

A Semantic Search Engine for Mathlib4

Guoxiong Gao^{1*}, Haocheng Ju^{2*}, Jiedong Jiang², Zihan Qin¹, Bin Dong^{2,3†}

¹School of Mathematical Sciences, Peking University,

²Beijing International Center for Mathematical Research, Peking University,

³Center for Machine Learning Research, Peking University

samggx@stu.pku.edu.cn hcju@pku.edu.cn

{emailboxofjjd,15261375660}@163.com dongbin@math.pku.edu.cn

Abstract

The interactive theorem prover Lean enables the verification of formal mathematical proofs and is backed by an expanding community. Central to this ecosystem is its mathematical library, mathlib4, which lays the groundwork for the formalization of an expanding range of mathematical theories. However, searching for theorems in mathlib4 can be challenging. To successfully search in mathlib4, users often need to be familiar with its naming conventions or documentation strings. Therefore, creating a semantic search engine that can be used easily by individuals with varying familiarity with mathlib4 is very important. In this paper, we present a semantic search engine¹ for mathlib4 that accepts informal queries and finds the relevant theorems. We also establish a benchmark for assessing the performance of various search engines for mathlib4.²

1 Introduction

Lean (de Moura et al., 2015; de Moura and Ullrich, 2021) is an interactive theorem prover built on dependent type theory, designed to verify mathematical proofs written in a formal language and thereby enhancing their rigor. It has a vibrant and supportive community, with its popularity growing among mathematicians. A prime example of the Lean community's collaborative spirit is mathlib4, the mathematical library for Lean 4. This library, regularly updated by contributors from around the globe, acts as a basis for the formalization of new mathematical theories. This eliminates the need to repeatedly formalize established results, as users can simply check if mathlib4 contains the necessary theorems referenced in informal proofs. However, locating these theorems is often challenging due to

the limitations of officially provided search tools, which struggle to find relevant theorems with informal queries.

There are primarily two methods to search for theorems in mathlib4: using the mathlib4 documentation³ and searching the GitHub repository of mathlib4⁴. The mathlib4 documentation allows users to search for theorems by their formal names. However, this feature can be challenging for beginners to utilize effectively, as they may not be familiar with the naming conventions. For example, Cauchy's Mean Value Theorem is named as `exists_ratio_deriv_eq_ratio_slope` in mathlib4, meaning a direct search for "Cauchy's Mean Value Theorem" yields no results. An alternative method involves searching within the GitHub repository of mathlib4, which allows for keyword-based searches across the source files, including formal statements, proofs, and documentation strings. This method can locate Cauchy's Mean Value Theorem as it is mentioned in the documentation string of `exists_ratio_deriv_eq_ratio_slope`. However, this approach faces two issues: many theorems in mathlib4 lack documentation strings, and semantically similar user queries that don't exactly match the theorems or documentation strings may lead to unsuccessful searches.

Consequently, neither method adequately supports finding theorems based on informal queries, leading to beginners spending significant time on this task. Discussions on Zulip⁵ have highlighted the need for creating a semantic search engine for mathlib4. Therefore, developing such a search engine for mathlib4 is highly desirable to improve the efficiency of theorem retrieval.

In this paper, we introduce a semantic search

*These authors contributed equally.

†Corresponding author

¹The search engine is available at <https://leansearch.net>.

²The data and the source code are available at <https://github.com/reaslab/LeanSearch>.

³https://leanprover-community.github.io/mathlib4_docs/

⁴<https://github.com/leanprover-community/mathlib4>

⁵<https://leanprover.zulipchat.com/#narrow/stream/219941-Machine-Learning-for-Theorem-Proving/topic/Semantic.20Search.20for.20Mathematics>

engine for mathlib4 that allows users to input an informal query and retrieve a list of relevant theorems from mathlib4. To construct this search engine, we translate the formal statements of mathlib4 theorems into informal ones and integrate these pairs into a database. Each database entry consists of a formal theorem statement and its informal version. Upon receiving a user query, we augment the query for improved context understanding and perform semantic search across the database to find relevant results. Additionally, we have established a benchmark to compare the effectiveness of various search engines for mathlib4.

The remaining part of the paper is organized as follows. In Section 2, we review the related works on text retrieval and mathematical information retrieval. Section 3 describes our approach to developing a semantic search engine for mathlib4. Section 4 presents the mathlib4 semantic search benchmark. Numerical results are discussed in Section 5. We conclude this paper in Section 6.

2 Related Work

Text Retrieval. Text retrieval is the task of finding relevant information within a corpus based on user queries. Early methods, such as BM25 (Robertson et al., 1995; Robertson and Zaragoza, 2009), used sparse vector representations for queries and documents, assessing relevance by comparing these vectors with certain weighting techniques (Salton and Buckley, 1988). While effective in measuring lexical similarity, these approaches fall short in capturing the semantic similarity between queries and documents.

To address this limitation, deep learning techniques (Huang et al., 2013; Xiong et al., 2017; McDonald et al., 2018; Liu et al., 2018) have been introduced. Utilizing deep neural networks, these techniques encode queries and documents into dense vectors, thereby assessing relevance based on semantic similarity. Further developments have been made in neural architectures and training paradigms. There are two main architectures: the cross-encoder (Qiao et al., 2019; Nogueira et al., 2020) and the bi-encoder (Karpukhin et al., 2020; Qu et al., 2021; Ni et al., 2022b,a). Cross-encoders take the concatenation of query and document as input and produce the final relevance of this query-document pair, while bi-encoders map the query and document into vectors separately, determining relevance through similarity between the two vec-

tors. Training paradigms have also evolved, with the Inverse Close Task (ICT) (Lee et al., 2019) initially proposed for dense retriever pre-training. Subsequently, other pre-training tasks have been developed, including Body First Selection and Wiki Link Prediction, both introduced in (Chang et al., 2020). Recent studies (Neelakantan et al., 2022; Wang et al., 2022; Su et al., 2023; Xiao et al., 2023) have explored large-scale unsupervised pre-training using contrastive loss, followed by fine-tuning on smaller, labeled datasets. Pre-training with extensive text pairs allows the language model to grasp textual semantics and the fine-tuning stage further enhances its performance across various retrieval tasks. However, the authors of (Wang et al., 2024) argue that a two-stage training approach might not be necessary, demonstrating that directly fine-tuning a decoder-only model on both synthetic and labeled data can yield competitive results.

Mathematical Information Retrieval. Mathematical Information Retrieval (MIR) differs from text retrieval in that it involves queries and documents that contain mathematical formulas. These formulas are highly structured, which distinguishes them from plain text. To effectively capture the semantics of math formulas, several representations are employed. The Symbol Layout Tree (SLT) (Zanibbi and Blostein, 2012) preserves the original layout of formulas, while the Operator Tree (OPT) (Gao et al., 2016) represents mathematical symbols as nodes, with edges denoting the relationships between operators and operands. Classical MIR methods (Fraser et al., 2018; Gao et al., 2016; Kristianto et al., 2016; Zanibbi et al., 2016; Kane et al., 2022; Ng et al., 2021, 2020) rely on structure search, identifying matching substructures across various features. Among these, the Approach0 structural search method (Zhong and Zanibbi, 2019; Zhong et al., 2020) has shown to be particularly effective. It indexes formulas through leaf-root paths in the OPT and utilizes subexpression matching to assess formula similarity. With advancements in deep learning, dense retrievers have been integrated into MIR, often combined with structural searches (Kane et al., 2022; Zhong et al., 2022a,b, 2023). A notable example is the Approach0 hybrid search (Zhong et al., 2022a), which combines Approach0 structure search with a bi-encoder dense retriever, CoBERT (Khattab and Zaharia, 2020). This combination not only facilitates effective for-

mula matching but also enhances understanding of context.

The preceding literature focuses on extracting mathematical content from a corpus composed of natural language texts and formulas formatted in markup languages. This contrasts with our objective of conducting searches within mathlib4, a formal mathematical library. The work most closely related to ours is Moogle⁶, a semantic search engine for mathlib4. However, its technical details have not been released. We will compare the performance of Moogle and our search engine in Section 5.

3 Methodology

In this section, we will describe the implementation of our semantic search engine for mathlib4⁷. This engine is designed to accept a user query in natural language and return a list of relevant theorems in mathlib4. Our approach involves converting formal theorems from mathlib4 into their informal counterparts, as illustrated in Figure 1. These informal-formal theorem pairs are then vectorized and stored in a vector database, a step that can be executed offline. When a user query is submitted, we augment the query to better grasp its context, vectorize the enriched query, and locate the corresponding theorems in the embedding space. In the following subsections, we will elaborate on the informalization of mathlib4, the design of the search engine, and the method of query augmentation.

3.1 Informalizing Mathlib4

Our strategy for informalizing mathlib4 involves employing a large language model (LLM). Central to this strategy is providing the LLM with sufficient context to accurately grasp the formal theorem’s exact meaning. To this end, we not only extract the theorem’s name, statement, and documentation string⁸ from the mathlib4 documentation but also include related definitions found in the theorem statements through hyperlinks. For example, as illustrated in Figure 2, we extract the definition of `Exists.choose` because it is referenced and linked in the theorem `Exists.choose_spec`. The gathered information is then fed into an LLM⁹ to gener-

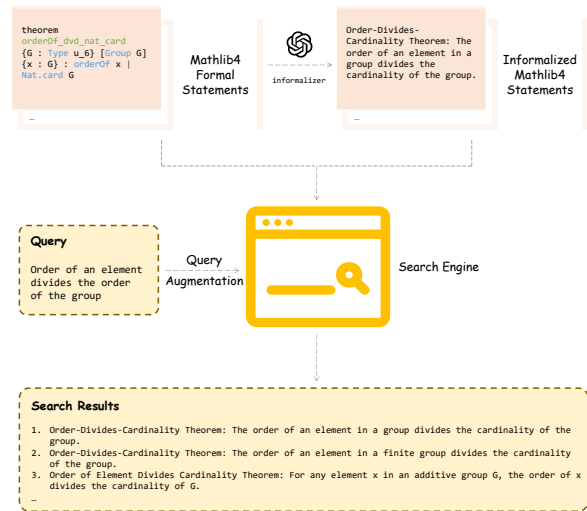


Figure 1: Overview of our method for creating a semantic search engine for mathlib4. We employ an informalizer to convert formal statements from mathlib4 into their informal counterparts. These informal-formal pairs are then stored in a vector database. When users input a query, the system augments the query and searches across the database, yielding a list of relevant theorems.

ate informal theorem names and statements. These are subsequently formatted as "theorem name: informal statement" and incorporated into an informal corpus. Figure 3 shows the prompt used for informalizing the theorem `Exists.choose_spec`, including the formal statement, documentation string, and the definition of `Exists.choose` as annotations to aid the LLM’s understanding.

Notably, the authors of (Jiang et al., 2023) also employed an LLM for informalizing mathlib4 statements, relying solely on the formal statements. We argue that providing the LLM with additional context, such as related definitions and documentation strings, enhances its ability to accurately interpret and convert formal theorems into their informal counterparts. An example comparing the informalized mathlib4 statement generated by (Jiang et al., 2023) to the one produced by our method is provided in Appendix A.

3.2 Semantic Search Engine for Mathlib4

After obtaining the informal corpus, we employ dense embedding models, which excel at capturing semantic information, to encode the informal-formal theorem pairs. Recent advancements in text embedding models have introduced the practice of integrating specific task instructions into either queries or documents, enhancing the model’s adaptability to diverse tasks and boosting performance in

⁶<https://www.moogle.ai/>

⁷We use mathlib4 with commit db04a978b67b2200691f0bd968f334c83261b66a.

⁸For many theorems in mathlib4 that lack a documentation string, we will extract only the theorem name and statement.

⁹We use gpt-3.5-turbo-16k for generating informal names and statements, setting the temperature to 0.

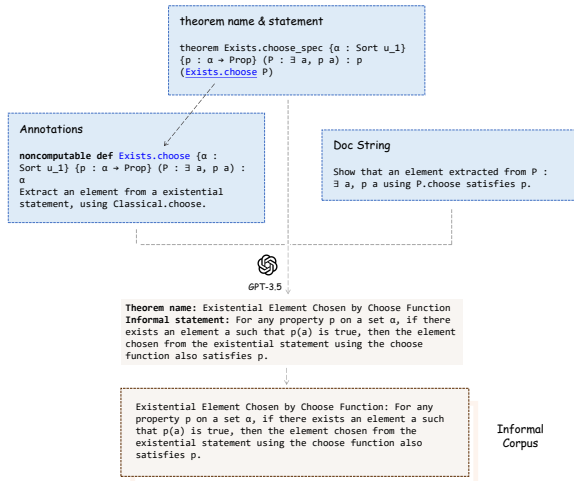


Figure 2: Our approach to informalizing mathlib4 theorems. We extract the theorem name, statement, and documentation string from the mathlib4 documentation. Moreover, we collect related definitions via the hyperlinks in the theorem statements. The gathered information is then inputted into GPT-3.5 to generate informal names and statements. These are then organized in the format "theorem name: informal statement" and added to an informal corpus.

zero-shot settings (Su et al., 2023; Li and Li, 2023; Wang et al., 2024). Consequently, for the purpose of theorem retrieval within mathlib4, we enrich our corpus documents with specific instructions. We adopt the following instruction template:

```
Instruct: Retrieve math theorems stated
in bilingual Lean 4 + natural language
that are mathematically equivalent to
the given one \n Doc:{document}
```

Here, {document} denotes

```
"{formal statement} \n {informal
name}:{informal statement}"
```

These task-specific instructions significantly influence the performance of embedding models. We will investigate the effect of varying task instructions on the overall retrieval performance in Appendix D.

Embedding the entire corpus, although a time-consuming process, is performed offline and does not require repetition for each use of the engine. We utilize Chroma DB to store the embeddings. Upon receiving a query, our system vectorizes it and retrieves theorems based on their cosine distance to the query in the embedding space. Chroma DB employs the Hierarchical Navigable Small World (HNSW) algorithm (Malkov and Yashunin, 2020),

Input:

System message: As a mathematician and expert in Lean and Mathlib, your task is to translate the formal theorem provided below into an informal statement that is more accessible to mathematicians. Please utilize the provided doc string and annotations to better understand the formal theorem. Additionally, please express the mathematical formulas in LaTeX when necessary.

Formal theorem:
`theorem Exists.choose_spec {α : Sort u_1} {p : α → Prop} (P : ∃ a, p a) : p (Exists.choose P)`

Doc string:
 Show that an element extracted from P : ∃ a, p a using P.choose satisfies p.

Annotations:
`noncomputable def Exists.choose {α : Sort u_1} {p : α → Prop} (P : ∃ a, p a) : α`
 Extract an element from an existential statement, using Classical.choose.

Output:

Theorem name: Existential Element Chosen by Choose Function
Informal statement: For any property p on a set α, if there exists an element a such that p(a) is true, then the element chosen from the existential statement using the choose function also satisfies p.

Figure 3: Prompt for informalizing mathlib4 statements with documentation strings.

an efficient approximate nearest neighbor search method, ensuring rapid retrieval from the corpus.

3.3 Query Augmentation

To enhance query clarity and achieve a more accurate embedding, our system incorporates a query augmentation step. This process involves prompting an LLM¹⁰ to transform a brief, vague query into a detailed statement that includes both informal and formal statements, ensuring mathematical equivalence with the original query. We guide the LLM with specific principles for query augmentation, emphasizing the importance of precision, the use of LaTeX for mathematical expressions, and the clarification of ambiguous inputs. Additionally, we provide examples of query augmentations to improve the LLM’s comprehension of the task. A detailed prompt for query augmentation is provided in Appendix B. Although Lean 4 code generation by LLMs may occasionally introduce inaccuracies due to the limited presence of mathlib4 in their training data, this approach effectively enriches the query with additional contextual information.

Following augmentation, the enriched query is structured as "{formal statement} \n {informal name}:{informal statement}", matching the structure of our database. In a similar manner to adding task instructions in document processing, we enrich the query with specific instructions, utilizing a template as follows:

```
Instruct: Retrieve math theorems stated
in bilingual Lean 4 + natural language
that are mathematically equivalent to
the given one \n Query:{formal statement}
\n {informal name}:{informal statement}"
```

This formatted query is then vectorized and utilized

¹⁰We use gpt-4-0125-preview for query augmentation.

by the search engine to retrieve theorems based on the cosine distance in the embedding space.

4 Mathlib4 Semantic Search Benchmark

To rigorously assess and compare the efficacy of various retrieval methods, we have established the Mathlib4 Semantic Search Benchmark. This benchmark encompasses a curated set of queries, relevance labels of each mathlib4 theorem for these queries, and a collection of performance metrics. In this section, we provide a detailed explanation of the benchmark. Section 4.1 describes the composition of the query set. In Section 4.2, we explain the relevance criteria for our benchmark and the labeling procedure. Section 4.3 lists the performance metrics used in our benchmark.

4.1 Composition of the Query Set

The query set of our benchmark has 50 distinct queries, spanning various mathematical disciplines including calculus, abstract algebra, linear algebra, number theory, algebraic number theory, set theory, and mathematical logic. This selection aims to cover a broad spectrum of topics and complexities. We argue that this is a typical size for the query sets in mathematical information retrieval datasets, as demonstrated by the ARQMath1, ARQMath2, and ARQMath3 databases from the MIR field, which contain 77, 71, and 78 queries, respectively, for their answer retrieval tasks (Mansouri et al., 2022).

To optimize labeling efforts, we organize queries with identical search intents into 18 distinct, non-overlapping groups, each containing at least two distinct queries. This approach assumes that all queries within a group have the same relevance score for any document, significantly reducing the need for repetitive labeling, as each document is evaluated only once per query group. Furthermore, to provide a more detailed assessment and mitigate the impact of duplicate document labels, we consider four prevalent forms of mathematical queries: natural language descriptions, LaTeX formulas, theorem names, and Lean 4 term descriptions. In each query group, we strive to include as many different description forms as possible. Table 1 provides statistics and examples of the query set. However, not all query groups contain all four representation forms, as some theorems may not have an official name or their representation in a LaTeX formula might be redundant.

Category	Count	Example 1	Example 2
Natural Description	18	If there exist injective maps of sets from A to B and from B to A , then there exists a bijective map between A and B .	If p implies q , then not q implies not p .
LaTeX Formula	15	If there exist $f : A \rightarrow B$ injective, $g : B \rightarrow A$ injective, then there exists $h : A \rightarrow B$ bijective.	$(p \rightarrow q) \rightarrow (\neg q \rightarrow \neg p)$
Theorem Name	7	Schroeder Bernstein Theorem	Modus Tollens
Lean 4 Term	10	$\{f : A \rightarrow B\} \{g : B \rightarrow A\} (hf : \text{Injective } f) (hg : \text{Injective } g) : \exists h, \text{Bijective } h$	$(p \rightarrow q) \rightarrow (\neg q \rightarrow \neg p)$

Table 1: Statistics and examples of query groups in our benchmark.

4.2 Relevance Judgments

During the process of collecting relevance labels for query-document pairs, we adopt an approach similar to ARQMath (Mansouri et al., 2022). Initially, we establish a carefully crafted relevance assessment criteria, elaborated in Table 2. While performing the labeling, assessors are instructed to evaluate whether a given search result facilitates their mathematical formalization workflow. Instead of merely considering similarity in topic, presented form, formula structure, or mathematical induction relationship, the relevance we focus on here is deeply aligned with the needs of Lean 4 experts.

Rating	Label	Score	Definition
Exact Match	2	1	Exact match to the query or being a stronger statement
Relevant	1	0.3	Useful in locating where the corresponding statement should be
Irrelevant	0	0	Not expected to be useful in mathematics formalization workflow

Table 2: The relevance assessment criteria for our benchmark.

For each set of queries grouped by identical search intentions, assessors are presented with the top 50 theorems retrieved by an intermediate version of our search engine. In addition to evaluating the provided theorems, assessors are tasked with identifying and adding any relevant theorems that may have been omitted from the initial list by inspecting the files where "Exact Match" items in the list are located, with particular attention to those with "Exact Match" rating.

Given the structured organization of mathlib4, where related theorems are often located within the same file, it is reasonable to assume that any items not in the list are irrelevant to this query group. Additionally, all queries are guaranteed to have at least one exact match in mathlib4. These assumptions support the performance metrics we use in the benchmark, as described in the following subsection.

4.3 Performance Measures

To compare different retrieval paradigms' performance on our labeled dataset, three commonly used metrics are adopted in our benchmark. Precision@k calculates the average relevant document proportion in top-k retrieved results among the query set \mathcal{Q} :

$$\text{Precision@k} = \frac{1}{k|\mathcal{Q}|} \sum_{i=1}^{|\mathcal{Q}|} \sum_{j=1}^k \mathbb{I}(i, d_i^j), \quad (1)$$

where d_i represents the list of retrieved theorems for the i -th query, and $\mathbb{I}(i, d_i^j) = 1$ if and only if the j -th retrieved result of i -th query is "Exact Match", otherwise $\mathbb{I}(i, d_i^j) = 0$. Assuming that all unlabeled theorems are irrelevant, we can also calculate Recall@k:

$$\text{Recall@k} = \frac{1}{|\mathcal{Q}|} \sum_{i=1}^{|\mathcal{Q}|} \frac{1}{\sigma_i} \sum_{j=1}^k \mathbb{I}(i, d_i^j), \quad (2)$$

where σ_i is the number of "Exact Match" theorems for i -th query.

The third metric, nDCG (normalized Discounted Cumulative Gain), incorporates the position of the retrieved result. A decaying weight is allocated to each position:

$$\text{DCG}_i@k = \sum_{j=1}^k \frac{s(i, d_i^j)}{\log_2(j+1)}, \quad (3)$$

where $s(i, d_i^j)$ represents the score of the j -th retrieved theorem for the i -th query, as defined in Table 2. Our scoring system is not proportional to the label number, emphasizing the importance of exact matching in the mathlib4 retrieval task. We further define IDCG _{i} @k as the highest possible DCG _{i} @k by properly arranging the retrieved theorems, and finally:

$$\text{nDCG@k} = \frac{1}{|\mathcal{Q}|} \sum_{i=1}^{|\mathcal{Q}|} \frac{\text{DCG}_i@k}{\text{IDCG}_i@k} \quad (4)$$

To sum up, nDCG@k measures both the retrieval and ranking ability of the given engine in a unified and detailed way. Meanwhile, P@10 (Precision@10) and R@10 (Recall@10) focus solely on the retrieval effectiveness. In our benchmark, all three metrics will be reported for both the entire query set and each query category.

5 Experiments

In this section, we evaluate the performance of various theorem retrieval methods using our benchmark. The experimental setup is detailed in Section 5.1. We compare the performances of different retrieval methods in Section 5.2. Additionally, the ablation study on document preparations and query augmentations is presented in Section 5.3.

5.1 Experiment Setup

We have considered four embedding models in our experiments: text-embedding-ada-002 and text-embedding-3-large from OpenAI, UAE-Large-V1 (Li and Li, 2023), and E5_{mistral-7b} (Wang et al., 2024). For the baseline methods, we compare BM25 (Robertson and Zaragoza, 2009), MoogLe, and the four embedding models applied to the original Lean 4 formal corpus and unaugmented query. The same four embedding models equipped with formal + informal query augmentation and formal + informal document corpus are also tested on our benchmark to demonstrate the efficacy of our approach. The implementation details are provided in Appendix C.

5.2 Main Results

Table 3 and 4 present the results of BM25, MoogLe, and four embedding models. The model E5_{mistral-7b}, when integrated with our retrieval pipeline, achieves the best performance across three overall metrics and significantly outperforms all other methods, including its own performance on a formal corpus with unaugmented queries. We observe that all retrieval methods benefit from using an augmented corpus and queries, as augmentation expands concrete mathematics and Lean 4 terms, and these embedding models typically perform better on symmetrical retrieval tasks (Wang et al., 2024).

Examining the results averaged by category, we find that E5_{mistral-7b} attains the best results in nearly all categories, with the exception of nDCG@20 for Lean 4 Terms. This exception is likely due to the augmented informal information negatively impacting retrieval performance, as it is not essential in formal-formal retrieval (formal queries, formal results) which focuses on lexical matches. Notably, our method significantly improves performance in the Theorem Name category. The enhancement is due to the expansion of the statement for a given theorem name by our query augmentation, which

is essential for successful retrieval in this category. Furthermore, the combined corpus of formal and informal statements facilitates the retrieval process.

Model	Corpus	Query aug.	nDCG@20	P@10	R@10
Baselines					
Mooglet [†]	\	\	0.365 [†]	0.092 [†]	0.513 [†]
BM25	F	\	0.024	0.004	0.030
OpenAI v2	F	\	0.312	0.078	0.405
OpenAI v3	F	\	0.493	0.128	0.622
UAE-Large-V1	F	\	0.233	0.066	0.307
E5 _{mistral-7b}	F	\	0.593	0.132	0.687
Our Methods					
OpenAI v2	F+IF	F+IF	0.553	0.144	0.707
OpenAI v3	F+IF	F+IF	0.691	0.178	0.837
UAE-Large-V1	F+IF	F+IF	0.368	0.084	0.440
E5 _{mistral-7b}	F+IF	F+IF	0.733	0.196	0.913

Table 3: Results on our benchmark, averaged across all queries in our dataset. Here "aug.", F, IF, P@10 and R@10 represent augmentation, formal, informal, Precision@10, and Recall@10, respectively. The terms OpenAI v2 and v3 refer to the text-embedding models ada-002 and 3-large, respectively. Mooglet[†], unlike other retrieval systems, not only fetches theorems but also definitions, structures, instances, etc., making it incomparable under our performance metrics directly. For the purpose of this analysis, all non-theorem items retrieved are considered irrelevant, given the explicit theorem-searching intent of our queries. This approach, however, might advantage our theorem retrieval systems over Mooglet, as non-theorem items occupy potential slots in the retrieval list. The notation [†] is used to denote this adjustment.

Model	nDCG@20				P@10			
	ND	LF	TN	LT	ND	LF	TN	LT
Baselines								
Mooglet [†]	0.369 [†]	0.324 [†]	0.333 [†]	0.441 [†]	0.083 [†]	0.107 [†]	0.071 [†]	0.100 [†]
BM25	0.000	0.000	0.000	0.119	0.000	0.000	0.000	0.020
OpenAI v2	0.276	0.379	0.000	0.498	0.061	0.107	0.000	0.120
OpenAI v3	0.479	0.553	0.235	0.610	0.122	0.160	0.057	0.140
UAE-Large-V1	0.301	0.216	0.004	0.298	0.078	0.067	0.000	0.090
E5 _{mistral-7b}	0.576	0.633	0.294	0.774	0.139	0.140	0.043	0.170
Our Methods								
OpenAI v2	0.536	0.533	0.571	0.600	0.111	0.160	0.186	0.150
OpenAI v3	0.681	0.657	0.772	0.703	0.167	0.180	0.200	0.180
UAE-Large-V1	0.415	0.337	0.371	0.329	0.100	0.080	0.071	0.070
E5 _{mistral-7b}	0.748	0.712	0.855	0.654	0.194	0.200	0.214	0.180

Table 4: Results on our benchmark averaged by category. Here ND, LF, TN and LT stands for Natural Description, LaTeX Formula, Theorem Name and Lean 4 Term respectively. Corpus type and query augmentation type remain the same as Table 3 and are omitted here due to space limitations. [†] indicates a different performance calculation rule, detailed in captions of Table 3.

5.3 Ablation Studies

In this subsection, we perform ablation studies on document preparation and query augmentation. The results of these studies are presented in Sections 5.3.1 and 5.3.2, respectively. Additionally,




Figure 4: nDCG@20 performance of E5_{mistral-7b} across formal, informal and formal + informal corpus on our benchmark, averaged by query groups. These evaluations are conducted using non-augmented queries.

the ablation study on task instructions is provided in Appendix D.

5.3.1 Ablation of Document Preparations

Table 5 presents the effects of changing corpus type with and without query augmentation. We use the original Lean 4 statement and its informalized version to create the formal and informal corpus. Without query augmentation, results indicate that both OpenAI v3 large and E5_{mistral-7b} underperform with incomplete corpus components, yielding lower overall scores on all metrics. Specifically, the formal + informal corpus enhances performance on Lean 4 Terms by incorporating formal data and improves score on the Natural Description category by providing hybrid information. For the Theorem Name category, E5_{mistral-7b} significantly benefits from the formal + informal corpus by utilizing the unfolded mathematical descriptions provided by the informalized statement, whereas OpenAI v3 large shows diminished results, likely due to the lack of training on mixed-domain texts and absence of adaptive instructions. We note that the formal corpus outperforms the formal+informal corpus in the Lean 4 Terms category for E5_{mistral-7b}, as explained in Section 5.2: retrieval on Lean 4 Terms relies heavily on lexical rather than semantic information, and the presence of informal information negatively impacts retrieval performance. Similar results are observed with augmented queries.

A more detailed examination of the results for each query group without query augmentation is shown in Figure 4. In most instances, the "formal + informal" configuration avoids the abrupt performance decline observed in solely formal or informal settings. The informalized statement provided in the corpus mitigates the lexical mismatch between the query and the formal statement, which justifies our choice of a combined corpus.

Model & Corpus	nDCG@20					P@10	R@10
	ND	LF	TN	LT	All	All	All
without query augmentation							
OpenAI v3 large							
formal corpus	0.479	0.553	0.235	0.610	0.493	0.128	0.622
informal corpus	0.591	0.527	0.365	0.451	0.512	0.134	0.628
formal+informal corpus	0.607	0.549	0.267	0.622	0.545	0.140	0.677
E5_{mistral-7b}							
formal corpus	0.576	0.633	0.294	0.774	0.593	0.132	0.687
informal corpus	0.696	0.636	0.616	0.604	0.648	0.174	0.773
formal+informal corpus	0.749	0.687	0.698	0.612	0.701	0.184	0.845
with query augmentation							
OpenAI v3 large							
formal corpus	0.637	0.617	0.565	0.647	0.623	0.160	0.747
informal corpus	0.603	0.650	0.810	0.575	0.640	0.166	0.783
formal+informal corpus	0.681	0.657	0.772	0.703	0.691	0.178	0.837
E5_{mistral-7b}							
formal corpus	0.688	0.685	0.631	0.694	0.680	0.178	0.847
informal corpus	0.711	0.646	0.824	0.670	0.699	0.184	0.877
formal+informal corpus	0.748	0.712	0.855	0.654	0.733	0.196	0.913

Table 5: Results of ablation studies on document preparation and query augmentation. "All" indicates the performance averaged across the entire query set. In the query augmentation section, all queries are augmented to match the corpus type. Here, E5_{mistral-7b} uses different task instructions for the document and query in the non-augmented query setting and the same task instructions for both in the query augmentation setting. Detailed task instructions can be found in Table 6.

Query aug. & Doc Type	Side	Task Instructions
None & Informal	Query	Given a math search query, retrieve theorems mathematically equivalent to the query
	Doc	Represent the given math theorem statement for retrieving related statement by natural language query
None & F+IF	Query	Given a math search query, retrieve theorems stated in bilingual Lean 4 + natural language that mathematically match the query
	Doc	Represent the given formal math statement written in Lean 4 concatenated with its natural language explanation for retrieving related statement by natural language query
Formal & Formal	Query & Doc	Retrieve math theorems stated in Lean 4 that are mathematically equivalent to the given one
IF & IF	Query & Doc	Retrieve math theorems that are mathematically equivalent to the given one

Table 6: Task Instructions for E5_{mistral-7b} in ablation studies, supplementing Table 7. The abbreviation "aug." denotes augmentation, while "F" and "IF" represent formal and informal, respectively.

5.3.2 Ablation of Query Augmentations

The effect of query augmentation is evident when comparing the results in the upper and lower sections of Table 5. As shown, augmented queries enhance the performance across all compared methods. Specifically, query augmentation significantly improves the ability of all methods to identify Theorem Names by expanding terms and providing richer context. Notably, the performance of E5_{mistral-7b} on the formal + informal corpus without query augmentation is comparable to its performance on the same corpus with query augmen-

tation. This makes it a cost-effective option for deploying our semantic search engine, as it eliminates the need to prompt an LLM with every query.

6 Conclusion

In this paper, we introduce a semantic search engine designed to enable users to locate theorems in mathlib4 using informal queries. Specifically, we translate the formal statements of mathlib4 theorems into informal versions and develop our search engine to work with a corpus of informal-formal theorem pairs. Additionally, we construct a dataset to facilitate evaluation. Our comprehensive experiments on this dataset reveal that the best theorem retrieval performance is attained by augmenting the user’s query appropriately and embedding the content of the corpus simultaneously. Consequently, our system employs a strategy that first augments the query, followed by a semantic search, thereby precisely aligning with the users’ search intentions.

Our future research will focus on three primary directions. First, in terms of informalizing mathlib4, we aim to design guidelines for translation and provide examples of converting mathlib4 statements to informal language for LLMs, enhancing the informal corpus’s quality. Second, regarding the mathlib4 semantic search benchmark, we plan to continually enlarge the query set, aiming for a more comprehensive benchmark. Lastly, for the semantic search engine itself, we intend to fine-tune a text embedding model on the task of theorem retrieval, aiming to improve the search engine’s performance.

7 Limitations

When translating mathlib4 theorems into their informal versions, we extract the related definitions through the hyperlinks in the theorem statements. This approach to extracting dependencies can sometimes be inaccurate and may omit certain dependencies. A more precise method is to extract the dependencies of a theorem by interacting with Lean’s language server.

Another limitation of our work is the use of E5_{mistral-7b}, a 7B embedding model, in our search engine. This model may pose challenges when deploying the search engine on resource-constrained devices like laptops. In future work, we aim to improve the performance of smaller embedding models for theorem retrieval tasks to enable broader accessibility.

References

- Wei-Cheng Chang, Felix X. Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2020. [Pre-training Tasks for Embedding-based Large-scale Retrieval](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Leonardo de Moura and Sebastian Ullrich. 2021. [The Lean 4 Theorem Prover and Programming Language](#). In *Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021, Proceedings*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer.
- Leonardo Mendonça de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. [The Lean Theorem Prover \(System Description\)](#). In *Automated Deduction - CADE-25 - 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings*, volume 9195 of *Lecture Notes in Computer Science*, pages 378–388. Springer.
- Dallas Fraser, Andrew Kane, and Frank Wm. Tompa. 2018. [Choosing Math Features for BM25 Ranking with Tangent-L](#). In *Proceedings of the ACM Symposium on Document Engineering 2018, DocEng '18, New York, NY, USA*. Association for Computing Machinery.
- Liangcai Gao, Ke Yuan, Yuehan Wang, Zhuoren Jiang, and Zhi Tang. 2016. [The Math Retrieval System of ICST for NTCIR-12 MathIR Task](#). In *Proceedings of the 12th NTCIR Conference on Evaluation of Information Access Technologies, National Center of Sciences, Tokyo, Japan, June 7-10, 2016*. National Institute of Informatics (NII).
- Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. [Learning deep structured semantic models for web search using clickthrough data](#). In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13*, page 2333–2338, New York, NY, USA. Association for Computing Machinery.
- Albert Q. Jiang, Wenda Li, and Mateja Jamnik. 2023. [Multilingual Mathematical Autoformalization](#). *CoRR*, abs/2311.03755.
- Andrew Kane, Yin Ki Ng, and Frank Wm. Tompa. 2022. [Dowsing for Answers to Math Questions: Doing Better with Less](#). In *Proceedings of the Working Notes of CLEF 2022 - Conference and Labs of the Evaluation Forum, Bologna, Italy, September 5th - 8th, 2022*, volume 3180 of *CEUR Workshop Proceedings*, pages 40–62. CEUR-WS.org.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. [ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, page 39–48, New York, NY, USA. Association for Computing Machinery.
- Giovanni Yoko Kristianto, Goran Topic, and Akiko Aizawa. 2016. [MCAT Math Retrieval System for NTCIR-12 MathIR Task](#). In *Proceedings of the 12th NTCIR Conference on Evaluation of Information Access Technologies, National Center of Sciences, Tokyo, Japan, June 7-10, 2016*. National Institute of Informatics (NII).
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent Retrieval for Weakly Supervised Open Domain Question Answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Xianming Li and Jing Li. 2023. [Angle-optimized Text Embeddings](#). *CoRR*, abs/2309.12871.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2018. [Entity-Duet Neural Ranking: Understanding the Role of Knowledge Graph Semantics in Neural Information Retrieval](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2395–2405, Melbourne, Australia. Association for Computational Linguistics.
- Yury A. Malkov and Dmitry A. Yashunin. 2020. [Efficient and Robust Approximate Nearest Neighbor Search Using Hierarchical Navigable Small World Graphs](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(4):824–836.
- Behrooz Mansouri, Vít Novotný, Anurag Agarwal, Douglas W. Oard, and Richard Žanibbi. 2022. [Overview of ARQMath-3 \(2022\): Third CLEF Lab on Answer Retrieval for Questions on Math \(Working Notes Version\)](#). In *Conference and Labs of the Evaluation Forum*.
- Ryan McDonald, George Brokos, and Ion Androutsopoulos. 2018. [Deep Relevance Ranking Using Enhanced Document-Query Interactions](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1849–1860, Brussels, Belgium. Association for Computational Linguistics.
- Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong Wook Kim, Chris Hallacy, Johannes Heidecke, Pranav Shyam, Boris Power, Tyna Eloundou Nékoul, Girish Sastry, Gretchen

- Krueger, David Schnurr, Felipe Petroski Such, Kenny Hsu, Madeleine Thompson, Tabarak Khan, Toki Sherbakov, Joanne Jang, Peter Welinder, and Lilian Weng. 2022. [Text and Code Embeddings by Contrastive Pre-Training](#). *CoRR*, abs/2201.10005.
- Yin Ki Ng, Dallas J. Fraser, Besat Kassaie, George Labahn, Mirette S. Marzouk, Frank Wm. Tompa, and Kevin Wang. 2020. [Dowsing for Math Answers with Tangent-L](#). In *Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum, Thessaloniki, Greece, September 22-25, 2020*, volume 2696 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Yin Ki Ng, Dallas J. Fraser, Besat Kassaie, and Frank Wm. Tompa. 2021. [Dowsing for Answers to Math Questions: Ongoing Viability of Traditional MathIR](#). In *Proceedings of the Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, Bucharest, Romania, September 21st - 24th, 2021*, volume 2936 of *CEUR Workshop Proceedings*, pages 63–81. CEUR-WS.org.
- Jianmo Ni, Gustavo Hernandez Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer, and Yinfei Yang. 2022a. [Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1864–1874, Dublin, Ireland. Association for Computational Linguistics.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, and Yinfei Yang. 2022b. [Large Dual Encoders Are Generalizable Retrievers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9844–9855, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. [Document Ranking with a Pre-trained Sequence-to-Sequence Model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718, Online. Association for Computational Linguistics.
- Yifan Qiao, Chenyan Xiong, Zhenghao Liu, and Zhiyuan Liu. 2019. [Understanding the Behaviors of BERT in Ranking](#). *CoRR*, abs/1904.07531.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. [RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5835–5847, Online. Association for Computational Linguistics.
- Stephen Robertson, S. Walker, S. Jones, M. M. Hancock-Beaulieu, and M. Gatford. 1995. [Okapi at TREC-3](#). In *Overview of the Third Text REtrieval Conference (TREC-3)*, pages 109–126. Gaithersburg, MD: NIST.
- Stephen Robertson and Hugo Zaragoza. 2009. [The Probabilistic Relevance Framework: BM25 and Beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Gerard Salton and Christopher Buckley. 1988. [Term-weighting approaches in automatic text retrieval](#). *Information Processing & Management*, 24(5):513–523.
- Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2023. [One Embedder, Any Task: Instruction-Finetuned Text Embeddings](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1102–1121, Toronto, Canada. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing Data using t-SNE](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. [Text Embeddings by Weakly-Supervised Contrastive Pre-training](#). *CoRR*, abs/2212.03533.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. [Improving Text Embeddings with Large Language Models](#). *CoRR*, abs/2401.00368.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighof. 2023. [C-Pack: Packaged Resources To Advance General Chinese Embedding](#). *CoRR*, abs/2309.07597.
- Chenyan Xiong, Zhuyun Dai, Jamie Callan, Zhiyuan Liu, and Russell Power. 2017. [End-to-End Neural Ad-hoc Ranking with Kernel Pooling](#). In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17*, page 55–64, New York, NY, USA. Association for Computing Machinery.
- Richard Zanibbi and Dorothea Blostein. 2012. [Recognition and retrieval of mathematical expressions](#). *Int. J. Document Anal. Recognit.*, 15(4):331–357.
- Richard Zanibbi, Kenny Davila, Andrew Kane, and Frank Wm. Tompa. 2016. [Multi-Stage Math Formula Search: Using Appearance-Based Similarity Metrics at Scale](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, page 145–154, New York, NY, USA. Association for Computing Machinery.
- Wei Zhong, Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2023. [One Blade for One Purpose: Advancing Math Information Retrieval using Hybrid Search](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 141–151, New York, NY, USA. Association for Computing Machinery.

Wei Zhong, Shaurya Rohatgi, Jian Wu, C. Lee Giles, and Richard Zanibbi. 2020. [Accelerating Substructure Similarity Search for Formula Retrieval](#). In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part I*, page 714–727, Berlin, Heidelberg. Springer-Verlag.

Wei Zhong, Yuqing Xie, and Jimmy Lin. 2022a. [Applying Structural and Dense Semantic Matching for the ARQMath Lab 2022, CLEF](#). In *Proceedings of the Working Notes of CLEF 2022 - Conference and Labs of the Evaluation Forum, Bologna, Italy, September 5th - to - 8th, 2022*, volume 3180 of *CEUR Workshop Proceedings*, pages 147–170. CEUR-WS.org.

Wei Zhong, Jheng-Hong Yang, Yuqing Xie, and Jimmy Lin. 2022b. [Evaluating Token-Level and Passage-Level Dense Retrieval Models for Math Information Retrieval](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 1092–1102. Association for Computational Linguistics.

Wei Zhong and Richard Zanibbi. 2019. [Structural Similarity Search for Formulas Using Leaf-Root Paths in Operator Subtrees](#). In *Advances in Information Retrieval - 41st European Conference on IR Research, ECIR 2019, Cologne, Germany, April 14-18, 2019, Proceedings, Part I*, volume 11437 of *Lecture Notes in Computer Science*, pages 116–129. Springer.

A Example of Informalized Mathlib4 Statements

In this section, we provide an example comparing the informalized mathlib4 statement generated by (Jiang et al., 2023), which used GPT-4 without any additional context, to the statement produced by our approach that incorporates additional context.

Formal Statement: $\text{theorem dite_eq_iff } \{ \alpha : \text{Sort } u_2 \} \{ P : \text{Prop} \} [\text{Decidable } P] \{ c : \alpha \} \{ A : P \rightarrow \alpha \} \{ B : \neg P \rightarrow \alpha \} : \text{dite } P \ A \ B = c \leftrightarrow (\exists h, A \ h = c) \vee \exists h, B \ h = c$

Informal Statement (Jiang et al., 2023)¹¹: dite P A B equals c if and only if there exists a proof h such that A h equals c, or there exists a proof h such that B h equals c.

Informal Statement (ours): Dependent if-then-else equivalence: For any proposition P that is decidable, and any elements c, A, and B, the expression ‘if P then A else B’ is equal to c if and only if either there exists a proof h such that A h is equal to c, or there exists a proof h such that B h is equal to c.

In this example, the informalized statement provided by (Jiang et al., 2023) retains the "dite" notation, which may be difficult for readers unfamiliar with mathlib4 naming conventions. In contrast, our informalized statement translates "dite" into "dependent if-then-else," making the statement more accessible. This translation was achieved by extracting the definition of "dite" and its documentation string "Dependent if-then-else..." in the prompt, demonstrating how additional context enhances the interpretation and conversion of formal theorems into informal language.

B Prompt for Query Augmentation

In this section, we present the prompt template for query augmentation, as illustrated in Figure 5. We outline specific principles for query augmentation, highlighting the importance of providing precise information, using LaTeX for mathematical expressions, and clarifying ambiguous inputs. Additionally, we offer examples of query augmentations to enhance understanding of the task.

Input:

System message: Suppose you are an expert in mathematics and Lean4. Your task is to rewrite a short, vague input description of a mathematical theorem into a detailed, explicit statement of the theorem, first written in natural mathematical language, then in the formal language of Lean4.

Output Format

Your output must contain exactly two parts. The first part is a natural mathematical language statement, starting with a theorem name, followed by the detailed statement of the theorem. The second part is a Lean4 code block, write the formal theorem statement directly without `import` before or `proof` following. NO other paragraph is allowed to appear in the output.

Principles of Output

1. You should provide as much information and be as precise as possible.
2. The natural language part should be accessible to mathematicians. Additionally, please express the mathematical formulas in LaTeX when necessary.
3. Even in the case that the input is very vague, you should guess a most similar mathematical theorem and write it out explicitly.
4. The formal statement can assume certain definitions in hypothesis is already formalized, in the case you are not sure whether or not.

Example1

Input description

The image of a compact set under a continuous function is compact.

Output

Theorem (Image of a Compact Set Under a Continuous Function):
 Let $f: X \rightarrow Y$ be a continuous function between two topological spaces X and Y . If K is a compact subset of X , then the image $f(K)$ is a compact subset of Y .

`theorem image_compact_of_compact {X Y : Type*} [TopologicalSpace X] [TopologicalSpace Y] (f : X → Y) (K : Set X) (hf : Continuous f) (hK : IsCompact K) : IsCompact (f '' K) :=`

<...other examples...>

Now, the statement for you to rewrite is:

Input description

degree of product of two polynomials equals to sum of degree

Output:

Theorem (Degree of the Product of Two Polynomials):
 Let $P(x)$ and $Q(x)$ be two non-zero polynomials over a field F . The degree of the product $P(x)Q(x)$ is equal to the sum of the degrees of $P(x)$ and $Q(x)$. Formally, if $\deg(P) = m$ and $\deg(Q) = n$, then $\deg(PQ) = m + n$.

`theorem degree_product_eq_sum_degree (F : Type*) [Field F] (P Q : Polynomial F) (hP : P ≠ 0) (hQ : Q ≠ 0) : Polynomial.degree (P * Q) = Polynomial.degree P + Polynomial.degree Q :=`

Figure 5: Prompt for query augmentations.

¹¹<https://github.com/albertqjiang/MMA/tree/main/data>

C Implementation Details

In this section, we provide the implementation details of our experiments.

We use default task instructions in UAE-Large-V1. For $E5_{\text{mistral-7b}}$, we employ different task instructions for the document and query in the non-augmented query setting and the same task instructions for both in the query augmentation setting, as shown in Table 7. Notably, the original implementation of $E5_{\text{mistral-7b}}$ omits document-side task instructions to reduce computational costs during document indexing in multiple retrieval tasks. Given the complexity of the theorem retrieval task, we have modified this approach to use two-sided prompts.

During evaluation, inputs to the embedding model exceeding 4096 characters were truncated due to GPU memory limitations. For all tested models, the corpus embedding process was completed within three hours on a single Nvidia V100 GPU with 16GB of GPU memory.

Query aug. & Doc Type	Side	Input with Task Instructions
None & Formal	Query	"Instruct: Given a math search query, retrieve theorems stated in Lean 4 that mathematically match the query \n Query:{F+IF augmented query}"
	Doc	"Instruct: Represent the given formal math statement written in Lean 4 for retrieving related statement by natural language query \n Doc:{Formal statement}"
F+IF & F+IF	Query	"Instruct: Retrieve math theorems stated in bilingual Lean 4 + natural language that are mathematically equivalent to the given one \n Query:{query}"
	Doc	"Instruct: Retrieve math theorems stated in bilingual Lean 4 + natural language that are mathematically equivalent to the given one \n Doc:{F + IF statement}"

Table 7: Task instructions used in $E5_{\text{mistral-7b}}$. Here "aug." stands for augmentation, and F and IF stands for formal and informal respectively. The "None & Formal" setting is used as baseline, and "F+IF & F+IF" is used in our method.

D Ablation Study on Task Instructions

In this section, we investigate the impact of task instructions on the theorem retrieval performance. We present a visualization of the embeddings of all 50 queries, generated by four embedding engines in our benchmark, using t-SNE(van der Maaten and Hinton, 2008). We use the default task instruction for UAE-Large-V1. For $E5_{\text{mistral-7b}}$, we test three instructions: an empty instruction, a mathematics retrieval instruction ("Given a math search query, retrieve theorems mathematically equivalent to the query"), and a Lean 4 retrieval instruction ("Given a

math search query, retrieve Lean 4 written theorems that mathematically match the query").

As illustrated in Figure 6, clusters of theorem names appear in the results produced by all four embedding models. However, these clusters vanish when we provide $E5_{\text{mistral-7b}}$ with mathematics-aware task instructions. Given that we assume all queries within the same query group (denoted by identical colors in Figure 6) have the same search intent, they should be proximal on the graph, as t-SNE maintains the relative distance relationships between vectors. Thus, the more effective the embedding engine, the more closely the dots of the same color group together. A comparison of subfigures 6c and 6d with 6a and 6b reveals that text-embedding-3-large and $E5_{\text{mistral-7b}}$ demonstrate superior performance in discerning mathematical search intentions, aligning with our main results (Table 3).

The presence of theorem name clusters in these subfigures, however, contradicts this principle, suggesting that the embedding models perceive these theorem names as more similar to each other rather than correctly associating them with their respective query groups. While this might be acceptable for other retrieval tasks, it is inappropriate for our mathematical information retrieval context. In contrast, the absence of theorem clusters in subfigures 6e and 6f suggests that these embedding methods successfully group theorem names with their corresponding statements, despite their notable differences in appearance and structure. This indicates their proficiency in comprehending mathematical theorems and recognizing search intents. In summary, appropriate task instructions enhance the embedding models' sensitivity to the search intents of mathematical concepts, thereby improving retrieval effectiveness on mathlib4.

E Models for Query Augmentation

In this section, we compare the performance of two LLMs, DeepSeek-V2-Chat and GPT-4, for query augmentation. Table 8 presents the performance of both models using the informal+formal corpus setting, with results averaged over 10 runs.

API	nDCG@20	P@10	R@10	Cost per 1M input tokens	Cost per 1M output tokens
gpt-4-0125-preview	0.721	0.195	0.911	\$10.00	\$30.00
deepseek-chat	0.749	0.202	0.933	\$0.14	\$0.14

Table 8: Performance of DeepSeek-V2-Chat and GPT-4 for query augmentation.

As shown, DeepSeek-V2-Chat not only outper-




Figure 6: t-SNE visualization of the embeddings of all queries in our benchmark. Each dot denotes a distinct query, with queries within the same query group sharing identical colors. The shape of the markers differentiates the four distinct query categories. Theorem name clusters are emphasized with circles.

forms GPT-4 in terms of effectiveness but also offers a significantly lower cost. The results suggest that DeepSeek-V2-Chat is a more efficient and cost-effective choice for query augmentation tasks.