Example 1 (Verified Ring Buffer). A ring buffer is a fixed-size data structure that treats memory as a circular array. Ring buffers are widely used in streaming, networking, and real-time systems [10, 16]. Fig. 1 shows the Rust implementation for a ring buffer, including the Verus annotations necessary for verifying key properties of the implementation. \square

Generally speaking, deductive program verification augments programs with *logical annotations*. These annotations connect directly to correctness: the verifier generates proof obligations from the code and annotations; if all obligations are proved, then the implementation is correct with respect to the annotations. Otherwise, the verifier produces an error message. Below, we describe the different categories of annotations available in the Verus program verification framework.

Specification Functions. In Verus, functions are classified into two categories: *specification functions* and *executable functions*. Specification functions, marked with the `spec` keyword (e.g., `view` and `inv` in Fig. 1), are pure and cannot modify any global or mutable state. They are simply logical definitions of pure functions without side effects and are called exclusively from specifications, invariants, and proofs. In contrast, executable functions (e.g., `new`, `is_full`, `enqueue`, and `dequeue`) define concrete program behavior and may modify state. When writing

annotations, only specification functions can be called, since annotations must remain side-effect-free and semantically independent of execution.

Pre- and Postconditions. Annotations include preconditions and postconditions for each function. The implicit program requirement is that: if a function is invoked with a program state satisfying the precondition, then it will return (or terminate) in a state satisfying the postcondition. In Verus, one uses the keywords **requires** (Lines 38–39) and **ensures** (Lines 40–41, 48–49) for specifying preconditions and postconditions, respectively. The verifier checks that every call site meets the callee’s preconditions and that every function body establishes its postconditions from the assumed preconditions.

Proof Blocks. Implementation code may include *proof blocks* embedded in the code. These blocks provide hints to the verifier to help it discharge difficult proof obligations (Lines 51–53 in Fig. 1).

Verus includes two additional constructs that are helpful when verifying data structure modules: views and type invariants.

View Traits. The *view trait* in Verus establishes a bridge between the concrete implementation of a data structure and its abstract mathematical representation used in specifications. It requires implementing a **view()** function, which defines how an instance of a data structure maps to a logical object that the verifier can reason about. While one could theoretically encode a data structure’s state as a Cartesian product of all its fields, defining a custom view offers better abstraction, clarity, and proof efficiency, allowing specifications and proofs to focus on the logical behavior rather than the implementation details.

Once the view is implemented, specifications can refer to the logical abstraction using **a.view()** or its shorthand **a@**. Verus provides built-in specification types such as **nat**, **Seq<T>**, **Set<T>**, and **Map<K,V>** to represent natural numbers, sequences, sets, and maps, respectively. These types serve as good building blocks for implementing views.

Example 2. In Fig. 1, the **View** for **RingBuffer** (Lines 8–28) abstracts the buffer as a pair (**Seq<T>**, **usize**). The first element in the pair is a mathematical sequence (**Seq<T>**) that reflects the elements currently stored in the buffer in order, starting from the element at **head** and continuing up to (but not including) the element at **tail**. The second element in the pair is the capacity of the ring buffer. To implement the view function, we slice and concatenate segments of the underlying ring according to the positions of **head** and **tail**, while hiding **head** and **tail** themselves.

This logical view abstracts the circular implementation, enabling specifications for buffer operations—like **enqueue** and **dequeue**—to be written as simple sequence transformations (e.g., **appending** or **removing** elements), without dealing with low-level index arithmetic or memory layout details.

Type Invariants. A *type invariant* is a logical formula that all instances of a data structure must satisfy. In Verus, one declares a type invariant using the

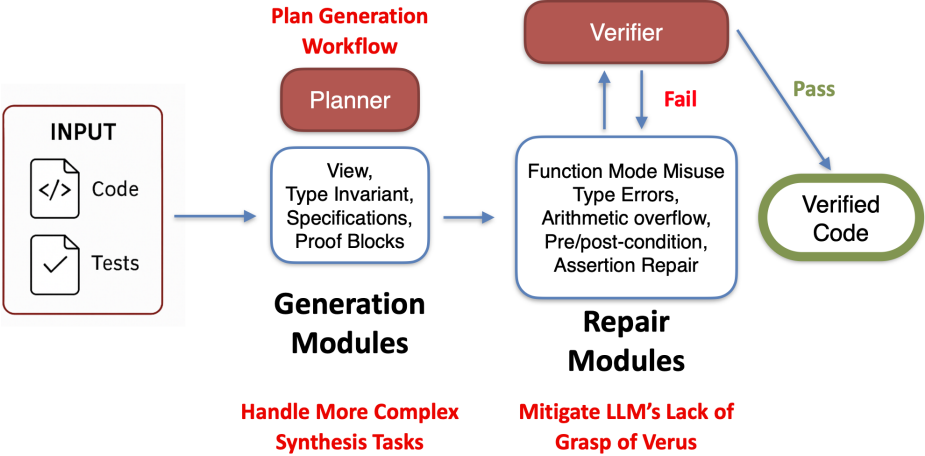


Fig. 2. Workflow of VERISTRUCT, where rounded rectangles denote modules

`#[verifier::type_invariant]` attribute. We can call `use_type_invariant` (Line 51) within a proof block to make invariant facts available in the current proof context.

Example 3. In Fig. 1 (Lines 29–36), the ring buffer invariant states that `head` and `tail` are valid indices and that the ring has positive capacity.

3 The VERISTRUCT Workflow

In this section, we give an overview of the workflow used in VERISTRUCT using the `RingBuffer` (Example 1) as a running example.

The input to VERISTRUCT consists of two parts: (1) a user-provided Rust implementation of the code to be verified (e.g., the code in Fig. 1 but *without* any annotations—the green-highlighted lines would be absent), and (2) a unit-test suite that conveys typical usage patterns and rules out trivial specifications. The output is a fully annotated version of the code that passes the Verus verifier (e.g., the full `RingBuffer` code as shown in Fig. 1). Care is taken to ensure that both the code and the tests remain unchanged throughout the workflow.

The overall workflow is shown in Fig. 2. VERISTRUCT builds on the general approach of the AutoVerus system [39], in that it consists of a network of modules that collaborate to complete program verification tasks. However, as noted in the challenge discussion of Sect. 1, verifying data structures requires techniques and components not found in AutoVerus or any prior work [20, 26, 29, 37, 39].

VERISTRUCT applies a two-stage pipeline as shown in Algorithm 1 to synthesize annotations and perform verification.

1. Stage 1: Generate the initial draft of annotations through the function `GenAnnos`. The detailed process is described in Sect. 4.

Algorithm 1: The Outermost Algorithm of VERISTRUCT

```

1 Input. : the data structure code to be verified; test: unit test suite
2 Output. verified version of 
3 Function VerifyModule(, test)
4    ← GenAnnos(, test) // Stage 1: Generate the Annotations
5   return RepairAnnos(, test) // Stage 2: Repair Annotations
```

Algorithm 2: The Generation Stage

```

1 Parameter. n: Number of samples generated in each module invocation.
2 Input. : the data structure code to be verified; test: unit test suite
3 Output.  with annotations generated
4 Function GenAnnos(, test)
5   plan ← ExecPlanner()
6   foreach  $M_i \in \mathbf{plan}$  do
7      ← ExecWithSampling( $M_i$ , , test, n)
8   return 
```

2. Stage 2: Invoke **RepairAnnos** to iteratively detect and repair incorrect annotations. The details are presented in Sect. 5.

Due to space limitations, we only present sketches of our prompts, but the full details of our prompts are publicly available [1].

4 Stage 1: Generating the Annotations

The generation stage creates annotations. As mentioned above, we may need to generate different kinds of annotations, so VERISTRUCT uses four dedicated modules, one for each type of annotation: (M_1) a view module, (M_2) a type invariant module, (M_3) a specifications module (for specification functions, preconditions, and postconditions), and (M_4) a proof blocks module, which generates both proof blocks and invariants.

By examining the dependencies among these kinds of annotations, we observe that it should typically be sufficient to invoke the modules in order: $\langle M_1, M_2, M_3, M_4 \rangle$. At the same time, it is not always necessary to execute every module. For example, some data structures do not require non-trivial type invariants; in such cases, generating an invariant may incur: (1) more instability in the language model (and thus affect the robustness of our framework); (2) unnecessary computational cost, and (3) more runtime overhead. To mitigate these issues, we employ a *planner* module that determines which generation modules should be executed.

Planner. The generation process begins with the planner agent (Line 5). The function **ExecPlanner** accepts the target code as input, invokes the planner agent, and outputs the modules to execute. The planner agent instructs the LLM to select only the necessary generation modules to be executed. Its prompt

comprises four blocks: (i) a concise description of the role of the planner; (ii) a catalog of the available generation modules: there are four in our current framework; (iii) background knowledge on Verus annotation generation; and (iv) a machine-parseable output format for the execution sequence.

The prompt is designed with several guiding principles:

1. Invoke M_1 if the data structure can be represented by a sequence, set, or map.
2. Invoke M_2 when non-trivial relationships exist among fields (e.g., range constraints or arithmetic relations).
3. If M_2 is invoked, also execute M_4 to add type invariants into the proof context.

For the ring buffer, the planner emits the full execution sequence, namely $\langle M_1, M_2, M_3, M_4 \rangle$.

Module Invocation. Following the generated plan, the framework sequentially invokes each module M_i (Lines 6–10). The function **ExecWithSampling** takes as input: (1) a module to be executed; (2) the code to be verified; and (3) the unit test suite. It executes the module, and returns the annotated code. To improve robustness, it asks for n samples during each module invocation [28]. It then selects the result yielding the maximum number of successfully verified functions. We describe the view generation module in detail below, and then describe some differentiating aspects of the other modules, which are otherwise similar.

View Generation (and Refinement) Module. To generate a view, our prompt (Fig. 3) is structured as follows:

1. *Objective.* This part articulates the task of generating a **View** implementation.
2. *Verus Guidelines.* This part provides conceptual foundations for generating a valid **View** implementation. As discussed in Sect. 1, LLMs often struggle with Verus’ syntax and verification-specific semantics. To mitigate this, our framework incorporates guidelines on Verus syntax and semantics imported from the Verus standard library documentation [33] and the official Verus tutorial [32]. For each different module in VERISTRUCT, a specific set of guidelines is added to the prompt. For **View** generation, for example, we include:
 - An overview of the **View** trait (as introduced in Sect. 2).
 - Guidelines on how to manipulate Verus’ logical types, which serve as foundational building blocks for providing a **View** implementation.
 - A chapter on Verus’ specialized logical syntax.
3. *Step-by-Step Instructions.* This part provides step-by-step procedural instructions on how to systematically generate a view.
4. *Examples* [6]. This part provides examples of **View** implementations for other verified data structures (e.g., doubly-linked lists). These examples are kept the same across (and are disjoint from) all benchmarks.
5. *Code with Tests.* Finally, the prompt adds the code to be verified code along with the unit test suite.

```

[Objective]
You are an expert in Verus (verifier for rust). Your task is to
generate a View trait for the given data-structure module...
[Verus Guidelines]
Background on View and relevant Verus features...
[Step-by-Step Instructions]
1. Infer what should be contained in the View trait.
2. Generate the view based on the inferred information.
[Examples]
The View trait for other data structures, e.g., doubly-linked list.
[Code and Tests]
Code to be generated a View trait...

```

Fig. 3. The Prompt for View Generation

For the RingBuffer example, LLMs sometimes produce a trivial Cartesian-product view, which includes all of the data fields in RingBuffer (see Fig. 4), i.e., the view function simply maps the buffer content `self.ring` to its logical abstraction `self.ring.view()` and converts `head` and `tail` into logically defined natural numbers. This view function, while syntactically valid, fails to abstract away the circular structure of the ring buffer. As a result, the subsequent specification generation becomes significantly more complex: we have to explicitly reason about the arithmetic relationship between `head` and `tail` when specifying the behavior of operations such as `enqueue` and `dequeue`.

```

1  impl View for RingBuffer<T> {
2      type V = (Seq<T>, nat, nat);
3      closed spec fn view(&self) -> Self::V {
4          (self.ring.view(), head as nat, tail as nat)
5      }
6  }

```

Fig. 4. Result of View Generation for RingBuffer

To address this issue, we introduce an additional *view refinement* step that prompts the LLM to reconsider the generated `View` and reconstruct it with fewer, more abstract components. This module encourages the model to focus on the conceptual essence of the data structure rather than on reproducing its concrete implementation. The refinement process substantially simplifies subsequent specification and proof generation tasks.

The *view refinement* prompt follows the same structure as Fig. 3, but emphasizes abstraction, minimality, and logical coherence over completeness. After applying this refinement, the model typically is able to produce an improved `View` implementation, like the one shown in Fig. 1.

Other Modules. The prompts for the other modules share a similar structure to that shown in Fig. 3. Each module prompt begins with a task description,

followed by module-specific Verus guidelines. These guidelines are specialized to the individual modules.

The *Type Invariant Module* is focused on generating type invariants, which were introduced in Sect. 2. In this module, the LLM is instructed to focus on common patterns such as range constraints, capacity checks, and arithmetic relationships between fields.

The *Specification Module* is focused on generating specification functions, preconditions and postconditions. We add mutability guidelines to ensure compatibility with Rust’s ownership and mutability system in **requires** and **ensures** clauses. For example, the LLM is prompted to use `old(·)` in **requires** to reference a variable’s value prior to function invocation.

The *Proof Block Module* handles proof blocks and loop invariants. The LLM is directed to construct the appropriate proof blocks and loop invariants, while explicitly incorporating type invariants into the proof context or applying relevant lemmas when necessary.

5 Stage 2: Repairing the Annotations

Unfortunately, a single generation run rarely yields fully verified code. LLMs often struggle with Verus’ syntax and rules. For example, the following is an *incorrect* specification that was generated for the `enqueue` method of the `RingBuffer` example:

```

1 pub fn enqueue(&mut self, val: T) -> (succ: bool)
2 ensures
3     succ == !old(self).is_full(),
4     ... // omitted
5 { /* omitted */ }
```

The intent is reasonable: `enqueue` succeeds if and only if the original buffer is not full. However, `is_full` is an *executable* function (see Line 46 of Fig. 1) and therefore cannot be called from a specification context. The verifier reports:

```
Error: cannot call the executable function is_full in annotation.
```

To handle such issues, VERISTRUCT executes an iterative repair loop. In each iteration, it: (i) selects an error reported by the verifier; (ii) applies a corresponding predefined repair module; and (iii) reruns verification. The process continues until either all errors are eliminated or a preset iteration budget is exhausted. VERISTRUCT provides a suite of repair modules, each targeting a specific class of errors. Error messages are routed to the appropriate module via lightweight pattern matching over the verifier’s error message.

For the error reported above, VERISTRUCT invokes a dedicated mode-repair module, which provides instructions for fixing this mode misuse (Figure 5). After applying this module, a new specification is generated. For example, one possible solution (one we actually observed) is to create a specification version of the `is_full` function and use it to replace all occurrences of `is_full`. With this change, the error no longer occurs, and the workflow can continue.

```
1 spec fn is_full_spec(&self) -> bool { self@0.len() == self@().1 - 1 }
```

```
[Objective]
Fix the error for the following code. The error indicates that an
executable function is called from annotations or vice versa...
[Relevant Background]
Background on mode and relevant Verus Features
[Code with Error, and Tests] ...
[Instructions] Make one of these changes:
1. Adjust the function to be compatible with the calling context
...
```

Fig. 5. The Prompt for Repairing the Misuse of Specification and Executable Functions

Algorithm 3: The Repair Stage

```
1 Parameter. n: Number of samples generated in each module invocation.
2 Parameter. m: Maximum number of iterations in the repair loop.
3 Given. rep_modules: A pre-defined set of repair modules.
4 Input. code: the code after generating annotations; test: unit test suite
5 Output. code with annotations repaired
6 Function RepairAnnos(code, test)
7   for r ← 1 to m do
8     err ← TryVerify(code, test)
9     if err is None then
10      return Success, code
11     cur_module ← default_module
12     foreach (pattern, module) ∈ rep_modules do
13       if err matches pattern then
14         cur_module ← module
15         break
16     code ← ExecWithSampling(cur_module, code, test, n)
17   return Failure, code
```

Algorithm 3 shows the repair algorithm in detail. VERISTRUCT provides a predefined set of repair modules, denoted as `rep_modules`. Each module is associated with a regular expression `pattern` specifying the class of error messages that it addresses. In the current implementation, VERISTRUCT includes repair modules for: (1) misuse between specification and executable functions (as demonstrated above), (2) inconsistencies with Rust’s mutability type system, (3) violations of preconditions and postconditions, (4) arithmetic overflow or underflow, (5) mismatches between logical and Rust types, (6) test assertion failures, and (7) a fallback module, `default_module`, which prompts the LLM to attempt a repair based directly on the provided error message.

Note that the repair module for test assertion failures requires the LLM to do interprocedural analysis, something largely absent from previous work. The repair module inspects the failing assertion, identifies the method invoked immediately before it, and then instructs the LLM to strengthen that method’s postcondition so as to ensure the assertion holds.

The repair process proceeds iteratively, with a maximum number of iterations defined by the parameter `m`. In each iteration, the system first attempts to verify the current implementation via **TryVerify** (Line 8), which takes the code and test suite as input and returns the highest priority error message `err` produced by Verus. If no error is reported, verification succeeds and the code is returned (Lines 9–10). Otherwise, `err` is matched against all module patterns; if a match is found, the corresponding repair module is executed (Lines 13–15). If no match is found, the `default_module` is invoked. Each module execution includes `n` sampling attempts to improve robustness (Line 16). If the implementation remains unverified after `m` iterations, the system reports failure and returns the final unverified result (Line 17).

6 Evaluation

In this section, we describe an evaluation of VERISTRUCT on a set of Rust benchmarks.

Benchmarks. We evaluate VERISTRUCT on a benchmark set comprising eleven Rust data structure modules drawn from the Verus GitHub repository [13] as well as other open-source repositories. A list of the benchmarks is shown in Table 1. Although the original sources are publicly available, our benchmarks have been substantially modified with more complete specifications, additional methods, and unit tests. We verified that the versions used in our evaluation do not appear verbatim in the `o1-2024-12-17` training snapshot; for example, no form of our fully verified `RINGBUFFER` appears in the dataset. Furthermore, even if older versions existed in pretraining data, our baseline results demonstrate that LLMs still struggle to synthesize correct Verus specifications without structured guidance—underscoring the necessity of VERISTRUCT’s workflow. The benchmarks cover a set of commonly-used data structures, including a ring buffer, a vector implementation with set abstractions and common algorithms, a polymorphic option type, tree-based modules for standalone nodes and a treemap, a bitmap with 64-bit blocks, and concurrent data structures such as a read–write lock and an invariant framework for shared state. Each benchmark contains 5–21 functions to verify, where the number of functions refers to the total number of member methods in the data structure module and the corresponding test functions. In total, our benchmark set includes 129 functions to be verified across all modules.

Baseline. To demonstrate the effectiveness of VERISTRUCT, we compare it against a baseline that iteratively invokes an LLM to generate annotations, without employing the systematic generation-and-repair workflow described in Sect. 3. In each iteration, the baseline calls a simple `BaselineModule` that invokes

Table 1. Benchmark set consisting of eleven data structures.

Benchmark	Description	#Functions
ATOMICS	An atomic counter in concurrent programming.	11
BITMAP	Bitmap implemented over 64-bit words.	14
TREEMAP	BST-based map with ordered key/value bindings.	21
INVARIANTS	Concurrent routines showing reusable invariants.	7
NODE	Implementing the node class of the binary search tree.	12
OPTION	Polymorphic option wrapper with safe accessor methods.	15
RINGBUFFER	Implementing a ring buffer.	13
RWLOCKVSTD	Read/write lock implementation.	5
SETFROMVEC	Set abstraction constructed from backing vectors.	10
TRANSFER	Account transfer routines that enforce balanced updates.	5
VECTORS	Implementing a vector with basic algorithms.	16

Table 2. The Comparison between VERISTRUCT and Baseline.

	#Solved	#Functions
VERISTRUCT	10 (↑ 150.0%)	128 (↑ 146.2%)
Claude Code	8	102
BASELINE	4	52

the LLM once to generate all annotations simultaneously. If the Verus verifier reports an error, we record the failing code and the associated error message, which are then injected into the prompt for subsequent iterations. The baseline prompt includes: (1) the code to be verified, (2) the unit test suite, and (3) the previously failing code along with any corresponding error message.

We also compare against Claude Code [2], a state-of-the-art coding agent that can autonomously invoke external tools such as the Verus verifier. We configure Claude Code (using Claude Sonnet 4.5) to iteratively generate annotations and respond to Verus error messages in an agentic loop, analogous to our baseline but with the added capability of tool invocation and autonomous iteration.

Procedure. We run VERISTRUCT and the baseline on every benchmark individually. For each benchmark we first execute VERISTRUCT to completion and then record how many LLM invocations it required. We observe that the maximum number of LLM invocations for VERISTRUCT is 13, so we run the baseline on each benchmark with 13 LLM invocations. All experiments are conducted on a server running Ubuntu 22.04.5 LTS with a 24-core Intel(R) Core(TM) i9-12900K CPU and 64 GB RAM. We use OpenAI o1 [23] as the back-end model for both VERISTRUCT and the baseline. We instantiate VERISTRUCT with $n=3$ and $m=5$; that is, the number of samples generated per module is 3, and the maximum number of repair rounds is set to 5.

Results. We summarize the results in Tables 2 and 3. VERISTRUCT successfully solves 10 out of 11 benchmarks. For the only unsolved benchmark (NODE), VERISTRUCT still verifies 11 out of 12 functions. In contrast, the baseline solves

only 4 out of 11 benchmarks. In terms of the total number of functions verified, VERISTRUCT verifies 128 out of 129 (99.2%) functions, whereas the baseline verifies only 52. Claude Code, despite using a more recent model (Claude Sonnet 4.5) and having the ability to autonomously invoke Verus, verifies 102 functions across 8 benchmarks—substantially better than the single-prompt baseline but still short of VERISTRUCT’s performance. Notably, Claude Code consumes approximately 24k tokens per benchmark on average, while VERISTRUCT uses only 22k tokens, demonstrating that our structured workflow achieves better results with comparable or lower token cost. In summary, VERISTRUCT solves substantially more benchmarks and verifies significantly more functions, demonstrating that our approach improves the capabilities of LLMs to generate correct annotations.

To provide a more fine-grained comparison, we inspect the verification trajectories of VERISTRUCT across all benchmarks. For each integer k , and for each benchmark, we compute the total number of functions fully verified within the first k LLM invocations. We then sum the numbers across all benchmarks for that value of k , plot the result, and then increase k . Plotting these points creates a visualization that shows the verification progress as a result of increasing LLM calls. We show the result in Fig. 7, where the x-axis denotes the LLM invocation index k , and the y-axis shows the cumulative number of functions verified within the first k invocations across all benchmarks.

Although the graph primarily trends upward, there are brief dips where the number of verified functions decreases. These temporary decreases occur during specification generation, when none of the generated specifications across the requested samples is syntactically correct. In such cases, VERISTRUCT must adopt a candidate specification that inadvertently breaks previously verified functionality. Our pipeline deliberately retains these specifications – even when they momentarily reduce coverage –so that subsequent repair modules can diagnose the faulty contracts once again and make progress towards verification. Typically, later stages recover from these dips and continue to progress toward full coverage.

We can observe the consistent dominance of VERISTRUCT over the baseline throughout the verification process. In particular, VERISTRUCT achieves the largest gain in verified functions during the second invocation (the first is for the planner), highlighting the effectiveness of our planner in selecting a highly effective initial module that verifies a large portion of functions early on. Overall, the experimental results clearly demonstrate that VERISTRUCT substantially outperforms the baseline, validating the effectiveness of our workflow design. Fig. 6 further shows the per-benchmark verification progression, where each benchmark makes steady progress toward its ground-truth target through iterative refinement.

Case Study: Bitmap. Interestingly, we find that the LLM produces a solution for the BITMAP benchmark that is fundamentally different from the ground-truth solution. In the ground-truth implementation written by human experts, the `View` trait models the bitmap as a two-dimensional array: each `u64` block is first abstracted into a bit sequence of length 64, and then the entire bitmap data structure is represented as a 2D array of these blocks. To manipulate this

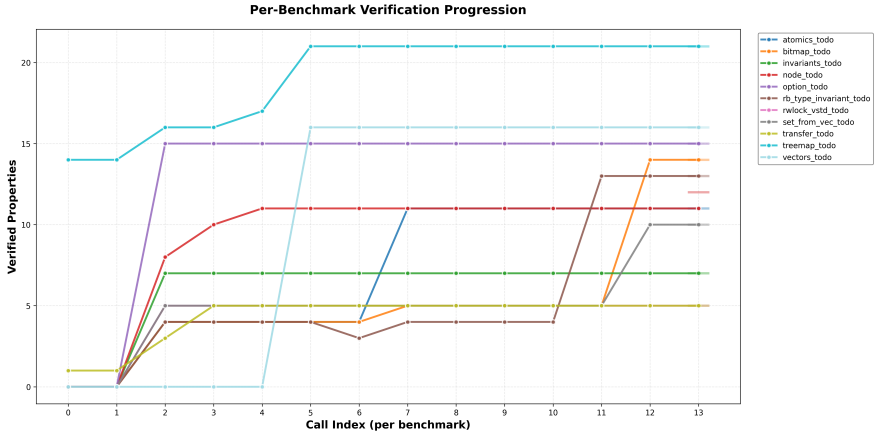


Fig. 6. Per-benchmark verification progression. Each benchmark makes steady progress toward its ground-truth target through iterative refinement, demonstrating consistent improvement across diverse verification tasks.

representation, the ground truth defines several auxiliary functions. In contrast, the LLM adopts a simpler abstraction—it models the entire bitmap as a single array, thereby eliminating the need for auxiliary functions and reasoning directly with Verus’ built-in APIs for the `Seq` type.

7 Related Work

Towards AI-assisted automated program verification, there is a long research line focusing on verifying the correctness of either an individual function [9, 20, 37–39] or the type invariant for a given module [29]. In contrast, this paper targets the verification of data structure modules, which requires the synthesis of mathematical abstractions, structural invariants, and specifications and proof blocks spanning multiple methods. Consequently, none of the previous work can be applied in our setting.

Traditional formal-methods tools (e.g., Frama-C [5], Why3 [7], Dafny [15], Coq [31], and Isabelle [21]) provide strong correctness guarantees but require developers to manually construct detailed specifications and proofs. Thus, they impose substantial annotation overhead and demand significant expertise. In contrast, VERISTRUCT leverages large language models to automatically synthesize annotations. As a result, VERISTRUCT bridges the gap between fully manual formal verification and prior LLM-based assistants, achieving greater automation without sacrificing rigor.

8 Conclusion and Future Work

This paper introduced VERISTRUCT, a novel framework that combines large language models with Verus to automatically verify Rust data-structure modules.

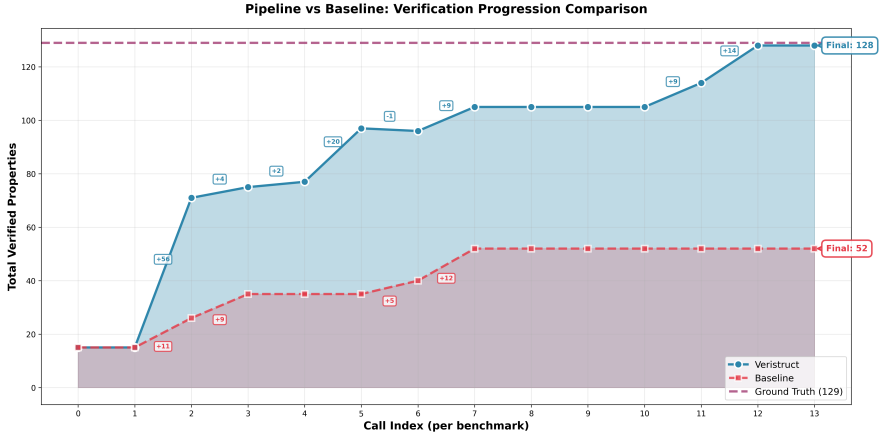


Fig. 7. Aggregated verification progress across all benchmarks. VERISTRUCT rapidly approaches the ground-truth total of 129 functions, while the single-shot baseline plateaus well below full coverage.

VERISTRUCT systematically generates `View` implementations, type invariants, specifications, and proof code under the scheduling of an extra planner module. To mitigate the issue that LLMs do not fully grasp Verus’ annotation syntax and verification-specific semantics, we not only inject syntax guidelines in the prompt, but also introduce an additional repair stage to fix errors.

We evaluated VERISTRUCT on eleven benchmarks drawn from the Verus GitHub repository. Across eleven benchmark modules, VERISTRUCT successfully verifies ten of them and verifies 128/129 (99.2%) of functions in our benchmark. VERISTRUCT shows a significant improvement over the baseline, demonstrating the framework’s effectiveness.

We conclude with several promising directions for future work. A natural next step is to extend VERISTRUCT to support more complex verification tasks. First, we could integrate retrieval-augmented generation (RAG) [17], which could help identify useful lemmas from the Verus standard library [14], strengthening proof synthesis. Second, we could additionally support the synthesis of Verus’ resource algebra library [12] for verifying complex concurrent data structures. Finally, applying constrained decoding techniques [8] could alleviate syntax errors in generated annotations.

Automatic Unit Test Generation. At present, our approach requires users to manually provide a comprehensive unit test suite. However, constructing such suites, especially those that achieve high coverage and include corner cases, can be both time-consuming and error-prone. Recent studies have shown that large language models are highly capable of generating high-quality test cases, including those targeting rare or edge conditions [18, 35]. Building on these advances, we plan to incorporate LLM-based unit test generation into our framework, thereby

Table 3. Detailed Statistics of VERISTRUCT, where #Funcs refers to the number of functions to be verified in each benchmark.

Benchmark	#Funcs	Time (minutes)	#LLM Calls	Solved?
ATOMICS	11	2.1	8	✓
BITMAP	14	12.2	13	✓
TREEMAP	21	0.3	6	✓
INVARIANTS	7	1.5	4	✓
NODE	12	10.4	12	✗ (11/12 Verified)
OPTION	15	2.6	3	✓
RINGBUFFER	13	4.2	12	✓
RWLOCKVSTD	5	6.4	3	✓
SETFROMVEC	10	9.1	13	✓
TRANSFER	5	1.2	4	✓
VECTORS	16	3.1	6	✓

further reducing manual effort and enhancing the completeness of the verification process.

Reinforcement Learning for Verification. Another promising direction is to apply reinforcement learning (RL) to specification inference and proof generation. The Verus verifier provides a natural reward signal: successful verification yields positive feedback, while failed verification attempts with diagnostic messages can guide policy updates. By training models to optimize for verification success, RL could enable more effective exploration of the specification and proof search space, potentially discovering annotation strategies that are difficult to obtain through prompting alone.

References

1. Artifact of veristruct (2025), <https://github.com/ChuyueSun/VeriStruct>
2. Anthropic: Claude code. <https://claude.ai/claude-code> (2025), accessed: 2025
3. Barrett, C., Boyd, B., Bursztein, E., Carlini, N., Chen, B., Choi, J., Chowdhury, A.R., Christodorescu, M., Datta, A., Feizi, S., Fisher, K., Hashimoto, T., Hendrycks, D., Jha, S., Kang, D., Kerschbaum, F., Mitchell, E., Mitchell, J., Ramzan, Z., Shams, K., Song, D., Taly, A., Yang, D.: Identifying and mitigating the security risks of generative ai. *Foundations and Trends in Privacy and Security* **6**(1), 1–52 (2023). <https://doi.org/10.1561/33000000041>, <http://dx.doi.org/10.1561/33000000041>
4. Barrett, C., Sebastiani, R., Seshia, S., Tinelli, C.: Satisfiability modulo theories. In: Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T. (eds.) *Handbook of Satisfiability, Second Edition*, *Frontiers in Artificial Intelligence and Applications*, vol. 336, chap. 33, pp. 825–885. IOS Press (Feb 2021), <http://theory.stanford.edu/~barrett/pubs/BSST21.pdf>

5. Cuoq, P., Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B.: Frama-C: A software analysis perspective. In: Eleftherakis, G., Hinchey, M., Holcombe, M. (eds.) *Proceedings of the 10th International Conference on Software Engineering and Formal Methods (SEFM 2012)*. Lecture Notes in Computer Science, vol. 7504, pp. 233–247. Springer, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33826-7_16
6. Dong, Q., Li, L., Dai, D., Zheng, C., Ma, J., Li, R., Xia, H., Xu, J., Wu, Z., Liu, T., et al.: A survey on in-context learning. arXiv preprint arXiv:2301.00234 (2022)
7. Filiâtre, J.C., Paskevich, A.: Why3. <https://why3.lri.fr> (2024), version 1.6.0
8. Fu, Y., Baker, E., Ding, Y., Chen, Y.: Constrained decoding for secure code generation. arXiv preprint arXiv:2405.00218 (2024)
9. Kamath, A., Senthilnathan, A., Chakraborty, S., Deligiannis, P., Lahiri, S.K., Lal, A., Rastogi, A., Roy, S., Sharma, R.: Finding inductive loop invariants using large language models. arXiv preprint arXiv:2311.07948 (2023)
10. Knuth, D.E.: *The Art of Computer Programming, Volume 1: Fundamental Algorithms*. Addison-Wesley (1997)
11. Lahiri, S.K.: Evaluating llm-driven user-intent formalization for verification-aware languages. In: *CONFERENCE ON FORMAL METHODS IN COMPUTER-AIDED DESIGN–FMCAD 2024*. p. 142 (2024)
12. Lattuada, A., Hance, T., Bosamiya, J., Brun, M., Cho, C., LeBlanc, H., Srinivasan, P., Achermann, R., Chajed, T., Hawblitzel, C., Howell, J., Lorch, J.R., Padon, O., Parno, B.: Verus: A practical foundation for systems verification. In: *Proceedings of the ACM SIGOPS 30th Symposium on Operating Systems Principles*. p. 438–454. SOSP '24, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3694715.3695952>, <https://doi.org/10.1145/3694715.3695952>
13. Lattuada, A., Hance, T., Cho, C., Brun, M., Subasinghe, I., Zhou, Y., Howell, J., Parno, B., Hawblitzel, C.: Verus: Verifying Rust programs using linear ghost types. *Proceedings of the ACM on Programming Languages* (2023). <https://doi.org/10.1145/3586037>
14. Lattuada, A., Parno, B., Bosamiya, J., Hawblitzel, C., Hance, T., et al.: vstd: Verus standard library. <https://github.com/verus-lang/verus/tree/main/vstd> (Jun 2024), version 0.0.0, MIT License
15. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2010)*. Lecture Notes in Computer Science, vol. 6355, pp. 348–370. Springer, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17511-4_20
16. Levy, A., Campbell, B., Ghena, B., Giffin, D.B., Pannuto, P., Dutta, P., Levis, P.: Multiprogramming a 64kb computer safely and efficiently. In: *Proceedings of the 26th Symposium on Operating Systems Principles*. p. 234–251. SOSP '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3132747.3132786>, <https://doi.org/10.1145/3132747.3132786>
17. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., Riedel, S., Kiela, D.: Retrieval-augmented generation for knowledge-intensive nlp tasks. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS '20, Curran Associates Inc., Red Hook, NY, USA (2020)
18. Liu, K., Chen, Z., Liu, Y., Zhang, J.M., Harman, M., Han, Y., Ma, Y., Dong, Y., Li, G., Huang, G.: Llm-powered test case generation for detecting bugs in plausible programs (2025), <https://arxiv.org/abs/2404.10304>

19. Miltner, A., Padhi, S., Millstein, T., Walker, D.: Data-driven inference of representation invariants. In: Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation. p. 1–15. PLDI 2020, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3385412.3385967>, <https://doi.org/10.1145/3385412.3385967>
20. Misu, M.R.H., Lopes, C.V., Ma, I., Noble, J.: Towards ai-assisted synthesis of verified Dafny methods. *Proc. ACM Softw. Eng.* **1**(FSE) (Jul 2024). <https://doi.org/10.1145/3643763>, <https://doi.org/10.1145/3643763>
21. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle. <https://isabelle.in.tum.de> (2024), version 2023
22. O’Hearn, P.W.: Separation logic. *Communications of the ACM* **62**(2), 86–95 (Feb 2019). <https://doi.org/10.1145/3211968>
23. OpenAI: Openai o1 system card. Technical report, OpenAI (2024), includes safety evaluations and red teaming results for the o1 and o1-mini models
24. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL ’05). pp. 247–258. Association for Computing Machinery, New York, NY, USA (2005). <https://doi.org/10.1145/1040305.1040327>
25. Perry, N., Srivastava, M., Kumar, D., Boneh, D.: Do users write more insecure code with AI assistants? In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS). pp. 2785–2799. ACM (2023). <https://doi.org/10.1145/3576915.3623157>
26. Poesia, G., Loughridge, C., Amin, N.: dafny-annotator: AI-assisted verification of Dafny programs. *CoRR abs/2411.15143* (2024), <https://arxiv.org/abs/2411.15143>, arXiv:2411.15143
27. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science (LICS 2002). pp. 55–74. IEEE Computer Society, Los Alamitos, CA, USA (2002). <https://doi.org/10.1109/LICS.2002.1029817>
28. Sun, C.E., Gao, S., Weng, T.W.: Breaking the barrier: Enhanced utility and robustness in smoothed drl agents. *ICML* (2024)
29. Sun, C., Agashe, V., Chakraborty, S., Taneja, J., Barrett, C.W., Dill, D.L., Qiu, X., Lahiri, S.K.: ClassInvGen: Class invariant synthesis using large language models. *CoRR abs/2502.18917* (2025), <https://arxiv.org/abs/2502.18917>
30. Sun, C., Sheng, Y., Padon, O., Barrett, C.: Clover: Closed-loop verifiable code generation. In: *AI Verification (SAIV 2024)*. Lecture Notes in Computer Science, vol. 14846, pp. 134–155. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-65112-0_7
31. The Coq Development Team: The Coq proof assistant. <https://coq.inria.fr> (2024), version 8.19.0
32. Verus Contributors: Verus tutorial and reference (2025), <https://verus-lang.github.io/verus/guide/>
33. Verus Contributors: vstd: Verus standard library api documentation (2025), <https://verus-lang.github.io/verus/verusdoc/vstd/>
34. Verus Project: Verilib. <https://verilib.org/>, accessed: 2025-02-17
35. Wang, Z., Liu, K., Li, G., Jin, Z.: Hits: High-coverage LLM-based unit test generation via method slicing. In: Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering. p. 1258–1268. ASE ’24, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3691620.3695501>, <https://doi.org/10.1145/3691620.3695501>

36. Wang, Z., Zhou, Z., Song, D., Huang, Y., Chen, S., Ma, L., Zhang, T.: Towards understanding the characteristics of code generation errors made by large language models. In: Proceedings of the 47th IEEE/ACM International Conference on Software Engineering (ICSE) (2025), to appear
37. Wen, C., Cao, J., Su, J., Xu, Z., Qin, S., He, M., Li, H., Cheung, S.C., Tian, C.: Enchanting program specification synthesis by large language models using static analysis and program verification. In: Computer Aided Verification (CAV 2024). Lecture Notes in Computer Science, vol. 14682, pp. 302–328. Springer (2024). https://doi.org/10.1007/978-3-031-65630-9_16
38. Wu, H., Barrett, C., Narodytska, N.: Lemur: Integrating large language models in automated program verification. In: The Twelfth International Conference on Learning Representations
39. Yang, C., Li, X., Misu, M.R.H., Yao, J., Cui, W., Gong, Y., Hawblitzel, C., Lahiri, S., Lorch, J.R., Lu, S., Yang, F., Zhou, Z., Lu, S.: AutoVerus: Automated proof generation for Rust code. CoRR [abs/2409.13082](https://arxiv.org/abs/2409.13082) (2024), <https://arxiv.org/abs/2409.13082>

Data Availability Statement. The source code of VERISTRUCT, including all prompts and the benchmark suite used in our evaluation, is publicly available at <https://github.com/ChuyueSun/VeriStruct>.

Open Access. This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits any noncommercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if you modified the licensed material. You do not have permission under this license to share adapted material derived from this chapter or parts of it.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

