

Formally Verified AI Co-Reasoning: An Education Case Study

Arnav Mehta*, Chenjun Guo*, Thomas Lu*, Manooshree Patel†,
Niels Voss†, Rayna Bhattacharyya, Peter Donovan, Gireeja Ranade

University of California, Berkeley

{arnavmehta, chenjun_guo, thomaslu, manooshreepatel}@berkeley.edu
{niels_voss, rayna_b, peterdonovan}@berkeley.edu, ranade@eecs.berkeley.edu

Abstract

Proof assistants such as Lean 4 offer machine-checkable verification of mathematical reasoning, but their adoption is limited because formal languages have a steep learning curve and can be challenging to use. Large language models offer a way to lower this barrier to entry through *natural-language co-reasoning systems*: proof assistants that users interact with entirely in natural language, and which use formal verification to provide feedback to users who submit incomplete and incorrect proofs. This paper explores such a co-reasoning system in an educational setting.

We explore metrics for three properties that are important for natural language co-reasoning systems — *faithfulness*, *accuracy* and *appropriate granularity*. We introduce an **evaluation framework** for these properties in the context of educational proofs. Second, we present a **case study** on an educational natural-language co-reasoning system for math proofs, evaluating its autoformalization harness on student-written proofs from an undergraduate linear algebra class using these three properties.

1 Introduction

Proof assistants like Lean 4 (de Moura and Ullrich, 2021) and Isabelle (Paulson, 1994) have the ability to assist math research and education because they provide a mechanism to scalably and reliably check mathematical proofs. They have been adopted by some mathematicians in large-scale formalization projects such as Mathlib (mathlib Community, 2020), but adoption of proof assistants in mainstream mathematics remains limited. Formal languages pose challenges to users due to syntax and readability (Shi et al., 2025), and proofs in formal languages require a much greater level of detail and justification than those in natural language (Patel et al., 2026).

*These authors contributed equally to this work.

†These authors contributed equally to this work.




Figure 1: NL co-reasoning system architecture

Recent advances in LLMs have shown promise in bridging the gap between natural and formal language. LLMs now assist users of proof assistants with tactic suggestion (Song et al., 2024; Welleck and Saha, 2023), premise selection (Yang et al., 2023), proof repair (Thakur et al., 2025; Lu et al., 2024c), conjecturing (Onda et al., 2025), informalization (Gao et al., 2024), proof generation (Ren et al., 2025; Chen et al., 2025a; Lin et al., 2025), and autoformalization (Wu et al., 2022; Lu et al., 2024b). These tools make it easier for users to interact with a proof assistant, but they still require knowledge of the formal language.

We introduce the notion of a **natural language co-reasoning** system (*NL co-reasoning*): a proof assistant that a user interacts with entirely in natural language. Such a system allows users to iteratively develop their informal proofs: they may submit incomplete or incorrect proofs, receive formally verified natural-language feedback, and revise their reasoning accordingly. A user’s workflow is as follows (see Figure 1): (1) the user inputs their reasoning in free-form natural language; (2) the input is autoformalized into a formal language such as Lean 4; (3) the resulting formal program is compiled; (4) the generated code and compiler output are informalized back into natural language as verification feedback; and (5) the user uses this feedback to revise and resubmit their reasoning. Steps (2)–(4) are hidden from the user, who is never required to read or write formal code.

Existing autoformalization approaches use formal proof primarily as a post-hoc check—to verify correctness of a natural language proof after it is considered complete (Wu et al., 2022; Lu et al.,

2024b). NL co-reasoning instead integrates formal verification into the proof discovery process itself, using it to surface errors to the user at intermediate steps.

Co-reasoning systems that users interact with in natural language are valuable because they can deliver formally verified feedback without disrupting a user’s existing workflow. In mathematics for instance, both expert mathematicians proving new conjectures and students learning to write proofs can use NL co-reasoning systems to receive per-step verification and feedback that is trustworthy (Ospanov et al., 2025; Avigad, 2019; Lodder et al., 2021; Keenan and Omar, 2024; Patel et al., 2026). Beyond math, legal and tax professionals must reason within the constraints of statutes and tax codes (Jurayj et al., 2026; PramaanaLabs), and clinicians must produce evidence-based justifications for their diagnostic conclusions (Hekmatnejad et al., 2019; Singh et al., 2026). In each case, a co-reasoning system can productively surface inconsistencies, with no additional burden to the user.

Our main contributions are: (1) desiderata for natural language co-reasoning systems; (2) an evaluation framework for faithfulness, accuracy and appropriate granularity in the context of education; and (3) a case study applying this framework to an educational NL co-reasoning system, whose autoformalization harness is evaluated on student-written proofs from an undergraduate linear algebra course.

2 Natural language co-reasoning desiderata

In earlier work, we identified desiderata (Patel et al., 2026) for an educational tutor with a formal backend. We build on and expand this list for the case of a natural language co-reasoning system. In the next section, we propose quantitative measures for each desideratum discussed here. We do not claim that this is a complete list. For example, from an interaction perspective, a co-reasoning system should also be transparent, allowing users to interpret the system’s understanding of their proof, and should support incremental editability. However, evaluating these desiderata requires user feedback and assessment, so we leave them for future work.

1. **Faithfulness:** the formal proof accurately reflects both the structure (Cabral et al., 2025) and semantics (Murphy et al., 2024) of the natural-language proof provided as input.

2. **Accuracy:** the system provides mathematically accurate feedback for each reasoning step, i.e. it identifies correct steps as correct and incorrect steps as incorrect.
3. **Appropriate granularity:** the system must require users to justify their reasoning at a level of detail appropriate for the target application. Formal languages such as Lean require fully rigorous formal proofs. This can impose unnecessary burdens on users, who may naturally think at a higher level of granularity (Guo et al., 2025). The level of granularity required depends on the context; an advanced mathematical course may allow students to assume lemmas that compress substantial logical jumps, while an introductory course will require more justification. Thus, it must be possible to calibrate an NL co-reasoning system to the needs of the context it is deployed in.

The educational setting is a particularly interesting case study for NL co-reasoning and provides a natural environment to measure the above desiderata. Students learning to write proofs naturally generate diverse informal arguments, including incomplete and incorrect proofs. Instructors implicitly or explicitly communicate a notion of sufficiently rigorous or justified proofs (that is specific to their classrooms) (Yackel and Cobb, 1996).

Evaluating these properties presents a significant challenge. Faithfulness is not typically considered when judging autoformalization success, as a compiling proof is often treated as a successful output (Azerbayev et al., 2023; Wang et al., 2024; Ying et al., 2024; Liu et al., 2025b; Huang et al., 2025). Much of the prior work also assumes well-formed informal statements (Lu et al., 2024b; Liu et al., 2025a), while a co-reasoning system must handle proofs that are incomplete, ambiguous, or incorrect.

3 Metrics

3.1 Faithfulness

We attempt to measure faithfulness by considering two proxy metrics—*information preservation* and *non-dilution*. Essentially, these measure “did the autoformalizer retain everything the user included?” and “did the autoformalizer refrain from adding anything the user did not include?”

3.1.1 Information Preservation and Non-Dilution

Let a *natural-language proof* Π_{NL} and a *formal-language proof* Π_{FL} each be decomposed into a set of content-bearing pieces, which we call *units*, and we will say more about units below. Let u_{NL} and u_{FL} be a unit of natural and formal language each. We assume a relation $\sim$ between NL units and FL units, where $u_{NL} \sim u_{FL}$ indicates that the two units express the same mathematical information.

Definition (Information preservation) Π_{FL} *preserves the information* of Π_{NL} iff every unit in the natural-language proof has a corresponding unit (under $\sim$) in the formal proof — i.e., nothing got dropped. We measure this as:

$$IP = \frac{\# \text{ NL units matched in FL under } \sim}{\# \text{ NL units}}$$

Definition (Non-dilution) Π_{FL} is *non-diluting* with respect to Π_{NL} iff every unit in the formal proof has a corresponding unit in the natural-language proof—nothing was added. We measure this as:

$$ND = \frac{\# \text{ FL units matched in NL under } \sim}{\# \text{ FL units}}$$

Together, the two properties say the formal proof is no more and no less than what was said by the user in natural language. Both scores are in the range $[0, 1]$.

We note that “unit” is intentionally left abstract in these definitions. The choice of unit determines the resolution at which we evaluate information preservation and non-dilution. We call this a *resolution choice* and fix it before measuring. For example, we may examine a Lean 4 proof at the level of tactics, which gives a high-resolution view, or at the level of have statements, which gives a lower-resolution view. The matching relation $\sim$ is also fixed as part of the resolution choice. The definition of a “unit” in both natural and formal-language proofs is fundamental to this definition. We provide working definitions for both.

3.1.2 Units in natural language

A natural language proof consists of two kinds of reasoning steps (Guo et al., 2025; Sieg, 1990): *declarations*, which introduce new variables, objects, or hypotheses (e.g. “let x be the zero vector”), and *assertions*, which make a mathematical claim (e.g. “the matrix M is symmetric”).

We adopt the unitization procedure of Guo et al. (2025), which decomposes a natural-language proof into a sequence of *declaration* and *assertion* steps exactly as written by the user. Each resulting step is one natural-language (NL) unit.

3.1.3 Units in formal language:

The choice of proof-unit resolution is application-dependent. In educational settings, the coarsest unit may be a single theorem, since students are typically asked to prove one theorem at a time. In other settings, such as formalizing a legal document, the coarsest unit may instead be an entire document comprising multiple theorems. To evaluate faithfulness using the metrics above, one must first specify the proof units that the system is designed to support. In our case, we evaluate co-reasoning systems at three levels of resolution, ordered from coarse to fine. These reflect common levels at which Lean 4 proofs are structured, although this list is not intended to be exhaustive. The proposed metrics are agnostic to the choice of resolution and apply to any decomposition that defines meaningful proof units.

Low-resolution: Top-level statements. A Lean 4 file typically contains named theorems and lemmas, each of which asserts a proposition. The unit is the *initial goal state* at the start of each such declaration. This corresponds to the resolution used in statement autoformalization benchmarks (Azerbayev et al., 2023; Ying et al., 2024; Gulati et al., 2024; Poiroux et al., 2024).

Medium resolution: Intermediate sub-claims. Within a proof, tactics such as *have*, *suffices*, and *show* introduce intermediate propositions; the unit is the goal state following each such statement. This granularity corresponds to proof-sketch and sub-goal decomposition approaches in the literature (Jiang et al., 2022; Ren et al., 2025; Cabral et al., 2025).

High resolution: Tactic states. In a tactic-mode proof, Lean 4 exposes a goal state after every tactic. The unit is each such goal state. A block of n tactics typically yields $n + 1$ goal states. This is the granularity used by tactic-prediction models, which learn to predict the next tactic given the current proof state (Yang et al., 2023; Welleck and Saha, 2023).

3.2 Measuring Appropriate Granularity

We wish to identify when a co-reasoning system operates at an ‘appropriate granularity’. Here it is easier to define what failure looks like than to characterize success. We do not want users to provide extensive justification for mathematical facts that may seem trivial to humans but require long proofs in Lean 4. On the other hand, particularly in educational contexts, we do not want to allow students to make large logical leaps and have gaps in their reasoning be filled in by a powerful autoformalizer.

Evaluating whether the first of the undesirable behaviors exists in the system requires extensive user studies that are out of scope for the current paper. We focus on detecting the second kind of undesirable behavior in this work.

Autograder-based evaluation: Evaluating whether a proof step contains a large jump is tricky, since the required ground truth is difficult to obtain. Directly measuring step-wise large-jump detection would require labels indicating whether each student step contained a large jump. Such labels are not common in natural-language proof benchmarks and would require additional human annotation. In contrast, in typical instructional settings, past exam records more often contain instructors’ rubric grades for student proofs. We therefore measure the success in detecting large jumps through looking at the system performance as an autograder for the proofs.

Proof autograders (Brown and Michel, 2010; Moore, 2016; Miller et al., 2018) automatically assign scores to student proofs. A central part of that task is identifying *logical leaps*: steps that are stated by a student without sufficient justification (Fallis, 2003; Weber, 2008; Inamdar et al., 2025). A co-reasoner that can detect insufficiently justified steps should, in principle, provide information useful for grading. We use this idea to define our proxy metric.

Our evaluation compares two grading modes. In both modes, we first associate each human-written rubric item with the formal proof steps relevant to evaluating that item. We then classify each step as a large jump or not.

We compare two modes to show the advantage of our system. In the *lenient* mode, the student gets credit for all steps, whether they are appropriately justified or not. In the *strict* mode, steps flagged as insufficiently justified do not get any credit. We then ask whether strict mode agrees better with

human graders than lenient mode.

Let C denote the set of student submissions, and let A_m^i denote the agreement rate on rubric item i under grading mode m :

$$A_m^i = \frac{1}{|C|} \sum_{c \in C} \mathbf{1}[g_m(c)_i = h(c)_i],$$

where $h(c)_i$ and $g_m(c)_i$ are the human and system credit decisions for rubric item i on submission c , and $m \in \{\text{strict}, \text{lenient}\}$.

We define the granularity score for rubric item i as the gain in agreement from using strict mode:

$$\Delta^i = A_{\text{strict}}^i - A_{\text{lenient}}^i.$$

The overall granularity score averages this gain across rubric items:

$$\Delta = \frac{1}{|R|} \sum_{i \in R} \Delta^i.$$

We weight all rubric items equally, regardless of the number of points assigned to each. A large positive Δ means that withholding credit for insufficiently justified steps brings the system’s grading closer to the human grader’s decisions. This is preliminary evidence that the detector identifies the gaps a human grader would penalize.

This measure is indirect and has three limitations. First, human graders are imperfect (Thomas et al., 2025), yet we treat them as ground truth. Second, if the dataset is such that most proofs do not contain any large jumps, we may still have a small Δ , regardless of the quality of the system. Third, the metric assumes that the formalization of the rubric has been done accurately.

3.3 Measuring Accuracy

We define accuracy as whether the system correctly identifies the mathematical validity of each natural-language proof unit, as defined in Section 3.1.2, under the problem context. We report true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN) at the unit level.

4 A case study: An educational co-reasoning system for math proofs

We developed a proof-of-concept NL co-reasoning system for proof tutoring. The system comprises an autoformalizer and an auto-informalizer. It accepts free-form student proofs and allows users to formally verify their proofs, solely in natural




Figure 2: Autoformalization architecture of our system

language. This tool was deployed in an undergraduate linear algebra course. In our case study, we assess the autoformalizer’s **faithfulness**, **appropriate granularity**, and **accuracy**. Figure 2 illustrates the autoformalizer architecture, which consists of the proof scaffolder, lemma prover, and proof assembler.

4.1 Autoformalization harness

Our system preserves the sub-claim structure of a student’s proof by converting each sub-claim into an independent Lean 4 lemma, proving those lemmas, and assembling them into a proof of the original theorem. To faithfully capture the student’s reasoning, we maintain the local context established by the proof. Whenever a student introduces new assumptions or variables, these are added to the context and inherited by all subsequent sub-claims until they go out of scope. For example, if a student writes “let u be a unit vector,” then all later sub-claims within that scope are checked under the corresponding hypotheses that u is defined and is a unit vector. The pipeline has three stages (Figure 2):

Proof scaffolder: Extracts an ordered list of sub-claims from the student’s proof and converts each into a self-contained Lean 4 lemma statement. We call these scaffold lemmas.

Lemma prover: Proves each scaffold lemma via a compiler-error repair loop augmented with a curated library of *supporting mini lemmas*—routine facts that can be invoked during proof construction (Jiang et al., 2022)

Proof assembler: Stitches the scaffold lemmas into a proof of the original theorem.

4.2 Stage 1: Proof scaffolding

The scaffolder turns each student-written natural language sub-claim into a Lean 4 lemma statement (with proof left as sorry) in three steps:

NL preprocessing: We build on the preprocessing system of Guo et al. (2025), which decomposes natural language proofs into minimal independent units at the user-specified granularity, in three steps: (1) resolve referential expressions such as “they” or “by similar reasoning” to their referents; (2) classify each step as a declaration or assertion (Sieg, 1990); (3) for assertions, extract antecedent and consequent. The output is a list of disambiguated steps, with implication assertions labeled with its antecedent and consequent.

Assertion scaffolding: An LLM then converts each assertion step into a Lean 4 proposition — the type after the colon in

lemma name (hypotheses): prop := by ...

For implication assertions, the antecedent is lifted into a hypothesis ($h_{\text{antecedent}} : \dots$) and the consequent becomes the proposition. The proof is left as sorry; the hypothesis signatures for the lemmas are completed in the declaration scaffold step.

Declaration scaffolding: To recover the hypotheses in scope at each assertion, we prompt an LLM to translate each student declaration into a Lean 4 tactic. We assemble a partial proof by placing these tactics in order and inserting `extract_goal` (mathlib Community, 2020) at every assertion position. `extract_goal` is a Mathlib tactic that serializes the current goal state as a standalone theorem whose hypotheses are exactly those in scope at the call site, giving us the hypotheses to attach to the subsequent assertion’s lemma (Figure 3).

Type checking and error correction: Lean 4 enforces strict type correctness, so each lemma statement is type-checked before proof synthesis. If a generated lemma statement fails to compile, it is passed into a loop in which an LLM attempts to revise the statement. Importantly, the LLM is provided with explicit instructions to correct only formalization errors and never alter the mathematical content of the assertion. However, this is not always enforceable in practice. A compilation error could be caused by a mathematical error made by the student: for instance, a matrix multiplication with incompatible dimensions will trigger a type error in Lean 4 even when the assertion faithfully reflects the student’s claim. In such cases, the correction loop cannot fix the compilation error without editing the mathematical content, which we identify as a limitation of this approach.




Figure 3: using `extract_goals` to get the local context for each assertion




Figure 4: Lemma proving stage: scaffolded lemma is proved with a one-shot attempt followed by branch for proof diversity. prover concludes if one of the branches succeed or every branch exhausts 5 attempts

4.3 Stage 2: Lemma proving

The prover attempts to prove the generated Lean 4 lemma statements. Real-time feedback demands speed and low cost (Patel et al., 2026; Barnes and Stamper, 2008; Lodder et al., 2021), yet formal proof search is expensive and slow (Yang et al., 2023; Xin et al., 2024b; Liu et al., 2026). We address this with two design choices. First, we augment the prover with *supporting mini lemmas* that supply problem-specific facts, saving the prover the burden of deriving low-level results from scratch (Jiang et al., 2022). Scaffold lemmas corresponding to earlier assertions in the proof are also available to the prover. Second, a staged pipeline (Figure 4) tries a one-shot proof (Xin et al., 2024a), then falls back to parallel sampling for diversity (Xin et al., 2024b; Lin et al., 2025), with a compiler-driven repair loop of up to 5 attempts per branch (Thakur et al., 2025; Zhang et al., 2025). Together these shrink the search space enough that a lightweight model such as Gemini Flash (Comanici et al., 2025) suffices as the prover.

4.4 Supporting mini lemma generation

A *supporting mini lemma* is a small auxiliary result that the prover may invoke when proving sub-claims in a students’ proof. Supporting mini lemmas serve two purposes. First, they provide reusable facts that allow the prover to focus on higher-level reasoning. Second, they provide a mechanism to calibrate and detect the “size” of

sub-claims.

A stronger mini-lemma set permits larger inferential steps, while a weaker set requires more intermediate reasoning to be made explicit.

Rather than retrieving from a large library such as Mathlib (mathlib Community, 2020; Yang et al., 2023), we construct problem-specific mini lemmas tailored to the target theorem. In this paper, the mini lemmas are written by humans, though automated generation is a promising direction for future work. This design allows us to maintain a consistent proof granularity across problems (Gusakov et al., 2025).

Main and side lemmas. Mini lemmas differ in mathematical significance, a distinction used by the granularity detector in the next section. Each is annotated by the authors as either *main* or *side*. A *main* lemma is a key mathematical fact for the problem relative to the course rubric, and its appearance indicates substantive reasoning; a *side* lemma is auxiliary Lean 4 machinery — type coercions, definitional unfolding, and other boilerplate without significant mathematical content. The annotation is made once per problem and reused across submissions; see Appendix G for an example.

4.5 Granularity Detection

Each *main* mini lemma corresponds to a rubric-scorable step (Section 4.4). The number of main mini lemmas invoked in the proof of a scaffold lemma serves as a proxy for the amount of mathematical reasoning compressed into that student claim. If the prover invokes more than one main mini lemma to discharge a single scaffold lemma, we flag the claim as a **large jump**. This is the systems’ *appropriate granularity detector* as introduced in Section 3.2.

4.6 Stage 3: Proof assembler

The proof assembler’s purpose is to verify that the student’s top-level argument is logically valid: given the sub-claims as stated, does the main theorem follow? Given (i) the theorem statement, (ii) the ordered scaffold lemmas, (iii) the Lean 4 statements for the declarations, and (iv) the student’s NL proof, an LLM is prompted to produce a Lean 4 proof invoking the lemmas and declaration tactics in the order mentioned in the student’s proof. Similar to Thakur et al. (2025), compilation failures are returned to the LLM for revision, up to 5 attempts.




Figure 5: System accuracy evaluation. Left: per-problem step-level agreement between the system and the LLM evaluator. Right: aggregate agreement and disagreement counts.

The assembler runs whether or not every scaffold lemma was closed. Unresolved lemmas are left as sorry and treated as assumed facts during assembly, analogous to sub-goal decomposition in recent auto-provers (Ren et al., 2025; Zhou et al., 2025).

The assembler is prompted to close the main theorem using the proven scaffold lemmas, the declaration statements, and routine tactics: `rw`, arithmetic normalization (`ring`, `linarith`, `nlinarith`) and positivity. It is instructed not to introduce substantial mathematical facts.

Calibrating how much mathematical content the assembler is allowed to add is domain-dependent and remains open: if the assembler is too conservative, it may fail even when the student’s argument is correct; if it is too permissive, it may fill gaps the student did not provide. In future work, we hope to replace the LLM assembler with a custom Lean 4 tactic that makes this boundary explicit.

5 Evaluations

We evaluate the system on 100 natural-language proofs from real student homework submissions: 10 proof problems with 10 student solutions each. All the student solutions have been anonymized by the researchers. For each proof, we run the student’s natural-language steps through the system pipeline and record step-level verification outcomes.

5.1 Accuracy

Accuracy is defined as whether the system’s verification outcome matches the mathematical validity of each proof step under the problem context. As a baseline, an LLM evaluator judges each natural language proof step. We manually audited a random 10% of the LLM judge’s decisions and found that 34 of 35 steps (97%) were judged correctly. Appendix H reports the evaluator prompt. The system’s outputs are then compared with the baseline judgments (Figure 5), and all disagreements are

	Information Preservation	Non-Dilution
All Proofs (100)		
Top-Level	0.6305	0.7616
Sub-Claims	0.6704	0.3968
Tactic States	0.3645	0.1000
Complete Proofs Only (74)		
Top-Level	0.7194	0.7926
Sub-Claims	0.7625	0.4051
Tactic States	0.4096	0.1040

Table 1: Information preservation and non-dilution scores for formalized proofs, by granularity level.

Human vs. LLM Disagreement		
Top-Level	Sub-Claims	Tactic States
6.19%	3.09%	19.59%

Table 2: Human–LLM judge disagreement rate on the faithfulness mapping, by granularity level, over a manually audited 10% sample (10 proofs, 97 NL steps).

manually reviewed. The main source of false positive and false negative is the scaffolder generating unfaithful lemma statements. We manually review all cases and report the analysis in Appendix D.

We conduct an ablation study on all ten problems (five student solutions per problem) by withholding supporting mini-lemmas from the lemma prover (Section 4.3). Appendix E shows that this ablation reduces accuracy between 14% to 48% across problems.

5.2 Faithfulness

We measure faithfulness as defined in Section 3.1, defining the correspondence relation $\sim$ to be LLM-as-a-judge that decides whether a natural language unit and a informalized formal-language unit convey the same mathematical content. We randomly picked and manually audited 10% of these decisions to validate them. Table 2 shows the human and LLM disagreement report. Appendix B details our methodology.

We report information preservation and non-dilution across the 100 natural-language proofs at varying granularity (Table 1). For 26 of the 100 solutions, the system does not generate a complete formal-language proof, due to lemma statements that failed to compile. For example, in Figure 6, a student step involving a matrix with mismatched dimensions fails the Lean 4 type checker, as do all steps that follow. Information preservation naturally suffers here, since fewer formal statements




Figure 6: An example student solution where the formal lemma statement does not compile. The solution contains a dimension error (red): the student writes Σ as an $m \times m$ matrix when it should be $m \times n$.

are produced; we therefore report separate results for the 74 problems that yield a complete formal-language proof.

Our results show a trade-off between non-dilution and information preservation as unit resolution changes: finer units expose more FL content, raising information preservation ($0.72 \rightarrow 0.76$) but lowering non-dilution ($0.79 \rightarrow 0.40$) through additional low-level scaffolding. None of the resolution levels tested approached perfect information preservation and non-dilution. We also approached faithfulness measurement using cosine similarity in a vector embedding space (Appendix C), but found it unreliable: semantic similarity does not track mathematical equivalence. A verbose student step and its clean informalized equation can be equivalent yet score low, while a minor auto-correction (e.g., an added absolute value) can alter the mathematical meaning yet score high.

5.3 Appropriate Granularity

We measure granularity using the metric defined in Section 3.2. To do this, we run the pipeline on one problem (P10 in Appendix A), using 50 student submissions. Each submission is compared to the ground truth grades assigned by a TA using a 7-item rubric worth 8 points in total (Appendix G.1). Each rubric item corresponds to one or a combination of mini-lemmas defined in Section 4.4. We explain the mapping between rubric items and mini-lemmas in Appendix F.

For each student proof, we collect the *main mini lemmas* used by its successfully verified scaffold lemmas (Section 4.4). The detector then flags a mini lemma as insufficiently justified when it appears *only* in large-jump scaffold steps (Section 4.5). Across the 50 proofs, submissions invoke 4.1 main mini lemmas on average, and 71 of the 540 scaffold lemmas in the corpus (13%) were flagged as large-jumps. *Strict mode* removes steps flagged as large jumps before grading, whereas




Figure 7: Per-item agreement between our system and the human grader under strict mode (left) and lenient mode (right) for problem P10 ($n = 50$ submissions). Each bar shows the fraction of submissions where our system exactly matched the human (green), over-credited (red), or under-credited (blue). Strict mode achieves higher exact-match rates on most rubric items.

lenient mode keeps them.

5.4 Granularity Results

We hypothesized that the strict mode would agree more closely with human graders than lenient mode, and our experiment confirms this. Strict mode agrees with the human grader on 62.1% of item-level decisions versus 57.1% for lenient mode (an increase of 5.0%). Strict mode matches or exceeds lenient on every rubric item and is strictly better on four of seven (Figure 7). These results provide preliminary evidence that the granularity detector is correctly identifying insufficiently justified steps.

A limitation of this evaluation is that rubric items are scored solely based on supporting mini lemmas. However, human graders also award credit for intermediate results, correct use of hypotheses, and other forms of proof progress that may not be reflected in supporting mini lemmas. Extending the rubric representation to capture these forms of evidence is an important direction for future work.

6 Related Work

Faithful Autoformalization While many prior works on autoformalization focus on translating informal mathematical statements into Lean propositions (Ying et al., 2024; Gao et al., 2024; Wu et al., 2022; Azerbayev et al., 2023), whole-proof autoformalization is comparatively more difficult (Jiang et al., 2022; Huang et al., 2024; Patel et al., 2026; Tarrach et al., 2024). A natural desiderata for whole-proof autoformalization is faithfulness – the requirement that a formal proof accurately reflects both the structure and semantics of its natural-language counterpart. Metrics to evaluate faithfulness assess syntactic correctness, semantic similarity, and structural equivalence (Jiang et al., 2023;

Moore and Shah, 2025; Lu et al., 2024a; Liu et al., 2025a; Li et al., 2024) between informal and autoformalized proofs. We extend these criteria by defining information preservation and non-dilution, which extend structural equivalence by requiring a one-to-one pairing between formal-language and natural-language units.

Several approaches have been proposed for faithful proof-level autoformalization. LeanTutor (Patel et al., 2026) formalizes informal proofs step-by-step to preserve semantic equivalence, and ProofFlow (Cabral et al., 2025) constructs an explicit dependency graph of proof steps before formalizing each node as a lemma, preserving structural fidelity. Quan et al. (2025) and Guo et al. (2025) parse the structure of informal proofs to reduce downstream ambiguity. Despite this progress, faithfulness remains an open challenge: LLMs have been observed to mistranslate informal proofs, fabricate or omit axioms, and produce logically inconsistent formal statements (Kim et al., 2026; Moore and Shah, 2025).

Appropriate Granularity The role of granularity—the level of detail at which a system operates relative to the user’s reasoning—has been acknowledged but not fully studied in the autoformalization literature. LeanTutor (Patel et al., 2026) assumes a one-to-one mapping between Lean 4 tactics and informal steps, but observes that more complex operations in Lean 4 lack direct natural-language correspondents. Nevertheless, following the user’s level of granularity is valuable; Tarrach et al. (2024) iteratively elaborates informal proof steps when formalization fails, and Wang et al. (2025) generates formal proof states aligned with the informal proof’s logical structure. When evaluating formalization faithfulness, ProofFlow (Cabral et al., 2025) acknowledges that multiple levels of formal-language granularity may exist and be valid, but does not attempt to further categorize them. Our work can be viewed as a development in the spirit of Guo et al. (2025), which identifies the “unspoken logic” implicit in informal proofs and develops a preprocessing pipeline that surfaces this latent structure before formalization; we explicitly define and study the different levels of formal language granularity, and argue that granularity should be *calibratable to a specific use case*; a system grading proofs in an introductory course must demand more justification at each step than one used in a graduate course.

Accuracy Common automated theorem proving and autoformalization evaluations often report compile success, proof success, or $\text{pass}@k$ (Azerbaiyev et al., 2023; Ying et al., 2024; Wu et al., 2022; Yang et al., 2023; Xin et al., 2024a; Lin et al., 2025; Ren et al., 2025; Chen et al., 2025b). Our step-level framing is closer to work on natural-language proof verification and proof grading (Hu et al., 2025; Patel et al., 2026; Ma et al., 2025; Mahdavi et al., 2025; Gonzalez et al., 2026; Mahinpei et al., 2026).

7 Limitations

This work evaluates the autoformalization harness of an educational NL co-reasoning system for math proofs, not a full end-to-end tutoring system. We do not evaluate the natural-language feedback shown to students or whether using the system improves proof-writing ability; we leave this to future work. Our results should therefore be read as evidence about the formalization and verification pipeline rather than about educational effectiveness.

Our experiments are limited to undergraduate linear algebra proofs from one course, so performance may differ in other mathematical domains or courses with different justification standards.

The supporting mini-lemmas are also problem-specific and human-authored, so the reported results do not reflect the cost or reliability of constructing them automatically.

The accuracy evaluation uses an LLM-based reference evaluator with manual review of disagreements. This provides a useful proxy, though some student steps are ambiguous or valid only under course-specific conventions, making step-level correctness inherently difficult to label in edge cases.

Our faithfulness metrics are proxy measures that depend on proof unitization, LLM informalization, embedding similarity, and one-to-one matching. These choices may not capture every valid correspondence between informal and formal steps, and do not fully encode dependency structure or proof order.

The granularity evaluation is indirect: we measure whether strict grading improves agreement with a human rubric rather than directly measuring large-jump detection. As it is also run on a single problem, these results should be treated as preliminary.

8 Acknowledgment

The authors would like to acknowledge support from: NSF CAREER grant ECCS-2240031, the 2023 CITRIS Institute Seed Grant, the UC Berkeley Instructional Technology and Innovation MicroGrant Program, Google and the Renaissance Philanthropy AI for Math 2025 fund.

The authors would also like to thank Aditya Chandrasekhar, Dakota Feldkamp, and Krishiv Kothari for their work on the human auditing of the evaluation results.

References

- Jeremy Avigad. 2019. [Learning logic and proof with an interactive theorem prover](#). In Gila Hanna, David A. Reid, and Michael de Villiers, editors, *Proof Technology in Mathematics Research and Teaching*, volume 14 of *Mathematics Education in the Digital Era*, pages 277–290. Springer, Cham.
- Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W Ayers, Dragomir Radev, and Jeremy Avigad. 2023. Proofnet: Autoformalizing and formally proving undergraduate-level mathematics. *arXiv preprint arXiv:2302.12433*.
- Tiffany Barnes and John Stamper. 2008. [Toward automatic hint generation for logic proof tutoring using historical student data](#). In *Intelligent Tutoring Systems (ITS 2008)*, volume 5091 of *Lecture Notes in Computer Science*, pages 373–382, Berlin, Heidelberg, Springer.
- David E. Brown and Shayla Michel. 2010. Assessing proofs with rubrics: The RVF method. In *Proceedings of the 13th Conference on Research in Undergraduate Mathematics Education*.
- Rafael Cabral, Tuan Manh Do, Xuejun Yu, Wai Ming Tai, Zijin Feng, and Xin Shen. 2025. ProofFlow: A dependency graph approach to faithful proof autoformalization. *arXiv preprint arXiv:2510.15981*.
- Jiangjie Chen, Wenxiang Chen, Jiacheng Du, Jinyi Hu, Zhicheng Jiang, Allan Jie, Xiaoran Jin, Xing Jin, Chenggang Li, Wenlei Shi, Zhihong Wang, Mingxuan Wang, Chenrui Wei, Shufa Wei, Huajian Xin, Fan Yang, Weihao Gao, Zheng Yuan, Tianyang Zhan, Zeyu Zheng, Tianxi Zhou, and Thomas Hanwen Zhu. 2025a. [Seed-prover 1.5: Mastering undergraduate-level theorem proving via learning from experience](#).
- Luoxin Chen, Jinming Gu, Liankai Huang, Wenhao Huang, Zhicheng Jiang, Allan Jie, Xiaoran Jin, Xing Jin, Chenggang Li, Kaijing Ma, et al. 2025b. Seed-Prover: Deep and broad reasoning for automated theorem proving. *arXiv preprint arXiv:2507.23726*.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Leonardo de Moura and Sebastian Ullrich. 2021. The Lean 4 theorem prover and programming language. In *International Conference on Automated Deduction*, pages 625–635. Springer.
- Don Fallis. 2003. [Intentional gaps in mathematical proofs](#). *Synthese*, 134(1–2):45–69.
- Guoxiong Gao, Yutong Wang, Jiedong Jiang, Qi Gao, Zihan Qin, Tianyi Xu, and Bin Dong. 2024. Herald: A natural language annotated lean 4 dataset. *arXiv preprint arXiv:2410.10878*.
- Santiago Gonzalez, Alireza Amiri Bavandpour, Peter Ye, Edward Zhang, et al. 2026. QEDBench: Quantifying the alignment gap in automated evaluation of university-level mathematical proofs. *arXiv preprint arXiv:2602.20629*.
- Aryan Gulati, Devanshu Ladsaria, Shubhra Mishra, Jasdeep Sidhu, and Brando Miranda. 2024. [An evaluation benchmark for autoformalization in Lean4](#). In *ICLR 2024 Tiny Papers*.
- Chenjun Guo, Manooshree Patel, Bjoern Hartmann, JD Zamfirescu-Pereira, Sarah Chasins, and Gireeja Ranade. 2025. Unspoken logic: Understanding and bridging the gap between free-form and llm-interpretable natural language mathematical proofs. In *The 5th Workshop on Mathematical Reasoning and AI at NeurIPS 2025*.
- Alena Gusakov, Peter Nelson, and Stephen M. Watt. 2025. [Structuring definitions in mathematical libraries](#). *arXiv preprint arXiv:2509.10828*.
- Mohammad Hekmatnejad, Andrew M. Simms, and Georgios Fainekos. 2019. [Model checking clinical decision support systems using SMT](#). ArXiv preprint arXiv:1901.04545 (v2, revised 4 March 2019).
- Xiaolin Hu, Qinghua Zhou, Bogdan Grechuk, and Ivan Y. Tyukin. 2025. [Stepproof: Step-by-step verification of natural language mathematical proofs](#).
- Yanxing Huang, Xinling Jin, Sijie Liang, Peng Li, and Yang Liu. 2025. [Formarl: Enhancing autoformalization with no labeled data](#).
- 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. *arXiv preprint arXiv:2402.08957*.
- Amogh Inamdar, Uzay Macar, Michel Vazirani, Michael Tarnow, Zarina Mustapha, Natalia Dittren, Sam Sadeh, Nakul Verma, and Ansaf Salheb-Aouissi. 2025. [LogicLearner: A tool for the guided practice of propositional logic proofs](#).

- Albert Q Jiang, Wenda Li, and Mateja Jamnik. 2023. Multilingual mathematical autoformalization. *arXiv preprint arXiv:2311.03755*.
- Albert Q Jiang, Sean Welleck, Jin Peng Zhou, Wenda Li, Jiacheng Liu, Mateja Jamnik, Timothée Lacroix, Yuhuai Wu, and Guillaume Lample. 2022. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. *arXiv preprint arXiv:2210.12283*.
- William Jurayj, Nils Holzenberger, and Benjamin Van Durme. 2026. [Language models and logic programs for trustworthy tax reasoning](#).
- Matthew Keenan and Cyrus Omar. 2024. [Learner-centered design criteria for classroom proof assistants](#). In *Workshop on Human Aspects of Types and Reasoning Assistants (HATRA)*.
- Kyuhee Kim, Auguste Poiroux, and Antoine Bosselut. 2026. Do LLMs game formalization? Evaluating faithfulness in logical reasoning. In *ICLR 2026 Workshop on Verified AI (VerifAI-2)*.
- Zenan Li, Yifan Wu, Zhaoyu Li, Xinming Wei, Fan Yang, Xian Zhang, and Xiaoxing Ma. 2024. Autoformalize mathematical statements by symbolic equivalence and semantic consistency. *Advances in Neural Information Processing Systems*, 37:53598–53625.
- Yong Lin, Shange Tang, Bohan Lyu, Ziran Yang, Jui-Hui Chung, Haoyu Zhao, Lai Jiang, Yihan Geng, Jiawei Ge, Jingruo Sun, et al. 2025. Goedel-prover-v2: Scaling formal theorem proving with scaffolded data synthesis and self-correction. *arXiv preprint arXiv:2508.03613*.
- Junqi Liu, Zihao Zhou, Zekai Zhu, Marco Dos Santos, Weikun He, Jiawei Liu, Ran Wang, Yunzhou Xie, Junqiao Zhao, Qiufeng Wang, et al. 2026. Numina-Lean-Agent: An open and general agentic reasoning system for formal mathematics. *arXiv preprint arXiv:2601.14027*.
- Qi Liu, Xinhao Zheng, Xudong Lu, Qinxiang Cao, and Junchi Yan. 2025a. [Rethinking and improving autoformalization: Towards a faithful metric and a dependency retrieval-based approach](#). In *The Thirteenth International Conference on Learning Representations*.
- Xiaoyang Liu, Kangjie Bao, Jiashuo Zhang, Yunqi Liu, Yu Chen, Yuntian Liu, Yang Jiao, and Tao Luo. 2025b. Atlas: Autoformalizing theorems through lifting, augmentation, and synthesis of data. *arXiv preprint arXiv:2502.05567*.
- Josje Lodder, Bastiaan Heeren, Johan Jeuring, and Wendy Neijenhuis. 2021. [Generation and use of hints and feedback in a hilbert-style axiomatic proof tutor](#). *International Journal of Artificial Intelligence in Education*, 31(1):99–133.
- Jianqiao Lu, Yingjia Wan, Yinya Huang, Jing Xiong, Zhengying Liu, and Zhijiang Guo. 2024a. [Formalalign: Automated alignment evaluation for autoformalization](#). *arXiv preprint arXiv:2410.10135*.
- Jianqiao Lu, Yingjia Wan, Zhengying Liu, Yinya Huang, Jing Xiong, Chengwu Liu, Jianhao Shen, Hui Jin, Jipeng Zhang, Haiming Wang, et al. 2024b. Process-driven autoformalization in lean 4. *arXiv preprint arXiv:2406.01940*.
- Minghai Lu, Benjamin Delaware, and Tianyi Zhang. 2024c. Proof automation with large language models. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering (ASE)*.
- Wenjie Ma et al. 2025. Reliable fine-grained evaluation of natural language math proofs. *arXiv preprint arXiv:2510.13888*.
- Hamed Mahdavi, Pouria Mahdavinia, Samira Malek, Pegah Mohammadipour, Alireza Hashemi, Majid Daliri, Alireza Farhadi, Amir Khasahmadi, Niloo-far Mireshghallah, and Vasant Honavar. 2025. Ref-Grader: Automated grading of mathematical competition proofs using agentic workflows. *arXiv preprint arXiv:2510.09021*.
- Romina Mahinpei, Sofia Druchyna, and Manoel Horta Ribeiro. 2026. When LLMs help – and hurt – teaching assistants in proof-based courses. *arXiv preprint arXiv:2602.23635*.
- The mathlib Community. 2020. [The lean mathematical library](#). In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, POPL ’20*, page 367–381. ACM.
- David Miller, Nicholas Infante, and Keith Weber. 2018. How mathematicians assign points to student proofs. *Journal of Mathematical Behavior*, 49:24–34.
- Hayden Moore and Asfahan Shah. 2025. Evaluating autoformalization robustness via semantically similar paraphrasing. *arXiv preprint arXiv:2511.12784*.
- Robert C. Moore. 2016. [Mathematics professors’ evaluation of students’ proofs: A complex teaching practice](#). *International Journal of Research in Undergraduate Mathematics Education*, 2(2):246–278.
- Logan Murphy, Kaiyu Yang, Jialiang Sun, Zhaoyu Li, Anima Anandkumar, and Xujie Si. 2024. Autoformalizing euclidean geometry. *arXiv preprint arXiv:2405.17216*.
- Naoto Onda, Kazumi Kasaura, Yuta Oriike, Masaya Taniguchi, Akiyoshi Sannai, and Sho Sonoda. 2025. LeanConjecturer: Automatic generation of mathematical conjectures for theorem proving. *arXiv preprint arXiv:2506.22005*.
- Azim Ospanov, Zijin Feng, Jiacheng Sun, Haoli Bai, Xin Shen, and Farzan Farnia. 2025. [HERMES: Towards efficient and verifiable mathematical reasoning in LLMs](#). *arXiv preprint arXiv:2511.18760*.

- Manooshree Patel, Rayna Bhattacharyya, Thomas Lu, Arnav Mehta, Niels Voss, Narges Norouzi, and Gireeja Ranade. 2026. [Leantutor: Towards a verified AI mathematical proof tutor](#). In *Proceedings of the Fortieth AAAI Conference on Artificial Intelligence*, AAAI '26, pages 40670–40678. AAAI Press.
- Lawrence C. Paulson. 1994. Isabelle: A generic theorem prover. In *Proceedings of the 12th International Conference on Automated Deduction (CADE-12)*, volume 814 of *Lecture Notes in Computer Science*, pages 37–41. Springer.
- Auguste Poiroux et al. 2024. [Reliable evaluation and benchmarks for statement autoformalization](#).
- PramaanaLabs. 2026. Pramaana labs. <https://www.pramaanalabs.ai/>. Accessed: 2026-05-08.
- Xin Quan, Marco Valentino, Louise A Dennis, and André Freitas. 2025. Faithful and robust llm-driven theorem proving for nli explanations. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 17734–17755.
- ZZ Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanxia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, et al. 2025. Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. *arXiv preprint arXiv:2504.21801*.
- Jessica Shi, Cassia Torczon, Harrison Goldstein, Benjamin C. Pierce, and Andrew Head. 2025. [Qed in context: An observation study of proof assistant users](#). *Proceedings of the ACM on Programming Languages*, 9(OOPSLA1):1–27.
- Wilfried Sieg. 1990. *Proof theory*. Carnegie Mellon University.
- Vikash Singh, Debargha Ganguly, Haotian Yu, Chengwei Zhou, Perna Singh, Brandon Lee, Vipin Chaudhary, and Gourav Datta. 2026. [Toward guarantees for clinical reasoning in vision language models via formal verification](#).
- Peiyang Song, Kaiyu Yang, and Anima Anandkumar. 2024. Towards large language models as copilots for theorem proving in lean. *arXiv preprint arXiv:2404.12534*.
- Guillem Tarrach, Albert Q Jiang, Daniel Raggi, Wenda Li, and Mateja Jamnik. 2024. More details, please: Improving autoformalization with more detailed proofs. In *AI for Math Workshop@ ICML 2024*.
- Amitayush Thakur et al. 2025. Apollo: Automated llm and lean collaboration for advanced formal reasoning. *arXiv preprint arXiv:2505.05758*.
- Danielle R. Thomas, Conrad Borchers, and Kenneth R. Koedinger. 2025. Beyond agreement: Rethinking ground truth in educational AI annotation. *arXiv preprint arXiv:2508.00143*.
- Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi, and Tong Zhang. 2024. Theoremllama: Transforming general-purpose llms into lean4 experts. *arXiv preprint arXiv:2407.03203*.
- Ziyu Wang, Bowen Yang, Chenyi Li, Yuan Zhang, Shihao Zhou, Bin Dong, and Zaiwen Wen. 2025. Translating informal proofs into formal proofs using a chain of states. *arXiv preprint arXiv:2512.10317*.
- Keith Weber. 2008. How mathematicians determine if an argument is a valid proof. *Journal for Research in Mathematics Education*, 39(4):431–459.
- Sean Welleck and Rahul Saha. 2023. LLMSTEP: LLM proofstep suggestions in Lean. *arXiv preprint arXiv:2310.18457*.
- Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus Rabe, Charles Staats, Mateja Jamnik, and Christian Szegedy. 2022. Autoformalization with large language models. *Advances in Neural Information Processing Systems*, 35:32353–32368.
- Huajian Xin, Daya Guo, Zhihong Shao, Zhizhou Ren, Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li, and Xiaodan Liang. 2024a. Deepseek-prover: Advancing theorem proving in llms through large-scale synthetic data. *arXiv preprint arXiv:2405.14333*.
- Huajian Xin, Z.Z. Ren, Junxiao Song, Zhihong Shao, Wanxia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, et al. 2024b. Deepseek-prover-v1.5: Harnessing proof assistant feedback for reinforcement learning and monte-carlo tree search. *arXiv preprint arXiv:2408.08152*.
- Erna Yackel and Paul Cobb. 1996. Sociomathematical norms, argumentation, and autonomy in mathematics. *Journal for research in mathematics education*, 27(4):458–477.
- Kaiyu Yang, Aidan M Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan Prenger, and Anima Anandkumar. 2023. LeanDojo: Theorem proving with retrieval-augmented language models. In *Advances in Neural Information Processing Systems*, volume 36.
- Huaiyuan Ying, Zijian Wu, Yihan Geng, Jiayu Wang, Dahua Lin, and Kai Chen. 2024. Lean workbook: A large-scale lean problem set formalized from natural language math problems. *Advances in Neural Information Processing Systems*, 37:105848–105863.
- Xingguang Zhang et al. 2025. Leanabell-prover-v2: Verifier-integrated reasoning for formal theorem proving via reinforcement learning. *arXiv preprint arXiv:2507.08649*.
- Yichi Zhou, Jianqiu Zhao, Yongxin Zhang, Bohan Wang, Siran Wang, Luoxin Chen, Jiahui Wang, Haowei Chen, Allan Jie, Xinbo Zhang, et al. 2025. Solving formal math problems by decomposition and iterative reflection. *arXiv preprint arXiv:2507.15225*.

A Problem Statements

The following ten problems from EECS 127 were used in our evaluation. Each problem was attempted by $n = 10$ students.

1. **Q1: Eigenvector Orthogonality**

Let $A \in S^n$ (i.e., the set of $n \times n$ symmetric matrices) and $(\lambda_1, \mathbf{u}_1), (\lambda_2, \mathbf{u}_2)$, with $\lambda_1 \neq \lambda_2$, be distinct eigen-pairs of A . Show that $\mathbf{u}_1^\top \mathbf{u}_2 = 0$, i.e., eigenspaces corresponding to distinct eigenvalues are mutually orthogonal.

2. **Q2: Norm Inequality:** $\|\cdot\|_\infty \leq \|\cdot\|_2$

Show that the following inequality holds for any $\vec{x} \in \mathbb{R}^n$:

$$\|x\|_\infty \leq \|x\|_2.$$

3. **Q3: Norm Inequality:** $\|\cdot\|_1 \leq \sqrt{n} \|\cdot\|_2$

Show that the following inequality holds for any $\vec{x} \in \mathbb{R}^n$:

$$\|x\|_1 \leq \sqrt{n} \cdot \|x\|_2.$$

4. **Q4: Frobenius Norm: Zero Iff**

Let P be an $m \times n$ matrix. Show that the Frobenius norm satisfies $\|P\|_F = 0$ if and only if $P = 0$.

5. **Q5: Frobenius Norm: Homogeneity**

Let P be an $m \times n$ matrix. Show that for all $a \in \mathbb{R}$, $\|aP\|_F = |a| \|P\|_F$.

6. **Q6: Frobenius Norm: Triangle Inequality**

Let P and Q be $m \times n$ matrices. Show that $\|P + Q\|_F \leq \|P\|_F + \|Q\|_F$.

7. **Q7: Frobenius Norm via Singular Values**

Let P be an $m \times n$ matrix with singular values $\sigma_1, \dots, \sigma_n$. Show that $\|P\|_F^2 = \sum_{i=1}^n \sigma_i^2$.

8. **Q8: PSD Implies Non-negative Diagonal**

A symmetric matrix $M \in \mathbb{R}^{n \times n}$ is positive semidefinite if $\vec{x}^\top M \vec{x} \geq 0$ for all $\vec{x} \in \mathbb{R}^n$. Assume $A \in \mathbb{R}^{n \times n}$ is symmetric. Show that if $A \succeq 0$ then $A_{ii} \geq 0$ for all i .

9. **Q9: $Q^\top Q$ is PSD**

Show that for any matrix $Q \in \mathbb{R}^{m \times n}$, $A = Q^\top Q$ is positive semidefinite.

10. **Q10: Constrained Eigenvalue Optimization**

Let $A \in \mathbb{R}^{n \times n}$ be symmetric with eigenvalue-eigenvector pairs

$(\lambda_1, \vec{v}_1), \dots, (\lambda_n, \vec{v}_n)$ where $\lambda_1 > \dots > \lambda_n$.

Consider

$$p^* = \max_{\vec{x} \in \mathbb{R}^n} \vec{x}^\top A \vec{x}$$

$$\text{s.t. } \|\vec{x}\|_2^2 = 1 \quad \vec{x}^\top \vec{v}_1 = 0.$$

Show that $p^* = \lambda_2$.

B Description of Faithfulness Evaluation

Given a formal proof in Lean 4, we extract content units at the three levels of granularity given in Section 3.1.3. At all three levels, our units consist of Lean 4 goal states, in order to fully disambiguate variable names and definitions. For example, a top-level unit could consist of the goal state immediately following the statement of a theorem. All formal units are then informalized into natural language using Claude Sonnet 4.6. Below, we give the prompt used at each level, together with an example of an original formal statement and its informalized version. We observed that the LLM judge is less accurate at the tactic level than at the other levels. We hypothesize the formatting of “Before... After...” could confuse the LLM judge.

Top level. Prompt:

```
f"""--- ENTRY {idx} ---
Kind: {entry.kind}
Name: {entry.name}
Lean source:
{entry.lean_code}

Source-level proposition, if explicit:
{entry.proposition}

Lean-derived checked proposition:
{checked_prop}

Lean goal state to informalize:
{goal_state}
"""
)

return f"""PROBLEM CONTEXT
-----
This is {problem_label}. The overall theorem/
problem context is:

{lean_theorem}

NOTATION CONVENTIONS USED IN THIS PROBLEM
-----
{custom_instructions}

LEAN GOAL STATES TO INFORMALIZE ({len(batch)})
-----
{chr(10).join(items)}

TASK
----
```

For each entry, informalize the mathematical claim represented by the Lean goal state to informalize.

Usually this means informalizing the `...` line. For inferred `have` statements such as `have := term`, the displayed goal state has been synthesized so that the newly-created hypothesis type appears as the goal; informalize that created fact.

Do not informalize the unchanged outer goal when the entry is an inferred `have` that merely adds a hypothesis.

Return JSON with a single field `informalizations`, an array of length $\{\text{len}(\text{batch})\}$ in the same order as the entries above. Each entry has one field: `informalization`.

Input example:

```
lemma step_3_transpose_mul (A : Matrix (Fin n) (Fin n) ) ... (h_A_symm : A.IsSymm) :
  A * v1 = lamb1 v1
goal fed to the model: A * v1 = lamb1 v1
```

Output informalization:

$$A^T v_1 = \lambda_1 v_1$$

Sub-claims level. Prompt:

```
f"""--- ENTRY {idx} ---
Kind: {entry.kind}
Name: {entry.name}
Lean source:
{entry.lean_code}

Source-level proposition, if explicit:
{entry.proposition}

Lean-derived checked proposition:
{checked_prop}

Lean goal state to informalize:
{goal_state}
"""

    )

    return f"""PROBLEM CONTEXT
-----
This is {problem_label}. The overall theorem/
problem context is:

{lean_theorem}

NOTATION CONVENTIONS USED IN THIS PROBLEM
-----
{custom_instructions}

LEAN GOAL STATES TO INFORMALIZE ({len(batch)})
-----
{chr(10).join(items)}

TASK
----
```

For each entry, informalize the mathematical claim represented by the Lean goal state to informalize.

Usually this means informalizing the `...` line. For inferred `have` statements such as `have := term`, the displayed goal state has been synthesized so that the newly-created hypothesis type appears as the goal; informalize that created fact.

Do not informalize the unchanged outer goal when the entry is an inferred `have` that merely adds a hypothesis.

Return JSON with a single field `informalizations`, an array of length $\{\text{len}(\text{batch})\}$ in the same order as the entries above. Each entry has one field: `informalization`.

Input example:

```
have h_left : (A * v1) v2 = lamb1 * (v1 v2) :=
  step_5_symm_apply
goal fed to the model: A * v1 v2 = lamb1 * v1
v2
```

Output informalization:

$$(Av_1) \cdot v_2 = \lambda_1(v_1 \cdot v_2)$$

Tactic level. Prompt:

```
PROBLEM CONTEXT
-----
This is {problem_label}. The overall theorem we
are proving is:

{lean_theorem}

NOTATION CONVENTIONS USED IN THIS PROBLEM
-----
{custom_instructions}

PROOF STEPS TO INFORMALIZE ({len(batch)}) steps
-----
{steps_block}

TASK
----
For EACH step above, informalize ONLY the goal (
the `` line(s)) of the BEFORE
and AFTER states into one short, human-style
English statement of what is being
claimed. Use LaTeX where helpful. Do NOT list
all the hypotheses; assume the
standard problem context. If a post-state is "no
goals", set its
informalized_end to "Proof complete." (or
similar).

Return JSON with a single field `
informalizations`, an array of length
{len(batch)} in the same order as the steps
above. Each entry has two fields:
informalized_start, informalized_end.
```

Input example:

```

exact hv1_eigen
goal fed to the model:
BEFORE: A * v1 = lamb1 v1
AFTER: A * v1 v2 = lamb1 * v1 v2

```

Output informalization:

Before this step, the goal is: $Av_1 = \lambda_1 v_1$
After this step, the goal is: $(A^T v_1) \cdot v_2 = \lambda_1(v_1 \cdot v_2)$

C Description of Faithfulness Evaluation (embedding model)

The embeddings of the informalizations are computed using BGE-M3. The natural-language units, produced using the unitization procedure of (Guo et al., 2025), are embedded using the same model, and we compute the formal-language natural-language pairing that maximizes total cosine similarity using the Hungarian method. Our pairing is one-to-one; each unit can be paired to at most one other unit, and only pairings with a cosine similarity of 0.6 or greater are permitted. Based on the number of paired and unpaired units, we compute information preservation and non-dilution scores following Section 3.1.

D Accuracy examples

There are 86 cases where the LLM evaluator marks a step as correct but the system marks it as incorrect. Manual review shows that 6 are LLM evaluator errors, and 4 are lemma-prover failures, where the prover fails to prove a faithful scaffolded lemma (Sections 4.3 and 4.2). The remaining cases occur when **the scaffolder generates an unfaithful or invalid lemma because the student writes ambiguous or under-specified steps**. Appendix ?? gives three examples in which students omit declarations, types, or assumptions that humans can infer, but the system produces false negatives due to unfaithful lemma statements.

Example 1: Ambiguous assumption markers.

In Figure 8 (bottom), the proof requires both directions of an iff statement. For the forward direction, the system must add $\|P\|_F = 0$ as a hypothesis before proving $P = 0$. However, the student begins with “Given $\|P\|_F = 0$,” which the preprocessor classifies as a restatement of known context rather than as an assumption for the current proof goal. As a result, the generated Lean4 statement omits the needed hypothesis (`h_norm_zero : P.frobeniusNorm = 0`), making the claim false.

Example 2: Undeclared variables and implicit conventions.

In Figure 8 (top), the student writes $aP = a \sum_i \sigma_i u_i v_i^\top = \sum_i a \sigma_i u_i v_i^\top$. The variables i , σ_i , u_i , and v_i are not declared, but a human reader infers them from notation used in lectures and textbooks, where U and V commonly refer to the left and right singular-vector matrices in an SVD. A faithful formalization would need either an explicit SVD-expansion hypothesis for P , or explicit declarations of σ_i , u_i , and v_i with their types and index ranges.

Example 3: Intuitively correct but formally invalid steps.

In Figure 8 (center), the student writes $\langle P, P \rangle + 2\langle P, B \rangle + \langle B, B \rangle = \|P\|_F^2 + 2\|PB\|_F^2 + \|B\|_F^2$. The term $\|PB\|_F$ is undefined because P and B are both $m \times n$ matrices, so the system cannot prove the step. However, because the intended expansion is clear, a human grader or LLM evaluator may not focus on notation error and mark the step as correct.

Other false negatives arise from two cases: (1) the student asserts a correct but irrelevant step, or a step that uses a non-trivial fact, absent from our mini helper lemma library; and (2) The proof of a scaffold lemma depends on an earlier scaffold lemma whose proof has failed.

For the 24 false positive cases, we identify three causes. First, despite the scaffold lemma correction loop being instructed never to alter mathematical content (Section ??), we observe it doing so for some type errors — a design flaw to be addressed in future work. Second, some student steps are mathematically true only in context, e.g., `expressionA * a = expressionB * a` holds trivially when `a = 0`, but an LLM judging the claim in isolation marks it false. Third, the preprocessor (Section 4.2) occasionally misclassifies an assertion as a declaration; since declarations are always treated as valid, this produces a false positive.

E Mini-Lemma Ablation

We evaluate the contribution of supporting mini-lemmas on all ten problems, using five student solutions per problem. In the ablated setting, the lemma prover has no access to these mini-lemmas. As shown in Table 3, removing them reduces accuracy by 14–48 percentage points across questions.




Figure 8: examples for student steps which lead to unfaithful autoformalization

F Mapping of mini-lemmas to rubric items

We use the mini lemmas in a student proof to assign rubric grades. A rubric item may correspond to one or more mini lemmas. Therefore we assign a logical condition over the set of mini lemmas to each rubric item; the item earns credit when that condition is satisfied.

1. $\text{any_of}(L)$: the student's proof contains at least one mini lemma from the set L ; we use this when there are multiple admissible mathematical facts which equivalently demonstrate the rubric item.
2. $\text{all_of}(L)$ the student's proof contains all mini lemma from the set L ; this is used when multiple mathematical facts are necessary to demonstrate the rubric item.
3. $\text{all_any}(L_1, \dots, L_k)$: the student's proof

Problem	Ablated (%)	Full (%)	Gain (%)
Q1	45	84	39
Q2	50	83	33
Q3	40	88	48
Q4	49	70	21
Q5	39	71	32
Q6	26	70	44
Q7	27	62	35
Q8	49	63	14
Q9	30	63	33
Q10	57	73	16

Table 3: Lemma-prover accuracy with and without supporting mini-lemmas. Gain denotes the improvement of the full system over the ablated setting in percentage points.

contains at least one mini lemma from each of the sets $L_1, \dots, L_k$; this is used when multiple mathematical facts are necessary and each of them may be demonstrated in several equivalent ways.

G Supporting mini lemma example

Problem statement:

```

theorem Problem4_ineq1
  {n : ℕ}
  (hn : 0 < n)
  (x : Fin n → ℝ) :
  1 / Real.sqrt ↑n * ||WithLp.toLp 2 x|| ≤
  ↪ ||WithLp.toLp T x|| :=
---- AXIOMS START
-- @side
-- max x_i ≥ 0 when all x_i ≥ 0
theorem max_is_positive {n : ℕ} (x : Fin n → ℝ)
(h : ∀ i : Fin n, 0 ≤ x i) : 0 ≤ ⌊ i : Fin n, x
↪ i := by
  sorry

-- @side
-- ||x||_2 = √(∑ |x_i|^2)
theorem def_l2_norm {n : ℕ} (x : Fin n → ℝ)
: ||WithLp.toLp 2 x|| = √(∑ i : Fin n, |x i|^2) := by
  ↪ 2) := by
  sorry

-- @side
-- ||x||_∞ = max |x_i|
theorem def_linfty_norm {n : ℕ} (x : Fin n → ℝ)
: ||WithLp.toLp T x|| = ⌊ i : Fin n, |x i| := by
  sorry

-- @side
-- x ≤ y iff x^2 ≤ y^2 for non-negative x, y
theorem square_bs_ineq (x y : ℝ) (hx : 0 ≤ x)
(hy : 0 ≤ y) : x ≤ y ↔ x^2 ≤ y^2 := by
  sorry

-- @side
-- |x|^2 = x^2
lemma abs_sq_eq_sq (x : ℝ) : |x|^2 = x^2
:= by
  sorry

```

```

-- @side
--  $(\max |x_i|)^2 = \max |x_i|^2$ 
theorem move_squares_into_max {n : ℕ} (x : Fin n
  → ℝ)
: (⋒ i : Fin n, |x i|) ^ 2
= ⋒ i : Fin n, |x i| ^ 2 := by
  sorry

--  $y/z \leq x$  iff  $y \leq x*z$  for positive  $z$ 
theorem division_ineq (x y : ℝ)
{z : ℝ} (hz : 0 < z) : y / z ≤ x ↔ y ≤ x * z
  → := by
  sorry

--  $x \leq y$  implies  $\sqrt{x} \leq \sqrt{y}$  for non-negative  $x$ ,
  → y
theorem sqrt_ineq (x y : ℝ)
(hx : 0 ≤ x) (hy : 0 ≤ y) : x ≤ y → √x ≤
  → √y := by
  sorry

-- @side
-- sum of a constant is the constant times the
  → number of elements in the sum
recall Finset.sum_const : ∀ {α : Type u_1} {M :
  → Type u_4} {s : Finset α}
[AddCommMonoid M] (b : M), ∑ _x ∈ s, b = s.card
  → • b

-- @side
--  $\sum c = n * c$ 
theorem me_Finset_sum_const {n : ℕ}
(c : ℝ) : ∑ _ : Fin n, c = ↑n * c := by
  simp

--  $f_i \leq g_i$  for all  $i$  implies  $\sum f_i \leq \sum g_i$ 
theorem me_Finset_sum_le_sum {n : ℕ}
(f g : Fin n → ℝ) (h : ∀ i : Fin n, f i ≤ g i) :
  → ∑ i : Fin n, f i ≤ ∑ i : Fin n, g i := by
  sorry

-- @side
recall le_ciSup : ∀ {α : Type u_1} {ι : Sort u_4}
  → [ConditionallyCompleteLattice α]
{f : ι → α},
  BddAbove (Set.range f) → ∀ (c : ι), f c ≤ iSup
  → f

--  $f_i \leq \max f_j$ 
theorem me_le_ciSup {n : ℕ}
(f : Fin n → ℝ) (i : Fin n) : f i ≤ ⋒ i : Fin
  → n, f i := by
  sorry

-- @side
--  $(1/\sqrt{n}) * \sqrt{n*x} = \sqrt{x}$  for positive  $n$  and
  → non-negative  $x$ 
theorem cancel_sqrt_n {n : ℕ}
(hn : 0 < n) (x : ℝ) (hx : 0 ≤ x) :
  1 / Real.sqrt n * Real.sqrt (n * x) = Real.sqrt
  → x := by
  sorry

-- @side
--  $1 \leq a*(b/c)$  iff  $c \leq a*b$  for positive  $c$ 
theorem mul_div_le_iff (a b c : ℝ)
(hc : 0 < c) : 1 ≤ a * (b / c) ↔ c ≤ a * b := by
  rw [div_eq_mul_inv, ← mul_assoc,
  → le_mul_inv_iff0 hc, one_mul]

```

```

-- @side
--  $\sum 1 = n$ 
theorem sum_one_over_n {n : ℕ} :
  (∑ _ : Fin n, 1 = ↑n) := by
  simp

-- @side
recall norm_smul {α : Type u_1} {β : Type u_2}
[Norm α] [Norm β] [SMul α β] [NormSMulClass α
  → β] (r : α) (x : β) :
  ||r • x|| = ||r|| * ||x||

-- @side
recall Real.norm_of_nonneg {r : ℝ}
(hr : 0 ≤ r) : ||r|| = r

-- @side
theorem one_div_sqrt_eq_sqrt_one_div (x : ℝ) :
  (1 / Real.sqrt x) = Real.sqrt (1 / x)
  := by
  rw [Real.sqrt_div]
  simp
  positivity

-- @side
theorem mul_one_div_sqrt (a n : ℝ) :
  a * (1 / Real.sqrt n) = a / Real.sqrt n := by
  field_simp

---- AXIOMS END

```

G.1 Rubric

These rubric items were crafted by the course instructor.

Second Principal Component (8 pts) *Shows that $\bar{x} = \bar{v}_2$ is achievable.*

1. +2 pts: Concludes that the feasible point $\bar{x} = \bar{v}_2$ attains $\bar{v}_2^\top A \bar{v}_2 = \bar{v}_2^\top \lambda_2 \bar{v}_2 = \lambda_2$.

2. +2 pts: Shows that $\bar{x} = \bar{v}_2$ is feasible: $\bar{v}_1^\top \bar{v}_2 = 0$ and $\|\bar{v}_2\|_2 = 1$.

Proves that $p^* \leq \lambda_2$.

1. +1 pt: Claims $\bar{x} \in \text{span}(\bar{v}_2, \dots, \bar{v}_n)$, e.g., $\bar{x} = \sum_{k=2}^n y_k \bar{v}_k$.

2. +1 pt: Substitutes into the objective and derives

$$\bar{x}^\top A \bar{x} = \sum_{i=2}^n \lambda_i y_i^2.$$

3. +1 pt: Uses

$$\lambda_i y_i^2 \leq \lambda_2 y_i^2 \quad (i \geq 2).$$

4. +1 pt: Concludes

$$\bar{x}^\top A \bar{x} \leq \lambda_2 \sum_{i=2}^n y_i^2 = \lambda_2.$$

5. +0 pts: Incorrect but non-blank.

6. +0 pts: Blank.

H Prompt for LLM evaluator

System prompt:

You are an expert mathematics professor evaluating whether individual proof steps are correct given the proof question. Respond with valid JSON only, no extra text.

User prompt:

QUESTION: question

STUDENT PROOF STEPS (OCR transcription, 1-indexed): steps

Evaluate the student's proof: - List the 1-based step indices of any mathematically INCORRECT steps in "incorrect_step" (use [] if every step is correct). - If "incorrect_step" is non-empty, set "incomplete" to "NA". - If "incorrect_step" is empty, judge whether the proof is INCOMPLETE — i.e. missing steps or cases needed to fully establish the result — and set "incomplete" to "Yes" or "No". - "reason": one concise explanation of your judgment.

Respond with JSON only: <"incorrect_step": [<ints>], "incomplete": "Yes"|"No"|"NA", "reason": "<text>">