


Evaluating LLM-driven User-Intent Formalization for Verification-Aware Languages

Shuvendu K. Lahiri 

Microsoft Research, Redmond, USA

shuvendu@microsoft.com

Abstract—Verification-aware programming languages such as Dafny and F* provide means to formally specify and prove properties of a program. Although the problem of checking an implementation against a specification can be defined mechanically, there is no algorithmic way of ensuring the correctness of the *user-intent formalization for programs* — that a specification adheres to the user’s intent behind the program. This is because intent or requirement is expressed *informally* in natural language and the specification is a formal artefact. However, the advent of large language models (LLMs) has made tremendous strides bridging the gap between informal intent and formal program implementations in the last couple of years, driven in large parts due to benchmarks and automated metrics to evaluate different techniques.

Recent work has developed a framework for evaluating and benchmarking the *user-intent formalization* problem for mainstream programming languages [12]. However, as we argue in this paper, such an approach does not readily extend to verification-aware languages that support rich specifications (using quantifiers and ghost variables) that cannot be evaluated through dynamic execution. Previous work also required generating program mutants using LLMs to create the benchmark. We advocate an alternate, perhaps simpler approach of *symbolically testing specifications* to provide an intuitive metric for evaluating the quality of specifications that can be easily instantiated with most verification-aware languages. We demonstrate that our automated metric agrees closely on a human-labeled dataset of Dafny specifications for the popular MBPP code-generation benchmark, yet demonstrates cases where the human labeling is not perfect. We also outline formal verification challenges that need to be addressed to apply the technique more widely. We believe our work provides a stepping stone to enable the establishment of a benchmark and research agenda for the problem of user-intent formalization for programs.

Index Terms—formal verification, specifications, large language models

I. INTRODUCTION

Formal verification is only as good as the specification it verifies. A formal specification unambiguously defines the (possibly partial) intent behind a program, often in an declarative manner. Although there has been decades of research in advancing the state-of-the-art in automating the problem of verifying an implementation against a specification, relatively less attention has been spent on how to aid the generation and evaluation of specifications or formal requirements. It is well-known that the lack of formal specifications is a significant impediment to deployment of formal verification in production code [10]. On the other hand, although a software is often

accompanied by informal intent expressed in natural language comments and API documentation, these intents are seldom formally enforced on the underlying implementation.

Large language models (LLMs) have recently demonstrated potential to bridge the gap between informal intent and formal artefact such as code [9], [5], by performing code generation from natural language. Such progress has been spurred in large parts through the creation of crowd-sourced benchmarks such as HumanEval [9] and Mostly Basic Python Programs (MBPP) [5] with automated program-semantics based metrics (such as tests), unlike NLP metrics such as BLEU scores [25]. For code generation, the quality of an implementation generated from informal intent is measured through the set of *hidden* validation tests. We term the tests as “hidden” since they are not available to the model at the time of code generation. The benchmarks also provide a reference code in Python for each problem, which are also hidden from the code generation model. These approaches rely heavily on the presence of a high-quality set of validation tests to exercise most corner cases. These benchmarks allow a community to evaluate their models and techniques on a common set of benchmarks automatically and measure progress. One can hope that the establishment of similar benchmarks and metrics for specification generation (decoupled from code generation) can enable synthesizing useful specifications from informal intent in practice.

Motivated by such a need, Endres *et al.* [12] describe the problem of generating formal declarative specifications (namely, method postconditions) from informal intent using LLMs and automated metrics for evaluating them — we refer to this as *user-intent formalization* problem in this paper. For user-intent formalization, Endres *et al.* re-purpose code-generation benchmarks (such as HumanEval) and introduce two metrics (a) *correctness* and (b) *completeness*, with respect to the set of (hidden) validation tests and the (hidden) reference code. Correctness captures that the specification satisfies the reference code for all the validation tests; completeness captures the strength of the specification to discriminate against buggy mutations of the reference code (under the set of tests). Further, these code mutants are generated through LLMs and are grouped by the subset of tests that fail them. They demonstrate that these automatic metrics closely resemble the quality of specifications as determined through manual analysis. As with code generation, the approach relies on the

presence of a high-quality test-suite.

A. Summary

In this paper, we investigate creating benchmarks (dataset and associated automatic metrics) for user-intent formalization for *verification-aware languages* such as F* [27], Dafny [21] and Verus [20]. A verification-aware language supports a rich program logic for expressing specifications and offers the ability to verify them statically using automated theorem provers. However, for most non-trivial programs, the verification requires manually decomposing the problem through the use of ghost variables, intermediate lemmas, invariants and assertions.

We focus on the problem of evaluating specifications automatically. We discuss why prior approaches do not readily apply to evaluating specifications. Then we propose the use of automatic program verification to symbolically “test” the specifications to determine their quality. We have developed a prototype and applied it to a dataset of Dafny specifications from prior work [23]. We demonstrate that the quality of the specifications obtained through the automated means aligns well with the human-labeling of the specifications in most cases. Finally, we describe the unique challenges (such as quantifier instantiation) that need to be addressed to apply the technique to a larger class of problems.

B. Running example

Consider a snippet of the JSON specification of an example from the MBPP dataset for code generation from natural language in Python [9].

```
"prompt": "Write a function to find the shared elements
from the given two lists.",
"code": "def similar_elements(test_tup1, test_tup2):\n
    ...",
"test_list": [
    "assert set(similar_elements((3, 4, 5, 6), (5, 7, 4, 10)
    )) == set((4, 5))",
    ...
]
```

The dataset comes with a set of problems, each containing a natural language prompt, the reference code as well as a set of 3 tests in `test_list`. Misu *et al.* [23] port these requirements to Dafny (call it MBPP-DFY), including slight change to the prompt, explicitly representing the Dafny method signature (with a slightly changed name), as well as the test cases [1].¹

```
"task_description": "Write a method in Dafny to find the
shared elements from the given two array.",
"method_signature": "method similarElements (arr1:array<int>
>, arr2:array<int>) returns (res: array<int>)",
"test_cases": {
    "test_1": "var a1:= new int[] [3, 4, 5, 6];\nvar a2:=
new int[] [5, 7, 4, 10];\nvar e1:= new int[] [4,
5];\nvar res1:=similarElements(a1,a2);\nassert
arrayEquals(res1,e1);",
    ...
}
```

¹We take the snapshot of the repository at COMMIT a57ce24.

Misu *et al.* use GPT-4 [4] and other language models (with sophisticated prompting) to generate a pair of Dafny specification and implementation, and retain generations where the generated specification is provable for the generated implementation. Below we show the specification component of the GPT-4 generated Dafny artefact for this example (with slightly altered method name and signature over JSON) [3]. The two postconditions (marked with `ensures`) define the specification and use an auxiliary predicate `InArray`.

```
predicate InArray(a: array<int>, x: int)
reads a
{ exists i :: 0 <= i < a.Length && a[i] == x }

method SharedElements(a: array<int>, b: array<int>)
    returns (result: seq<int>)
    ensures forall x :: x in result ==>
        (InArray(a, x) && InArray(b, x))
    ensures forall i, j :: 0 <= i < j < |result| ==>
        result[i] != result[j]
{
    ...Dafny implementation...
}
```

Although the implementation satisfies the specification, there is no evidence that the specification indeed captures what the user intended (and specified implicitly in the hidden test cases). The authors perform *manual* review of the specifications and mark them as either `{WRONG_SPEC, WEAK_SPEC, STRONG_SPEC}` to denote if the specification is respectively, inconsistent, weakly or strongly consistent with the intent expressed in the natural language and tests [2]. Therefore, although this work addresses generating verified Dafny programs from informal intent, it relies completely on a human to ensure that the specification matches the intent. In particular, the framework allows for the Dafny code to only satisfy a vacuous specification (such as `ensures true`) or worse, an incorrect one. This aspect renders this dataset unsuitable as an automated benchmark for specification generation in verification-aware languages. In other words, if a language model generates a specification that is non-equivalent to the ground-truth specification labeled by the user, we cannot determine if it is incorrect without requiring a user. This also makes the evaluation subjective. Since verification of (generated) Dafny code is only as good as the specification, this makes the dataset unsuitable for an automated benchmark for verified code generation as well.

C. Proposal: Symbolically testing specifications

Our objective in this work is to define metrics that can be automatically evaluated to determine the quality of a specification. At the same time, we would also like to establish that these metrics correspond well to what a manual labeling would establish. Further, we decouple the problem of specification generation from code generation, to be able to (a) use the specifications to find bugs in underlying or generated code [12], or (b) use the validated specifications to refine the prompts for code generation.

We first argue that the none of the prior approaches are readily applicable to yield an automatic metric: (i) First, running the tests on the implementation without specifications will

not rule out vacuous specifications. (ii) Second, one cannot apply the approach of Endres *et al.* for rich specifications that contain ghost state or constructs that cannot be evaluated at runtime (such as a universal/existential quantifier). (iii) One can ignore tests and attempt to statically verify the synthesized specification directly against the hidden reference code. This can still allow vacuous specifications such as `true`. Further, it is unlikely that such verification would be automatic for non-trivial specifications given the need to infer intermediate lemmas and invariants. (iv) One can instead provide a hidden reference specification and check the candidate specification for semantic equivalence. However, such a metric would be too strict as it would not be able to distinguish weak and vacuous specifications (such as `true`) from strong (yet incomplete) specifications.

Next, we propose a method for evaluating user-intent formalization for verification-aware languages that leverages the validation tests and symbolic verification capabilities of these languages. Given a set of tests consisting of a set of *consistent* input-output pairs, we define the correctness and completeness metrics purely over the tests, without the need for the reference code. Of course, just as in the case of code generation, the technique for determining specification quality assumes a high-quality set of validation tests.

Consider a method with signature $m(x):y$ denoting the name m , input parameters x and output parameters y , a candidate postcondition/summary specification $\phi(x, y)$ and a set of input-output tests T .

1) *Correctness*: A postcondition $\phi(x, y)$ is correct (or sound) with respect to T if it is consistent with all the input-output pairs in T . In other words, for each $(i, o) \in T$ the following Hoare-triple [15] holds.

$$\models \{\text{true}\} x := i; y := o; \{\phi(x, y)\}$$

2) *Completeness*: The completeness measure for a specification ϕ given T is the fraction of output mutations of the tests in T that ϕ is inconsistent with. Let $T_1 \doteq \bigcup_{(i,o) \in T} \bigcup_{o' \neq o} \{(i, o')\}$ be a finite set of mutations of T that mutate the output values for the given inputs. In this work, we restrict the set of mutants per input i to a fixed number (5 for this paper). Let $T_2 \subseteq T_1$ be the largest subset such that for each $(i, o') \in T_2$, the following Hoare-triple *does not* hold:

$$\not\models \{\text{true}\} x := i; y := o'; \{\phi(x, y)\}$$

Then the completeness measure of ϕ with respect to T is $|T_2|/|T_1|$ (this is inspired by kill-set in *mutation-testing* literature [16]). For interested readers, we also contrast these with a plausible proposal in Appendix B.

Let us demonstrate an implementation of these Hoare triples as Dafny programs for our running example. Consider the first test (see JSON input in Section I-B) that asserts that `SharedElements` should return the set $\{4, 5\}$ for input arrays $[3, 4, 5, 6]$ and $[5, 7, 4, 10]$. For correctness against this test, we create the following Dafny program by providing a definition of `SharedElements` that performs the soundness check:

```

predicate InArray(a: array<int>, x: int)
reads a
{ exists i :: 0 <= i < a.Length && a[i] == x }

method SharedElements(a: array<int>, b: array<int>)
returns (result: seq<int>)
ensures forall x :: x in result ==>
(InArray(a, x) && InArray(b, x))
ensures forall i, j :: 0 <= i < j < |result| ==>
result[i] != result[j]
{
var a1 := new int[] [3, 4, 5, 6];
var a2 := new int[] [5, 7, 4, 10];
assume {:axiom} a[..a.Length] == a1[..a1.Length];
assert a[0] == a2[0] && ... && a[3] == a2[3];
assume {:axiom} b[..b.Length] == a2[..a2.Length];
assert b[0] == a2[0] && ... && b[3] == a2[3];
result := [4, 5];
}

```

The Dafny program above is identical to the program in Section I-B, except for the body of `SharedElements` specified within the curly braces $\{\dots\}$. Instead of the original implementation of the method, we model the Hoare-triple for correctness described above. The precondition of the Hoare-triple (`true`) translates to `assume true` which is dropped. The input assignment $x := i$; is modeled as assignments of the input i to temporary variables `a1`, `a2`, followed by constraining the actual parameters `a`, `b` respectively. For Dafny, this amounts to saying that two arrays `a` (respectively, `b`) and `a1` (respectively, `a2`) are equal on all elements up to their lengths (which implies that the lengths are identical as well). The redundant asserts are needed for the verifier to trigger the quantifiers used in `InArray` to enable the proof. Finally, the output parameter `result` is assigned one of the expected values (`[5,4]` is also an acceptable value). The Dafny program symbolically checks (using Satisfiability Modulo Theories solvers [6]) that the specification (provided by the `ensures` statements) holds for the specific input and output.

However, the above verification only proves that the specification is *correct* for the test; a vacuous specification `ensures true` would also be verified. For completeness, we mutate the output value in `result` in several ways and check that the verification fails. In this case, mutating the value in `result` to `[6]` would fail the first postcondition, since `6` is not present in both the input arrays.

For this program, the above specification is marked as `STRONG_SPEC` by the authors of MBPP-DFY. However, our automated test harness coupled with mutation discovers that the specification is *incomplete* (score 0.6). When we check the above program with either of the two mutations `result := [4]` or `result := [5]`, the specification still verifies. This is because the implication `==>` in the first postcondition only checks that values in `result` is present in both the arrays; it does not check that all such common values are present in `result!` Thus our automated metric is able to assign the above specification a lower score than the full functional specification for this example where the `==>` is replaced by `<==>`.

Note: One may caution against using a verification failure (as used for the completeness checks) as a means to show that the corresponding Hoare-triple does not hold. In other words, given that deductive verifiers are sound, but not nec-

essarily complete, a verification failure is typically interpreted as an unknown outcome. However, note that we check for completeness for a specification ϕ against an input-output example (i, o') only if ϕ is *proved correct* for all the tests in T (including the original test (i, o)). Since the two verification conditions differ only in a concrete value (between o and o'), we conjecture that the underlying reasoning (including quantifiers instantiated) would be quite identical in most cases, and the verification failure strongly indicates that the Hoare-triple does not hold.

II. IMPLEMENTATION AND EVALUATION

We report ongoing work in implementing the two metrics for the MBPP-DFY dataset mentioned earlier in Section I-B. Our tool (a 400 line Python script) consumes the method signature, test cases and the candidate specification from the JSON and Dafny files, and creates Dafny programs for verifying the correctness and completeness for each test. It then invokes the Dafny verifier, and reports the aggregate correctness as well completeness score for each specification, averaged over the different test cases; it finally compares the metrics against the labels provided by authors. Of the 153 problems with specifications (a subset written by humans), we have managed to apply our tool to check 64 of the specifications (at the time of writing). In other words, for these 64 examples, our tool is able to verify the soundness of the specifications over the set of validation tests. We have released the scripts, dataset and outputs at the website <https://github.com/microsoft/nl-2-postcond>.

A. Mutating values

Each of these examples contain 3 test cases in Dafny format, and we consider up to 5 distinct mutants of the output values for each test. We currently have a simple mutation scheme for the output values. For Booleans, we flip the value between `true` and `false`. For integers, we choose a random value between 1 and 10, and randomly add or subtract it. For strings, we choose a random character and either replace one of the characters or append it to the string. For arrays (we only restrict to integer arrays and sequences), we choose between dropping an element or inserting a random value at a randomly chosen index. For our running example of `SharedElements`, this allows us to create a mutant test case with the array value `[4]` that demonstrates that the GPT-4 generated specification (marked as `STRONG_SPEC`) is not the most precise specification for this problem.

B. Results

We briefly report some details of the evaluation, and outline the challenges to handle the remaining examples.

We find that for large majority of examples, the manual labels align with the metrics computed by our tool. That is, for the incorrect specifications, our tool reports a verification failure for correctness check, and report a high completeness score (usually > 0.66) for precise specifications. In other words, barring a few exceptions below, all the specifications

that are marked `STRONG_SPEC` have a completeness score above 0.66, and all the `WEAK_SPEC` have a score below 0.66.

This provides evidence that the automated metrics serve as a good proxy for the user label. One such incorrect specification generated by GPT-4 (and correctly labeled) is the problem of `RemoveDuplicates` ("task_id" 572) (also the motivating example by Endres *et al.* [12] demonstrating ambiguity of natural language) where the task is to remove all elements with duplicates, but GPT-4 interprets it as the problem of retaining a single copy of each value. Similarly, for `countSubstrings` ("task_id" 61), that counts the number of substrings whose sum equals their length, our completeness check gives a low score to the specification (correctly labeled `WEAK_SPEC`) that only ensures count is a non-negative number.

In addition to the `SharedElements` ("task_id" 2), we found at least 2 more cases where the specification (labeled `STRONG_SPEC`) is weaker than the most precise one. These include the cases for `maxAbsDiff` ("task_id" 145), and `removeElements` ("task_id" 161). `maxAbsDiff` is expected to compute the maximum difference between any two elements in the list; the provided specification only ensures that the result upper bounds the difference between any two elements, but is not the precise bound. `removeElements` is expected to remove the elements in one array from another; the specification only ensures that any element in the result array is present in the first array and not the second but fails to ensure that the result array is the precise difference (can be ensured by changing `==>` by `<==>`).

Interestingly, we also discovered a few cases where the correctness check failed for specifications labeled `STRONG_SPEC`. These include examples with "task_id"s 234, 240, and 445. For these examples, the authors had accidentally introduced bugs while copying from the test cases in Python. For instance, for "task_id" 234 that computes the cube of a number, the authors mistakenly copied the value 25 instead of 125 for the cube of 5. We noticed that a couple of these errors have been fixed in the latest commit at the time of writing (fb4f53e), but a few still remain (e.g., 445 where the value 32 is replaced by 31).

Our preliminary results demonstrates that the automatic metrics not only helps to add objectivity to the manual labeling, argues for having a quantitative metric for completeness (instead of a Boolean `WEAK_SPEC` vs `STRONG_SPEC`), but also clearly demonstrates the potential to serve as a metric for evaluating specifications for a benchmark.

C. Limitations

Our implementation currently has a few limitations that precludes analyzing all the 153 specifications correctly. A large class of them stem from inaccurate parsing through a simple regex based parsing of the Dafny files. We expect these (e.g., handling of 2D arrays) to be addressed as we improve the parsing through Dafny AST-based analysis. However, a few fundamental challenges (related to failure to perform proofs of correct specifications) remain that requires further investigation. We note a couple of them here:

a) *Recursive predicates.*: Specifications for tasks such as 105 (counting number of true Booleans in an array) require the use of recursive functions `countTo` that expresses the count of an array in terms of `countTo` of the tail of the array. Dafny uses a notion of "fuel" to control the level of unfolding of such predicates, which suffices for inductive proofs. However, for concrete arrays in test cases, these recursive functions need to be unrolled proportional to the length of the array.

b) *Quantifier instantiation.*: Specifications for tasks such as 3 (that classifies a number as a non-prime) uses existential quantifiers. For example, checking a number n as non-prime requires checking for the existence of a smaller number k that divides n . Unlike their usage in inductive proofs where a loop establishes the inductive hypothesis to only instantiate the quantifier on a few expressions, testing a symbolic specification against a concrete value is sometimes non-trivial. For example, to show that the value 97 is a prime, the quantifier needs to be instantiated on all values up to 96.

We have already automated a few failed proofs due to quantifier instantiation by adding intermediate redundant assertions equating two arrays on every index (see `assert` statements in the body of `SharedElements` in Sec I-B). In future work, we plan to automate the inference of fuel amount and quantifier instantiation witnesses when ranging over the indices of an array.

III. RELATED WORK

Although the literature on specification mining is fairly rich, it has historically focused on techniques that infer specifications of an implementation through dynamic [13] and recently neural techniques [26]. The objective of these techniques is to learn invariants from a few runs that generalize to unseen executions. Static [14], [19], [24] and more recently neural approaches [28], [17] have also been applied to the problem (inductive) invariant generation to aid program proofs. Finally, recent works have explored the generation of programs, specifications and proofs from natural languages [23], [29], [22]. On the other hand, work on user-intent formalization (i.e., translating natural language comments to specifications) [7], [12] synthesize the intended specification, even in the absence of an implementation. The problem of generating *test oracle* assertions have also been investigated [11], but they only apply to single input or test prefix. TiCoder [18] introduced the term "user-intent formalization" over test cases for interactive code generation, and alluded to metrics for measuring quality of weak specifications (namely, tests) in terms of correctness (user acceptance) and completeness (prunes buggy codes). The work on *autoformalization* [30] for translating natural language comments to mathematical theorems is closest to our work. However, these techniques do not use semantic checks (such as tests and symbolic verification) to ensure the quality of the generated formal mathematical statements; instead they leverage textual similarity metrics such as BLEU.

IV. CONCLUSION

In this paper, we motivate the problem of evaluating user-intent-formalization for verification-aware languages. We demonstrate that the idea of symbolically testing specifications against validation tests can provide an automated metric for a benchmark. We plan to curate a benchmark with a large fraction of examples from MBPP-DFY dataset using the above metric. We hope such benchmarks will accelerate the research on specification generation from informal intent in verification-aware languages. This when coupled with work on program and proof synthesis (given a specification) can greatly lower the cost required to create formally verified modules in the future [8].

APPENDIX A EXTENDED VERSION OF THE JSON EXAMPLES

```
"prompt": "Write a function to find the shared elements
from the given two lists.",
"code": "def similar_elements(test_tup1, test_tup2):\n
    .....",
"test_list": [
    "assert set(similar_elements((3, 4, 5, 6), (5, 7, 4, 10)
    )) == set((4, 5))",
    "assert set(similar_elements((1, 2, 3, 4), (5, 4, 3, 7))
    ) == set((3, 4))",
    "assert set(similar_elements((11, 12, 14, 13), (17, 15,
    14, 13))) == set((13, 14))"
]

```

```
"task_description": "Write a method in Dafny to find the
shared elements from the given two array.",
"method_signature": "method similarElements (arr1:array<int
>, arr2:array<int>) returns (res: array<int>)",
"test_cases": {
    "test_1": "var a1:= new int[] [3, 4, 5, 6];\nvar a2:=
new int[] [5, 7, 4, 10];\nvar e1:= new int[] [4,
5];\nvar res1:=similarElements(a1,a2);\nassert
arrayEquals(res1,e1);",
    "test_2": "var a3:= new int[] [1, 2, 3, 4];\nvar a4:=
new int[] [5, 4, 3, 7];\nvar e2:= new int[] [3,
4];\nvar res2:=similarElements(a3,a4);\nassert
arrayEquals(res2,e2);",
    "test_3": "var a5:= new int[] [11, 12, 14, 13];\nvar
a6:= new int[] [17, 15, 14, 13];\nvar e3:= new
int[] [13, 14];\nvar res3:=similarElements(a5,a6)
;\nassert arrayEquals(res3,e3);"
}

```

APPENDIX B ALTERNATE PROPOSAL

A reader may wonder if the following check achieves a similar objective for checking completeness:

$$\models \{x == i \wedge \phi(x, y)\} \text{ skip } \{y == o\}$$

There are two issues with this formulation: (a) this only provides a Boolean metric that only rewards a precise specification that constrains y to a unique value o . For example, it would award the strong (yet imprecise) specifications of `SharedElements` to 0. Secondly, (b) this formulation is also too strong when the precise functional specification allows for non-determinism in the output. Consider the running example of `SharedElements`, where the specification allows for the output to be one of either $\{[4, 5], [5, 4]\}$ for the given inputs.

REFERENCES

- [1] mbpp-san-dfy-228-all-task-test.json. <https://github.com/Mondego/dafny-synthesis/blob/a57ce24/MBPP-san-DFY-228/mbpp-san-dfy-228-all-task-test.json>.
- [2] rq3-dynamic-few-shot-prompting-gpt-4-temp_0.5-verified-unverified-tagged.json. https://github.com/Mondego/dafny-synthesis/blob/a57ce24/RQs/RQ3-%5BDynamic-Few-Shot-Prompting%5D/rq3-dynamic-few-shot-prompting-GPT-4-temp_0.5-verified-unverified-tagged.json.
- [3] task_id_2.dfy. https://github.com/Mondego/dafny-synthesis/blob/a57ce24/MBPP-DFY-153/test/task_id_2.dfy.
- [4] Gpt-4 technical report, 2024.
- [5] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le, and C. Sutton. Program synthesis with large language models, 2021.
- [6] C. W. Barrett, L. M. de Moura, S. Ranise, A. Stump, and C. Tinelli. The SMT-LIB initiative and the rise of SMT - (HVC 2010 award talk). In S. Barner, I. G. Harris, D. Kroening, and O. Raz, editors, *Hardware and Software: Verification and Testing - 6th International Haifa Verification Conference, HVC 2010*, volume 6504 of *Lecture Notes in Computer Science*, page 3. Springer, 2010.
- [7] A. Blasi, A. Goffi, K. Kuznetsov, A. Gorla, M. D. Ernst, M. Pezzè, and S. D. Castellanos. Translating code comments to procedure specifications. In *Proceedings of the 27th ACM SIGSOFT international symposium on software testing and analysis*, pages 242–253, 2018.
- [8] S. Chakraborty, G. Ebner, S. Bhat, S. Fakhoury, S. Fatima, S. Lahiri, and N. Swamy. Towards neural synthesis for smt-assisted proof-oriented programming, 2024.
- [9] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- [10] D. Craigen, S. Gerhart, and T. Ralston. Formal methods reality check: industrial usage. *IEEE Transactions on Software Engineering*, 21(2):90–98, 1995.
- [11] E. Dinella, G. Ryan, T. Mytkowicz, and S. K. Lahiri. Toga: A neural method for test oracle generation. In *Proceedings of the 44th International Conference on Software Engineering*, pages 2130–2141, 2022.
- [12] M. Endres, S. Fakhoury, S. Chakraborty, and S. K. Lahiri. Can large language models transform natural language intent into formal method postconditions? In *Proceedings of the Foundations of Software Engineering 2024 (FSE'24)*, 2024.
- [13] M. D. Ernst, J. Cockrell, W. G. Griswold, and D. Notkin. Dynamically discovering likely program invariants to support program evolution. In *Proceedings of the 21st international conference on Software engineering*, pages 213–224, 1999.
- [14] C. Flanagan and K. R. M. Leino. Houdini, an annotation assistant for `esc/java`. In J. N. Oliveira and P. Zave, editors, *FME 2001: Formal Methods for Increasing Software Productivity*. Springer Berlin Heidelberg, 2001.
- [15] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, oct 1969.
- [16] Y. Jia and M. Harman. An analysis and survey of the development of mutation testing. *IEEE transactions on software engineering*, 37(5):649–678, 2010.
- [17] A. Kamath, A. Senthilnathan, S. Chakraborty, P. Deligiannis, S. K. Lahiri, A. Lal, A. Rastogi, S. Roy, and R. Sharma. Finding inductive loop invariants using large language models, 2023.
- [18] S. K. Lahiri, A. Naik, G. Sakkas, P. Choudhury, C. von Veh, M. Musuvathi, J. P. Inala, C. Wang, and J. Gao. Interactive code generation via test-driven user-intent formalization. *arXiv preprint arXiv:2208.05950*, 2022.
- [19] S. K. Lahiri, S. Qadeer, J. P. Galeotti, J. W. Vounq, and T. Wies. Intra-module inference. In A. Bouajjani and O. Maler, editors, *Computer Aided Verification*, pages 493–508, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [20] A. Lattuada, T. Hance, C. Cho, M. Brun, I. Subasinghe, Y. Zhou, J. Howell, B. Parno, and C. Hawblitzel. Verus: Verifying rust programs using linear ghost types. *Proc. ACM Program. Lang.*, 7(OOPSLA1):286–315, 2023.
- [21] K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In E. M. Clarke and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 16th International Conference, LPAR-16*, volume 6355 of *Lecture Notes in Computer Science*, pages 348–370. Springer, 2010.
- [22] C. Loughridge, Q. Sun, S. Ahrenbach, F. Cassano, C. Sun, Y. Sheng, A. Mudide, M. R. H. Misu, N. Amin, and M. Tegmark. Dafnybench: A benchmark for formal software verification, 2024.
- [23] I. M. Md Rakib Hossain Misu, Cristina V. Lopes and J. Noble. Towards ai-assisted synthesis of verified dafny methods, 2024.
- [24] O. Padon, J. R. Wilcox, J. R. Koenig, K. L. McMillan, and A. Aiken. Induction duality: primal-dual search for invariants. *Proc. ACM Program. Lang.*, 6(POPL):1–29, 2022.
- [25] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, page 311–318, USA, 2002. Association for Computational Linguistics.
- [26] K. Pei, D. Bieber, K. Shi, C. Sutton, and P. Yin. Can large language models reason about program invariants? 2023.
- [27] J. Protzenko, J. K. Zinzindohoue, A. Rastogi, T. Ramananandro, P. Wang, S. Zanella-Béguelin, A. Delignat-Lavaud, C. Hrițcu, K. Bhargavan, C. Fournet, and N. Swamy. Verified low-level programming embedded in f^* . In *22nd International Conference on Functional Programming (ICFP 2017)*. ACM SIGPLAN, May 2017.
- [28] X. Si, A. Naik, H. Dai, M. Naik, and L. Song. Code2inv: A deep learning framework for program verification. In S. K. Lahiri and C. Wang, editors, *Computer Aided Verification - 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21-24, 2020, Proceedings, Part II*, volume 12225 of *Lecture Notes in Computer Science*, pages 151–164. Springer, 2020.
- [29] C. Sun, Y. Sheng, O. Padon, and C. Barrett. Clover: Closed-loop verifiable code generation, 2024.
- [30] Y. Wu, A. Q. Jiang, W. Li, M. N. Rabe, C. Staats, M. Jamnik, and C. Szegedy. Autoformalization with large language models, 2022.