

000 001 002 003 004 005 PROOFBRIDGE: AUTO-FORMALIZATION OF NATURAL 006 LANGUAGE PROOFS IN LEAN VIA JOINT EMBEDDINGS 007 008

009 **Anonymous authors**
 010
 011
 012
 013
 014
 015

016 Paper under double-blind review
 017
 018
 019
 020
 021
 022
 023
 024
 025
 026
 027
 028
 029
 030
 031

ABSTRACT

032 Translating human-written mathematical theorems and proofs from natural language (NL) into formal languages (FLs) like Lean 4 has long been a significant challenge for AI. Most state-of-the-art methods either focus on theorem-only NL-to-FL auto-formalization or on FL proof synthesis from FL theorems. In practice, auto-formalization of both theorem and proof still requires human intervention, as seen in AlphaProof’s silver-medal performance at the 2024 IMO, where problem statements were manually translated before automated proof synthesis.

033 We present PROOFBRIDGE, a unified framework for automatically translating entire NL theorems and proofs into Lean 4. At its core is a joint embedding model that aligns NL and FL (NL-FL) theorem-proof pairs in a shared semantic space, enabling cross-modal retrieval of semantically relevant FL examples to guide translation. Our training ensures that NL-FL theorems (and their proofs) are mapped close together in this space if and only if the NL-FL pairs are semantically equivalent. PROOFBRIDGE integrates retrieval-augmented fine-tuning with iterative proof repair, leveraging Lean’s type checker and semantic equivalence feedback to ensure both syntactic correctness and semantic fidelity. Experiments show substantial improvements in proof auto-formalization over strong baselines (including GPT-5, Gemini-2.5, Kimina-Prover, DeepSeek-Prover), with our retrieval-augmented approach yielding significant gains in semantic correctness (SC, via proving bi-directional equivalence) and type correctness (TC, via type-checking theorem+proof) across pass@k metrics on MINIF2F-TEST-PF, a dataset we curated. In particular, PROOFBRIDGE improves cross-modal retrieval quality by up to $3.28\times$ Recall@1 over all-MiniLM-L6-v2, and achieves +31.14% SC and +1.64% TC (pass@32) compared to the baseline Kimina-Prover-RL-1.7B.

1 INTRODUCTION

034 In mathematics, ensuring the correctness of proofs is a crucial yet inherently difficult task. Traditionally, mathematicians rely on the peer-review process for *proof verification*, yet as proofs grow 035 increasingly complex, even careful human scrutiny can overlook subtle errors. For instance, in 1989, 036 Kapranov and Voevodsky published a proof connecting ∞ -groupoids and homotopy types, which 037 was later disproven by Carlos Simpson in 1998; more recently, while formalizing his 2023 paper 038 (Tao, 2023) on the Maclaurin-type inequality, Terence Tao discovered a non-trivial bug. To 039 mitigate challenges of verifying complex proofs, proof assistants and formal mathematical languages 040 like Coq (Barras et al., 1999), Isabelle (Nipkow et al., 2002), HOL Light (Harrison, 2009), Metamath 041 (Megill & Wheeler, 2019), Lean 4 (Moura & Ullrich, 2021), and Peano (Poesia & Goodman, 042 2023) have been developed, offering a way to create *computer-verifiable formal proofs*. Such formal 043 language (FL) proofs, defined by strict syntax and symbolic logic, enable reliable automated 044 verification guarantees that resolve the inherent ambiguity of natural language (NL) proofs. However, 045 constructing FL proofs is time-intensive and demands both deep mathematical expertise and detailed 046 knowledge of the language and its libraries, making the process challenging even for experienced 047 mathematicians and limiting the wider adoption of such theorem provers and FL proofs.

048 To simplify the task of writing proofs in FL, two key research directions have emerged: *auto-*
 049 *formalization* and *automated formal proof synthesis (AFPS)*. *Auto-formalization* targets NL-to-FL
 050 translation, but most prior works (Wang et al., 2025; Wu et al., 2025; Jiang et al., 2024; Gao et al.,
 051 2024b) focus only on formalizing theorems (statements), not proofs. In contrast, *automated for-*
 052 *mal proof synthesis* (Ren et al., 2025; Wang et al., 2025) aims to generate FL proofs given an FL
 053 theorem. Proof auto-formalization is relatively less explored, with Draft-Sketch-Prove (Jiang et al.,
 2022) for Isabelle and FormL4 (Lu et al., 2024) for Lean serving as notable examples. In practice,

054

055

056

057

058

059

060

061

062

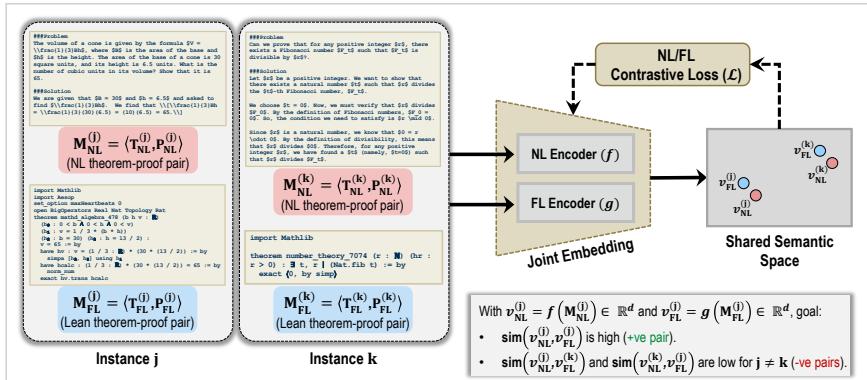
063

064

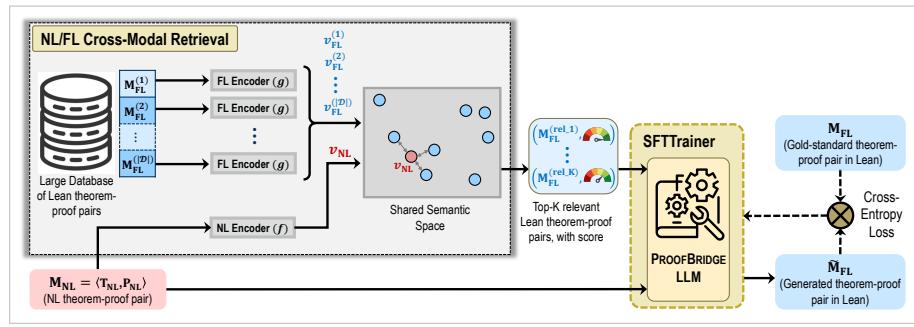
065

066

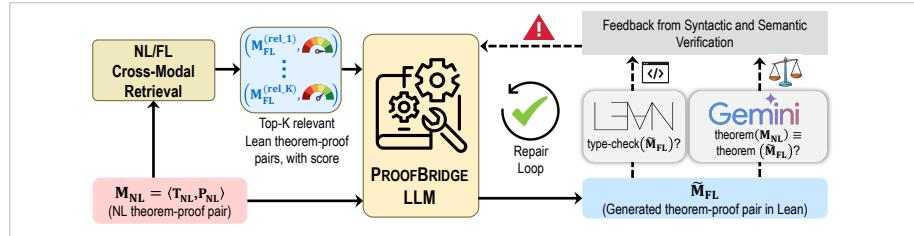
067



(a) Joint embedding of NL and FL (Lean) theorems and proofs into shared semantic space



(b) Retrieval-augmented Supervised Fine-Tuning (SFT) of PROOFBRIDGE with NL/FL cross-modal retrieval



(c) Inference phase of retrieval-augmented proof auto-formalization with iterative repair

Figure 1: **Pipeline of PROOFBRIDGE for proof auto-formalization.** We first train a joint embedding model for NL and FL via contrastive learning, enabling cross-modal retrieval of semantically related FL theorem-proof pairs for a given NL input. An LLM is then fine-tuned on NL-to-Lean translations, conditioned on retrieved proofs and relevance scores. At inference, the system retrieves related Lean proofs and applies an iterative repair loop to the generated FL theorem-proof pair.

formalizing an entire NL proof requires first performing *theorem-only auto-formalization* to translate the NL theorem into FL, followed by AFPS to generate the FL proof from the FL theorem. AlphaProof (Deepmind, 2024), which achieved silver-medal standard in the 2024 International Mathematical Olympiad, followed this two-step process: problems were first manually translated into formal mathematical language, then formal proofs were synthesized. Thus, in practice, pipelines still require manual formalization of the theorem before proof synthesis, even though SoTA theorem-only auto-formalization and AFPS tools exist. This illustrates the broader challenge that current systems often rely on human intervention to ensure semantic correctness of proof auto-formalization.

Contemporary LLMs face several challenges that limit their effectiveness for proof auto-formalization in Lean 4. **First**, large-scale datasets pairing NL theorems with Lean 4 proofs are scarce. Most existing resources (Goedel-Pset-v1 (Lin et al., 2025), Herald statements (Gao et al., 2024b), Lean Workbook (Ying et al., 2024), MMA (Jiang et al., 2024)) cover only theorems, while those with proofs (Herald proofs, Lean Workbook proofs (Lin et al., 2025), and FormL4 (Lu et al., 2024)) are much smaller and do not align with the popular miniF2F (Zheng et al., 2021) benchmark in the same Lean 4 version. Lean versions are not backward compatible, so cross-version evalua-

108 tion often fails. **Second**, most fine-tuned LLMs for Lean 4 target either theorem auto-formalization
 109 or proof synthesis. Proof auto-formalization is harder, as it requires both translating the NL theo-
 110 rem and constructing the corresponding FL proof. **Third**, Lean 4 has an effectively *infinite action
 111 space* (Poesia & Goodman, 2023), with proofs using complex *tactics* that reuse prior theorems or
 112 introduce new variables. Prior work generates FL directly from NL, ignoring semantic relations like
 113 shared tactics and DAG dependencies, causing LLMs to often violate Lean 4’s strict syntactic and
 114 semantic constraints and produce hallucinated or invalid proofs (Jha et al., 2025; Jana, 2024; Ugare
 115 et al., 2024). **Fourth**, automated evaluation is a major bottleneck. Lean’s type-checker verifies the
 116 FL proof but cannot ensure semantic equivalence. Existing methods often type-check only the the-
 117 orem (leaving proofs incomplete using placeholders like `sorry`) or rely on proxies such as BLEU,
 118 which are unreliable (Jiang et al., 2024; Lu et al., 2024; Wu et al., 2022; Ying et al., 2024).

119 **Key Insight.** In this paper, we address the task of *proof auto-formalization*, focusing on Lean 4 as
 120 the FL, via a combination of a joint embedding model, an LLM, and Lean for verification. It takes
 121 as input an NL theorem-proof pair and produces the corresponding FL theorem-proof pair in Lean 4.
 122 The key insight behind our approach is to treat proof auto-formalization as *learning from demon-
 123 strations*: the LLM is guided not only by the NL proof but also by FL proofs retrieved using an
 124 NL/FL joint embedding model that leverages contrastive learning and encodes linear DAG trav-
 125 ersals of Lean proofs. Rather than relying on randomly chosen few-shot examples, these retrieved
 126 proofs capture far richer *reusable formalization patterns* (tactic choices, DAG structures), providing
 127 grounded signals that guide generation toward Lean-verifiable proofs, as illustrated in Figure 1.

128 Contributions:

129 **The PROOFBRIDGE Auto-formalization Method and Tool:** We present PROOFBRIDGE, an
 130 LLM-based, retrieval-augmented proof auto-formalization framework. At its core is an *NL/FL joint
 131 embedding model* that maps semantically equivalent NL and FL theorem-proof pairs to nearby points
 132 in a shared space, enabling effective cross-modal retrieval of related FL proofs. We then fine-tune the
 133 SoTA LLM Kimina-Prover-RL-1.7B (Wang et al., 2025) to perform NL-to-Lean 4 proof translation,
 134 conditioned on the retrieved FL proofs and their relevance scores. During inference, generation is
 135 refined with an iterative verifier-guided repair loop combining Lean type-checking with LLM-based
 136 bi-directional equivalence proving to ensure syntactic correctness and semantic fidelity. (Section 4)

137 **NUMINAMATH-LEAN-PF Dataset:** We curate NUMINAMATH-LEAN-PF, a large-scale dataset
 138 of 38.9k NL \leftrightarrow Lean 4 theorem-proof pairs, specialized for *proof auto-formalization*. Each Lean
 139 theorem-proof pair is type-checked and paired with an NL counterpart. Additionally, we release
 140 MINIF2F-TEST-PF, a Lean v4.15.0 version of miniF2F-Test with 244 instances tailored for proof
 141 auto-formalization, enabling a consistent pipeline in the same Lean version. (Section 5.1)

142 **Extensive Experimental Evaluation:** Compared to the baseline encoder all-MiniLM-L6-v2,
 143 PROOFBRIDGE’s cross-modal NL \rightarrow FL retrieval achieves $3.28\times$ higher Recall Rate@ K at $K=1$
 144 and $2.74\times$ MRR, with top- K retrieved embeddings 23% closer and non-retrieved 104% farther. We
 145 evaluate PROOFBRIDGE against 13 SoTA LLMs, including foundation models (Gemini-2.5, GPT-
 146 5-mini) and automated proof synthesis LLMs (DeepSeek-Prover, STP, Leanabell-Prover, Kimina-
 147 Prover), using verifier-grounded metrics: *type correctness* (TC) and *semantic correctness* (SC, a new
 148 metric based on Lean bidirectional equivalence proofs). Built on Kimina-Prover-RL-1.7B, PROOF-
 149 BRIDGE achieves +31.14% SC and +1.64% TC (pass@32) on MINIF2F-TEST-PF. (Section 5)

150 2 RELATED WORK

151 Our work lies at the intersection of three key AI-for-Math research areas: automated formal proof
 152 synthesis, auto-formalization, and retrieval-augmented learning for mathematical reasoning. We
 153 focus on the most relevant approaches and highlight differences from our unified framework.

154 **Auto-Formalization.** Auto-formalization translates NL mathematics into FL, but most existing
 155 work focuses on theorem formalization rather than proofs. **Theorem-only approaches** include
 156 Herald-translator (Gao et al., 2024b), which extracts FL theorems from Mathlib4 and trains on
 157 informal counterparts, and Kimina-Autoformalizer (Wang et al., 2025), which fine-tunes models
 158 with expert iteration. These excel at translating theorems but cannot handle proofs. **Proof auto-
 159 formalization** has received limited attention. Draft-Sketch-Prove (Jiang et al., 2022) converts NL
 160 proofs into formal sketches in Isabelle with open conjectures, then fills gaps using predefined tactics
 161 and tools like Sledgehammer (Paulsson & Blanchette, 2012). FormL4 (Lu et al., 2024) trains on
 GPT-4 informalized Mathlib proofs with process-supervised step-level Lean compilation feedback.

162 *Key Differences:* Our approach is the first to jointly learn representations for NL and FL theorem-
 163 proof pairs, enabling cross-modal retrieval to guide formalization. Unlike prior work on isolated
 164 proof generation, we leverage semantic relationships of NL and FL proofs for contextual guidance.
 165

166 **Automated Formal Proof Synthesis.** Automated formal proof synthesis (AFPS) takes FL theo-
 167 rems as input and generates FL proofs. Current approaches fall into two categories: *next-tactic*
 168 *prediction* and *whole-proof generation*. **Next-Tactic Prediction (NTP)** methods train models to
 169 predict single proof steps from current proof states. Representative systems include GPT-f (Polu &
 170 Sutskever, 2020) for Metamath, LIsa (Jiang et al., 2021) for Isabelle, and PACT (Han et al., 2022) for
 171 Lean. These use tree search over proof states, prioritizing tactics by cumulative probability. While
 172 NTP ensures stepwise correctness via interactive theorem-prover verification, it suffers from long-
 173 horizon dependencies and computational overhead from such repeated interactions. **Whole-Proof**
 174 **Generation (WPG)** methods generate complete FL proofs in single passes, offering computational
 175 efficiency but risking cascading errors. Recent advances include DeepSeek-Prover-v1 (Xin et al.,
 176 2024a), which combines SFT with expert iteration, and TheoremLlama (Wang et al., 2024b), which
 177 improves in-context learning through curriculum-based training. DeepSeek-Prover-v2 (Ren et al.,
 178 2025) integrates NL reasoning with formal proof generation, while Kimina-Prover (Wang et al.,
 179 2025) applies reinforcement learning with compilation-based rewards (Jana et al., 2024).
 180

181 *Key Differences:* Unlike AFPS approaches that assume FL theorems as input, our work addresses
 182 the more challenging task of generating theorem-proof pairs in FL from an NL input.
 183

184 **Retrieval-Augmented Learning for Mathematics.** Recent work has explored retrieval-
 185 augmented approaches for mathematical reasoning, though not specifically for auto-formalization.
 186 TLAPS (Zhou, 2025) retrieves verified proofs to assist proof generation, COPRA (Thakur et al.,
 187 2023) selects relevant lemmas to guide proof search, and REAL-Prover (Shen et al., 2025) re-
 188 trievals Mathlib theorems for next-tactic prediction. These methods rely on plain-text encoding.
 189 LeanSearch (Gao et al., 2024a), HERALD (Gao et al., 2024b), and RAuformalizer (Liu et al.,
 190 2025) also use plain text encoders for FL theorem retrieval; however, as shown in Section 5.5, this
 191 does not extend to the more demanding task of FL theorem+proof pair retrieval.
 192

193 *Key Differences:* Our joint embedding enables NL/FL cross-modal retrieval of theorem+proof pairs
 194 and encodes the DAG structure of Lean proofs, unlike plain-text encoders. Capturing proof-structure
 195 semantics is essential for proof auto-formalization and is not addressed by existing tool-chains.
 196

197 **Positioning our contributions.** PROOFBRIDGE makes several novel contributions relative to existing
 198 approaches: (a) *Unified Proof Auto-Formalization*: We address complete translation (theorem
 199 + proof) rather than treating theorem formalization and proof synthesis separately. (b) *Joint Se-
 200 mantic Embedding*: Our contrastive learning framework for aligning NL and FL proofs is novel,
 201 enabling effective cross-modal retrieval. (c) *Retrieval-Augmented Translation*: We are the first
 202 to apply retrieval-augmented fine-tuning and generation to auto-formalization, leveraging seman-
 203 tic relationships between FL proofs to guide translation. (d) *Rigorous Evaluation*: We introduce
 204 systematic metrics for proof auto-formalization, including type correctness via bi-directional equiv-
 205 alence rather than proxy measures. This combination of joint embedding, retrieval augmentation,
 206 and unified translation distinguishes our approach from prior work.
 207

208 3 PRELIMINARIES: TACTIC-STYLE PROOFS IN LEAN

209 Lean (Moura & Ullrich, 2021) is a functional programming language and interactive theorem prover
 210 that is based on the propositions-as-types principle, where proving a proposition is equivalent to
 211 constructing a term of the corresponding type. Rather than building these terms manually, users
 212 write proofs in a tactic language, which provides high-level steps to guide term construction. Lean 4
 213 (henceforth Lean) represents tactic-style proofs as directed acyclic graphs (DAGs) of *proof states*
 214 and *tactics*, automatically generating the corresponding proof term in the background. The kernel
 215 then verifies the term, ensuring correctness by enforcing the axiomatic foundations of Lean’s logic,
 216 the Calculus of Inductive Constructions. This combination of a formal system and a small, trusted
 217 kernel provides strong confidence in the validity of proofs. In the DAG (Figure 2) of a Lean proof:
 218

- 219 • Each **proof state** $S_i \equiv [G_1, \dots, G_n]$ consists of a sequence of zero or more *open goals*. Initial
 220 state S_0 has one goal, the theorem $T_{FL} \equiv pr \vdash cn$ itself. Leaf-level states have no open goal.
 221 • Each **open goal** $G_i \equiv pr_i \vdash cn_i$ of a proof state represents a proposition cn_i that needs to be
 222 proven, given a set of premises pr_i .
 223

216 • Each **tactic** tac_i represents a proof step. It is a high-level command (rooted in metaprogramming)
 217 applied to an open goal G_i , producing a new proof state. If the resulting proof state has no open
 218 goal, it directly resolves the current goal. A parent goal is resolved once all subgoals are resolved.
 219

220 Tactic-style proof synthesis in Lean follows a *sequential decision process*. Lean provides an interactive REPL (Leanprover, 2025) that applies tactics step by step to manipulate
 221 proof states. An FL proof is a sequence of tactics, and at each step, the REPL updates the proof state if the tactic is
 222 valid or returns an error identifying the faulty one. Each tactic advances the proof by breaking the current goal into
 223 simpler subgoals, similar to the ‘*suffices to show*’ construct.
 224

225 **Proof Auto-formalization as a Learning Problem.** Given
 226 an NL theorem-proof pair $M_{NL} = \langle T_{NL}, P_{NL} \rangle$, the goal is
 227 to learn a function $f: M_{NL} \mapsto M_{FL}$ that produces a corresponding Lean theorem-proof pair $M_{FL} = \langle T_{FL}, P_{FL} \rangle$.
 228 Here, $T_{FL} \equiv pr \vdash cn$ denotes the formal theorem, and
 229 $P_{FL} \equiv (tac_0, \dots, tac_{n-1})$ is the proof as a sequence of
 230 tactics. The generated pair must satisfy:

231 (a) *Type correctness*: $M_{FL} = \langle T_{FL}, P_{FL} \rangle$ passes Lean type-checking, ensuring that P_{FL} proves T_{FL}
 232 with no open goals in the DAG.
 233 (b) *Semantic correctness*: FL theorem is semantically equivalent to the NL one, i.e., $T_{FL} \equiv T_{NL}$.

238 4 OUR APPROACH AND TOOL ARCHITECTURE

239 4.1 JOINT EMBEDDING OF NL AND LEAN PROOFS FOR CROSS-MODAL RETRIEVAL

240 A core component of our framework is the *joint embedding model*, which learns to represent
 241 NL theorem-proof pairs and their FL (Lean) counterparts in a shared semantic space. The goal
 242 is to align these modalities so that cross-modal retrieval between NL and FL becomes possible.
 243 Formally, given an NL theorem-proof pair M_{NL} and a large database of FL theorem-proof pairs
 244 $\mathcal{D} = \{M_{FL}^{(i)} = \langle T_{FL}^{(i)}, P_{FL}^{(i)} \rangle\}$, the model retrieves subset $\mathcal{R}(M_{NL}, \mathcal{D}) \subseteq \mathcal{D}$ of size $K \ll |\mathcal{D}|$, that
 245 serve as in-context demonstrations to guide downstream auto-formalization.

246 During training the joint embedding model, we start with NL-FL pairs $(M_{NL}^{(i)}, M_{FL}^{(i)})$, which are
 247 encoded into vectors using two modality-specific encoders. Each encoder is initialized with a pre-
 248 trained model, and a subset of parameters is subsequently fine-tuned. Given $M_{NL}^{(i)}$, the NL encoder f
 249 produces an embedding $v_{NL}^{(i)} = f(M_{NL}^{(i)}, \theta_f \parallel \phi_f) \in \mathbb{R}^d$, and given $M_{FL}^{(i)}$, the FL encoder g produces
 250 $v_{FL}^{(i)} = g(M_{FL}^{(i)}, \theta_g \parallel \phi_g) \in \mathbb{R}^d$, where θ denotes frozen parameters, ϕ denotes trainable parameters,
 251 and d is the dimension of the shared semantic space. The details of each encoder are as follows:

252 • **NL encoder $f(M_{NL}^{(i)}, \theta_f \parallel \phi_f)$:** To encode $M_{NL}^{(i)}$, we use all-MiniLM-L6-v2 (Reimers &
 253 Gurevych, 2020), a lightweight model (22.7M parameters) that effectively captures semantic sim-
 254 ilarity in NL. It encodes $M_{NL}^{(i)}$ into 384-dimensional embeddings, thereby projected into the joint
 255 embedding space of dimension $d = 512$ via a linear layer included in the trainable set ϕ_f .
 256 • **FL encoder $g(M_{FL}^{(i)}, \theta_g \parallel \phi_g)$:** Given $M_{FL}^{(i)} = \langle T_{FL}^{(i)}, P_{FL}^{(i)} \rangle$, we first extract the linearized DAG
 257 traversal of tactics from $P_{FL}^{(i)}$ using Lean REPL (Leanprover, 2025). This traversal is represented
 258 as an ordered sequence of proof-state transformations induced by successive tactic applications:
 259 $S_0 \xrightarrow{tac_0} S_1 \xrightarrow{tac_1} \dots \xrightarrow{tac_{H-1}} S_H$, where $S_0 \equiv T_{FL}^{(i)}$, each $S_h \equiv [G_1, \dots, G_l]$ denotes a
 260 proof state consisting of zero or more open goals, and tac_{h-1} is the tactic applied at step h .
 261 This sequence captures the entire proof as an ordered series of state transformations. To create
 262 embeddings for the full proof, we first encode each state S_h using LeanDojo’s ByT5 proof-state
 263 encoder (Yang et al., 2023) (218M parameters), producing embeddings of size 1,472 per state.
 264 We then obtain a single embedding for the entire proof via mean-pooling, which is subsequently
 265 projected into a shared semantic space of dimension $d = 512$ using a linear layer included in the
 266 267 268 269

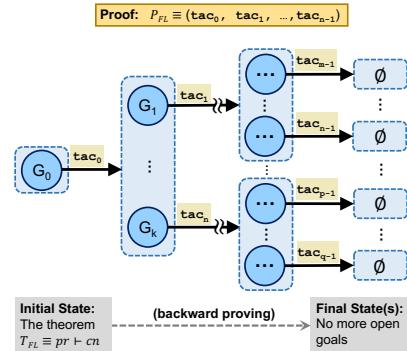


Figure 2: **Tactic-style Proof.** A Lean proof represented as a DAG of tactics.

270 trainable parameters ϕ_g . The intuition behind this approach is to ensure that semantically similar
 271 proofs (those with similar DAG structures of proof-states and tactics) produce similar embeddings.
 272

273 **Contrastive Learning.** To enable cross-modal retrieval between NL and FL representations, it is
 274 essential to align the two modalities in the embedding space. Specifically, for each positive pair
 275 $(M_{\text{NL}}^{(i)}, M_{\text{FL}}^{(i)})$, we aim for their embeddings $(v_{\text{NL}}^{(i)}, v_{\text{FL}}^{(i)})$ to exhibit high cosine similarity, while em-
 276 beddings of mismatched pairs are pushed apart. Denoting ℓ_2 -normalization by $\hat{v} = v/\|v\|_2$ and
 277 defining the cosine similarity between two embeddings u and w as $[\hat{u}, \hat{w}]$, we adopt the following
 278 symmetric contrastive loss for a mini-batch $\mathcal{B} = \{(M_{\text{NL}}^{(i)}, M_{\text{FL}}^{(i)})\}_{i=1}^n$ of NL-FL pairs:
 279

$$\mathcal{L}(\mathcal{B}) = -\frac{1}{2n} \sum_{i=1}^n \left[\log \left(\frac{\exp([\hat{v}_{\text{NL}}^{(i)}, \hat{v}_{\text{FL}}^{(i)}]/\tau)}{\sum_{j=1}^n \exp([\hat{v}_{\text{NL}}^{(i)}, \hat{v}_{\text{FL}}^{(j)}]/\tau)} \right) + \log \left(\frac{\exp([\hat{v}_{\text{FL}}^{(i)}, \hat{v}_{\text{NL}}^{(i)}]/\tau)}{\sum_{j=1}^n \exp([\hat{v}_{\text{FL}}^{(i)}, \hat{v}_{\text{NL}}^{(j)}]/\tau)} \right) \right] \quad (1)$$

282 where $\tau > 0$ is a temperature hyperparameter. This loss encourages each NL embedding to be
 283 closest to its corresponding FL embedding, and vice versa, using other in-batch embeddings as
 284 negatives. The negatives are sampled randomly for each mini-batch.
 285

286 **NL/FL Cross-Modal Retrieval.** We precompute the normalized embeddings $\{\hat{v}_{\text{NL}}^{(i)}\}_{i=1}^{|\mathcal{D}|}$ and
 287 $\{\hat{v}_{\text{FL}}^{(j)}\}_{j=1}^{|\mathcal{D}|}$ for all NL and FL theorem-proof pairs respectively in our database \mathcal{D} , which enables
 288 efficient cross-modal retrieval. Given a query theorem-proof pair in either source modality (NL or
 289 FL), we encode it into the shared semantic space (yielding \hat{q}_{NL} or \hat{q}_{FL}) and compute cosine similar-
 290 ities with all items in the target modality, producing the set $\{[\hat{q}_{\text{NL}}, \hat{v}_{\text{FL}}^{(j)}]\}_{j=1}^{|\mathcal{D}|}$ or $\{[\hat{q}_{\text{FL}}, \hat{v}_{\text{NL}}^{(i)}]\}_{i=1}^{|\mathcal{D}|}$,
 291 depending on the retrieval direction. The top- K nearest neighbors from these sets are then selected
 292 as demonstrations, reflecting similar proof structures, patterns, and mathematical domains.
 293

294 4.2 RETRIEVAL-AUGMENTED FINE-TUNING FOR PROOF AUTO-FORMALIZATION

295 We fine-tune an LLM to translate NL theorem-proof pairs into FL (Lean), conditioned on retrieved
 296 FL demonstrations that provide rich contextual knowledge. For each training instance, we construct
 297 a prompt containing (a) input NL theorem-proof pair M_{NL} and (b) top- K retrieved FL theorem-
 298 proof pairs: $\mathcal{R}(M_{\text{NL}}, \mathcal{D}) = \{M_{\text{FL}}^{(k)}\}_{k=1}^K$ with relevance scores $\{r^{(k)}\}_{k=1}^K$. The retrieved examples
 299 demonstrate how similar mathematical concepts and proof strategies are formalized in Lean. We
 300 include relevance scores to help the model weight the importance of each retrieved example.
 301

302 **Training Objective.** We fine-tune Kimina-Prover-RL-1.7B (Wang et al., 2025) using supervised
 303 learning on our NUMINAMATH-LEAN-PF dataset (details in Section 5.1). The model is trained to
 304 generate an FL theorem-proof pair $\widetilde{M}_{\text{FL}}$ given the input context. This retrieval-augmented approach
 305 allows the LLM to learn from similar formalization patterns rather than generating formal theorems
 306 in isolation. As illustrated in Figure 1b, we use the standard auto-regressive language modeling loss:
 307

$$\mathcal{L}_{\text{CE}} = -\frac{1}{|\mathcal{T}|} \sum_{t=1}^{|\mathcal{T}|} \log P_{\theta}(\tau_t \mid \tau_{<t}, \mathcal{C}) \quad (2)$$

308 where $\mathcal{T} = \widetilde{M}_{\text{FL}}$ is the generated formalization tokenized as sequence $(\tau_1, \dots, \tau_{|\mathcal{T}|})$, \mathcal{C} represents
 309 the input context (NL theorem-proof + retrieved FL examples), and θ are the LLM parameters. This
 310 corresponds to the cross-entropy loss between $\widetilde{M}_{\text{FL}}$ and the gold-standard formalization M_{FL} .
 311

312 4.3 ITERATIVE PROOF REPAIR WITH VERIFIER FEEDBACK

313 During inference, we perform retrieval-augmented proof auto-formalization with the fine-tuned
 314 LLM (Figure 1c). However, LLM being a stochastic model may still generate FL theorem-proof
 315 pair that contain syntactic errors or semantic misalignments with the input NL theorem-proof. To
 316 address, we implement an iterative repair mechanism that combines Lean’s type checker with se-
 317 mantic equivalence verification. For an input NL theorem-proof pair $M_{\text{NL}} = \langle T_{\text{NL}}, P_{\text{NL}} \rangle$ the LLM
 318 generates an FL counterpart $\widetilde{M}_{\text{FL}} = \langle \widetilde{T}_{\text{FL}}, \widetilde{P}_{\text{FL}} \rangle$, on which we perform two types of verification:
 319

- 320 1. **Syntactic Verification:** We compile $\widetilde{M}_{\text{FL}}$ using Lean’s type checker. If compilation fails, we
 321 extract the specific error message and location from Lean’s diagnostic output.
 322 2. **Semantic Verification:** We assess whether the generated theorem $\widetilde{T}_{\text{FL}}$ accurately represents the
 323 original NL theorem T_{NL} using an LLM-based equivalence judge.

Repair Process. When either syntactic or semantic verification fails, we initiate an iterative repair process. The procedure terminates once both checks succeed or the maximum number of repair attempts ($R_{\max} = 5$) is reached. This bounded, iterative strategy improves the reliability of proof auto-formalization by catching and correcting common errors while maintaining computational efficiency. The overall process is described in Algorithm 1.

5 EXPERIMENTAL EVALUATION

5.1 DATASETS AND PREPARATION: NUMINAMATH-LEAN-PF AND MINIF2F-TEST-PF

For training, we construct NUMINAMATH-LEAN-PF from NuminaMath-LEAN (Wang et al., 2025), containing 104,155 competition-level problems in algebra, geometry, number theory, combinatorics, and calculus. Each instance pairs an NL theorem T_{NL} with a human-written Lean v4.15.0 theorem T_{FL} ; 38,951 include FL proofs P_{FL} (30% human-written, rest by KiminaProver), forming $\{T_{\text{NL}}, \langle T_{\text{FL}}, P_{\text{FL}} \rangle\}$. Next, we prepare NUMINAMATH-LEAN-PF via the following steps:

Formal Verification and Repair. Each FL pair is type-checked in Lean (Leanprover, 2025). About 6% (2,337) failed due to syntax, library mismatches, or incomplete proofs. These were automatically repaired via Gemini-2.5-Pro: error messages and locations are extracted from Lean, used to prompt the LLM for corrections, and re-verified iteratively up to five times. All errors were successfully fixed, showing most issues were syntactic rather than mathematical.

NL Proof Generation. NuminaMath-LEAN provides only theorems, so we generate NL proofs in two stages. First, *solution sketch retrieval* matches T_{NL} to NuminaMath 1.5 (Li et al., 2024) (896 problem-solution pairs) via exact string matching, retrieving sketches (median 79 words) for 25,792 instances (66%). Second, *FL-to-NL informalization* uses Gemini-2.5-Pro to translate verified P_{FL} into detailed, human-readable NL proofs. Each instance uses $P_{\text{FL}}, \langle T_{\text{FL}}, T_{\text{NL}} \rangle$, and sketches when available, producing 38,951 NL-FL theorem-proof pairs $\{\langle T_{\text{NL}}, P_{\text{NL}} \rangle, \langle T_{\text{FL}}, P_{\text{FL}} \rangle\}$.

For validation, we curate MINIF2F-TEST-PF by combining two versions of miniF2F-test (Zheng et al., 2021), a widely-used auto-formalization benchmark. It contains 244 Olympiad-level problems from the AIME, AMC, IMO, and undergraduate courses in algebra, number theory, and inequalities. We use the Lean v4.15.0 version (NuminaMath, 2025) and add missing NL proofs from Yang (2025).

5.2 EVALUATION METRICS

Metrics for NL/FL Cross-Modal Retrieval. We evaluate cross-modal alignment of our joint embedding model in two directions. $\text{NL} \rightarrow \text{FL}$ measures retrieval of FL theorem-proof pairs given an NL input, which is relevant for proof auto-formalization, while $\text{FL} \rightarrow \text{NL}$ assesses the reverse. For a test pair $(M_{\text{NL}}, M_{\text{FL}})$, a retrieval in the $\text{NL} \rightarrow \text{FL}$ direction is deemed successful if the model retrieves the FL counterpart M_{FL} given M_{NL} , and unsuccessful otherwise; the $\text{FL} \rightarrow \text{NL}$ direction is evaluated analogously. We assess retrieval performance using five metrics. **(i) Recall Rate @ K** measures the percentage of queries for which the query’s cross-modal counterpart appears among the top- K retrieved results. We report $K = 1, 5, 10, 20, 50$. **(ii) Mean Reciprocal Rank (MRR)** is the average reciprocal rank of the retrieved cross-modal counterpart for each query, $\text{MRR} = \frac{1}{N} \sum_{q=1}^N \frac{1}{\text{rank}_q}$, indicating how highly it is ranked. **(iii) Cosine Similarity of Top- K Retrieved** measures the cosine similarity between the query embedding and those of the top- K retrieved instances. For each query, we sort these scores in ascending order and record three statistics: median (**M**), 25th percentile (**Q1**), and 75th percentile (**Q3**), and report their average over the test set. **(iv) Cosine Similarity of Non-Retrieved** applies the same procedure to all non-retrieved instances and reports the median (**M**), 25th percentile (**Q1**), and 75th percentile (**Q3**) averaged over the test set. **(v) mean Median Gap (mMG)** measures the difference between the median (**M**) cosine similarity of top- K retrieved and that of non-retrieved instances, averaged over $K = 1, 5, 10, 20, 50$.

Metrics for Proof Auto-Formalization. Given $M_{\text{NL}} = \langle T_{\text{NL}}, P_{\text{NL}} \rangle$, an auto-formalizer LLM/tool generates an FL version $M_{\text{FL}} = \langle T_{\text{FL}}, P_{\text{FL}} \rangle$. We evaluate performance using two metrics: **(i) Type Correctness (TC)** measures whether M_{FL} is accepted by Lean’s type-checker, i.e., P_{FL} proves T_{FL} without using `sorry`. **(ii) Semantic Correctness (SC)** is evaluated only for type-correct generations

378 and measures whether \tilde{T}_{FL} is definitionally equal ¹ to the gold-standard T_{FL} by prompting Gemini-
 379 2.5-Pro up to five times to produce a Lean proof of the bi-directional equivalence $\tilde{T}_{\text{FL}} \leftrightarrow T_{\text{FL}}$ using
 380 a restricted set of tactics (`rfl`, `simp`, `ring`, etc.; details in Appendix A.4). Although relying on an
 381 LLM judge, it is based on the Lean proof of equivalence rather than the LLM’s judgment alone. TC
 382 and SC are computed as *pass@k*, i.e., over the top- k generated candidates, for $k = 1, 2, 4, 8, 16, 32$.
 383

384 5.3 STATE-OF-THE-ART BASELINES

385 **SoTA for NL/FL Cross-Modal Retrieval.** To our knowledge, no existing model jointly embeds
 386 theorems and proofs in NL and FL. Pre-trained encoders alone do not yield embeddings suitable
 387 for meaningful cross-modal retrieval. To illustrate, we evaluate two *SoTA Encoders*: Qwen3-
 388 Embedding-8B (Zhang et al., 2025b) and E5-Mistral-7B-Instruct (Wang et al., 2024a). Qwen3
 389 allows user-defined output dimensions up to 4096, and we use 512 to match our joint embedding
 390 model, while E5-Mistral (used by LeanSearch-PS (Shen et al., 2025)) has a fixed dimension of 4096.
 391 We also include a *Baseline Encoder*, all-MiniLM-L6-v2 (Reimers & Gurevych, 2020), used for the
 392 NL encoder in PROOFBRIDGE, producing 384-dim embeddings. All these encoders treat theorem-
 393 proof pairs as plain text, ignoring the DAG structure of FL proofs, which PROOFBRIDGE explicitly
 394 leverages. Retrieval is performed via cosine similarity in the respective embedding spaces.
 395

396 **SoTA Tools for Proof Auto-Formalization.** Existing proof auto-formalization tools include
 397 DSP (Jiang et al., 2022) which supports only Isabelle; since we target Lean 4, direct comparison
 398 is not feasible. Another recent tool, FormL4 (Lu et al., 2024), has not released its trained model.
 399 We therefore compare against three categories: foundation models, SoTA automated formal proof
 400 synthesis (AFPS) LLMs, and SoTA theorem auto-formalization LLMs. The four *foundation models*
 401 we include are GPT-5-mini (OpenAI, 2025) and the Gemini-2.5 (Comanici et al., 2025) variants
 402 Flash-Lite, Flash, and Pro. We evaluate seven *AFPS LLMs*, all trained to generate full FL proofs
 403 with tactics, the same output space targeted in proof auto-formalization: DeepSeek-Prover-V1.5-
 404 RL (Xin et al., 2024b), STP_model_Learn_0320 (Dong & Ma, 2025), Goedel-Prover-SFT (Lin et al.,
 405 2025), Leanabell-Prover-V2-KM (Zhang et al., 2025a), and three Kimina-Prover variants (Wang
 406 et al., 2025) (RL-1.7B, Distill-8B, 72B). Lastly, we evaluate two *auto-formalization LLMs*: Kimina-
 407 Autoformalizer-7B (Wang et al., 2025) and Herald-Translator (Gao et al., 2024b).
 408

409 5.4 TRAINING AND EVALUATION SETUP

410 We train our joint embedding model on 90% (35,056 instances) of NUMINAMATH-LEAN-PF and
 411 evaluate it on the remaining 3,895 instances. The split is domain-stratified across all mathematical
 412 areas, ensuring hard negatives in the test set. The train split serves as a database \mathcal{D} of FL theorem-
 413 proof pairs. PROOFBRIDGE, built on Kimina-Prover-RL-1.7B, is SFT-tuned for NL-to-FL transla-
 414 tion using $(M_{\text{NL}}, M_{\text{FL}})$ from NUMINAMATH-LEAN-PF, with the joint embedding model retrieving
 415 the top-5 relevant FL proofs from \mathcal{D} for retrieval-augmented SFT and inference. Inference-time
 416 iterative proof repair is applied, and the model is evaluated on MINIF2F-TEST-PF. See Appen-
 417 dices A.2–A.3 for implementation and training details, and Appendix A.5 for an example inference.
 418

419 The *SoTA Tools* are evaluated in three settings: (a) *zero-shot*, with no in-context I/O examples; (b)
 420 *random few-shot*, with five randomly selected in-context examples; and (c) *text-based retrieval few-
 421 shot*, where the top-5 FL theorem-proof pairs are retrieved from \mathcal{D} via Qwen3-Embedding-8B and
 422 paired with their NL counterparts as in-context examples. Further, we evaluate a *SoTA Two-Step*
 423 setting: a theorem-only auto-formalizer (T1) first translates T_{NL} to \tilde{T}_{FL} , and an AFPS LLM (T2)
 424 then generates $\langle \tilde{T}_{\text{FL}}, \tilde{P}_{\text{FL}} \rangle$ from \tilde{T}_{FL} , both in zero-shot. For *pass@k*, we sample the top- k from T1,
 425 select one that is TC and judged equivalent to T_{NL} by Gemini-2.5-Pro (SC is not used as the gold
 426 T_{FL} is withheld until the pipeline finishes), and then generate the top- k proof candidates from T2.
 427

428 5.5 EXPERIMENTAL RESULTS

429 **NL/FL Cross-Modal Retrieval.** Table 1 compares PROOFBRIDGE’s joint embedding model with
 430 the two *SoTA Encoders* and the *Baseline Encoder*. Since PROOFBRIDGE is obtained by contrastively
 431 training the NL encoder together with an FL encoder, all improvements are reported relative to the
 432 original NL encoder (the Baseline Encoder). PROOFBRIDGE achieves consistently higher recall
 433 rates across all top- K values: for NL \rightarrow FL, it yields $3.28 \times$ gain for NL \rightarrow FL and $1.94 \times$ for FL \rightarrow NL

¹We admit some propositional equalities in addition to definitional ones, see details in Appendix A.4

432 **Table 1: NL/FL Cross-Modal Retrieval Performance.** Retrieval performance of *SoTA* and *Baseline*
 433 encoders versus PROOFBRIDGE’s joint embedding. (**Q1**: 25th %tile, **M**: median, **Q3**: 75th %tile)

Method	Retrieval Direction	Recall Rate @ K (%) ↑					MRR ↑	Cos. Similarity of top-K Retrieved ↑					Cos. Similarity of NOT Retrieved ↓					Gap ↑ (mMG)
		K=1	K=5	K=10	K=20	K=50		K=1	K=5	K=10	K=20	K=50	K=1	K=5	K=10	K=20	K=50	
Qwen3-Embedding-8B (<i>SoTA Encoder</i> , 8B params)	NL → FL	46.75	71.96	82.49	87.04	93.34	0.567	Q1: 0.74	0.67	0.62	0.60	0.57	0.317	0.317	0.317	0.317	0.317	0.29
	FL → NL	44.68	68.26	79.79	83.78	87.36		Q1: 0.74	0.68	0.63	0.62	0.58	0.362	0.362	0.362	0.362	0.362	
E5-Mistral-7B-Instruct (<i>SoTA Encoder</i> , 7B params)	NL → FL	35.60	53.22	60.22	67.82	77.18	0.441	Q1: 0.76	0.66	0.63	0.62	0.59	0.309	0.309	0.309	0.309	0.308	0.10
	FL → NL	30.27	41.93	46.41	50.83	57.88		Q1: 0.76	0.68	0.65	0.64	0.62	0.418	0.418	0.418	0.418	0.417	
all-MiniLM-L6-v2 (<i>Baseline Encoder</i> , 22.7M params)	NL → FL	16.06	30.93	38.63	47.31	60.95	0.237	Q1: 0.86	0.83	0.82	0.82	0.81	0.711	0.711	0.711	0.711	0.710	0.31
	FL → NL	26.35	45.12	54.28	63.75	75.54		Q1: 0.86	0.84	0.83	0.82	0.82	0.749	0.749	0.749	0.749	0.749	
PROOFBRIDGE (<i>Proposed</i> , 22.7M + 218M + 1M params)	NL → FL	52.83	79.81	87.06	91.49	95.08	0.650	Q1: 0.87	0.85	0.84	0.84	0.83	0.704	0.704	0.704	0.704	0.704	0.65
	FL → NL	51.23	78.77	86.18	90.50	94.83		Q1: 0.87	0.86	0.85	0.85	0.84	0.766	0.760	0.760	0.760	0.759	

450 at $K = 1$. MRR also improves by $2.74 \times$ and $1.79 \times$ for NL → FL and FL → NL, respectively. This
 451 indicates that PROOFBRIDGE more frequently retrieves the correct cross-modal counterpart among
 452 the highest-ranked results. For the NL and FL embeddings by PROOFBRIDGE, the median (**M**) co-
 453 sine similarity of retrieved cross-modal instances averages 0.64 across top- K ($K = 1, 5, 10, 20, 50$)
 454 in both directions, showing that select cross-modal theorem–proof pairs are tightly clustered. In con-
 455 trast, non-retrieved instances are much farther apart, averaging -0.01 . Compared to the Baseline,
 456 retrieved-pair similarities increase by 23%, while non-retrieved similarities decrease by 104%.

457 PROOFBRIDGE outperforms the *SoTA Encoders* in recall and MRR, with a much higher mMG
 458 despite being $32 \times$ smaller, showing a clearer separation between retrieved and non-retrieved items.
 459 For NL → FL, E5-Mistral-7B-Instruct attains an mMG of 0.10, while Qwen3-Embedding-8B reaches
 460 0.29. We believe these low values arise because: *first*, these QA-oriented encoders capture coarse
 461 domain-level signals rather than fine-grained mathematical semantics, so most mathematical texts
 462 cluster together; *second*, as plain-text, non-DAG-aware encoders, they rely on superficial lexical
 463 cues (keywords), but in Lean many proofs share the same tactics, making keyword overlap non-
 464 discriminative. In contrast, PROOFBRIDGE leverages the DAG to distinguish proofs, achieving an
 465 mMG of 0.65. Overall, encoding Lean proofs via linearized DAG traversals and contrastive align-
 466 ment with a DAG-aware FL encoder yield an effective joint embedding space, where equivalent
 467 NL-FL pairs cluster tightly and inequivalent ones remain well separated. This enables reliable re-
 468 trieval of the most relevant FL demonstrations to condition the LLM during auto-formalization.

469 **Proof Auto-Formalization.** Table 2 reports the proof auto-formalization performance of 13 SoTA
 470 LLM-based tools. *Theorem auto-formalization LLMs* achieve 0% TC and SC across all pass@k.
 471 These models are designed to formalize theorem statements only, leaving proofs as `sorry`. They
 472 lack knowledge of proof DAGs and tactics, making them unsuitable for end-to-end proof auto-
 473 formalization. Among *foundation models*, Gemini-2.5-Flash-Lite achieves 2.87% SC and 21.31%
 474 TC at pass@32, which increase to 4.10% SC, 18.44% TC for Flash and 8.61% SC, 31.56% TC
 475 for Pro. GPT-5-mini attains 9.02% TC and 34.84% SC at pass@32. They struggle with the strict
 476 syntax and semantics of specialized FLs like Lean, which are underrepresented in their training
 477 data. The *SoTA Two-Step* achieves 43.44% SC and 59.43% TC, but it is prone to cascading errors:
 478 an incorrect FL theorem from the first model causes the second to produce a semantically incorrect
 479 theorem-proof pair. Among the *SoTA AFPS LLMs*, Kimina-Prover-72B achieves the strongest zero-
 480 shot performance at pass@32, with 46.31% SC and 79.51% TC. We build PROOFBRIDGE on top of
 481 a smaller variant, Kimina-Prover-RL-1.7B, by retrieving five relevant FL proofs via NL/FL cross-
 482 modal retrieval and using them for retrieval-augmented SFT and inference, along with iterative proof
 483 repair. PROOFBRIDGE surpasses the zero-shot performance of Kimina-Prover-RL-1.7B by +22.54%
 484 SC and +20.49% TC, and its random few-shot performance by +31.14% SC and +1.64% TC.

485 In Figure 3, we present the pass@32 performance across mathematical domains. Following the
 486 taxonomy by Zheng et al. (2021), the benchmark includes 6 induction/sequence (2.46%), 69 number-
 487 theory (28.28%), 90 algebra (36.89%), and 79 contest problems (32.38%) sourced from AIME,
 488 AMC, and IMO. PROOFBRIDGE performs best on number theory, achieving over 85% SC, while
 489 contest problems remain the most challenging, reaching only about 35% SC.

486
487 **Table 2: Proof Auto-Formalization Performance.** Comparison of LLM-based tools on *Semantic*
488 *Correctness (SC)* and *Type Correctness (TC)* across pass@ k metrics ($k \in \{1, 2, 4, 8, 16, 32\}$).

Setting	LLM/Tool	Semantic Correctness (SC) (%) ↑						Type Correctness (TC) (%) ↑					
		pass@1	pass@2	pass@4	pass@8	pass@16	pass@32	pass@1	pass@2	pass@4	pass@8	pass@16	pass@32
SoTA Tools (zero-shot)	Kimina-Prover-RL-1.7B	9.02	13.93	22.54	30.33	35.25	40.16	26.23	41.80	56.15	62.70	68.03	75.00
	Kimina-Prover-Distill-8B	10.66	18.85	23.77	32.38	37.70	41.80	27.05	43.03	58.20	63.52	72.95	75.82
	Kimina-Prover-72B	12.70	21.31	25.00	34.84	38.52	43.03	30.33	45.08	61.07	69.26	75.41	79.51
SoTA Tools (random few-shot)	Kimina-Autoformalizer-7B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Herald-Translator	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Gemini-2.5-Flash-Lite	0.00	0.00	0.00	0.00	1.23	2.87	0.82	4.51	9.02	13.93	18.03	21.31
	Gemini-2.5-Flash	0.00	0.82	2.05	2.86	3.28	4.10	2.45	4.92	7.38	12.30	15.98	18.44
	Gemini-2.5-Pro	1.23	1.23	3.28	4.92	6.97	8.61	9.84	13.52	18.85	23.77	29.10	31.56
	GPT-5-mini	0.41	1.23	4.51	6.97	7.38	9.02	4.92	9.43	20.08	28.28	32.38	34.84
	DeepSeek-Prover-V1.5-RL	3.69	6.15	8.61	9.02	11.07	12.30	8.20	14.34	19.67	24.18	28.68	35.66
	STP.model.Lean.0320	4.51	6.56	8.20	9.84	11.48	13.11	12.70	18.03	23.36	28.69	33.61	39.34
	Goedel-Prover-SFT	4.92	5.33	7.38	8.20	12.70	16.39	13.52	17.21	25.41	31.56	36.88	42.21
	Leanabell-Prover-V2-KM	6.97	9.43	10.66	13.52	15.16	18.03	16.80	21.31	27.87	37.30	41.39	50.41
SoTA Tools (text-based retrieval few-shot)	Kimina-Prover-RL-1.7B	6.15	12.30	17.62	22.13	27.46	31.56	26.23	42.21	60.66	74.18	88.11	93.85
	Kimina-Prover-Distill-8B	7.38	11.89	16.39	24.18	28.69	32.38	24.59	44.26	61.89	75.00	85.25	89.34
	Kimina-Prover-72B	10.25	13.93	18.85	25.00	31.35	37.30	30.74	45.49	62.70	77.87	86.89	91.39
SoTA Two-Step	Kimina-Prover-RL-1.7B	6.15	12.70	18.85	22.54	28.28	32.78	26.63	43.39	58.68	70.66	86.36	89.75
	Kimina-Prover-Distill-8B	8.61	13.52	21.72	28.28	34.02	36.07	28.28	46.28	59.92	71.90	83.47	86.07
	Kimina-Prover-72B	12.29	14.34	24.59	29.51	37.70	44.26	31.15	45.08	64.88	75.62	86.78	88.93
Our Tool	Herald-Translator → Kimina-Prover-Distill-8B	14.75	19.26	27.05	33.20	38.52	43.44	30.33	32.79	43.03	48.36	54.10	59.43
	PROOFBRIDGE (SFT only)	6.97	13.52	19.26	24.18	29.92	34.84	27.87	45.90	60.66	66.39	72.13	78.69
	PROOFBRIDGE (Retrieval-augm. SFT)	13.11	20.90	27.87	35.66	47.95	55.33	29.92	46.31	60.25	71.31	83.20	89.75
	PROOFBRIDGE (Retrieval-augm. SFT + Repair)	16.39	25.41	29.51	37.70	50.41	62.70	32.79	47.13	64.75	78.69	90.16	95.49

5.6 ABLATION STUDIES

To assess the effect of in-context examples, Table 2 reports zero-shot, random few-shot, and text-based retrieval few-shot performance for three Kimina-Prover variants (72B, Distill-8B, and RL-1.7B). From the pass@32 results, Kimina-Prover-RL-1.7B achieves 40.16% SC and 75.00% TC in the zero-shot setting. When random examples are added, SC drops to 31.56% while TC rises to 93.85%, with similar trends across variants. This indicates that random examples improve TC (syntax) but hurt SC by causing the model to hallucinate semantically misaligned proofs. With text-based retrieval via Qwen3-Embedding-8B, SC rises to 32.38% but TC declines, likely because QA-based retrieval favors proofs with similar tactics, reducing tactic diversity. This highlights the need for retrieving semantically relevant examples via a DAG-aware encoder, as in PROOFBRIDGE.

To quantify the contribution of each component of PROOFBRIDGE, we perform an ablation over three variants. **PROOFBRIDGE (SFT)**, fine-tuned on labeled NL-FL pairs and evaluated in the few-shot setting with semantically relevant examples via our joint-embedding model, improves SC by +2.06% but reduces TC by -11.06% relative to Kimina-Prover-RL-1.7B (text-based retrieval few-shot). Next, in **PROOFBRIDGE (Retrieval-augmented SFT)**, we fine-tune the LLM with semantically relevant FL proofs included in the input, achieving +22.55% SC. **PROOFBRIDGE (Retrieval-augmented SFT + Repair)**, adding iterative proof repair, yields the best results: +29.92% SC and +5.74% TC over the same baseline. Relative to Kimina-Prover-RL-1.7B (random few-shot), the improvements are +31.14% SC and +1.64% TC.

6 CONCLUSION

We present PROOFBRIDGE, a unified framework for NL-to-Lean proof auto-formalization that translates both theorems and proofs end-to-end². At its core is a joint embedding model of NL and FL that encodes Lean proof DAGs, capturing tactic sequences and dependency structures. It enables highly effective cross-modal retrieval of semantically relevant FL proofs. These retrieved proofs act as demonstrations, guiding retrieval-augmented fine-tuning of an LLM. An iterative verifier-guided repair loop further refines generated proofs by combining Lean type-checking with semantic equivalence checking to ensure correctness. Evaluated on MINIF2F-TEST-PF, PROOFBRIDGE significantly outperforms state-of-the-art LLMs in both semantic correctness (by bi-directional equivalence proving) and type correctness, demonstrating that integrating structured embeddings, retrieval guidance, and verifier feedback leads to more reliable proof auto-formalization.

²**Reproducibility:** System details (Appendix A.2), code and datasets provided in the supplementary.

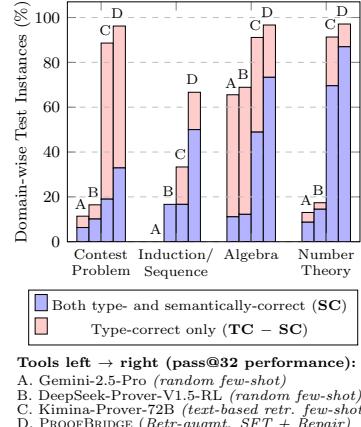


Figure 3: **Category-wise Results.**

Proof auto-formalization performance across mathematical domains.

540 REFERENCES
541

542 Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Yann Coscoy, David Delahaye,
543 Daniel de Rauglaudre, Jean-Christophe Filliâtre, Eduardo Giménez, Hugo Herbelin, et al. The
544 Coq Proof Assistant Reference Manual. *INRIA, version*, 6(11), 1999.

545 Kevin Buzzard. Formalising mathematics. Lecture notes from a course at Imperial
546 College London, 2022. URL <https://github.com/ImperialCollegeLondon/formalising-mathematics>.

547

549 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
550 Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the
551 frontier with advanced reasoning, multimodality, long context, and next generation agentic capa-
552 bilities. *arXiv preprint arXiv:2507.06261*, 2025.

553 Google Deepmind. AI achieves Silver-Medal Standard solving International Mathematical
554 Olympiad Problems, 2024. URL <https://deepmind.google/discover/blog/ai-solves-imo-problems-at-silver-medal-level/>. Accessed: Sep, 2025.

555

557 Kefan Dong and Tengyu Ma. Beyond Limited Data: Self-play LLM Theorem Provers with Iterative
558 Conjecturing and Proving. *arXiv preprint arXiv:2502.00212*, 2025.

559

560 Guoxiong Gao, Haocheng Ju, Jiedong Jiang, Zihan Qin, and Bin Dong. A Semantic Search Engine
561 for Mathlib4. *arXiv preprint arXiv:2403.13310*, 2024a.

562 Guoxiong Gao, Yutong Wang, Jiedong Jiang, Qi Gao, Zihan Qin, Tianyi Xu, and Bin Dong. Herald:
563 A Natural Language Annotated Lean 4 Dataset. *arXiv preprint arXiv:2410.10878*, 2024b.

564

565 Jesse Michael Han, Jason Rute, Yuhuai Wu, Edward Ayers, and Stanislas Polu. Proof Artifact Co-
566 Training for Theorem Proving with Language Models. In *International Conference on Learning
567 Representations*, 2022.

568

569 John Harrison. HOL Light: An Overview. In *International Conference on Theorem Proving in
570 Higher Order Logics*, pp. 60–66. Springer, 2009.

571

572 Prithwish Jana. NeuroSymbolic LLM for mathematical reasoning and software engineering. In
573 *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*,
574 pp. 8492–8493, 2024.

575

576 Prithwish Jana, Piyush Jha, Haoyang Ju, Gautham Kishore, Aryan Mahajan, and Vijay Ganesh.
577 CoTran: An LLM-based Code Translator using Reinforcement Learning with Feedback from
578 Compiler and Symbolic Execution. In *Proceedings of the 27th European Conference on Artificial
Intelligence (ECAI)*, pp. 4011–4018, 2024.

579

580 Piyush Jha, Prithwish Jana, Pranavkrishna Suresh, Arnav Arora, and Vijay Ganesh. RLSF: Fine-
581 tuning LLMs via Symbolic Feedback. In *Proceedings of the 28th European Conference on Arti-
582 ficial Intelligence (ECAI)*, 2025.

583

584 Albert Q Jiang, Sean Welleck, Jin Peng Zhou, Wenda Li, Jiacheng Liu, Mateja Jamnik, Timothée
585 Lacroix, Yuhuai Wu, and Guillaume Lampe. Draft, Sketch, and Prove: Guiding Formal Theorem
586 Provers with Informal Proofs. *arXiv preprint arXiv:2210.12283*, 2022.

587

588 Albert Q Jiang, Wenda Li, and Mateja Jamnik. Multi-language Diversity Benefits Autoformaliza-
589 tion. *Advances in Neural Information Processing Systems*, 37:83600–83626, 2024.

590

591 Albert Qiaochu Jiang, Wenda Li, Jesse Michael Han, and Yuhuai Wu. LISA: Language models of
592 ISAbelle proofs. In *6th Conference on Artificial Intelligence and Theorem Proving*, pp. 378–392,
593 2021.

594

595 Leanprover. leanprover-community/repl: A simple repl for lean 4, 2025. URL <https://github.com/leanprover-community/repl>. Accessed: Sep, 2025.

594 Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang,
 595 Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann
 596 Fleureau, Guillaume Lample, and Stanislas Polu. NuminaMath. [<https://huggingface.co/AI-MO/NuminaMath-1.5>] (https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024.

597

598

599 Yong Lin, Shange Tang, Bohan Lyu, Jiayun Wu, Hongzhou Lin, Kaiyu Yang, Jia Li, Mengzhou Xia,
 600 Danqi Chen, Sanjeev Arora, et al. Goedel-Prover: A Frontier Model for Open-Source Automated
 601 Theorem Proving. *arXiv preprint arXiv:2502.07640*, 2025.

602

603 Qi Liu, Xinhao Zheng, Xudong Lu, Qinxiang Cao, and Junchi Yan. Rethinking and Improving
 604 Autoformalization: Towards a Faithful Metric and a Dependency Retrieval-based Approach. In
 605 *The Thirteenth International Conference on Learning Representations*, 2025.

606

607 Jianqiao Lu, Yingjia Wan, Zhengying Liu, Yinya Huang, Jing Xiong, Chengwu Liu, Jianhao Shen,
 608 Hui Jin, Jipeng Zhang, Haiming Wang, et al. Process-driven Autoformalization in Lean 4. *arXiv*
 609 *preprint arXiv:2406.01940*, 2024.

610

611 Norman Megill and David A Wheeler. *Metamath: A Computer Language for Mathematical Proofs*.
 Lulu. com, 2019.

612

613 Leonardo de Moura and Sebastian Ullrich. The Lean 4 Theorem Prover and Programming Lan-
 614 guage. In *Automated Deduction–CADE 28: 28th International Conference on Automated Deduc-
 615 tion, Virtual Event, July 12–15, 2021, Proceedings 28*, pp. 625–635. Springer, 2021.

616

617 Tobias Nipkow, Markus Wenzel, and Lawrence C Paulson. *Isabelle/HOL: A Proof Assistant for
 Higher-order Logic*. Springer, 2002.

618

619 NuminaMath. minif2f.test. Hugging Face Dataset, September.23 2025. URL https://huggingface.co/datasets/AI-MO/minif2f_test. Version 1.0, Apache-2.0 Li-
 620 cense, 244 rows.

621

622 OpenAI. Introducing gpt-5, 2025. URL <https://openai.com/index/introducing-gpt-5/>. Accessed: 2025-09-23.

623

624 Lawrence C Paulsson and Jasmin C Blanchette. Three Years of Experience with Sledgehammer, a
 625 Practical Link Between Automatic and Interactive Theorem Provers. In *Proceedings of the 8th
 626 International Workshop on the Implementation of Logics (IWIL-2010), Yogyakarta, Indonesia.
 627 EPiC*, volume 2, 2012.

628

629 Gabriel Poesia and Noah D Goodman. Peano: Learning Formal Mathematical Reasoning. *Philoso-
 630 sophical Transactions of the Royal Society A*, 381(2251):20220044, 2023.

631

632 Auguste Poiroux, Gail Weiss, Viktor Kunčak, and Antoine Bosselut. Reliable Evaluation and Bench-
 633 marks for Statement Autoformalization. In *Proceedings of the 2025 Conference on Empirical
 634 Methods in Natural Language Processing*, pp. 17958–17980, 2025.

635

636 Stanislas Polu and Ilya Sutskever. Generative Language Modeling for Automated Theorem Proving.
 arXiv preprint arXiv:2009.03393, 2020.

637

638 Nils Reimers and Iryna Gurevych. Sentence-Transformers: Multilingual Sentence Embeddings
 639 using BERT and XLNet. <https://www.sbert.net/>, 2020.

640

641 ZZ Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia Zhao, Liyue Zhang,
 642 Zhe Fu, Qihao Zhu, Dejian Yang, et al. Deepseek-Prover-v2: Advancing Formal Mathe-
 643 matical Reasoning via Reinforcement Learning for Subgoal Decomposition. *arXiv preprint
 644 arXiv:2504.21801*, 2025.

645

646 Ziju Shen, Naohao Huang, Fanyi Yang, Yutong Wang, Guoxiong Gao, Tianyi Xu, Jiedong Jiang,
 647 Wanyi He, Pu Yang, Mengzhou Sun, et al. REAL-Prover: Retrieval Augmented Lean Prover for
 648 Mathematical Reasoning. *arXiv preprint arXiv:2505.20613*, 2025.

649

Terence Tao. A maclaurin type inequality. *arXiv preprint arXiv:2310.05328*, 2023.

648 Amitayush Thakur, Yeming Wen, and Swarat Chaudhuri. A Language-Agent Approach to Formal
 649 Theorem-Proving. 2023.

650

651 Shubham Ugare, Tarun Suresh, Hangoo Kang, Sasa Misailovic, and Gagandeep Singh. SynCode:
 652 LLM Generation with Grammar Augmentation. *arXiv preprint arXiv:2403.01632*, 2024.

653

654 Haiming Wang, Mert Unsal, Xiaohan Lin, Mantas Baksys, Junqi Liu, Marco Dos Santos, Flood
 655 Sung, Marina Vinyes, Zhenzhe Ying, Zekai Zhu, et al. Kimina-Prover Preview: Towards Large
 656 Formal Reasoning Models with Reinforcement Learning. *arXiv preprint arXiv:2504.11354*, 2025.

657

658 Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. Improv-
 659 ing Text Embeddings with Large Language Models. In *Proceedings of the 62nd Annual Meeting*
 660 *of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11897–11916,
 661 2024a.

662

663 Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi, and Tong Zhang. Theorem-
 664 Llama: Transforming General-purpose LLMs into Lean4 Experts. In *Conference on Empirical*
 665 *Methods in Natural Language Processing (EMNLP)*, pp. 11953–11974, 2024b.

666

667 Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus Rabe, Charles Staats, Mateja Jamnik, and
 668 Christian Szegedy. Autoformalization with Large Language Models. *Advances in Neural Infor-*
 669 *mation Processing Systems*, 35:32353–32368, 2022.

670

671 Yutong Wu, Di Huang, Ruosi Wan, Yue Peng, Shijie Shang, Chenrui Cao, Lei Qi, Rui Zhang, Zidong
 672 Du, Jie Yan, et al. Stepfun-formalizer: Unlocking the autoformalization potential of llms through
 673 knowledge-reasoning fusion. *arXiv preprint arXiv:2508.04440*, 2025.

674

675 Huajian Xin, Daya Guo, Zhihong Shao, Zhizhou Ren, Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li,
 676 and Xiaodan Liang. DeepSeek-Prover: Advancing Theorem Proving in LLMs through Large-
 677 Scale Synthetic Data. *arXiv preprint arXiv:2405.14333*, 2024a.

678

679 Kaiyu Yang. minif2f-lean4. GitHub repository, 2025. URL <https://github.com/yangky11/minif2f-lean4>.

680

681 Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil,
 682 Ryan Prenger, and Anima Anandkumar. LeanDojo: Theorem proving with retrieval-augmented
 683 language models. In *Neural Information Processing Systems (NeurIPS)*, 2023.

684

685 Huaiyuan Ying, Zijian Wu, Yihan Geng, Jiayu Wang, Dahua Lin, and Kai Chen. Lean Workbook:
 686 A Large-scale Lean Problem Set Formalized from Natural Language Math Problems. *Advances*
 687 *in Neural Information Processing Systems*, 37:105848–105863, 2024.

688

689 Jingyuan Zhang, Qi Wang, Xinguang Ji, Yahui Liu, Yang Yue, Fuzheng Zhang, Di Zhang, Guorui
 690 Zhou, and Kun Gai. Leanabell-prover: Posttraining scaling in formal reasoning. *arXiv preprint*
 691 *arXiv:2504.06122*, 2025a.

692

693 Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie,
 694 An Yang, Dayiheng Liu, Junyang Lin, et al. Qwen3 Embedding: Advancing Text Embedding and
 695 Reranking Through Foundation Models. *arXiv preprint arXiv:2506.05176*, 2025b.

696

697 Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. MiniF2F: A Cross-System Benchmark for
 698 Formal Olympiad-level Mathematics. *arXiv preprint arXiv:2109.00110*, 2021.

699

700 Yuhao Zhou. Retrieval-Augmented TLAPS Proof Generation with Large Language Models. *arXiv*
 701 *preprint arXiv:2501.03073*, 2025.

702 **A APPENDIX**
703704 **A.1 THE USE OF LARGE LANGUAGE MODELS (LLMs)**
705706 LLMs did *not* play a significant role in either the research ideation or the writing of this paper. Their
707 use was limited to correcting minor grammatical issues and typographical errors.
708709 **A.2 REPRODUCIBILITY**
710711 All implementations in this work, including dataset construction, model training, cross-modal re-
712 trieval, inference, and evaluation, use Python 3.12.10 and Lean v4.15.0. Experiments were con-
713 ducted on a high-performance AlmaLinux 9.5 (Teal Serval) cluster with a single Intel Xeon Plat-
714 inum 8480+ CPU (32 cores, 2.0-4.0 GHz), 251 GiB of RAM, and one NVIDIA H100 GPU. Our
715 full codebase, including scripts for dataset generation, model fine-tuning, and inference across both
716 GPU- and API-based setups, is shared through a supplementary .zip file. The repository is com-
717 plete, well-documented, and designed for full reproducibility: it includes instructions for creating a
718 Python virtual environment, and a comprehensive README outlining library dependencies, dataset
719 format, and step-by-step instructions to replicate the entire data pipeline and experimental workflow.
720721 **A.3 TRAINING AND INFERENCE HYPERPARAMETERS**
722723 **NL/FL Cross-Modal Retrieval.** We train two dense encoders to embed NL and FL
724 theorem+proof pairs into a shared semantic space. The NL encoder is initialized from
725 all-MiniLM-L6-v2 (Reimers & Gurevych, 2020), while the FL encoder builds on Lean-
726 Dojo’s (Yang et al., 2023) ByT5-based proof-state encoder, extended to process a linearized traversal
727 of the proof DAG. Each encoder is equipped with a projection head that maps representations into a
728 shared embedding space of dimension $d = 512$, with a dropout rate of 0.1. During fine-tuning, we
729 update only the top layers of each encoder to retain their pretrained linguistic and structural priors.
730 Specifically, we train the last 3 layers of the NL encoder and the last 2 layers of the FL encoder. We
731 set the maximum token length to 512 for both NL and FL sequences. The model is optimized using
732 our symmetric contrastive objective (Equation 1) with temperature $\tau = 0.07$, trained using AdamW
733 with a learning rate of 1×10^{-5} , weight decay of 0.01, and a batch size of 32. Training is run for 10
734 epochs with gradient accumulation steps set to 4. We also enable gradient checkpointing to reduce
735 memory usage during fine-tuning.
736737 **Proof Auto-Formalization.** PROOFBRIDGE builds on Kimina-Prover-RL-1.7B (Wang et al.,
738 2025), which we further fine-tune for NL \rightarrow FL translation using paired data (M_{NL}, M_{FL}) from
739 NUMINAMATH-LEAN-PF. During both SFT and inference, our NL/FL cross-modal retrieval model
740 gets the top-5 most relevant FL proofs from \mathcal{D} , which are provided as in-context demonstrations to
741 guide Lean proof synthesis. We use the HuggingFace Trainer for supervised fine-tuning with the
742 following settings: a per-device batch size of 8 and BF16 training enabled. Training is run for 5
743 epochs with a learning rate of 5×10^{-6} , cosine decay scheduling, and a warmup ratio of 0.05.
744745 For all SoTA baselines in Table 2, we compute pass@k using stochastic decoding. Specifically,
746 we run LLM inference with a temperature of 0.6 and top-p sampling of 0.95, ensuring sufficient
747 diversity across generated candidates.
748749 **A.4 SEMANTIC EQUIVALENCE OF LEAN THEOREMS**
750751 The task of autoformalization is to convert a mathematical theorem and proof from natural language
752 into a formal language, such as Lean. When evaluating the performance of such systems, we propose
753 two criteria for evaluation: *type correctness* and *semantic equivalence*. Type correctness, which
754 requires that the generated Lean proof is accepted by the Lean type-checker, is straightforward to
755 verify and serves as the standard evaluation metric in the field. However, semantic equivalence,
756 ensuring the FL theorem faithfully represents the meaning of the original NL theorem, presents a far
757 greater challenge. To the best of our knowledge, semantic equivalence has not been systematically
758 evaluated in prior work. This section introduces a novel methodology towards addressing this gap.
759760 While directly measuring the semantic alignment between a NL theorem and a Lean theorem is an
761 unsolved challenge, showing the logical equivalence of two Lean theorems is a tractable task. Our
762

756 training dataset, NUMINAMATH-LEAN-PF, contains pairs of $\langle T_{\text{NL}}, T_{\text{FL}} \rangle$ where most of the T_{FL}
 757 were manually created by experts at Numina. We treat these high-quality T_{FL} theorems as gold-
 758 standard references, assuming they are faithful translations of their T_{NL} counterparts. This allows us
 759 to reduce the intractable problem of verifying a model’s generated theorem \tilde{T}_{FL} against the original
 760 T_{NL} to the more tractable task of checking for logical equivalence between \tilde{T}_{FL} and the golden
 761 reference T_{FL} , which can be checked in Lean itself.

762 To be more specific, we enforce this semantic equivalence check by proving the logical biconditional
 763 $\tilde{T}_{\text{FL}} \leftrightarrow T_{\text{FL}}$ in Lean. Theorems like \tilde{T}_{FL} and T_{FL} are of type `Prop` in Lean. The following theorem
 764 from Mathlib states that for any two propositions, a logical biconditional between two propositions
 765 is itself logically equivalent to their propositional equality:

767 `theorem propext_iff{a b : Prop} :`
 768 `a = b ↔ (a ↔ b)`

770 The task thus converts to proving the equality $\tilde{T}_{\text{FL}} = T_{\text{FL}}$ within Lean. This requires clarifying the
 771 specific notion of equality being used, as Lean distinguishes between three primary types: *syntactic*,
 772 *definitional*, and *propositional* Buzzard (2022). Syntactic equality is the strictest form of equality in
 773 Lean, as it only admits expressions that are structurally identical according to their Abstract Syntax
 774 Trees, without any computation or reduction. Definitional equality is a more relaxed form of equality
 775 than syntactic equality, where two expressions are considered equal if they compute or reduce to the
 776 same normal form. Propositional equality is the weakest form of equality, and also the standard
 777 notion of equality used in mathematical theorems. Two terms `a`, `b` are propositionally equal in
 778 Lean if you can construct a proof term for the proposition `a = b`.

779 For our evaluation, we seek to measure how closely a \tilde{T}_{FL} matches the T_{FL} . The strictest criterion,
 780 syntactic equality, is too restrictive given the current state-of-the-art, as it would fail valid theorems
 781 with trivial notational differences. Conversely, full propositional equality can be too permissive; a
 782 proof of equivalence can be arbitrarily complex, making it difficult to automate and decide.

783 Therefore, we adopt a pragmatic compromise: we check for **definitional equality** supplemented by
 784 a form of **bounded propositional equality**. This means we primarily check if \tilde{T}_{FL} and T_{FL} reduce to
 785 the same normal form, but we also permit some propositional equality, provided they can be proven
 786 using a collection of tactics so that their proof complexity is bounded.

787 We then leverage Gemini 2.5 Pro as an automated equivalence checker. The model is prompted
 788 to synthesize a proof for the biconditional theorem $(\tilde{T}_{\text{FL}} \leftrightarrow T_{\text{FL}})$, with instructions limiting it to
 789 a specific subset of available tactics. This restricted set includes three powerful automated tactics,
 790 `rfl`, `simp`, and `ring`, each is designed to discharge a specific class of goals: `rfl` for definitional
 791 equality, `simp` for simplification, and `ring` for polynomial identities.

792 As definitional equality is our primary target, the equivalence checker first attempts to solve the goal
 793 with the `rfl` tactic. This single tactic should suffice for the majority of cases. If `rfl` fails, the
 794 checker then tries `simp`. This tactic performs additional simplifications by rewriting the goal using
 795 theorems from Mathlib that are tagged for its use. Critically, we use `simp` without any arguments.
 796 Providing explicit arguments would require a demanding search for the correct lemmas and could
 797 introduce unbounded complexity, violating our goal of a bounded proof search. Furthermore, the
 798 need for `simp` with arguments could imply that the required rewrite is non-trivial, since the default
 799 simplification set contains most of the trivial facts³. Since our goal is to ensure a close correspon-
 800 dence between \tilde{T}_{FL} and T_{FL} , a proof requiring such a targeted rewrite indicates a semantic distance
 801 that we classify as a mismatch. The `ring` tactic is a valuable complement to the previous tactics
 802 as it specializes in proving polynomial equalities. The `ring` tactic operates by reducing arithmetic
 803 expressions to a canonical normal form. This allows it to prove the equivalence of expressions that
 804 are algebraically identical but not definitionally so, such as x^2 and $x * x$, which `rfl` and default
 805 `simp` would otherwise fail to solve.

806
 807
 808 ³It is important to note that the default `simp` set intentionally excludes lemmas like associativity and com-
 809 mutativity, as they can cause the simplifier to loop indefinitely. However, since these lemmas primarily concern
 algebraic expressions, they can be handled by the `ring` tactic.

810 The three tactics discussed above cover most of the direct equivalences we aim to check. The re-
 811 remaining tactics in our instruction set are designed for a more nuanced case: proving the biconditional
 812 when two theorems differ only in their use of auxiliary variables. We observed that human experts
 813 and language models may make different but equally valid decisions on whether to introduce an
 814 auxiliary variable. We therefore classify such theorems as equivalent. For example, consider the
 815 following:

```
816
817 def Prop1 := (forall (b h v : ℝ), (0 < b ∧ 0 < h ∧ 0 < v) → (v = 1 / 3 * (b *
818   h)) → (b = 30) → (h = 13 / 2) → v = 65)
819 def Prop2 := (forall {B h : ℝ}, (B = 30) → (h = 6.5) → (1 / 3) * B * h = 65)
820 example : Prop1 ↔ Prop2 := by
821   constructor
822   . intro
823     simp
824     ring
825   . simp
826     intros
827     nlinarith
```

828 The main difference between the two propositions `Prop1` and `Prop2` is the presence of the auxiliary
 829 variable `v` in one. To prove that such theorems are equivalent, one must typically prove the bicon-
 830 ditional by separately proving the implications of both direction. This requires a step-by-step proof
 831 construction, and the tactics above are included in our instruction set.

832 Finally, we note that the LLM judge’s role is only to synthesize a biconditional proof under the
 833 bounded tactic set described above; the produced proof is then fully type-checked by the Lean
 834 kernel, so the SC decision ultimately depends on Lean’s verifier (making the metric conservative
 835 rather than prone to false positives).

836 **Comparison with Existing Semantic Correctness Metrics.** Prior work has proposed similar se-
 837 mantic correctness metrics, including BEq (Liu et al., 2025; Wu et al., 2025) and its extensions
 838 BEq+ (Poiroux et al., 2025). While the general idea behind our SC metric and BEq is similar, both
 839 aiming to establish bidirectional equivalence in Lean, the sets of allowed tactics differ. Because the
 840 notion of equivalence depends on the permitted tactics, these differences lead to meaningful distinc-
 841 tions between SC and BEq. BEq+ is a reference-based metric inspired by BEq and uses a set of
 842 tactics comparable to SC. However, BEq+ is deterministic and CPU-efficient, while SC relies on an
 843 LLM-based proof synthesizer. This creates a trade-off: the LLM can capture equivalences beyond
 844 the reach of the deterministic procedure, whereas BEq+ provides a reproducible evaluation.

845 Consider `Prop1` and `Prop2` above as an example. `Prop1` (gold-standard FL theorem from miniF2F-
 846 Test-PF) explicitly introduces an auxiliary variable `v` to denote volume, whereas `Prop2` (produced
 847 by Kimina-Prover-RL-1.7B) omits the auxiliary variable and substitutes the corresponding formula
 848 directly. The tactic set allowed by BEq is not expressive enough to establish equivalence in such
 849 cases, so these theorems would not be recognized as equivalent under BEq. Our SC metric, by
 850 contrast, was specifically designed to handle such variations, reflecting the fact that human experts
 851 may also differ in whether they introduce auxiliary variables. By explicitly handling these variations
 852 and using LLM-generated bidirectional proofs that are type-checked, SC provides an evaluation that
 853 is both more lenient and faithful in assessing the performance of auto-formalization models.

854 A.5 ILLUSTRATIVE EXAMPLE

855 We present an example of an NL theorem+proof pair from MINIF2F-TEST-PF and compare the
 856 `pass@1` output of proof auto-formalization generated by Kimina-Prover-RL-1.7B in the text-based
 857 retrieval few-shot setting with that produced by PROOFBRIDGE (using Retrieval-augmented SFT
 858 + Repair). We first show the retrievals of semantically relevant FL theorem+proof pairs from \mathcal{D} ,
 859 followed by the `pass@1` output proof auto-formalization generated by PROOFBRIDGE. In this
 860 example, the output by Kimina-Prover-RL-1.7B is type-correct (TC) but not semantically correct
 861 (SC), i.e., it is not bi-directionally equivalent to the gold-standard Lean proof. In contrast, the output
 862 by PROOFBRIDGE is both TC and SC.

864
865

Input NL theorem+proof pair (MINIF2F-TEST-PF)

866

<informal_theorem>

867

A point (x, y) on the coordinate plane with both coordinates negative is a distance of 6 units from the x -axis. It is a distance of 15 units from the point $(8, 3)$. It is a distance \sqrt{n} from the origin. What is n ? Show that it is 52.

868

</informal_theorem>

869

<informal_proof>

870

We know that $y = -6$ from the given information. By the distance formula, we have the equation $\sqrt{(x - 8)^2 + (-6 - 3)^2} = 15$. Solving, we have

871

$$\sqrt{(x - 8)^2 + (-6 - 3)^2} = 15$$

872

$$x^2 - 16x + 64 + 81 = 225$$

873

$$x^2 - 16x - 80 = 0$$

874

$$(x - 20)(x + 4) = 0$$

875

Thus, $x + 4 = 0$ or $x - 20 = 0$, so $x = -4$ or $x = 20$. $x = -4$ by the given conditions. Thus, our point is $(-4, -6)$ and is a distance of $\sqrt{(-4)^2 + (-6)^2} = \sqrt{52}$ units from the origin. $n = 52$.

876

</informal_proof>

877

878

879

880

881

882

883

Gold-standard Lean theorem (MINIF2F-TEST-PF)

884

```
import Mathlib
import Aesop
set_option maxHeartbeats 0
open BigOperators Real Nat Topology Rat

theorem mathd_algebra_288 (x y : ℝ) (n : NNReal) (h0 : x < 0 ∧ y < 0) (h1 : abs y = 6)
  (h2 : Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15)
  (h3 : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt n) : n = 52 := by
  sorry
```

885

886

887

Output Lean theorem+proof pair by Kimina-Prover-RL-1.7B

(text-based retrieval few-shot setting)

888

import Mathlib

889

/- A point (x, y) on the coordinate plane with both coordinates negative is a distance of 6 units from the x -axis. It is a distance of 15 units from the point $(8, 3)$. It is a distance \sqrt{n} from the origin. What is n ? Show that it is 52.-/

890

```
theorem algebra_60927 (n : N) (hn : n ≥ 1) (h : ∃ x y : ℝ, x < 0 ∧ y < 0 ∧ |y| = 6 ∧ √
  ((x - 8) ^ 2 + (y - 3) ^ 2) = 15 ∧ √(x ^ 2 + y ^ 2) = √n) : n = 52 := by
  rcases h with ⟨x, y, hx, hy, h1, h2, h3⟩
  have h1' : y = -6 := by
  cases' abs_cases y with hy2 hy2
  · linarith
  · linarith
  have h2' : (x - 8) ^ 2 + (y - 3) ^ 2 = 225 := by
  calc
    (x - 8) ^ 2 + (y - 3) ^ 2 = (√((x - 8) ^ 2 + (y - 3) ^ 2)) ^ 2 := by rw [Real.sq_sqrt];
    positivity
    _ = 15 ^ 2 := by rw [h2]
    _ = (225 : ℝ) := by norm_num
  rw [hy1] at h2'
  have h2x : (x - 8) ^ 2 = 144 := by nlinarith
  have hx1 : x = -4 := by
  nlinarith [h2x, hx]
  have h3' : x ^ 2 + y ^ 2 = (n : ℝ) := by
  calc
    x ^ 2 + y ^ 2 = (√(x ^ 2 + y ^ 2)) ^ 2 := by rw [Real.sq_sqrt]; positivity
    _ = (√n) ^ 2 := by rw [h3]
    _ = (n : ℝ) := by rw [Real.sq_sqrt]; positivity
  rw [hx1, hy1] at h3'
  have hn1 : (n : ℝ) = 52 := by nlinarith
  have hn2 : n = 52 := by
  exact_mod_cast hn1
  exact_mod_cast hn2
```

918 Note that, although the Lean theorem generated by Kimina-Prover-RL-1.7B is type-correct (TC), it
 919 differs semantically from the gold-standard theorem. The main difference lies in the quantification
 920 of variables. In the gold-standard theorem, the variables x, y, n are universally quantified as explicit
 921 arguments, and all hypotheses are stated as direct assumptions; this asserts that for any triple (x, y, n)
 922 satisfying the geometric constraints, $n = 52$. In contrast, the Kimina-Prover-RL-1.7B output univer-
 923 sally quantifies n but existentially quantifies x and y within the hypotheses. This more accurately
 924 reflects the intended geometric meaning: for a given n satisfying the distance constraints, there exists
 925 a point (x, y) realizing those constraints, and consequently $n = 52$. Therefore, while both theorems
 926 are syntactically valid in Lean, they encode slightly different logical statements. This difference pre-
 927 vents an LLM judge from producing a type-checkable proof of bi-directional equivalence between
 928 the two theorems. As a result, the Kimina-Prover-RL-1.7B's output is not semantically correct (SC).
 929

Comparison between the gold-standard theorem and Kimina-Prover-RL-1.7B's output

```

931 /- Gold-standard theorem -/
932 theorem mathd_algebra_288 (x y : ℝ) (n : NReal) (h0 : x < 0 ∧ y < 0) (h1 : abs y = 6)
933   (h2 : Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15)
934   (h3 : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt n) : n = 52 := by
935   sorry
936
937 /- Kimina-Prover-RL-1.7B output theorem -/
938 theorem algebra_60927 (n : N) (hn : n ≥ 1) (h : ∃ x y : ℝ, x < 0 ∧ y < 0 ∧ |y| = 6 ∧ √
939   ((x - 8) ^ 2 + (y - 3) ^ 2) = 15 ∧ √(x ^ 2 + y ^ 2) = √n) : n = 52 := by
940   sorry
  
```

938 Below, we present the relevant FL theorem+proof pairs (demonstrations) from \mathcal{D} retrieved by
 939 PROOFBRIDGE, along with their relevance scores.

Lean theorem+proof pairs retrieved by PROOFBRIDGE's NL/FL Cross-Modal Retrieval

Relevant Lean theorem+proof 1, with relevance score 0.786764 out of 1.0

```

941 import Mathlib
942
943 /- Find the distance between the points $(2,2)$ and $(-1,-1)$. -/
944 theorem algebra_13734 (p1 p2 : ℝ × ℝ) (hp1 : p1 = (2, 2)) (hp2 : p2 = (-1, -1)) :
945   Real.sqrt ((p1.1 - p2.1) ^ 2 + (p1.2 - p2.2) ^ 2) = 3 * Real.sqrt 2 := by
946   rw [hp1, hp2]
947   norm_num
948   ring
949   rw [Real.sqrt_eq_iff_sq_eq] <:> norm_num
950   ring
951   norm_num
  
```

Relevant Lean theorem+proof 2, with relevance score 0.768226 out of 1.0

```

952 import Mathlib
953 open Real
954
955 /- Prove that the angle (in degrees) between the vectors $(2,5)$ and $(-3,7)$ is $45$. -/
956 theorem calculus_17161 :
957   arccos ((2 * (-3) + 5 * 7) / (sqrt (2 ^ 2 + 5 ^ 2) * sqrt ((-3) ^ 2 + 7 ^ 2))) * 180
958   / π = 45 := by
959
960   have h1 : (2 * (-3) + 5 * 7 : ℝ) / (sqrt (2 ^ 2 + 5 ^ 2) * sqrt ((-3) ^ 2 + 7 ^ 2)) =
961   Real.sqrt 2 / 2 := by
962   have h2 : sqrt ((2 : ℝ) ^ 2 + (5 : ℝ) ^ 2) = Real.sqrt 29 := by
963   norm_num [Real.sqrt_eq_iff_sq_eq]
964
965   have h3 : sqrt ((-3 : ℝ) ^ 2 + (7 : ℝ) ^ 2) = Real.sqrt 58 := by
966   norm_num [Real.sqrt_eq_iff_sq_eq]
967
968   have h4 : (2 * (-3) + 5 * 7 : ℝ) = 29 := by norm_num
969
970   rw [h2, h3, h4]
971
972   have h5 : Real.sqrt 29 * Real.sqrt 58 = Real.sqrt 2 * (29 : ℝ) := by
973   calc
974     Real.sqrt 29 * Real.sqrt 58 = Real.sqrt (29 * 58 : ℝ) := by
975     rw [← Real.sqrt_mul (by norm_num)]
976     _ = Real.sqrt ((2 : ℝ) * (29 ^ 2 : ℝ)) := by norm_num
977     _ = Real.sqrt (2 : ℝ) * Real.sqrt ((29 : ℝ) ^ 2 : ℝ) := by
978     rw [Real.sqrt_mul (by norm_num)]
  
```

```

972
973     _ = Real.sqrt (2 : ℝ) * (29 : ℝ) := by
974     rw [Real.sqrt_sq (by norm_num)]
975
976     field_simp [h5]
977     <;> ring_nf <;> norm_num [Real.sqrt]
978
979     rw [h1]
980
981     have h5 : arccos (Real.sqrt 2 / 2) = Real.pi / 4 := by
982     have h6 : Real.sqrt 2 / 2 = Real.cos (Real.pi / 4) := by
983     rw [Real.cos_pi_div_four]
984     <;> ring_nf <;> norm_num
985     <;> ring
986     rw [h6]
987     have h7 : arccos (Real.cos (Real.pi / 4)) = Real.pi / 4 := by
988     apply arccos_cos
989     all_goals linarith [Real.pi_pos]
990     exact h7
991
992     rw [h5]
993
994     field_simp [Real.pi_pos]
995     <;> linarith [Real.pi_gt_three]
996
997     ### Relevant Lean theorem+proof 3, with relevance score 0.765285 out of 1.0
998
999     import Mathlib
1000
1001     /- Show that the sum of  $\sqrt{3x^2 + 2x + 1}$  and  $\sqrt{3x^2 - 4x + 2}$  is at least  $\sqrt{51}/3$  for all real  $x$ . -/
1002     theorem inequalities_201318 (x : ℝ) :
1003     Real.sqrt (3 * x^2 + 2 * x + 1) + Real.sqrt (3 * x^2 - 4 * x + 2) ≥
1004     Real.sqrt 51 / 3 := by
1005     set y := Real.sqrt (3 * x^2 + 2 * x + 1)
1006     set z := Real.sqrt (3 * x^2 - 4 * x + 2)
1007     have hy2 : y^2 = 3 * x^2 + 2 * x + 1 := by
1008     rw [Real.sqrt]
1009     nlinarith [sq_nonneg (x + 1 / 3)]
1010     have hz2 : z^2 = 3 * x^2 - 4 * x + 2 := by
1011     rw [Real.sqrt]
1012     nlinarith [sq_nonneg (x - 2 / 3)]
1013     have hy4_pos : 0 ≤ (3 * x^2 + 2 * x + 1 : ℝ) := by
1014     nlinarith [sq_nonneg (x * 3 + 1)]
1015     have hz4_pos : 0 ≤ (3 * x^2 - 4 * x + 2 : ℝ) := by
1016     nlinarith [sq_nonneg (x * 3 - 2)]
1017     have h11 : (Real.sqrt 51 / 3) ^ 2 = (51 / 9 : ℝ) := by
1018     calc
1019     (Real.sqrt 51 / 3) ^ 2 = (Real.sqrt 51) ^ 2 / 9 := by ring
1020     - = (51 / 9 : ℝ) := by
1021     rw [Real.sqrt_sq (by norm_num)]
1022     <;> ring
1023     have h50 : (y + z) ^ 2 ≥ (Real.sqrt 51 / 3) ^ 2 := by
1024     nlinarith [sq_nonneg (y - z), sq_nonneg (x - 2 / 3), sq_nonneg (x + 1 / 3),
1025     h11, Real.sqrt_nonneg 51, Real.sqrt_sq (show 0 ≤ (51 : ℝ) by norm_num),
1026     mul_nonneg (Real.sqrt_nonneg (3 * x^2 + 2 * x + 1)) (Real.sqrt_nonneg (3 * x^2 - 4 *
1027     x + 2)),
1028     sq_nonneg (y ^ 2 - z ^ 2), sq_nonneg (y * z - Real.sqrt ((3 * x^2 + 2 * x + 1) * (3
1029     * x^2 - 4 * x + 2)))]
1030
1031     have h51 : (y + z) ≥ 0 := by positivity
1032     have h54 : (Real.sqrt 51 / 3) ≥ 0 := by positivity
1033     have h52 : (y + z) ≥ (Real.sqrt 51 / 3) := by
1034     have h15 : (y + z) ^ 2 ≥ (Real.sqrt 51 / 3) ^ 2 := h50
1035     have h16 : (Real.sqrt 51 / 3) ≥ 0 := h54
1036     have h17 : (y + z) - (Real.sqrt 51 / 3) ≥ 0 := by
1037     nlinarith [sq_nonneg ((y + z) - (Real.sqrt 51 / 3)),
1038     sq_nonneg ((y + z) + (Real.sqrt 51 / 3)),
1039     Real.sqrt_pos.mpr (show (0 : ℝ) < 51 by linarith : (51 : ℝ) > 0)
1040     ]
1041     linarith
1042     linarith
1043
1044     ### Relevant Lean theorem+proof 4, with relevance score 0.764932 out of 1.0
1045
1046     import Mathlib
1047     open Real Set
1048     open scoped BigOperators
1049
1050     /- Given that  $\sqrt{(x-3)^2 + (y+4)^2} + \sqrt{(x+5)^2 + (y-8)^2} = 20$ , prove that the
1051     distance between  $(-5, 8)$  and  $(3, -4)$  is  $4\sqrt{13}$ . -/
1052

```

```

1026
1027 theorem algebra_10804 {x y : ℝ} (h : sqrt ((x - 3)^2 + (y + 4)^2) + sqrt ((x + 5)^2 + (y
1028 - 8)^2) = 20) :
1029   sqrt ((-5 - 3)^2 + (8 - (-4))^2) = 4 * sqrt 13 := by
1030   have h1 : sqrt ((-5 - 3)^2 + (8 - (-4))^2) = sqrt 208 := by
1031   norm_num
1032   rw [h1]
1033   have h2 : sqrt 208 = 4 * sqrt 13 := by
1034   rw [Real.sqrt_eq_iff_sq_eq] <; norm_num
1035   <;> ring_nf <;> norm_num
1036   rw [h2]
1037   all_goals norm_num
1038
1039 ### Relevant Lean theorem+proof 5, with relevance score 0.755091 out of 1.0
1040
1041 import Mathlib
1042
1043 /- Show that the square root of  $(-3)^2$  is  $3$ , that  $3$  is non-negative, and that  $3$ 
1044 squared equals  $(-3)$  squared. -/
1045 theorem algebra_8272 : Real.sqrt ((-3)^2) = 3 ∧ 0 ≤ 3 ∧ 3^2 = (-3)^2 := by
1046   have h1 : Real.sqrt ((-3 : ℝ) ^ 2) = 3 := by
1047   rw [Real.sqrt_sq_eq_abs]
1048   norm_num
1049   constructor
1050   · exact h1
1051   constructor
1052   · norm_num
1053   · norm_num
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

```

The proof auto-formalization generated by PROOFBRIDGE is as follows:

Output Lean theorem+proof pair by PROOFBRIDGE (Retrieval-augmented SFT + Repair)

```

1048 import Mathlib
1049
1050 /- A point (x,y) on the coordinate plane with both coordinates negative is a distance of
1051   6 units from the x-axis. It is a distance of 15 units from the point (8,3). It is a
1052   distance  $\sqrt{n}$  from the origin. What is n? Show that it is 52.-/
1053
1054 theorem algebra_53819 (x y : ℝ) (hx : x < 0) (hy : y < 0) (h1 : abs y = 6) (h2 :
1055   Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15) :
1056   Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt 52 := by
1057   have hy6 : y = -6 := by
1058   cases eq_or_eq_neg_of_abs_eq h1 with
1059   | inl h =>
1060     linarith
1061   | inr h =>
1062     linarith
1063   have h3 : (x - 8) ^ 2 + (y - 3) ^ 2 = 225 := by
1064   calc
1065     (x - 8) ^ 2 + (y - 3) ^ 2 = (Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2)) ^ 2 := by
1066     rw [Real.sq_sqrt]
1067     positivity
1068     _ = 15 ^ 2 := by
1069     rw [h2]
1070     _ = 225 := by
1071     norm_num
1072   rw [hy6] at h3
1073   have h4 : (x - 8) ^ 2 = 144 := by
1074   nlinarith
1075   have h5 : x - 8 = 12 ∨ x - 8 = -12 := by
1076   have h6 : (x - 8) ^ 2 - 144 = 0 := by linarith
1077   have h7 : (x - 8 - 12) * (x - 8 + 12) = 0 := by
1078   nlinarith
1079   cases (mul_eq_zero.mp h7) with
1080   | inl h8 =>
1081     left
1082     linarith
1083   | inr h9 =>
1084     right
1085     linarith
1086   cases h5 with
1087   | inl h6 =>
1088     have hx20 : x = 20 := by linarith
1089     linarith
1090   | inr h6 =>
1091     have hx4 : x = -4 := by linarith
1092

```

```

1080
1081     have h7 : x ^ 2 + y ^ 2 = 52 := by
1082         rw [hx4, hy6]
1083         norm_num
1084     have h8 : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt 52 := by
1085         rw [h7]
1086         exact h8

```

The Lean theorem generated by PROOFBRIDGE and the gold-standard theorem are semantically equivalent. Both correctly capture the intended geometric scenario: the point with negative coordinates satisfying the given distance constraints is at a distance of $\sqrt{52}$ from the origin. Using our set of restricted tactics, the LLM judge is able to generate a type-checkable proof establishing the bi-directional equivalence between the two theorems. Therefore, the PROOFBRIDGE output is semantically correct (SC).

Comparison between the gold-standard theorem and PROOFBRIDGE's output

```

/- Gold-standard theorem -/
theorem mathd_algebra_288 (x y : ℝ) (n : NNReal) (h0 : x < 0 ∧ y < 0) (h1 : abs y = 6)
  (h2 : Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15)
  (h3 : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt n) : n = 52 := by
  sorry

/- ProofBridge output theorem -/
theorem algebra_53819 (x y : ℝ) (hx : x < 0) (hy : y < 0) (h1 : abs y = 6) (h2 :
  Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15) :
  Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt 52 := by
  sorry

```

Lean proof of bi-directional equivalence between PROOFBRIDGE's output theorem and the gold-standard

```

import Mathlib.Analysis.SpecialFunctions.Sqrt
import Mathlib.Data.NNReal.Basic
open BigOperators Real Nat Topology Rat

-- We define the two statements as propositions to prove their equivalence.

def mathd_algebra_288_stmt :=
  ∀ (x y : ℝ) (n : NNReal) (h0 : x < 0 ∧ y < 0) (h1 : abs y = 6)
  (h2 : Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15)
  (h3 : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt n), n = 52

def algebra_53819_stmt :=
  ∀ (x y : ℝ) (hx : x < 0) (hy : y < 0) (h1 : abs y = 6)
  (h2 : Real.sqrt ((x - 8) ^ 2 + (y - 3) ^ 2) = 15),
  Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt 52

-- Proof of bidirectional equivalence
theorem equivalence : mathd_algebra_288_stmt ↔ algebra_53819_stmt := by
  constructor
  -- Proof that mathd_algebra_288_stmt implies algebra_53819_stmt
  · intro h_mathd_stmt
    intro x y hx hy h1 h2
    -- To use h_mathd_stmt, we need to provide a value for 'n' and a proof for 'h3'.
    -- Let's define 'n' as the value of 'x^2 + y^2'.
    -- First, we need to show 'x^2 + y^2' is non-negative to create an 'NNReal'.
    have h_nonneg : 0 ≤ x ^ 2 + y ^ 2 := by positivity
    let n : NNReal := (x ^ 2 + y ^ 2, h_nonneg)
    -- The hypothesis 'h3' now holds by definition of 'n'.
    have h3 : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt n := by
      simp [n, NNReal.coe_mk]
    -- The hypothesis 'h0' is equivalent to 'hx' and 'hy'.
    have h0 : x < 0 ∧ y < 0 := {hx, hy}
    -- Now we can apply the first statement.
    have n_eq_52 : n = 52 := h_mathd_stmt x y n h0 h1 h2 h3
    -- Our goal is 'Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt 52'.
    -- We can rewrite 'h3' using 'n_eq_52' to achieve the goal.
    rwa [n_eq_52] at h3

    -- Proof that algebra_53819_stmt implies mathd_algebra_288_stmt
    · intro h_algebra_stmt
      intro x y n h0 h1 h2 h3
      -- The hypotheses of the second statement are all available.

```

```

1134
1135 have h_sqrt_val : Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt 52 :=
1136   h_algebra_stmt x y h0.left h0.right h1 h2
1137   -- We are given 'h3': 'Real.sqrt (x ^ 2 + y ^ 2) = Real.sqrt n'.
1138   -- By transitivity, 'Real.sqrt n = Real.sqrt 52'.
1139 have sqrt_n_eq_sqrt_52 : Real.sqrt n = Real.sqrt 52 := by
1140   rw [← h3, h_sqrt_val]
1141   -- Since 'Real.sqrt' is injective on non-negative numbers, 'n' must equal '52'.
1142   -- We get the equality on 'ℝ' first.
1143 have n_val_eq_52 : (n : ℝ) = 52 :=
1144   (Real.sqrt_inj n.property (by norm_num)).mp sqrt_n_eq_sqrt_52
1145   -- Then we lift this equality to 'NNReal'.
1146 exact NNReal.eq n_val_eq_52
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

```