



# Cobblestone: A Divide-and-Conquer Approach for Automating Formal Verification

**Saketh Ram Kasibatla**

UC San Diego  
San Diego, California, USA  
skasibatla@ucsd.edu

**Arpan Agrawal**

University of Illinois  
Urbana-Champaign  
Urbana-Champaign, Illinois, USA  
arpan.agrawal94@gmail.com

**Yuriy Brun**

University of Massachusetts  
Amherst, Massachusetts, USA  
brun@cs.umass.edu

**Sorin Lerner**

Cornell University  
Ithaca, New York, USA  
sorin.lerner@cornell.edu

**Talia Ringer**

University of Illinois  
Urbana-Champaign  
Urbana-Champaign, Illinois, USA  
tringer@illinois.edu

**Emily First**

Rutgers University  
New Brunswick, New Jersey, USA  
emily.first@rutgers.edu

## Abstract

Formal verification using proof assistants, such as Coq, is an effective way of improving software quality, but requires significant effort and expertise. Machine learning can automatically synthesize proofs, but such tools are able to prove only a fraction of desired software properties. We introduce COBBLESTONE, a divide-and-conquer approach for proof synthesis. COBBLESTONE uses a large language model (LLM) to generate potential proofs, uses those proofs to break the problem into simpler parts, automatically identifies which of those parts were successfully proven, and iterates on the remaining parts to build a correct proof that is guaranteed to be sound, despite the reliance on unsound LLMs. We evaluate COBBLESTONE on four benchmarks of open-source Coq projects, controlling for training data leakage. Fully automatically, COBBLESTONE outperforms state-of-the-art non-LLM tools, and proves many theorems that other LLM-based tools cannot, and on many benchmarks, outperforms them. Each COBBLESTONE run costs only \$1.25 and takes 14.7 minutes, on average. COBBLESTONE can also be used with external input, from a user or another tool, providing a proof structure or relevant lemmas. Evaluated with such an oracle, COBBLESTONE proves up to 58% of theorems. Overall, our research shows that tools can make use of partial progress and external input to more effectively automate formal verification.

### ACM Reference Format:

Saketh Ram Kasibatla, Arpan Agrawal, Yuriy Brun, Sorin Lerner, Talia Ringer, and Emily First. 2026. Cobblestone: A Divide-and-Conquer Approach for Automating Formal Verification. In *2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26)*, April 12–18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3744916.3773178>

## 1 Introduction

Bugs in software systems can be costly and dangerous. In 2022, poor software quality cost the US economy \$2.41 trillion [35], and

bugs can bring down critical, global systems [45]. Formal verification using proof assistants, such as Coq [68] or Lean [14], is a promising method of improving software quality. Formal verification can mathematically prove the absence of entire classes of bugs, providing strong guarantees for the correctness of critical software systems. And formal verification is highly effective: A study [78] of C compilers found bugs in every tested compiler, including LLVM [38] and GCC [65], but not in the formally verified (in Coq) portions of CompCert [40].

But, formal verification requires specifying desired properties, writing mathematical proofs of the properties, and machine checking those proofs using the proof assistant. Writing these proofs requires significant expertise and manual effort. For example, the proofs verifying CompCert are 8 times longer than the functional code [39], and even small changes to the software can require heavy proof editing [57]. While hundreds of large software systems have been verified [57], like the sel4 microkernel [34, 49] and the CakeML compiler [36], and formal verification has seen industrial success, e.g., at Airbus France [64], Google, and Mozilla [16, 29], most software today is not verified due to the high manual cost.

Recent work has aimed to reduce the cost of formal verification by using machine learning to synthesize proofs [18–20, 53, 60, 61, 76]. However, these approaches, even when combined with an SMT-solver-based hammer [11] can only prove one third of the desired properties on a large benchmark of open-source Coq projects [18]. The advent of large language models (LLMs) has improved performance in this space [43, 66, 70].

In this paper, we aim to occupy the happy medium between two common approaches to proof synthesis. Some work uses LLMs to generate **whole proofs** at once, without using the intermediate proof state produced by partial proofs [20, 81]. For instance, PALM [43] samples an LLM for a whole proof attempt and then uses repair mechanisms. The success of whole-proof-based approaches relies on the powerful reasoning abilities of proprietary LLMs, such as OpenAI and Anthropic models, as they must construct proofs that follow viable high-level plans. However, these approaches fail to provide the LLM with access to the intermediate state, which contains important context to aid in reasoning.

Other approaches perform **tactic-by-tactic** search, building a proof one step at a time and allowing the LLM to condition the



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICSE '26, Rio de Janeiro, Brazil*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2025-3/26/04  
<https://doi.org/10.1145/3744916.3773178>

next step generation based on access to the intermediate state [6, 19, 30, 60, 76, 77]. However, navigating the complex search space of possible proofs is challenging since the local state steers the search without consideration for a high-level proof plan.

These approaches represent two extremes on the spectrum of proof synthesis granularity. Our work is a middle ground between these two extremes, incorporating the advantages of whole-proof generation with leveraging intermediate states where this generation fails. To achieve this middle ground, we make a fundamental insight: we can leverage the tree structure of whole-proof attempts to localize valid and invalid *subproofs*. We can keep valid subproofs and recurse on points of failure where we regain access to the state. This insight motivated the design of our tool COBBLESTONE, a novel LLM-based approach to proof synthesis.

We present COBBLESTONE, which uses a **divide and conquer** algorithm to combine an LLM’s high-level planning ability with the advantages of local exploration. COBBLESTONE starts by tackling the entire goal, automatically identifying working and broken parts of the proof, reducing the granularity of its attack from the whole proof to individual proof steps *by recursing*, if necessary.

COBBLESTONE uses an LLM to generate whole proofs, analyzes the proofs using its novel fail-safe execution to decompose the proof into subgoals, and determines which subgoals’ subproofs work and which fail. It then attempts to repair the failing subproofs, and, if necessary, recurses on the remaining, unproven subgoals to build a whole, correct proof. Because COBBLESTONE divides the goal into smaller subgoals, it is more effective at using automated theorem proving tools such as hammers; while a hammer may fail to prove the overall goal, it is more likely to succeed on smaller, decomposed goals. Unlike prior work POETRY [71], COBBLESTONE does not require fine-tuning an LLM to identify proof substructure and locations to recurse, making it easily extensible to use different LLMs (requiring only minor changes to API calls).

We implement COBBLESTONE for the Coq proof assistant, and evaluate it on GPT-4 and two other LLMs. We evaluate on subsets of four benchmarks: CoqGym [76], a set of open-source Coq projects from GitHub used for evaluating prior proof synthesis tools [18, 19, 60, 61, 76]; and coq-wigderson [55], coq-bb5 [50], and PnVRocqLib [56], three recent projects deliberately chosen to be after GPT-4’s training cutoff to control for the impact of pretraining data leakage on LLMs.

We compare COBBLESTONE to prior non-LLM-based proof synthesis tools, LLM-based tools, and LLM baselines of our own creation. Prior RNN-based proof synthesis tool, Proverbot9001 [60] proves 17% of the CoqGym subset and 10% of the coq-wigderson subset, while SMT-solver-based CoqHammer [11] proves 30% and 27%, respectively. We implement a baseline LLM-based approach that uses chain-of-thought reasoning and show that it proves 25% and 19%, respectively. Meanwhile, COBBLESTONE, automatically proves 48% of the CoqGym subset and 38% of the coq-wigderson subset. When prior LLM-based tool PALM [43] is run multiple times to match the token usage of our approach, COBBLESTONE outperforms it on the coq-wigderson, coq-bb5, and PnVRocqLib subsets, and is highly

complementary with it on the CoqGym subset. Therefore, COBBLESTONE complements prior work by proving different theorems they do not, so that together, they prove significantly more.

While completely automated, COBBLESTONE can also use external input, such as proof structure plans or relevant lemmas from other tools or a proof engineer, further improving proof-synthesis success. Combining automated COBBLESTONE with versions that have access to a proof structure and relevant lemmas can prove 58% of the CoqGym subset and 55% of the coq-wigderson subset.

The main contributions of our work are:

- COBBLESTONE, a novel LLM-based, divide-and-conquer proof-synthesis approach that uses partial successes of failing proofs to produce whole, correct proofs. This serves as a middle-ground strategy between whole-proof and tactic-by-tactic search.
- *Fail-safe mode*, a method for executing a proof to localize proof errors and separate working and failing parts of a proof.
- A COBBLESTONE evaluation on four Coq benchmarks and a comparison to state-of-the-art proof synthesis tools showing that COBBLESTONE consistently proves theorems that prior work cannot, often outright proving more than them.
- A data-based exploration of how external input from proof engineers or other tools can further improve COBBLESTONE’s success as part of interactive, semi-automated proof-synthesis.

To ensure reproducibility of our results and enable others to build on our work, we make all code, experimental scripts, and data publicly available [33].

## 2 The COBBLESTONE Approach

To formally verify software, in addition to writing the code for a system, an engineer needs to write a mathematical theorem formalizing the desired property, and a high-level proof (proof script) that the property holds. COBBLESTONE automatically writes the proof script. Given code and a theorem (e.g., `in_adj_exists` in Figure 2), COBBLESTONE synthesizes a proof using a divide and conquer approach, as shown in Figure 1. First, COBBLESTONE tries synthesizing a proof using CoqHammer (Section 2.3). If that fails, COBBLESTONE samples an LLM for candidate solutions (Section 2.4). If none of the candidates prove the theorem entirely, COBBLESTONE uses a novel fault localization method to isolate errors while noting which subgoals are successfully proven (Section 2.5). Next, COBBLESTONE keeps successful parts of candidate solutions and recurses to find proofs for the yet unproven subgoals (Section 2.6), ultimately combining them into a proof of the original theorem.

While COBBLESTONE is fully automated, it has extension points (labeled “API access point for external information” in Figure 1) that can be used to provide useful additional information. Section 3.5 will evaluate how this external information can increase COBBLESTONE’s proving power.

### 2.1 Illustrative Example

We illustrate our approach using a real-world theorem `in_adj_exists`.<sup>1</sup> Figure 2 lists a definition (lines 1–4), the theorem statement (lines 6–8), and a human-written proof (lines 9–15). This theorem states that for all graphs  $g$  and all nodes  $i, j$ , if  $i$  is adjacent to  $j$  in  $g$ , then  $g$ ’s structure must contain  $v$ .

<sup>1</sup>from `graph.v` in the `coq-wigderson` project

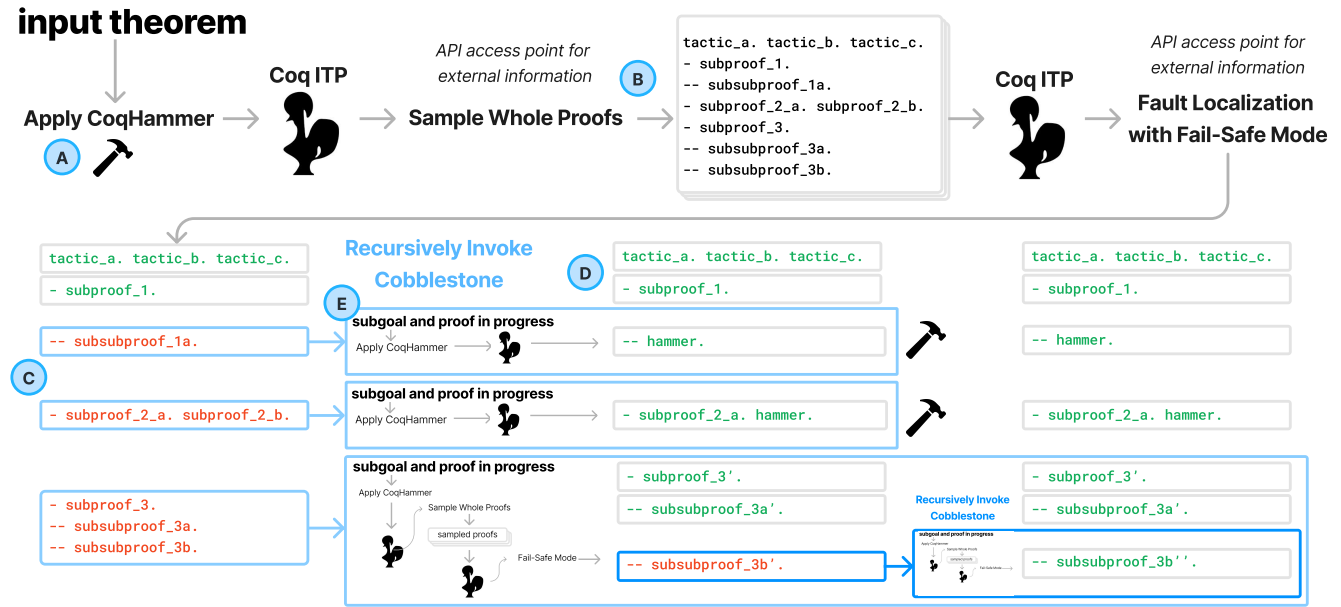


Figure 1: Given a software property to prove, COBBLESTONE first attempts to (A) use CoqHammer to prove that property. If that fails, it (B) generates a set of whole proofs, (C) localizes errors in those proofs, (D) extracts the working parts, and (E) recurses on the remaining unproven subgoals to assemble correct proofs. “API access points” provide external information as discussed in Section 3.5

```

1 Definition adj (g: graph) (i: node) : nodeset :=
2   match M.find i g with
3   | Some a => a
4   | None => S.empty end.
5
6 Lemma in_adj_exists : forall g i j,
7   S.In i (adj g j) ->
8     exists v, M.find j g = Some v /\ S.In i v.
9 Proof.
10  intros g i j H.
11  unfold adj in *.
12  destruct M.find eqn:E in *.
13  - hammer.
14  - rewrite SP.FM.empty_iff in H. contradiction.
15 Qed.

```

Figure 2: A sample theorem of the in\_adj\_exists property and its human-written proof (lines 9–15).

Coq proofs consist of high-level commands called *tactics*, like “try induction on n” or “try simplifying.” Tactics transform the *proof state*, which contains goals (statements that must be proven) and assumptions that can be used to prove the goals. The proof state starts with the initial theorem to be proven. When tactics transform the proof state to contain no more goals, the Qed command completes the proof.

Figure 3a shows the proof state before executing the “destruct M.find eqn:E in \*” tactic on line 12 of Figure 2, which performs case analysis. The proof state consists of a single goal and several assumptions, including H, which states that i is a member of adj g j.<sup>2</sup> Executing the tactic results in the proof state in Figure 3b.

<sup>2</sup>The unfold tactic on line 11 inlines adj in the proof state.

```

a
g: graph
i: S.elts
j: node
H: S.In i (
match M.find j g
with
| Some a => a
| None => S.empty
end)
(1/1)
exists v : nodeset,
M.find j g = Some
v /\
S.In i v

b
g: graph
...
E: M.find j g = Some n
H: S.In i n
---
(1/2)
exists v : nodeset, Some n = Some v
/\ S.In i v
=====
...
H: S.In i S.empty
---
(2/2)
exists v : nodeset, None = Some v
/\ S.In i v

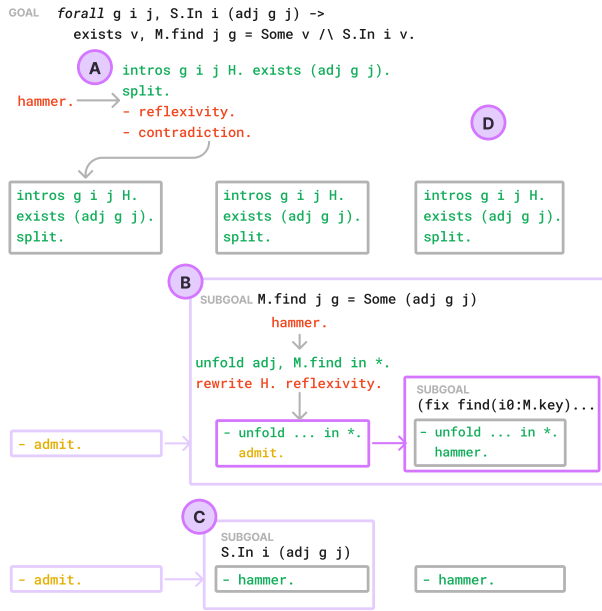
```

Figure 3: The proof states (a) before and (b) after executing destruct on line 12 of Figure 2.

The new proof state *decomposes* the preceding single goal into two *subgoals*. Each subgoal replaces “M.find” in the preceding goal with one of its two return values—“Some v” or “None”. If both subgoals are proven, then the original theorem is true.

The rest of the proof script is split into two *subproofs*, one for each subgoal, each starting with a bullet (-). The first bullet on line 13 invokes CoqHammer using the hammer tactic. CoqHammer [11] attempts to automatically prove a goal using SMT solvers and proof reconstruction procedures. If hammer succeeds, the goal is proven. Otherwise, the tactic produces an error.

While hammer is able to find a proof for the first subgoal, it has limitations. For example, CoqHammer cannot perform induction, a



**Figure 4: Proving `in_adj_exists` with COBBLESTONE.** (A) `hammer` fails, and an LLM-generated proof contains errors, resulting in recursive calls for each subgoal. (B) The first subgoal is proven using a new LLM completion and `hammer`. (C) The second subgoal is proven using `hammer`. (D) The final correct proof.

key step in many Coq proofs. In this case, `hammer` cannot prove the second subgoal. Instead, after the bullet on line 14 of Figure 2, the proof uses the `rewrite` tactic, which takes advantage of an already proven lemma (`SP.FM.empty_iff`) to change the hypothesis `H` to enable a proof by contradiction. After executing the rest of this proof script, which successfully proves the theorem, the proof state will contain no more goals, indicating that the theorem is proven.

## 2.2 COBBLESTONE on the Illustrative Example

We now demonstrate how COBBLESTONE proves `in_adj_exists` automatically. Figure 4 displays the steps COBBLESTONE takes to divide the proof into smaller pieces and prove each of them, resulting in a correct proof. First, COBBLESTONE attempts to prove `in_adj_exists` in one go, trying `hammer`, which fails, and then prompting an LLM to generate whole proofs, one of which is shown in (A). While the LLM’s proof starts off with reasonable structure, it contains errors (highlighted in red). Normally, Coq will not be able to identify all the errors; Section 2.5 describes COBBLESTONE’s novel fault localization mechanism. Though the proof in (A) is not correct as-is, it does decompose the initial goal into two subgoals. It may still be a useful part of a correct proof if these subgoals can be proven. However, its correctness can only be determined *retrospectively* after proving these subgoals.

Instead of naively trying to prove `in_adj_exists` again, as prior work has tried [20], COBBLESTONE recurses on each subgoal, attempting to prove it instead (B and C). In (B), COBBLESTONE tries using a `hammer`, which fails, and then prompts an LLM for a proof of the subgoal “`M.find j g = Some (adj g j)`”. This

results in a proof with an error, which COBBLESTONE temporarily fixes by adding the `admit` tactic. The proof with `admit`, a *proof in progress*, is passed as an argument to another recursive call. Here, COBBLESTONE changes the proof in progress to a working proof by replacing `admit` with the `hammer` tactic. In (C), COBBLESTONE is able to prove “`S.In i (adj g j)`” using only `hammer`.

With successful subproofs for each subgoal in hand, COBBLESTONE combines parts of the proofs in (A), (B), and (C) to form a correct proof of the whole theorem, shown in (D). Notably, by using LLM completions and its divide-and-conquer approach, COBBLESTONE can divide and simplify the goals, often reaching subgoals simple enough for `hammer` to prove. This allows the final proof to integrate portions of two separate LLM completions (for different subgoals), and two invocations of `hammer`, into a single proof.

## 2.3 Applying CoqHammer

COBBLESTONE combines neural-based proof synthesis and SMT-based proof synthesis by invoking CoqHammer [11]. In Figure 1, the `hammer` icon (🔨) shows how COBBLESTONE uses CoqHammer in two ways. First, at the start of COBBLESTONE’s execution, COBBLESTONE invokes CoqHammer on the goal, which sometimes produces a working proof. Second, in recursive calls, COBBLESTONE may be passed a *proof in progress*, which consists of working tactics followed by an “`admit`”. For example, “`unfold adj, M.find in *. admit.`” in Figure 4 (B) is a proof in progress. COBBLESTONE replaces `admit` with `hammer`, and checks if this proves the subgoal. Attempting to repair proofs in progress like this allows COBBLESTONE to use CoqHammer on a variety of goals, increasing the frequency of success.

## 2.4 Sampling LLM for Whole Proofs

Central to COBBLESTONE’s proof synthesis is using an instruction-tuned LLM [52] to sample whole proofs. COBBLESTONE’s prompt consists of two parts—a system message with high-level directions about the task the LLM should perform, and a user message<sup>3</sup> with details specific to the current theorem.

The system message directs the LLM to produce a whole proof for the provided theorem. The user message contains 5 types of information: (1) the theorem statement to prove, (2) the current proof state, (3) the definitions for all identifiers mentioned in the theorem statement and proof state that are not in the standard library, (4) contextual information, and (5) optional reasoning, which is used as part of chain-of-thought prompting. The user message is formatted as below, with each section preceded by a section header. COBBLESTONE constructs the user message fully automatically using the `coq-serapy` library [60] to interact with Coq.

```

[CURRENT THEOREM]
{{theorem statement from Coq ITP}}

[PROOF CONTEXT]
{{current proof context from Coq ITP}}

[DEFINITIONS]
{{definitions from Coq ITP}}

[OTHER PROVEN THEOREMS]
{{proven theorems from Coq ITP and optional oracle}}
  
```

<sup>3</sup>“user message” is terminology used by the OpenAI API. COBBLESTONE generates these prompts automatically, without a user.

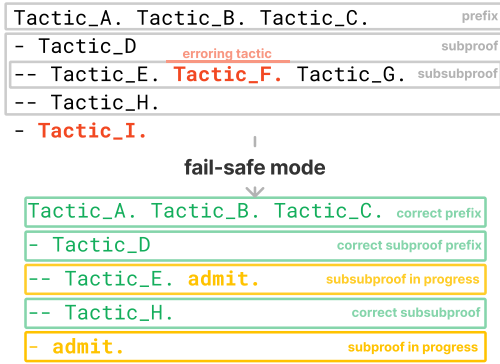


Figure 5: An execution of fail-safe mode on an input proof

```

[REASONING]
{{optional chain-of-thought reasoning}}

```

The contextual information in the prompt can either be empty or contain the theorem statements that were successfully proven before the given theorem in the file. If the information in the prompt is longer than the LLM’s token limit, the contextual information is left-truncated, removing information from the file that is further from the theorem statement. COBBLESTONE samples the LLM with this prompt, which results in candidate whole proof scripts.

Chain-of-thought [74], a popular prompting technique, prompts an LLM to provide natural language reasoning before generating its final response, and has been shown to help improve LLM performance. COBBLESTONE implements chain-of-thought by first prompting the LLM with a modified system message that instructs it to generate reasoning, and then prompting it a second time, with the generated reasoning added to the [REASONING] section of the user message. When COBBLESTONE samples an LLM without chain-of-thought prompting, this extra section is omitted.

Prior work has shown that the use of diverse data in models can increase the proving power of a proof-synthesis approach [18]. Accordingly, to increase variation in COBBLESTONE’s generated proofs, COBBLESTONE samples the LLM with 4 variations of the same prompt in 2 different dimensions—with and without contextual information and with and without chain-of-thought reasoning. COBBLESTONE benefits from this approach because the Coq theorem prover can serve as an oracle of proof correctness to pick a correct proof from among multiple samples. Our replication package [33] includes all prompts that were used as part of our evaluation.

## 2.5 Localizing Errors with Fail-Safe Mode

Once COBBLESTONE samples proofs from the LLM, it executes each of them in Coq. If a proof proves the goal, COBBLESTONE returns it as-is. If none of the sampled proofs are successful, COBBLESTONE localizes the errors in each proof, identifying which parts are correct, and which parts are not. Executing a generated proof naïvely using the theorem prover stops on the first error, which then misses correct parts that occur after that error. Instead, we create *fail-safe mode*, a novel way of executing a proof in the theorem prover. This approach localizes errors by recursing over the proof’s structure,

```

Proof →      Prefix Subproof*
Prefix →     Tactic+
Subproof →   Bullet Proof
Tactic →     split. | hammer. | ...
Bullet →     - | -- | --- | ...

```

Figure 6: The formal structure of a proof with bullets

making small edits to allow the execution of proof sections even after an error, and annotating sections as “correct” and “in progress.” These annotations are implemented via a parallel data structure that keeps track of which subproofs are correct and which are in progress.

Fail-safe mode relies on the fact that many proofs (43% of proofs from the ChainOfThought evaluation from Section 3.2) contain *subproofs*, sections of the proof organized using bullets. A proof consists of a *prefix* followed by 0 or more *subproofs*, each of which consists of a *bullet* followed by a nested proof. The last tactic in the prefix decomposes a single goal into multiple subgoals, one for each subproof (recall Section 2.1). Each subproof’s bullet focuses Coq on a single subgoal in the proof state; the proof that follows proves that subgoal. Figure 6 formalizes this structure.

Note that the definition of subproof is recursive. Each subproof can contain subproofs of its own (subsubproofs). In all figures, we indicate the level of nesting of the subproof by repeating the “-” character in bullets (i.e. “-” for a subproof, “--” for a subsubproof, etc.). A proof with no bullets can be expressed solely using a prefix and no subproofs. This means that all possible proofs can be expressed using this structure. (Section 3.6 will discuss fail-safe mode’s limitations in localizing errors in proofs consisting of only a prefix.)

### Algorithm 1 Fail-safe mode

```

Require: A prefix  $p$ , and 0 or more subproofs  $sps$ 
Ensure: the modified proof runs without errors, and the most deeply nested subproofs
           with errors are annotated as “in progress”
1: procedure RUNFAILSAFE( $p, sps$ )
2:   for all tactics  $t$  in  $p$  do                                     ▶ First, run  $p$ 
3:     Run( $t$ )
4:     if  $t$  results in an error then
5:       replace  $t$  and everything after it with admit.
6:       annotate  $p$  as “in progress”
7:       return
8:     annotate  $p$  as “correct”
9:   for all subproofs  $sp$  in  $sps$  do                               ▶ then, run  $sps$ 
10:     RunFailSafe( $sp.prefix, sp.subproofs$ )

```

Algorithm 1 describes fail-safe mode, which takes a proof script broken into prefixes and subproofs as input, modifies it to allow execution after errors, and annotates subproofs, subsubproofs, etc. as “correct” or “in progress.” First, it executes each tactic in the prefix. If a tactic results in an error, that tactic and everything after it is replaced with `admit`<sup>4</sup>, and the prefix is annotated as “in progress.”

If the prefix executes without an error, COBBLESTONE marks it as correct and fail-safe mode recurses on each of the subproofs, annotating them in the same way. Figure 5 shows the result of running fail-safe mode. Tactics that run successfully are colored

<sup>4</sup>Note that inserting `admit`s into proofs in progress is temporary; COBBLESTONE will remove these in a recursive call.

green, tactics that produce errors are colored red, and admits are colored yellow as they are temporary. The boxes around prefixes and subproofs in the output are annotations, colored yellow if a section is in progress and green if it is correct. Note that COBBLESTONE marks the subsubproof starting with `Tactic_E` as in progress, but not the subproof starting with `Tactic_D` that contains it. Because annotations occur as part of a depth-first search, fail-safe mode localizes errors to the most deeply nested subproof possible.

In addition to the description in Algorithm 1, fail-safe mode also handles two complications that arise when executing LLM-generated proofs. To understand these complications, consider the example proof that an LLM could generate below and to the left.

```
Proof.
  intros.
  destruct a.
  apply nonexistent.
  - assumption.
  - auto.
Qed.

Proof.
  intros.
  destruct a.
  - admit.
  - admit.
Qed.
```

The first complication is that the proof can encounter an error while executing a prefix tactic *after* it decomposes the original goal. Concretely, “`destruct a`” might successfully decompose the goal into subgoals, but then, “`apply nonexistent`” can fail. In this case, instead of just replacing the “`apply nonexistent`” with an `admit`, fail-safe mode searches for a shorter prefix that leads to multiple subgoals, which we call a *decomposing prefix*. For example, if the sequence “`intros. destruct a.`” leads to multiple subgoals, then fail-safe mode keeps it as the prefix and creates a bullet with an `admit` placeholder for each subgoal, as shown on the right.

The second complication is that a decomposing tactic can produce more or fewer subgoals than the number of bullets that follow. On the right, there are two bullets that follow “`destruct a`”. If “`destruct a`” produces fewer subgoals than two, fail-safe mode trims the number of bullets; if “`destruct a`” produces more than two, fail-safe mode adds the appropriate number of bullets with `admit` placeholders. This guarantees that the number of subproofs that appear syntactically in the proof matches the number of subgoals generated by the prefix.

## 2.6 Recursively Invoking COBBLESTONE

After executing fail-safe mode on each of the four sampled proofs, COBBLESTONE produces four annotated proofs like the one in Figure 5. This example contains two subproofs. The first one (starting with `Tactic_D`) has two subsubproofs. One is in progress, and the other succeeds. The second subproof is also in progress.

The overall theorem is not considered proven until COBBLESTONE generates and swaps in correct replacements for *all* the subproofs in progress for one of the four annotated proofs. To do this, COBBLESTONE loops over the annotated proofs. For each proof, COBBLESTONE makes one recursive call per subgoal in progress. It passes the subgoal and its corresponding subproof in progress (which contains `admit`) to the recursive call. Note that when fail-safe mode annotates a deeply nested subproof as in progress (e.g., the subsubproof that starts with `Tactic_E`), it only recurses on the corresponding subgoal, not on any of the subgoals containing it (e.g., the subgoal corresponding to `Tactic_D`). If none of the proofs are able to break the goal into subgoals, COBBLESTONE calls itself again, retrying with the original theorem as its goal. We

limit this search procedure using two hyperparameters—*maximum depth*, which limits the level of recursion, and *maximum invocations*, which limits the total number of times COBBLESTONE can be invoked. These settings are discussed further in Section 3.1.

Through this greedy depth-first recursion, COBBLESTONE uses divide-and-conquer, iteratively simplifying the goal and recursing on simpler subgoals, increasing the probability of success. This way, COBBLESTONE breaks down the problem of proving the original theorem into smaller chunks, and leverage progress made in the previous iterations to assist the subsequent ones.

## 3 Evaluation

We evaluate COBBLESTONE on 4 datasets of theorems from open-source Coq projects. We compare COBBLESTONE to 5 prior tools and two LLM baselines of our own creation. Our evaluation answers 5 research questions:

- RQ1:** How does COBBLESTONE compare to state-of-the-art proof synthesis methods?
- RQ2:** How much does CoqHammer contribute to COBBLESTONE’s performance?
- RQ3:** How much does COBBLESTONE’s search strategy contribute to its performance?
- RQ4:** How does external information affect COBBLESTONE’s performance?
- RQ5:** How do theorems COBBLESTONE proves and fails to prove differ?

### 3.1 Experimental Setup

**Benchmarks.** We construct our evaluation benchmarks from CoqGym’s test set of theorems [76], and from three other open-source projects—`coq-wigderson` [55], `coq-bb5` [50], and `PnVRocqLib` [56].

CoqGym is a widely used benchmark for evaluating proof-synthesis tools [18, 19, 61, 76], comprised of 68,501 theorems and their associated human-written proofs across 124 projects. CoqGym’s test set consists of 26 projects with 12,161 theorems, including `verdi-raft` (distributed software correctness), `goedel` (Gödel’s 1st incompleteness theorem), and `tree-automata` (verified algorithms). Evaluating using this dataset allows for a more fair comparison to prior tools, and assessing tool efficacy across a diverse set of projects. However, CoqGym was released in 2019, and its projects existed on GitHub before then. Because the pre-training cutoff for GPT-4 is September 2021, GPT-4 has potentially seen CoqGym in its pre-training.

To address this concern, we also evaluate using `coq-wigderson` [55], `coq-bb5` [50], and `PnVRocqLib` [56]. The first commit from each of these projects (March 2022, April 2024, and August 2024, respectively) is after the GPT-4 pre-training cutoff. `coq-wigderson` is a formal verification of Wigderson’s graph coloring algorithm. It consists of 174 theorems. `PnVRocqLib` contains 810 theorems about mathematics, including first order logic, Hilbert logic, and boolean algebra. `coq-bb5` contains 1446 theorems about busy-beaver values.

Because of the cost of performing inference on LLMs, full-scale evaluations on tens of thousands of theorems are impractical. In fact, the cost of just the LLM usage for our evaluation, including ablation studies, exceeded US\$8K. Instead, we create four 100-proof subsets of the four benchmark projects, sampling theorems at random. We call these subsets `CoqGym100`, `Wigderson100`, `PnVRocqLib100`,

and BB5100. To keep the cost of ablations low, we evaluate on PnVRocqLib100 and BB5100 only for RQ1 (Section 3.2). The total size of this dataset (400 theorems), is in line with other prior work which uses LLMs for proof synthesis [66, 71, 72], and 300 of the 400 theorems are published after GPT-4’s training cutoff.

**Comparisons to State-of-the-Art and Baselines.** We compare COBBLESTONE to 5 prior proof-synthesis tools—(1) Proverbot9001 [60], an RNN-based neural theorem prover; (2) CoqHammer [11], which uses SMT solvers and proof reconstruction procedures; (3) Tactician [37], which uses an online k-nearest-neighbor algorithm to select tactics during proof search; (4) PALM [43], which leverages retrieval augmentation (RAG), custom repair prompts, and a backtracking algorithm; and (5) Rango [70], which fine-tunes an LLM and uses RAG to include relevant lemmas and proof scripts. Proverbot9001 outperforms other ML-based proof synthesis tools that do not use LLMs, including ASTactic [76], TacTok [19], Diva [18], and Passport [61]. Proverbot9001 and CoqHammer prove 19.8% and 26.6% of theorems, respectively, on CoqGym, and prove 17% and 30% on CoqGym100. In a prior “informal comparison” on a subset of CoqGym, Tactician outperformed Proverbot9001 (except on one project) and sometimes outperformed CoqHammer [6].

For CoqHammer, we use the Z3 [12], CVC4 [4], Vampire [26], and E [63] SMT solvers, along with the default timeout settings (a 20 second prover timeout and a 5 second reconstruction timeout), following prior work [6, 60]. We evaluate Tactician’s ability to produce a proof fully automatically within a 10-minute timeout. Unlike Tactician, Proverbot9001 uses depth limits rather than a time limit, but usually finished its search within 10 minutes as well. We run PALM with gpt-4-0613 as its base model. One run of PALM on CoqGym100 used 3.7K tokens per theorem, 15 times less than the amount for COBBLESTONE (59K tokens per theorem). To correct for this, we run PALM 15 times for each dataset. Rango is run with the same setup noted in the paper—using a fine-tuned model with temperature 1.0 and a 10 minute timeout. Rango is run on an NVIDIA RTX 2080 GPU, along with a CPU with 16GB of RAM for proof checking. We do not compare against COPRA [66], another LLM-based tool, as it is too cost intensive to run in addition to our other experiments. In the COPRA paper, the authors only ran on a subset of theorems in CompCert that Proverbot9001 is not able to prove, citing budgetary constraints and high cost of evaluation.

We also compare against 2 new baselines, ChainOfThought, and LinearRegen. ChainOfThought prompts gpt-4-0613 20 times (the same as COBBLESTONE) in exactly the same way as COBBLESTONE does (using preceding lemmas and chain-of-thought prompting as described in Section 2.4) but only checks the proofs for correctness without further processing. LinearRegen also prompts gpt-4-0613 up to 20 times in exactly the same way as COBBLESTONE, but regenerates the proof after its first failure point. When a proof has a failure, LinearRegen keeps the prefix of the proof that succeeds, and continues from the last successful proof state.

**Metrics.** In line with prior evaluations [18–20, 61], we use two metrics: *success rate* and *added value*. The success rate of a tool is the fraction of all theorems the tool is able to prove. The added value of tool X over tool Y is the number of new theorems X proves that Y does not, divided by the number of theorems Y proves. To measure the number of theorems a tool can prove in a fixed number

of LLM samples, we define the *proven@k* metric. This metric is the sum of the *pass@k* [7] values for each proof in a dataset, and represents the expected number of theorems proven with k samples.

**Computing Resources.** Throughout our evaluation, we use OpenAI’s GPT-4 model (gpt-4-0613), which has a context length of 8,192 tokens. In Section 3.2, we also use GPT-3.5-turbo (gpt-3.5-turbo-0125) and Claude 3 Opus (claude-3-opus-20240229). GPT-3.5-turbo has a context length of 16,385. Claude 3 Opus has a context length of 200,000 tokens, but to manage inference cost, we capped its token limit at 10,000. We chose these 3 models because their APIs support function calling, a key part of our implementation. We call GPT-3.5-turbo and GPT-4 through the OpenAI Python library. We use the official Anthropic library to call Claude 3 Opus. For each call to an LLM, we sample with temperature 1.0. All experiments are run on machines with Intel Xeon Gold 6230 CPUs and 125GB of memory. We have access to 10 cores from these CPUs, as our experiments were run in a virtualized environment.

We run COBBLESTONE with a maximum depth of 5, and invoke it up to 5 times. Because each COBBLESTONE invocation samples an LLM 4 times, COBBLESTONE uses up to 20 LLM samples in total.

## 3.2 RQ1: How Does COBBLESTONE Compare to State-of-the-Art Proof Synthesis Methods?

Figure 7 shows success rates for non-LLM tools (CoqHammer, Proverbot9001, and Tactician), our newly created baselines (ChainOfThought and LinearRegen), LLM-based tools (PALM and Rango), and COBBLESTONE. Tactician and Rango cannot run on some datasets (reported with N/A) which use unsupported older Coq versions.

COBBLESTONE proves 48% of theorems on CoqGym100 and 38% of theorems on Wigderson100. It outperforms each of the non-LLM baselines as well as the *union* of these baselines, All non-LLM, adding 54.5% additional value in CoqGym100, and 25.0%, 40.0% and 14.6% added value in Wigderson100, PnVRocqLib100, and BB5100 respectively. Across all datasets, COBBLESTONE proves 46 theorems All non-LLM cannot, All non-LLM proves 14 theorems that COBBLESTONE cannot, and both prove an overlap of 127 theorems.

On CoqGym100, running COBBLESTONE costs \$1.46 and takes in 15.32 minutes, with successful runs costing an average of \$0.23 and completing in 2.54 minutes. On Wigderson100, running COBBLESTONE costs \$1.03 and completes in 14.21 minutes, and successful runs cost on average \$0.15 and complete in 3.66 minutes.

COBBLESTONE proves more theorems than ChainOfThought and LinearRegen in all datasets, adding value of up to 100.0% over ChainOfThought and up to 104.8% over LinearRegen. COBBLESTONE is also more token-efficient than LinearRegen. On CoqGym100 and Wigderson100, COBBLESTONE uses far fewer tokens (59.3K vs. 99.3K on CoqGym100 and 42.0K vs. 72.3K on Wigderson100). On PnVRocqLib100, COBBLESTONE uses 51.8K, vs. LinearRegen’s 83.7K, and on BB5100, COBBLESTONE uses 47.8K vs. LinearRegen’s 80.6K.

COBBLESTONE also consistently proves theorems that other LLM-based tools cannot in all datasets, outrightly proving more theorems than these tools in many cases. COBBLESTONE proves more theorems than a single run of PALM (PALM 1x in Figure 7). It also outperforms our PALM benchmark (run 15x) on Wigderson100, PnVRocqLib100, and BB5100, proving 8%, 9% and 2% more theorems (the difference in success rates) respectively. Furthermore,

	success added		success added		success added		success added	
	rate	value	rate	value	rate	value	rate	value
	CoqGym100		Wigderson100		PnVRocqLib100		BB5100	
CoqHammer	30%	66.7%	27%	44.4%	23%	91.3%	27%	59.3%
Proverbot9001	17%	194.1%	10%	280.0%	12%	275.0%	23%	87.0%
Tactician	N/A	N/A	13%	200.0%	23%	108.7%	22%	113.6%
All non-LLM	33%	54.5%	32%	25.0%	35%	40.0%	41%	14.6%
ChainOfThought	25%	92.0%	19%	100.0%	25%	84.0%	31%	48.4%
LinearRegen	31%	61.3%	21%	104.8%	33%	48.5%	22%	104.5%
PALM 1x	34%	52.9%	16%	143.8%	16%	175.0%	23%	104.3%
PALM	52%	17.3%	30%	40.0%	35%	42.9%	41%	26.8%
PALM + COBBLESTONE	61%		42%		50%		52%	
Rango	N/A	N/A	N/A	N/A	30%	56.7%	43%	25.6%
Rango + COBBLESTONE	N/A	N/A	N/A	N/A	47%		54%	
COBBLESTONE	48%		38%		44%		43%	
All Together	63%		48%		59%		61%	

**Figure 7: COBBLESTONE’s performance vs. baselines and other proof-synthesis tools. The “added value” columns report the percent of theorems COBBLESTONE proves over each of the other tools.**

CoqGym100						
	GPT-4		GPT-3.5		Claude	
	added value	added value	added value	added value	added value	added value
CoqHammer	30%	66.7%	30%	40.0%	30%	66.0%
ChainOfThought	25%	92.0%	13%	200.0%	28%	78.6%
ChainOfThought +CoqHammer	37%	35.1%	35%	25.7%	41%	24.4%
COBBLESTONE	48%		37%		49%	
Wigderson100						
	GPT-4		GPT-3.5		Claude	
	added value	added value	added value	added value	added value	added value
CoqHammer	27%	44.4%	27%	29.6%	27%	25.9%
ChainOfThought	19%	100.0%	8%	312.5%	22%	50.0%
ChainOfThought +CoqHammer	31%	25.0%	28%	25.0%	31%	12.9%
COBBLESTONE	38%		33%		32%	

**Figure 8: Performance of COBBLESTONE and ChainOfThought implemented with different LLMs. “added value” is COBBLESTONE’s value added relative to each row.**

in all datasets, COBBLESTONE proves many theorems that PALM cannot. COBBLESTONE alone prove 47 theorems, PALM alone proves 32 theorems, and both can prove an overlap of 126 theorems. On CoqGym100 and Wigderson100, COBBLESTONE uses a similar number of tokens to PALM, with an average of 59.3K vs. PALM’s 58.1K on CoqGym100 and 42.0K vs. PALM’s 41.8K on Wigderson100. On PnVRocqLib100, COBBLESTONE uses fewer tokens on average (51.8K vs. PALM’s 62.0K), while using more on BB5100 (47.8K vs. PALM’s 40.7K). COBBLESTONE and PALM’s complementarity indicates that COBBLESTONE may benefit from techniques such as PALM’s retrieval augmentation (RAG) and LLM-based proof repair.

COBBLESTONE proves the same number of theorems as Rango on BB5100, and significantly more on PnVRocqLib100. In both cases,

	CoqGym100		Wigderson100	
	success rate	added value	success rate	added value
CoqHammer	30%		27%	
ChainOfThought	25%		19%	
ChainOfThought $\cup$ CoqHammer	37%	48.0%	31%	63.2%
COBBLESTONE-NoHammer	25%		16%	
COBBLESTONE-NoHammer $\cup$ CoqHammer	38%	52.0%	30%	87.5%
COBBLESTONE	48%	96.0%	38%	137.5%

**Figure 9: CoqHammer’s use contributes significantly to COBBLESTONE and ChainOfThought. “added value” is the value of adding CoqHammer.**

COBBLESTONE and Rango also demonstrate significant complementarity. On BB5100, their union proves 54%, and on PnVRocqLib100, it proves 47%, more than either tool alone. Across all datasets, COBBLESTONE alone proves 28 theorems, Rango alone proves 15, and they both prove an overlap of 55 theorems. This indicates that COBBLESTONE may also benefit from leveraging Rango’s RAG. COBBLESTONE and Rango’s token consumption is not straightforward to compare, as Rango uses a 1.3 billion parameter LLM, which is not directly comparable to gpt-4-0613, a much larger model.

Since COBBLESTONE uses an LLM as a black box, it works with any LLM that supports tool use. We modified COBBLESTONE and ChainOfThought to use GPT-3.5-turbo and Claude 3 Opus, as shown in Figure 8. For all 3 models, COBBLESTONE outperforms ChainOfThought  $\cup$  CoqHammer. However, LLM choice can affect performance relative to GPT-4. GPT-3.5-turbo underperforms, both in ChainOfThought and in COBBLESTONE. Claude 3 Opus overperforms in ChainOfThought, with mixed results in COBBLESTONE.

**RA1:** COBBLESTONE outperforms non-LLM based proof generation methods and additional baselines. COBBLESTONE also performs comparably with LLM-based methods, complementing their results by proving theorems they cannot, and outperforming them in many cases. COBBLESTONE can be easily adapted to use LLMs other than GPT-4.

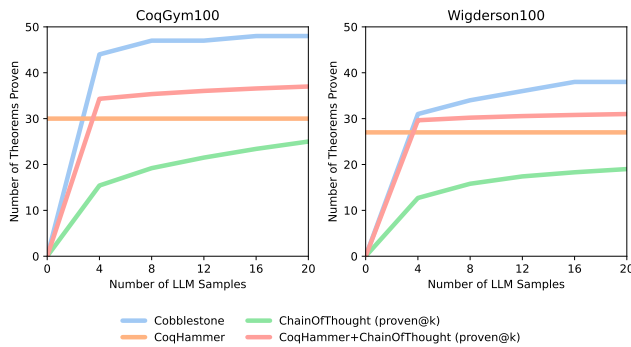
### 3.3 RQ2: How Much Does CoqHammer Contribute to COBBLESTONE’s Performance?

Figure 9 shows that CoqHammer plays an important role. On CoqGym100, COBBLESTONE-NoHammer (an ablated version of COBBLESTONE) proves 25% of theorems, while COBBLESTONE proves 48%. On Wigderson100, 16% vs. 38%. On both benchmarks, CoqHammer helps COBBLESTONE prove 96.0% and 137.5% more theorems.

Even running CoqHammer once improves ChainOfThought’s success rate—from 22% to 36% for CoqGym100, and from 17% to 31% for Wigderson100. However, COBBLESTONE outperforms this simpler combination of the two. Recall that COBBLESTONE uses CoqHammer in *each* invocation. Another simple combination is to invoke CoqHammer once per theorem, succeeding if it or COBBLESTONE-NoHammer succeeds. COBBLESTONE also outperforms this approach (“COBBLESTONE-NoHammer  $\cup$  CoqHammer”). On

	Wigderson100		CoqGym100	
	success rate	added value	success rate	added value
CoqHammer	27%		30%	
ChainOfThought	19%		25%	
Proverbot9001	10%		17%	
Tactician	13%		N/A	
All Prior Together	34%		39%	
TacticByTactic-NoHammer	8%	0.0%	11%	2.6%
COBBLESTONE-NoHammer	16%	0.0%	25%	7.7%
TacticByTactic	33%	11.8%	40%	20.5%
COBBLESTONE	38%	17.6%	48%	30.8%

**Figure 10: COBBLESTONE outperforms a simpler search strategy, both with and without CoqHammer. “added value” is the value each tool adds over “All Prior Together”**



**Figure 11: Theorems proven vs. LLM samples. The y values for “ChainOfThought” and “CoqHammer+ChainOfThought” use proven@k, detailed in Section 3.1.**

CoqGym100, COBBLESTONE adds 96.0% value over COBBLESTONE-NoHammer, compared to the simpler approach’s 52%. On Wigderson100, COBBLESTONE adds 137.5%, vs. the simpler approach’s 87.5%.

**RA2:** CoqHammer and COBBLESTONE are significantly complementary. CoqHammer contributes to COBBLESTONE’s performance, but COBBLESTONE outperforms CoqHammer and other simple combinations of the two. Even without CoqHammer, COBBLESTONE proves theorems other methods cannot.

### 3.4 RQ3: How Much Does COBBLESTONE’s Search Strategy Contribute to Its Performance?

COBBLESTONE’s search strategy is influenced by three key design decisions which we examine in this research question—(1) its use of whole-proof completions (2) its use of a recursive divide-and-conquer approach, and (3) its use of a sampling temperature of 1.0.

COBBLESTONE’s use of whole-proof completions sets it apart from many other proof synthesis systems. By contrast, most neural theorem provers synthesize proofs using tactic-by-tactic search, predicting the next tactic and using a tree search [18, 19, 60, 61, 76].

To measure the effectiveness of whole-proof completions, we implement another tool called TacticByTactic that, like COBBLESTONE, uses hammer (one for each search step), but only uses GPT-4 to

	CoqGym100		Wigderson100	
	success rate	added value	success rate	added value
ChainOfThought-temp0	10%		9%	
ChainOfThought	25%	150.0%	19%	111.1%
COBBLESTONE-temp0	48%		32%	
COBBLESTONE	48%	0.0%	38%	25.0%

**Figure 12: A temperature of 1 may improve performance, and does not harm performance. “added value” is value of changing from temperature 0 to 1.**

predict the next tactic at each search step. We run TacticByTactic, and TacticByTactic-NoHammer, a variant that does not use hammer, allowing up to 20 tactic predictions, a maximum proof depth of 20, and 3 attempts to predict the next tactic at each step. As each prediction attempt results in one LLM sample, these tools use the same number of samples as COBBLESTONE.

Figure 10 shows that without using CoqHammer, TacticByTactic-NoHammer underperforms all prior tools on Wigderson100 and CoqGym100. With calls to CoqHammer, TacticByTactic performs much better, proving 33% and 40% of the theorems in Wigderson100 and CoqGym100, respectively, and outperforms COBBLESTONE-NoHammer. Still, COBBLESTONE outperforms TacticByTactic and adds more value both in CoqGym100 and Wigderson100.

The second distinguishing feature of COBBLESTONE’s search strategy is its use of a divide-and-conquer approach. As described in Section 2.6, when a proof does not work, COBBLESTONE invokes itself recursively to prove *subgoals* of the original theorem, rather than retrying on the theorem itself. Figure 11 shows how this recursive approach contributes to COBBLESTONE’s performance.

One important datapoint to observe is what happens with a single invocation (which only uses 4 samples). To observe this datapoint, one should look at Figure 11, at the 4 point on the X axis. Even with a single invocation, COBBLESTONE proves more theorems than the combination of CoqHammer and ChainOfThought can in 20 samples. This is thanks to COBBLESTONE’s error localization (Section 2.5) and careful application of CoqHammer (Section 2.3).

This datapoint is also interesting to observe because it is an ablation that runs fail-safe mode without recursively invoking COBBLESTONE. This ends up being similar to PALM’s [43] backtracking algorithm (which also does not make recursive calls).

Invoking COBBLESTONE recursively further increases the number of theorems it can prove. On CoqGym100, COBBLESTONE is able to prove 44 out of 48 (91.6%) theorems within 1 invocation, proving an additional 4 (9.4%) by recursing. On Wigderson100, it is able to prove 31 out of 38 (81%) theorems within one invocation, proving 7 (19%) more by recursing. One possible explanation for this is test leakage in CoqGym100, which would make more samples of the original theorem “close enough” for COBBLESTONE’s to apply CoqHammer and error localization to produce a proof.

The third distinguishing feature of COBBLESTONE’s search strategy is its use of temperature 1.0 when sampling whole proofs. COBBLESTONE benefits from sampling diverse proofs (recall Section 2.4). A temperature of 1.0 contributes to this diversity. To examine the benefits of a higher temperature, we ran COBBLESTONE and ChainOfThought with the temperature set to 0. As shown in Figure 12, in both Wigderson100 and CoqGym100, ChainOfThought-temp0 proves a subset of the theorems that ChainOfThought is able to

	Wigderson100		CoqGym100	
	success rate	added value	success rate	added value
COBBLESTONE	38%		48%	
COBBLESTONE-PerfPremis	43%	21.1%	50%	12.5%
COBBLESTONE-PerfDecomp	52%	42.1%	47%	14.6%
All 3 COBBLESTONE versions	55%	44.7%	58%	20.8%

**Figure 13: External information in terms of relevant lemmas (COBBLESTONE-PerfPremis) and a breakdown of the proof into subgoals (COBBLESTONE-PerfDecomp) significantly improves COBBLESTONE’s proving power. “added value” is the value each tool adds over COBBLESTONE.**

prove. In Wigderson100, COBBLESTONE proves more theorems than COBBLESTONE-temp0, with a value added of 25%, and in CoqGym100, it proves the same number of theorems.

**RA3:** COBBLESTONE’s search strategy outperforms prior tools and an LLM-based tactic-by-tactic search. A single invocation of COBBLESTONE outperforms a combination of CoqHammer and ChainOfThought, and COBBLESTONE proves even more theorems when invoked recursively. A temperature of 1.0 can also improve performance.

### 3.5 RQ4: How Does External Information Affect COBBLESTONE’s Performance?

COBBLESTONE can use external information from another tool or a human. We created two oracles to provide such information (via the “API access points” in Figure 1). For each theorem, the *perfect premises* oracle (PerfPremis) knows which proven lemmas the human-written proof uses (e.g., with “rewrite” as described in Section 2.1). These are provided in addition to existing context in the [OTHER PROVEN THEOREMS] section of the prompt in Section 2.4. The *perfect decomposition* oracle (PerfDecomp) provides a decomposing prefix (recall Section 2.5) from the human-written proof, breaking the theorem into subgoals for recursive calls. Using the PerfPremis oracle is equivalent to asking “What lemmas are relevant to proving this theorem?” and using the PerfDecomp oracle is equivalent to asking “How would you break down this theorem into smaller goals?”

We evaluate COBBLESTONE with access to these oracles, calling the resulting variants COBBLESTONE-PerfPremis and COBBLESTONE-PerfDecomp. COBBLESTONE-PerfDecomp has access to both PerfPremis and PerfDecomp oracles. Figure 13 shows that COBBLESTONE-PerfPremis outperforms COBBLESTONE on both datasets, and COBBLESTONE-PerfDecomp sometimes does even better. Interestingly, there is some complementarity to the variants, with each proving some theorems none of the others do. Together, the variants prove 55% of CoqGym100 and 58% of Wigderson100. The additional information provided to COBBLESTONE by PerfDecomp does not always result in more theorems proven. COBBLESTONE-PerfPremis and COBBLESTONE are able to prove 6 theorems in CoqGym100 and 3 theorems in Wigderson100 that COBBLESTONE-PerfDecomp does not. This indicates that running COBBLESTONE several times, with different kinds of information in each run, may improve performance.

**RA4:** External information, such as useful already-proven lemmas or a decomposition of the theorem into subgoals can significantly increase COBBLESTONE’s proving power, suggesting a promising direction for future research.

### 3.6 RQ5: How do theorems COBBLESTONE proves and fails to prove differ?

To better understand where COBBLESTONE succeeds, we examine proofs that COBBLESTONE and its oracle-augmented variants generate. To better understand when it fails, we also examine human-written proofs of some theorems it cannot prove. All proofs referenced are available in our replication package[33].

**Successes:** COBBLESTONE’s successful proofs range from 1 to 24 tactics long, with the shortest consisting of a single invocation to CoqHammer. Successful proofs with length greater than 1 have an average length of 8.3 tactics, with bullet depth ranging from 1 (no bullets) to 4 (3 levels of bullets). These proofs can consist of 1 to 9 distinct CoqHammer invocations or LLM samples. Of the proofs that decomposed their goal, prefixes ranged from 1 to 4 tactics long, with an average of 1.6. Of these prefixes, 41% decomposed using `split`, 36% used `induction`, 18% used `destruct`, and 14% used `apply`.

The projects in our benchmarks broadly fall into 2 categories—mechanizations of mathematical proofs and formalizations of software. CoqGym100 consists of 12 mathematical projects and 8 software projects with 36 and 64 theorems respectively. COBBLESTONE proves 61.1% of CoqGym100’s mathematical theorems and 40.6% of its software theorems. `coq-wigderson` is a software project, and `coq-bb5` and `PnVRocqLib` are mathematical projects. Including these, our benchmarks contain 236 mathematical and 164 software theorems, with COBBLESTONE proving 46.2% and 39.0% respectively.

Together, COBBLESTONE-PerfPremis and COBBLESTONE-PerfDecomp generate 26 successful proofs for theorems unproven by COBBLESTONE (10 in CoqGym100 and 16 in Wigderson100). Interestingly, in one of these proofs<sup>5</sup>, the human-written proof relies on custom tactics in Coq’s Ltac language, while COBBLESTONE-PerfPremis generates a proof without custom tactics. In another<sup>6</sup>, COBBLESTONE asserts a proposition that helps prove the theorem.

**Failures:** Despite its improvements over the state-of-the-art, COBBLESTONE and its variants cannot prove 45 theorems from Wigderson100 and 42 from CoqGym100. We randomly sampled 14 such theorems (7 from each dataset) to examine manually.

Many of these theorems’ human-written proofs are much longer than COBBLESTONE’s proofs. COBBLESTONE’s average proof is 10 tactics long, whereas 5 of the 14 unproven theorems have a human-written proof of 20+ tactics. We also found that some theorems make it difficult for COBBLESTONE to make partial progress. For example, the ground-truth proof of `request_VoteReply_term_sanity_client_request`<sup>7</sup> uses tactics like `unfold` and `apply`, each of which helps prove the theorem, without decomposing the goal. It is difficult for COBBLESTONE to generate such proofs as it requires generating a working proof without recursing. Generalizing our

<sup>5</sup>`votesWithLog_update_elections_data_timeout` from `RefinementSpecLemmas.v` in `verdi-raft`

<sup>6</sup>`Mcardinal_Scardinal` from `graph.v` in `coq-wigderson`

<sup>7</sup>from `RequestVoteReplyTermSanityProof.v` in `verdi-raft`

approach to use partial progress that does not involve splitting a goal into subgoals may allow COBBLESTONE to prove such theorems.

**RA5:** COBBLESTONE’s proofs are sometimes shorter than their human-written counterparts, and often successfully leverage its divide and conquer approach. COBBLESTONE’s failures indicate an opportunity to leverage new kinds of partial progress.

### 3.7 Threats to Validity

All LLM evaluations may suffer from test data leaking into the pretraining dataset, reducing generalizability to unseen data. We mitigate this risk by using coq-wigderson, coq-bb5, and PnVRocqLib, whose first GitHub commits are after GPT-4’s publicly stated pretraining cutoff date (recall Section 3.1). We evaluate on 400 theorems, in line with other published work [66, 71, 72], to mitigate cost.

While our approach is general, our implementation is specific to Coq and uses GPT-4. We also evaluate with GPT-3.5-turbo and Claude 3 Opus to show generalizability. When evaluating against other LLM-based approaches such as PALM [43] and Rango [70], we give each tool similar resources *e.g.* by running PALM 15x to control for tokens. However, COBBLESTONE and Rango’s token consumption is not straightforward to compare, as Rango uses a model with fewer parameters and a different tokenization algorithm.

Finally, our evaluation only began exploring how external information can aid COBBLESTONE’s proof synthesis using oracles with access to human-written proofs. The intent of these oracles is to proxy how a human might interact with COBBLESTONE. But, their performance does not indicate how an interactive, semi-automated approach may perform. More research is necessary to study the effects of incomplete and noisy data, as well as user studies to show the potential impact of real, human-provided information.

## 4 Related Work

Recent work automating theorem proving in proof assistants has mostly explored three overarching approaches: hammers, machine learning techniques, and combinations of the two. Hammers such as CoqHammer [11] and Sledgehammer [54] call SMT solvers [13, 62] to construct low-level proofs. They cannot use induction, and thus are limited in what they can prove. Our evaluation shows that COBBLESTONE proves many theorems CoqHammer cannot.

Machine learning techniques (*i.e.* *neural theorem provers*) use a predictive model learned from existing proofs to predict next steps, and use these predictions to guide a search through the space of proofs [42]. They have been built using RNNs [27, 60], LSTMs [18, 19, 76], GNNs [3, 6], and LLMs [31, 77]. LLM-based methods prompt a pretrained model zero-shot or with few-shot examples [32, 81], fine-tune the model [20, 28], or use it as an agent [67, 72].

Hammers and machine learning techniques are complementary [18, 30]. Thor [30] fine-tunes an LLM to learn when to apply Sledgehammer vs. predict a tactic. By contrast, COBBLESTONE samples whole proofs rather than tactics and calls hammer without fine-tuning. Focusing on formalizing math, DraftSketchProve [32] uses an LLM to translate informal proofs into formal proof sketches with holes, filling the holes with Sledgehammer. Others extend this framework [72, 82, 83]. LEGO [72] decomposes theorems into

helper lemmas using an LLM, but provides no guarantee that proving them will prove the original theorem; by contrast, COBBLESTONE uses the theorem prover to decompose goals with such a guarantee.

LeanDojo [77] and Magnushammer [47] train models to select relevant premises (lemmas/definitions) at each proof step. In addition to premises, Rango [70] also selects similar proof scripts. COBBLESTONE uses the preceding lemmas in the file, but could benefit from a premise selection or a proof-script selection model in the same way it benefits from perfect premises provided by an oracle.

Proof engineers often repair previously working proofs [58]. Symbolic tools can automate some repair [44, 57]. Baldur fine-tunes an LLM to repair Isabelle proofs using error messages [20]. Unlike Baldur, which only attempts one repair, COBBLESTONE is an iterative, divide-and-conquer approach. PALM [43] samples from an LLM and uses repair mechanisms. By contrast, as detailed in Section 3.4, our approach is recursive in nature.

POETRY fine-tunes an LLM and estimates proof points on which to recurse [71]. By contrast, COBBLESTONE requires no tuning training data and uses fail-safe mode to precisely identify subgoals that recursion. Adapting COBBLESTONE to use more powerful models requires only minor changes to its API calls, whereas adapting POETRY is only possible with access to the model’s weights, and would require significant resources for fine-tuning.

Formally specifying system properties is complementary to proving them. Prior work has tackled formalizing natural language specifications [15, 24, 48, 80]. Once specified, some properties, like fairness [22], can be verified probabilistically [1, 23, 25, 46, 69].

Prompt engineering is key to effective LLM use [9, 17, 51]. Frameworks, such as chain-of-thought [75] and others [5, 8, 79], can elicit reasoning in LLMs [10, 21, 73, 75, 84]. LLMs can be used for quantitative reasoning tasks [2, 41] and programming tasks [7, 59]. These advances complement our efforts to generate proofs automatically.

## 5 Contributions

COBBLESTONE is a novel divide-and-conquer method for synthesizing formal verification proofs using LLMs. COBBLESTONE significantly outperforms prior tools and LLM-based baselines, proves more complex theorems, and complements other LLM-based tools. COBBLESTONE demonstrates a promising potential for collaborating with humans or other tools to synthesize even more proofs.

## Acknowledgments

We thank Zhanna Kaufman and Kyle Thompson for their help executing Proverbot9001 on the coq-wigderson dataset, and Rango on the PnVRocqLib and coq-bb5 datasets respectively. This work is supported by the National Science Foundation under grants no. NSF CCF-1955457, CCF-2210243, and CCF-2220892, and by the Defense Advanced Research Projects Agencies (DARPA) under Contract no. HR0011-24-2-0307.

## References

- [1] Aws Albarghouthi, Loris D’Antoni, Samuel Drews, and Aditya Nori. 2017. FairSquare: Probabilistic Verification for Program Fairness. In *OOPSLA*.
- [2] Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen Marcus McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2024. Llemma: An Open Language Model for Mathematics. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=4WnqRR915j>

- [3] Kshitij Bansal, Christian Szegedy, Markus N Rabe, Sarah M Loos, and Viktor Toman. 2019. Learning to Reason in Large Theories without Imitation. *arXiv preprint arXiv:1905.10501* (2019). arXiv:1905.10501
- [4] Clark W. Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanovic, Tim King, Andrew Reynolds, and Cesare Tinelli. 2011. CVC4. In *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6806)*, Ganesh Gopalakrishnan and Shaz Qadeer (Eds.). Springer, 171–177. doi:10.1007/978-3-642-22110-1\_14
- [5] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michał Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefler. 2024. Graph of thoughts: solving elaborate problems with large language models. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence (AAAI'24/IAAI'24/EAAI'24)*. AAAI Press, Article 1972, 9 pages. doi:10.1609/aaai.v38i16.29720
- [6] Lasse Blaauwbroek, Mirek Olšák, Jason Rute, Fidel Ivan Schaposnik Massolo, Jelle Piepenbrock, and Vasily Pestun. 2024. Graph2Tac: Online Representation Learning of Formal Math Concepts. In *Forty-first International Conference on Machine Learning*.
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, et al. 2021. Evaluating Large Language Models Trained on Code. arXiv:2107.03374 [cs.LG] <https://arxiv.org/abs/2107.03374>
- [8] Sijia Chen, Baochun Li, and Di Niu. 2024. Boosting of Thoughts: Trial-and-Error Problem Solving with Large Language Models. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=qBL04XXex6>
- [9] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* 24, 240 (2023), 1–113.
- [10] Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2024. Navigate through Enigmatic Labyrinth A Survey of Chain of Thought Reasoning: Advances, Frontiers and Future. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 1173–1203. doi:10.18653/v1/2024.acl-long.65
- [11] Łukasz Czajka and Cezary Kaliszzyk. 2018. Hammer for Coq: Automation for Dependent Type Theory. *Journal of Automated Reasoning* 61, 1-4 (2018), 423–453. doi:10.1007/s10817-018-9458-4
- [12] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- [13] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*. 337–340. doi:10.1007/978-3-540-78800-3\_24
- [14] Leonardo de Moura and Sebastian Ullrich. 2021. The Lean 4 Theorem Prover and Programming Language. In *Automated Deduction — CADE 28*. 625–635.
- [15] Madeline Endres, Sarah Fakhoury, Saikat Chakraborty, and Shuvendu K. Lahiri. 2024. Can Large Language Models Transform Natural Language Intent into Formal Method Postconditions? *Proceedings of the ACM Software Engineering (PACMSE)* 1, FSE (July 2024), 84:1–84:24. doi:10.1145/3660791
- [16] Andres Erbsen, Jade Philipoom, Jason Gross, Robert Sloan, and Adam Chlipala. 2019. Simple High-Level Code for Cryptographic Arithmetic — With Proofs, Without Compromises. In *IEEE Symposium on Security and Privacy (S&P)*. 1202–1219. doi:10.1109/SP.2019.00005
- [17] Hugo Touvron et al. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288
- [18] Emily First and Yuriy Brun. 2022. Diversity-Driven Automated Formal Verification. In *International Conference on Software Engineering (ICSE)* (22–27). 749–761. doi:10.1145/3510003.3510138
- [19] Emily First, Yuriy Brun, and Arjun Guha. 2020. TacTok: Semantics-Aware Proof Synthesis. *Proceedings of the ACM on Programming Languages (PACMPL) Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA) issue 4* (Nov. 2020), 231:1–231:31. doi:10.1145/3428299
- [20] Emily First, Markus Rabe, Talia Ringer, and Yuriy Brun. 2023. Baldur: Whole-Proof Generation and Repair with Large Language Models. In *ESEC/FSE* (6–8).
- [21] Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. Complexity-Based Prompting for Multi-step Reasoning. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=yfticZHC-19>
- [22] Sainyam Galhotra, Yuriy Brun, and Alexandra Meliou. 2017. Fairness Testing: Testing Software for Discrimination. In *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)* (6–8). 498–510. doi:10.1145/3106237.3106277
- [23] Stephen Giguere, Blossom Metevier, Yuriy Brun, Bruno Castro da Silva, Philip S. Thomas, and Scott Nielum. 2022. Fairness Guarantees under Demographic Shift. In *International Conference on Learning Representations (ICLR)* (25–29). <https://openreview.net/forum?id=wbP0BLm6ueA>
- [24] Alberto Goffi, Alessandra Gorla, Michael D. Ernst, and Mauro Pezzè. 2016. Automatic generation of oracles for exceptional behaviors. In *International Symposium on Software Testing and Analysis (ISSTA)*. 213–224. doi:10.1145/2931037.2931061
- [25] Austin Hoag, James E. Kostas, Bruno Castro da Silva, Philip S. Thomas, and Yuriy Brun. 2023. Seldonian Toolkit: Building Software with Safe and Fair Machine Learning. In *International Conference on Software Engineering (ICSE) Demo track* (14–20). 107–111. doi:10.1109/ICSE-Companion58688.2023.00035
- [26] Kryštof Hoder, Laura Kovács, and Andrei Voronkov. 2011. Invariant Generation in Vampire. In *Tools and Algorithms for the Construction and Analysis of Systems, Parosh Aziz Abdulla and K. Rustan M. Leino (Eds.)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 60–64.
- [27] Daniel Huang, Prafulla Dhariwal, Dawn Song, and Ilya Sutskever. 2018. GamePad: A Learning Environment for Theorem Proving. CoRR (2018). <https://arxiv.org/abs/1806.00608>
- [28] Yinya Huang, Xiaohan Lin, Zhengying Liu, Qingxing Cao, Huajian Xin, Haiming Wang, Zhenguo Li, Linqi Song, and Xiaodan Liang. 2024. MUSTARD: Mastering Uniform Synthesis of Theorem and Proof Data. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=8xliOUg9EW>
- [29] Kevin Jacobs and Benjamin Beurdouche. 2020. Performance Improvements via Formally-Verified Cryptography in Firefox. <https://blog.mozilla.org/security/2020/07/06/performance-improvements-via-formally-verified-cryptography-in-firefox/>.
- [30] Albert Jiang, Konrad Czechowski, Mateja Jamnik, Piotr Milos, Szymon Tworowski, Wenda Li, and Yuhuai Tony Wu. 2022. Thor: Wielding Hammers to Integrate Language Models and Automated Theorem Provers. In *NeurIPS*.
- [31] Albert Qiaochu Jiang, Wenda Li, Jesse Michael Han, and Yuhuai Wu. 2021. LISA: Language models of ISabelle proofs. In *Conference on Artificial Intelligence and Theorem Proving (AITP)*. 17.1–17.3.
- [32] Albert Qiaochu Jiang, Sean Welleck, Jin Peng Zhou, Timothee Lacroix, Jiacheng Liu, Wenda Li, Mateja Jamnik, Guillaume Lample, and Yuhuai Wu. 2023. Draft, Sketch, and Prove: Guiding Formal Theorem Provers with Informal Proofs. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=SMA9EAovKMC>
- [33] Saketh Ram Kasibatla, Arpan Agarwal, Yuriy Brun, Sorin Lerner, Talia Ringer, and Emily First. 2025. Cobblestone: A Divide-and-Conquer Approach for Automating Formal Verification Replication Package. <https://doi.org/10.5281/zenodo.17603489> <https://doi.org/10.5281/zenodo.17603489>
- [34] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. 2009. SeL4: Formal Verification of an OS Kernel. In *ACM SIGOPS 22nd Symposium on Operating Systems Principles (SOSP)*. 207–220. doi:10.1145/1629575.1629596
- [35] Herb Krasner. 2022. The Cost of Poor Software Quality in the US: A 2022 Report. <https://www.it-cisq.org/wp-content/uploads/sites/6/2022/11/CPSQ-Report-Nov-22-2.pdf>.
- [36] Ramana Kumar, Magnus O Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A verified implementation of ML. *ACM SIGPLAN Notices* 49, 1 (2014), 179–191.
- [37] Blaauwbroek Lasse, Josef Urban, and Herman Geuvers. 2020. The Tactician: A Seamless, Interactive Tactic Learner and Prover for Coq. In *International Conference on Intelligent Computer Mathematics (CICM)*. 271–277. doi:10.1007/978-3-030-53518-6\_17
- [38] Chris Latner and Nikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *International Symposium on Code Generation and Optimization (CGO)*. 75–86. doi:10.1109/CGO.2004.1281665
- [39] Xavier Leroy. 2006. Formal Certification of a Compiler Back-End or: Programming a Compiler with a Proof Assistant. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 42–54. doi:10.1145/1111037.1111042
- [40] Xavier Leroy. 2009. Formal verification of a realistic compiler. *Communications of the ACM (CACM)* 52, 7 (2009), 107–115. doi:10.1145/1538788.1538814
- [41] Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022. Solving quantitative reasoning problems with language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '22)*. Curran Associates Inc., Red Hook, NY, USA, Article 278, 15 pages.
- [42] Zhaoyu Li, Jialiang Sun, Logan Murphy, Qidong Su, Zenan Li, Xian Zhang, Kaiyu Yang, and Xujie Si. 2024. A Survey on Deep Learning for Theorem Proving. (2024). arXiv:2404.09939 [cs.AI] <https://arxiv.org/abs/2404.09939>
- [43] Minghai Lu, Benjamin Delaware, and Tianyi Zhang. 2024. Proof Automation with Large Language Models. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. Sacramento, CA, USA. doi:10.1145/3691620.3695521
- [44] Paolo Masci and Aaron Dutle. 2022. Proof Mate: An Interactive Proof Helper for PVS (Tool Paper). In *NASA Formal Methods Symposium*. Springer, 809–815.

- [45] Joseph Menn and Aaron Gregg. 2024. CrowdStrike blames global it outage on bug in system for checking updates. *The Washington Post* (July 2024). <https://www.washingtonpost.com/business/2024/07/24/crowdstrike-microsoft-crash-bug-report/>
- [46] Blossom Metevier, Stephen Giguere, Sarah Brockman, Ari Kobren, Yuriy Brun, Emma Brunskill, and Philip S. Thomas. 2019. Offline Contextual Bandits with High Probability Fairness Guarantees. In *Annual Conference on Neural Information Processing Systems (NeurIPS), Advances in Neural Information Processing Systems 32* (9–14). 14893–14904. <http://papers.nips.cc/paper/9630-offline-contextual-bandits-with-high-probability-fairness-guarantees>
- [47] Maciej Mikula, Szymon Antoniak, Szymon Tworkowski, Bartosz Piotrowski, Albert Qiaoju Jiang, Jin Peng Zhou, Christian Szegedy, Łukasz Kuciński, Piotr Miłoś, and Yuhuai Wu. 2024. Magnushammer: A Transformer-based Approach to Premise Selection. In *International Conference on Learning Representations (ICLR)*. <https://openreview.net/forum?id=oYjPk8mqAV>
- [48] Manish Motwani and Yuriy Brun. 2019. Automatically Generating Precise Oracles from Structured Natural Language Specifications. In *Proceedings of the 41st International Conference on Software Engineering (ICSE)* (29–31). 188–199. doi:10.1109/ICSE.2019.00035
- [49] Toby Murray, Daniel Matchuk, Matthew Brassil, Peter Gammie, Timothy Bourke, Sean Seefried, Corey Lewis, Xin Gao, and Gerwin Klein. 2013. seL4: From general purpose to a proof of information flow enforcement. In *IEEE Symposium on Security and Privacy (S&P)*, 415–429.
- [50] mxmlys. 2025. Coq-BB5 release v1.0.0. doi:10.5281/zenodo.17061968 <https://doi.org/10.5281/zenodo.17061968>
- [51] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774
- [52] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [53] Aditya Paliwal, Sarah M. Loos, Markus N. Rabe, Kshitij Bansal, and Christian Szegedy. 2020. Graph Representations for Higher-Order Logic and Theorem Proving. In *Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2967–2974.
- [54] Larry Paulson and Tobias Nipkow. 2023. The Sledgehammer: Let Automatic Theorem Provers write your Isabelle scripts! <https://isabelle.in.tum.de/website-isabelle2009-1/sledgehammer.html>
- [55] Siraphob Phipathananunth. 2023. Towards the Formal Verification of Wigderson’s Algorithm (SPLASH 2023). Association for Computing Machinery, 40–42. doi:10.1145/3618305.3623600
- [56] PnVDiscord. [n. d.]. *PnVRocqLib*.
- [57] Talia Ringer. 2021. *Proof Repair*. Ph.D. Dissertation. University of Washington.
- [58] Talia Ringer, Alex Sanchez-Stern, Dan Grossman, and Sorin Lerner. 2020. REPLica: REPL Instrumentation for Coq Analysis. In *International Conference on Certified Programs and Proofs (CPP)*, 99–113. doi:10.1145/3372885.3373823
- [59] 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
- [60] Alex Sanchez-Stern, Yousef Alhessi, Lawrence Saul, and Sorin Lerner. 2020. Generating correctness proofs with neural networks. In *ACM SIGPLAN International Workshop on Machine Learning and Programming Languages (MAPL)*, 1–10.
- [61] Alex Sanchez-Stern, Emily First, Timothy Zhou, Zhanna Kaufman, Yuriy Brun, and Talia Ringer. 2023. Passport: Improving Automated Formal Verification Using Identifiers. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 45, 2, Article 12 (June 2023), 12:1–12:30 pages. doi:10.1145/3593374
- [62] Stephan Schulz. 2013. System description: E 1.8. In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*. Springer, 735–743.
- [63] Stephan Schulz, Simon Cruanes, and Petar Vukmirović. 2019. Faster, Higher, Stronger: E 2.3. In *Proc. of the 27th CADE, Natal, Brasil (LNAI, 11716)*, Pascal Fontaine (Ed.), Springer, 495–507.
- [64] Jean Souyris. 2014. Industrial Use of CompCert on a Safety-Critical Software Product. [http://projects.laas.fr/IFSE/FMF/J3/slides/P05\\_Jean\\_Souyris.pdf](http://projects.laas.fr/IFSE/FMF/J3/slides/P05_Jean_Souyris.pdf)
- [65] Richard M. Stallman. 2012. *Using the GNU Compiler Collection*. Free Software Foundation.
- [66] Amitayush Thakur, George Tsoukalas, Yeming Wen, Jimmy Xin, and Swarat Chaudhuri. 2024. An In-Context Learning Agent for Formal Theorem-Proving. In *First Conference on Language Modeling*. <https://openreview.net/forum?id=V7HRrxXUhn>
- [67] Amitayush Thakur, George Tsoukalas, Yeming Wen, Jimmy Xin, and Swarat Chaudhuri. 2024. An In-Context Learning Agent for Formal Theorem-Proving. In *First Conference on Language Modeling*. <https://openreview.net/forum?id=V7HRrxXUhn>
- [68] The Coq Development Team. 2017. Coq, v.8.7. <https://coq.inria.fr>.
- [69] Philip S. Thomas, Bruno Castro da Silva, Andrew G. Barto, Stephen Giguere, Yuriy Brun, and Emma Brunskill. 2019. Preventing Undesirable Behavior of Intelligent Machines. *Science* 366, 6468 (22 November 2019), 999–1004. doi:10.1126/science.aag3311
- [70] Kyle Thompson, Nuno Saavedra, Pedro Carrott, Kevin Fisher, Alex Sanchez-Stern, Yuriy Brun, João F. Ferreira, Sorin Lerner, and Emily First. 2025. Rangor: Adaptive Retrieval-Augmented Proving for Automated Software Verification. In *Proceedings of the IEEE/ACM 47th International Conference on Software Engineering* (Ottawa, Ontario, Canada) (ICSE ’25). IEEE Press, 347–359. doi:10.1109/ICSE55347.2025.00161
- [71] Haiming Wang, Huajian Xin, Zhengying Liu, Wenda Li, Yinya Huang, Jianqiao Lu, Zhicheng YANG, Jing Tang, Jian Yin, Zhenguo Li, and Xiaodan Liang. 2024. Proving Theorems Recursively. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=yAa519t2TQ>
- [72] Haiming Wang, Huajian Xin, Chuanyang Zheng, Zhengying Liu, Qingxing Cao, Yinya Huang, Jing Xiong, Han Shi, Enze Xie, Jian Yin, et al. 2024. LEGO-Prover: Neural Theorem Proving with Growing Libraries. In *The Twelfth International Conference on Learning Representations*.
- [73] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=1PLNIMMrw>
- [74] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *arXiv [cs.CL]* (Jan. 2022). <http://arxiv.org/abs/2201.11903>
- [75] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [76] Kaiyu Yang and Jia Deng. 2019. Learning to prove theorems via interacting with proof assistants. In *International Conference on Machine Learning (ICML)*. <http://proceedings.mlr.press/v97/yang19a/yang19a.pdf>
- [77] Kaiyu Yang, Aidan M. Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J. Prenger, and Animashree Anandkumar. 2023. LeanDojo: Theorem Proving with Retrieval-Augmented Language Models. In *Conference on Neural Information Processing Systems (NeurIPS)*. [http://papers.nips.cc/paper\\_files/paper/2023/hash/4441469427094f8873d0fecb0c4e1ccc-Abstract-Datasets\\_and\\_Benchmarks.html](http://papers.nips.cc/paper_files/paper/2023/hash/4441469427094f8873d0fecb0c4e1ccc-Abstract-Datasets_and_Benchmarks.html)
- [78] Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. 2011. Finding and understanding bugs in C compilers. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 283–294. doi:10.1145/1993498.1993532
- [79] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NeurIPS)*. New Orleans, LA, USA, 517:1–517:14 pages.
- [80] Juan Zhai, Yu Shi, Minxue Pan, Guian Zhou, Yongxiang Liu, Chunrong Fang, Shiqing Ma, Lin Tan, and Xiangyu Zhang. 2020. C2S: Translating Natural Language Comments to Formal Program Specifications. In *European Software Engineering Conference and ACM SIGSOFT International Symposium on Foundations of Software Engineering (ESEC/FSE)*, 25–37. doi:10.1145/3368089.3409716
- [81] Shizhuo Dylan Zhang, Talia Ringer, and Emily First. 2023. Getting More out of Large Language Models for Proofs. *8th Conference on Artificial Intelligence and Theorem Proving* (Sept. 2023). arXiv:2305.04369
- [82] Xueliang Zhao, Wenda Li, and Lingpeng Kong. 2024. Subgoal-based Demonstration Learning for Formal Theorem Proving. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*, Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (Eds.), PMLR, 60832–60865. <https://proceedings.mlr.press/v235/zhao24h.html>
- [83] Chuanyang Zheng, Haiming Wang, Enze Xie, Zhengying Liu, Jiankai Sun, Huajian Xin, Jianhao Shen, Zhenguo Li, and Yu Li. 2024. Lyra: Orchestrating Dual Correction in Automated Theorem Proving. *Transactions on Machine Learning Research* (2024). <https://openreview.net/forum?id=Sw75k0tzs>
- [84] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. 2023. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=WZH7099tgfM>