

---

# Renaissance of Literate Programming in the Era of LLMs: Enhancing LLM-Based Code Generation in Large-Scale Projects

---

Wuyang Zhang<sup>\*1</sup> Yansong Li<sup>\*1,2</sup> Zeyu Dong<sup>3</sup> Yu Wu<sup>4</sup> Yingyao Zhou<sup>2</sup> Duolei Wang<sup>5</sup> Songsirou Xing<sup>6</sup>  
Chichun Zhou<sup>7</sup> Da Shen<sup>1</sup>

## Abstract

Large Language Models (LLMs) have helped programmers increase efficiency through code generation, comprehension, and repair. However, their application to large-scale projects remains challenging due to complex interdependencies and the extensive size of modern codebases. Although Knuth’s concept of Literate Programming (LP) combines code and natural language to convey logic and intent, its potential for enhancing relationships in large projects has not been fully explored. In this study, we introduce the idea of Interoperable LP (ILP), which leverages literate programming principles to enhance the development of both small-scale documents and large-scale projects with LLMs. We investigate how LLMs perform under ILP-style instructions for both document-oriented tasks and entire projects. Recognizing that many researchers rely on well-structured templates to guide LLMs, we propose a concise prompt engineering method to write LP documents so LLMs can better be involved in code generation. We also examine the capacity of various LLMs to generate Scheme and Python code on the RepoBench benchmark, illustrating the advantages of our approach. Our findings indicate that ILP with LLMs can enhance LLM-based code generation in large-scale project development.

## 1 Introduction

Generative machine learning models, particularly large language models (LLMs), have significantly enhanced how programmers approach coding tasks. These models excel at generating functional code snippets, adding detailed com-

ments, and performing tasks like code translation, repair, and test generation (Jiang et al., 2024; Chen et al., 2021). For example, LLMs can translate user prompts into working implementations while simultaneously offering clear explanations of the code’s logic, making programming more accessible to users with varying levels of expertise. This dual capability of generating and documenting code has positioned LLMs as indispensable tools for modern software development.

Currently, LLMs are widely used for two main purposes: efficient code generation (Jiang et al., 2024) and comprehensive code explanation (Nam et al., 2024). Traditional code generation focuses on creating single pieces of code based on user-provided prompts. These prompts can range from detailed, step-by-step instructions guiding the implementation process to high-level descriptions with sample inputs and outputs that define the behavior of an API. While effective for generating small and isolated pieces of code, this method is inherently limited in scope, as it treats each task as independent and does not account for broader project contexts. However, as software projects grow in scale and complexity, traditional approaches to code generation become increasingly inadequate (Dou et al., 2024).

Designing APIs or generating multiple interdependent components in large-scale projects demands more than isolated, prompt-based workflows. Such projects require a framework that manages complex interdependencies between APIs, ensures consistency across components, and supports interconnected development processes. To address these challenges, a new framework is necessary—one that facilitates bulk API generation handles dependencies, and enables the creation of components reliant on yet-to-be-developed APIs. This approach would provide the structure and coherence needed to tackle the intricacies of modern, large-scale software projects.

To improve LLM performance in generating and explaining code for large projects, we revisit the concept of Literate Programming (LP). LP, while established, still offers a foundation for organizing interdependencies and enforcing consistent standards across components. Tools like Jupyter Notebook and NoWEB demonstrate the benefits of combin-

---

<sup>\*</sup>Equal contribution <sup>1</sup>Liiv Network <sup>2</sup>University of Illinois Chicago <sup>3</sup>Stony Brook University <sup>4</sup>Rutgers University <sup>5</sup>The Chinese University of Hong Kong, Shenzhen <sup>6</sup>John Hopkins University <sup>7</sup>Dali University. Correspondence to: Yansong Li & Da Shen <da@liiv.pro>.

---

ing code and documentation within a single environment (Fig. 1). However, they lack universal standards and do not fully address the demands of modern software development. In particular, they do not offer robust methodologies for generating and explaining code across interdependent systems.

In this paper, we build upon LP to propose a new standard called Interoperable Literate Programming (ILP). ILP provides a structured approach tailored to LLM-based code generation and explanation. By defining explicit guidelines for code logic, ILP enhances accuracy and development efficiency. We also introduce two software tools—Mogan and Goldfish Scheme—that showcase practical applications of ILP. These tools help developers manage large projects, integrate documentation, and test code within a unified and secure environment.

This paper addresses the limitations of traditional code generation methods by introducing a framework based on ILP. Experiments are conducted in RepoBench (Liu et al., 2023) on Python code generation tasks. Our key contributions are:

- Our evaluations show that ILP reduces variability in code generation quality. The performance remains consistent across instruct-based models, including locally deployed pre-trained models and online models with internet search capabilities.
- ILP improves code generation efficiency, achieving speed increases that lead to higher productivity.
- ILP supports bulk API generation and resolves complex interdependencies, including those involving components yet to be developed. This approach addresses critical challenges in scaling code generation for large systems.
- We provide comprehensive benchmarks and evaluations, rather than a single metric of execution success, to demonstrate ILP’s effectiveness across different coding environments. These benchmarks compare code quality, speed, and completion rates across multiple models.

**Related works** *Literate programming supported editor*: In this paper, we introduce Mogan, a UTF-8-compatible literate programming editor that allows users to edit and export all project files within a single document, as shown in Figures 1 and 2. This unified approach aims to simplify the workflow for large-scale projects. Other literate programming editors are discussed in Section C.3.

*Prompt engineering*: Wei et al. (2022) propose the chain-of-thought (Chain of Thought (CoT)) method, guiding large language models (LLMs) to reason step by step. This improves both performance and interpretability. Building on

this idea, least-to-most (LtM) prompting (Zhou et al., 2022a) divides complex tasks into sequential subproblems, helping LLMs tackle more difficult scenarios. Recent work also focuses on increasing consistency and accuracy. For example, Wang et al. (2022) introduce self-consistency, a strategy that produces multiple reasoning paths from a single prompt and selects the most coherent outcome. Dhuliawala et al. (2023) present Chain-of-Verification (CoVE), a method that prompts LLMs to verify their outputs, reducing hallucinations. To reduce the manual effort in crafting prompts, researchers have explored automated prompt construction (Zhou et al., 2022b; Pryzant et al., 2023; Ye et al., 2023). These methods guide the LLM to examine its responses and propose improvements, as shown by Li et al. (2023), who refine the “think step by step” directive by prompting the LLM to identify its own limitations and update its solution accordingly.

*Scheme language*: Sussman & Steele Jr (1975) implement “Scheme”, a simple and lightweight functional programming language based on lambda calculus. Furthermore, several Scheme interpreters have been developed based on the Scheme standard, with the latest and most widely accepted being R7RS. Additionally, Scheme Requests for Implementation (SRFI) are community-driven specifications that extend the language with additional features and libraries. The Goldfish Scheme (Shen et al., 2024), built on S7 Scheme, supports the R7RS standard, parts of SRFI, and many functionalities inspired by the Python standard library. This simplicity, combined with several characteristics of Python, makes it easier for large language models to understand. More discussion about Scheme interpreters is in Sec. C.4.

## 2 Background


In this section, we provide background on Large Language Model (LLM)-based code generation, LP, and the Goldfish Scheme, a Scheme interpreter developed via the LP paradigm.

### 2.1 Challenges in LLM-based code generation


Large Language Models (LLMs) are generative models designed for natural language processing tasks, capable of performing functions such as text generation, translation, code generation, sentiment analysis, and more. More discussion on LLM-based code generation is attached in Sec. B.

Despite their capabilities, LLMs face limitations in code generation for large-scale projects. In this paper, we focus on the following challenges:

- **External bias**: LLMs often rely on pre-trained knowledge and external references rather than utilizing the APIs provided within a project. For instance, when



(a) Two files `packages/lp_in_packages.ipynb` and `lp.ipynb` are edited in different documents



```
The hello world in Python

lp.py
print("hello_world")

The sum in Python

packages/lp_in_packages.py
1 + 1

We can edit single file of code in different places inside the document.

lp.py
print("hello_world_again!")
```

(b) Two files `packages/lp_in_packages.py` and `lp.py` are edited in a single document. Also, a file `lp_in_packages.py` can be edited at multiple locations within a single document with other files.

Figure 1: Comparison: literate programming with Mogan and Jupyter Notebook programming

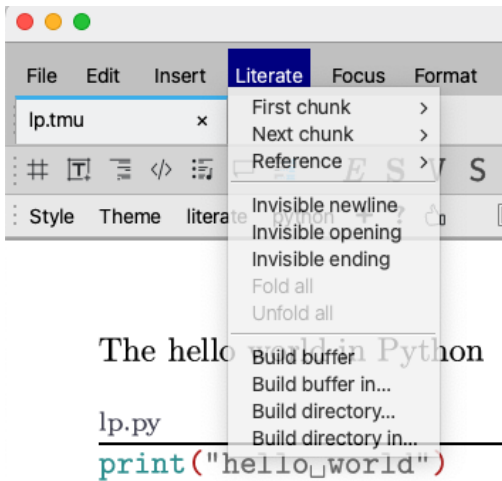


Figure 2: The whole project can be exported from the literate programming document with BUILD BUFFER button in Mogan.

tasked with implementing a backpropagation algorithm using a project’s APIs, LLMs are likely to bypass the given APIs and instead default to those in frameworks like PyTorch. Sec. B.2 provides a practical example of the external bias.

- **Inconsistent Solutions:** LLMs are highly sensitive to input prompts or context. For instance, changing an irrelevant token in the input can result in entirely different outputs. This inconsistency can lead to unreliable results, especially when stable and predictable solutions are essential for tasks involving complex logic or project-specific APIs.

The primary reasons why LLMs suffer from external bias

and inconsistent solutions are the abundance of training data for languages like Python and their widespread popularity. These factors result in an overwhelming number of pre-existing tools and libraries (“wheels”) for LLMs to choose from, leading to variability in the solutions generated during each run. Additionally, relying on wheels obscures the logical flow of the project. As a result, it becomes challenging for LLMs to deduce the project’s internal logic based solely on the provided codebase.

To enhance LLM-based code generation in large-scale projects, we reintroduce LP (Sec. 2.2) alongside the Scheme language (Sec. 2.3). LP organizes code and descriptive content in a logical manner, enabling LLMs to better comprehend the inherent logical flow of the project.

Additionally, we extend the concept of LP into ILP in Sec. 2.2, where the inherent logical flows between APIs are represented using a directed acyclic graph (DAG) structure within the LP document. In this paper, we use Scheme for these descriptions due to its simplicity and inherent  $\lambda$ -calculus foundation, which naturally aligns with the DAG representation. Optionally, Scheme can also be used as the target language, enabling LLMs to generate more consistent code without external bias, as Scheme lacks the extensive libraries and ‘wheels’ commonly associated with more popular programming languages. The detailed discussion of why the Scheme is suited for the ILP documents is in Sec. C.4.

## 2.2 Interoperable LP

There is no formal governing body or ISO-style standard for LP. There are only several widely agreed-upon practices. The closest to a “standard” we can find is Knuth’s original WEB system introduced in 1984, which established some foundational concepts:

```

[This is the original scheme sum function, which can only take integers as inputs]
old_packages/sum.scm
;; This is the original scheme sum function, which can only take integers as inputs
(define (add int1 int2) (+ int1 int2))

The new sum function we defined can add two rational numbers.
new_packages/sum.scm
;; The new sum function we defined can add two rational numbers.
;; The old sum is defined in old_packages/sum.scm

; Extend the + operator to handle both integers and rationals
(define (extended+ x y)
  (cond
    ((and (pair? x) (pair? y)) (add-rat x y)) ; If both are rationals, add as rationals
    ((pair? x) (add-rat x (make-rat y 1))) ; If x is rational, convert y to rational and add
    ((pair? y) (add-rat (make-rat x 1) y)) ; If y is rational, convert x to rational and add
    (else (+ x y)))) ; Otherwise, use standard integer addition

; Redefine add to use extended+
(define (add rat1 rat2)
  (extended+ rat1 rat2))

```

(a) Two idea-connected files edited in Mogan using literate programming; LLMs can directly read the connection.

(b) Connection between these two files are commented in new\_packages/sum.scm. Even with such comments, LLMs must read and interpret the paths mentioned within, locate the referenced file, and then analyze it to grasp the connection.

Figure 3: Comparison of a literate programming approach (left), which organizes code as a “web” of interconnected ideas, versus a typical Jupyter Notebook workflow (right) that does not follow literate programming principles.

- **Flexible code organization:** Code can be arranged in any logical order that suits the narrative, interwoven with descriptive text rather than mere comments.<sup>1</sup> This entire collection of code and text resides in a single document, as shown in Figure 3.
- **Tangling:** The combined code must be “tangled” to form a compilable source. While Jupyter does not support this process, Mogan includes tangling functionality, as seen in Figure 2.
- **Bidirectional workflows:** Tools and frameworks for LP often allow generating documentation from code (weaving) and extracting code from documentation (tangling) within the same environment. This “bidirectional” capability ensures that changes in the narrative are reflected in the compiled code and vice versa.

The focus of his design is on making programs comprehensible to humans first and executable by computers second. LP is particularly effective for algorithm-intensive, complex applications, such as large projects, where natural language enhances programmers’ understanding of the code. LP allows users to edit and export all project files within a single document, aligning with human logic, as illustrated in Fig. 1 and Fig. 2. Additionally, as shown in Fig. 3, LP facilitates connections between files. In contrast, in traditional methods, the logical flow between APIs is inherently hidden throughout the project files. LP overcomes this limitation by utilizing natural language to establish these connections.

In this paper, we use the term *documents* or *documentations* to refer to scripts that integrate code and descriptive content in a style consistent with LP principles. The term *files* denotes scripts consisting of code and comments. Documents in LP support a sequence that feels natural and coherent,

<sup>1</sup>Figure 3 illustrates how descriptive text differs from standard comments.

where each “*chunk*” of code contributes to an unfolding narrative, making the program easier to understand and compile.

Some engineers considered the Jupyter Notebook as an LP document. However, we do not consider a .ipynb file to be a *document* but rather a *file* in this paper. While a Notebook can mix text and code, it is not easily integrated or exported into a broader project structure in the way that LP requires. This lack of portability separates Notebook-based code from the type of unified narrative and codebase advocated by LP in this paper. Additionally, it does not allow files of codes to be structured as a flexible “web” of ideas in a single document. A more detailed discussion can be found in Sec. C.1. It is also important to note that code with detailed comments does not constitute LP; this topic is further explored in Sec. C.1.

To distinguish our definition of LP from Notebook-based coding, we introduce a modern form of LP called Interoperable Literate Programming (ILP). ILP incorporates all LP features we listed at the beginning of this subsection. More discussion on ILP is in Sec. C.1.

### 2.3 Goldfish Scheme as nature language compensation for descriptive contents in ILP

In this work, we introduce *Goldfish Scheme*, a Scheme interpreter with a Python-like standard library. As discussed in Sec. 2.1, Scheme is used to describe a directed acyclic graph (DAG) that represents the inherent logical flow between APIs in Python projects, leveraging Scheme’s simplicity and its  $\lambda$ -calculus foundation. Goldfish Scheme incorporates these features and includes an LLM interaction layer, making it highly suitable for LLM-based code generation. Additional features, such as its Python-like standard library for transfer learning, are further discussed in Sec. D.

**LLM interaction layer** The core idea behind this interaction layer is to supply machine-readable fields that describe

```

(define-with-docs quicksort
  #:pattern "divide-and-conquer"
  #:complexity "O(n log n)"
  #:stability "unstable"
  #:examples
  '((quicksort '(3 1 4 1 5 9 2 6 5 3))
    => (1 1 2 3 3 4 5 5 6 9))
  (lambda (lst)
    ;; Implementation follows...
  ))

```

Figure 4: Goldfish Scheme as descriptive contents in ILP.

key aspects of a function or library component. These annotations can specify algorithmic patterns, complexity metrics, stability characteristics, and concrete usage examples. By making this information explicit and accessible, the interaction layer allows LLMs to understand each piece of code not only at a syntactic level but also in terms of conceptual intent and practical implications. Consider the following `define-with-docs` macro:

Remember that the macro is not intended for execution but serves as an alternative to natural language in the descriptive content of an ILP document. The goal is not to generate or execute Scheme code but to facilitate the generation of Python code. In Fig. 4, the “divide-and-conquer” pattern informs the LLM of the underlying strategy behind the algorithm, making it easier for the model to relate this sorting routine to other similar functions and patterns it knows. By stating that the sort is “unstable” and complexity as  $O(n \log n)$ , the annotations reveal a critical property that distinguishes the function’s behavior, guiding the LLM toward more precise reasoning about potential use-cases and limitations. The inclusion of example tags allows the LLM to confirm its understanding of inputs, outputs, and expected transformations. Such an example can also serve as a reference point for generating further tests or variant solutions. More discussions are in Sec. D.

### 3 Interoperable literate programming documentation design

In Sec. 2.1, we identified multiple challenges in using LLMs for code generation in large-scale projects. We propose addressing these challenges by applying ILP with descriptive content that structures an API’s logic as a Directed Acyclic Graph (DAG). In this section, we present our methodology for creating the ILP document. Although we use Scheme code snippets to illustrate how the DAG is documented, these snippets are not intended for execution or compilation. Instead, they serve as placeholders to capture the logical flow of an API, allowing us to design and describe interdependencies without running the code itself.

A clear and logically structured documentation design is crucial for explaining API logic and guiding LLMs in generating correct code. We propose a step-wise framework to tackle the challenges from Sec. 2.1. This framework reduces reliance on external references, enforces logical consistency, and lowers the risk of errors.

We adopt a step-wise methodology inspired by mathematical induction and recursive algorithms, dividing large tasks into base (*zero-step*) and inductive (*successor-step*) parts. For instance, in recursive code, establishing a clear base case and carefully explaining the recursive step ensures consistency in the final output. Similarly, our method segments a task’s logic into smaller, sequential units, giving LLMs an incremental path instead of overwhelming them with entire program logic at once. By breaking down complex tasks into separate components, each step can be understood independently. Additional benefits are discussed in Sec. E. This structure inherently forms a DAG. Our approach relies on two core principles: *Scheme’s simplicity* and a *stepwise progression*. Scheme’s minimal syntax and functional style treat each part of the logic as a discrete unit. Furthermore, because Scheme is relatively uncommon in LLM training data, models must depend on the provided documentation, reducing the potential for errors caused by external biases (see Sec. 2.1).

Below, we outline the key components that enable this method to function effectively (additional details appear in Sec. E):

1. **Zero-step definition:** We start by specifying a *base case* or fundamental building block. In recursive functions, for example, this might be the condition under which recursion terminates. By designating this zero-step, we provide the LLM with a clear starting point for the broader logic.
2. **Successor-step explanation:** We then describe how one logical state transitions to the next. In inductive or recursive algorithms, this step outlines the “successor” or “next element,” ensuring that each additional step builds consistently on the zero-step.

To clarify these concepts, we offer several toy examples that target the Scheme language and use the R7RS document as a base for literate programming. We embed the following descriptive content in the R7RS document, converting it into an ILP document:

```

## take-right

### Zero-Step Logic
Return the list if it's empty.

```



---

```
### Succ-Step Logic
```

```
See below;
```

```
### Helper Function: `drop`
```

```
The helper function `drop` removes the first *n*  
↪ elements from a list. See Sec. 7.1.7 in this  
↪ document.
```

```
(define-with-docs take-right  
 #:pattern "divide-and-conquer"  
 #:complexity "O(n)"  
 #:stability "stable"  
 #:examples  
 '(take-right '(a b c d e) 2) => (d e)  
 (lambda (lst)  
   ;; Implementation follows...  
   ;; Base (zero-step):  
   ;; if n <= 0 or lst is empty, return lst as  
   ↪ is  
   ;; Succ-step:  
   ;; otherwise, (drop lst 1) and (take-right  
   ↪ lst (- n 1))  
 ))
```

We then prompt LLMs with:

```
Write the function (take-right flist i) using the zero-  
step/succ-step approach as described in the pdf  
file. The Scheme language grammar and necessary  
APIs can be found in the /opt/cachebox/r7rs/.
```

All toy-case experimental results appear in Sec. I.1 and Table 1. As shown in Table 1, GPT-4o generates correct code for all toy examples when guided by ILP. In the toy-case example above, the zero-step and successor-step align closely with the recursive implementation of `take-right`. However, not all functions rely on pure recursion. In some instances, the zero-step and successor-step capture the broader logical flow of APIs in the ILP document rather than a direct mapping to implementation details. Fig. 5 illustrates the DAG structure and corresponding Scheme description for the extended `add` function from Fig. 3a. Another example involving Python code generation from a Scheme-based description is provided in Sec. G.3.

Finally, the stepwise method itself is inherently recursive, and any code generated by LLMs reflects this pattern. Adding the generated code segments back into the ILP document produces an updated ILP document, reinforcing the same iterative design.

## 4 Experiments

In this section, we designed our experiments to observe the performance of LLMs’ Python code generation by learning ILP-style documents. Our experiments are conducted on two types of LLMs: local open-source LLMs and remote commercial LLMs. The machine details are in Sec. G.1.

### 4.1 RepoBench: benchmark for code generation

RepoBench (Liu et al., 2023) serves as our evaluation benchmark environment due to its realistic and comprehensive testing framework. Unlike benchmarks focused on isolated coding tasks, RepoBench emphasizes the interconnection between code components within projects. This design provides a practical environment for evaluating how LLMs handle dependencies and relationships across different parts of a codebase, reflecting the complexities of real-world software development.

RepoBench’s test cases are sourced from GitHub repositories, ensuring authenticity and diversity in the scenarios it presents. These real-world examples go beyond simplified toy problems, capturing the nuances and challenges typical of production-level software. This allows us to evaluate the performance of LLMs in contexts that closely mirror actual programming tasks.

Another critical feature of RepoBench is its diverse documentation patterns, which include comments in various programming languages. While most comments are not originally written in Scheme, this diversity presents a valuable opportunity to test the adaptability of our approach. To enable Scheme-based evaluations, we translated a selection of comments into Scheme, preserving their semantic meaning while aligning them with Scheme’s functional programming paradigm.

Further details about our adaptation of RepoBench for Scheme are discussed in Sec. F.1. In Sec. F.2, we provide a thorough comparison with alternative code generation benchmarks, highlighting why RepoBench is particularly well-suited for evaluating LLM-based code generation in our study.

### 4.2 Experiment: Why Scheme is ideal for LLM code generation?

In Sec. 2.3 and Sec. C.4, we discussed why Scheme is ideal as a description language for LLM code generation. The main reasons are: 1. LLMs are not extensively trained in Scheme. 2. Scheme is better suited for describing logical structures. In this section, we design an experiment to verify these conjectures.

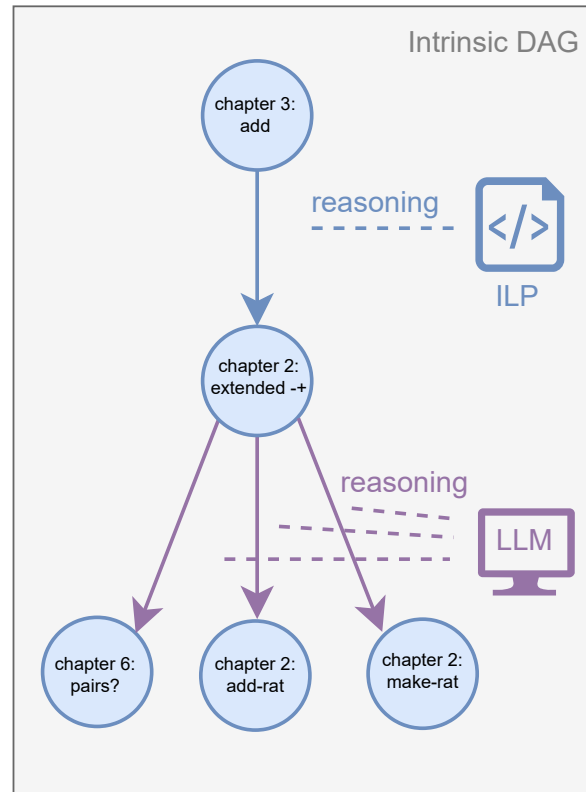
We first examined the extent of Scheme knowledge embedded in LLMs by prompting ChatGPT-4O with basic requests like “Generate a function called `xxx` in Scheme that outputs `xxx`.” The resulting code often contained syntax errors or flawed logic, suggesting that the models do not possess significant pre-trained Scheme patterns. Fig. 7 in Sec. C.4 provides an example where ChatGPT-4 fails to generate a simple `sort-int` function. This finding supports our hypothesis that Scheme’s limited presence in modern development contexts forces LLMs to engage in actual logical reasoning

```

;; The following are documents in Chapter. 3
(define-with-docs add
  #:pattern "api-calls"
  #:complexity "undefined"
  #:stability "stable"
  #:examples
  '(add (make-rat 1 2) (make-rat 1 2))=> 1
    ;; Implementation follows...
    ;; Base (zero-step):
    ;;   + in original Scheme
    ;; Succ-step:
    ;;   call extended-+ in Chapter.2
  ))
;; The following are documents in Chapter. 2,
↪ searched by LLMs.
(define-with-docs extended-+
  #:pattern "api-calls"
  #:complexity "undefined"
  #:stability "stable"
  #:examples
  '(extended-+ (make-rat 1 2) (make-rat 1
↪ 2))=> 1
    ;; Implementation follows...
    ;; Base (zero-step):
    ;;   + in the original Scheme
    ;; Succ-step:
    ;;   call add-rat and make-rat in Chapter.
↪ 2.
    ;; call pairs? in Chapter 6.
  ))
;; more define-with-docs follows

```

(a) Scheme representation for DAG: Note that these representations are located in different chapters of the ILP document. Only `(define-with-doc add)` needs to be written, as LLMs can recursively search for other descriptions previously written. Hyperlinks are included to reference the related chapters within the document.



(b) DAG representation: Each node corresponds to a Scheme description, with nodes distributed across different chapters. As shown in purple lines, LLMs shall trace the inherent logical flow of APIs through the descriptive contents in the document.

Figure 5: The DAG structure and the corresponding Scheme description of the extended add function are shown in Fig. 3a.

rather than recalling ready-made examples.

Next, we investigate whether Scheme’s mathematical underpinnings enhance logical reasoning. We compare code generation using standard R7RS and the ILP document based on R7RS discussed in Sec. 3.

As shown in Table 1, we test operations ranging from basic list manipulations to more advanced tasks like `delete-duplicates` and `split-at`. When comparing different prompts—R7RS-only descriptions, R7RS with detailed documentation, and R7RS with instructions to “fully depend on the file”—we observed that performance improved significantly with the ILP documentation. Even when we replaced standard function names with arbitrary terms, the models generated working code, implying that they were guided by the documented logic rather than default training data.

What’s more, to ensure LLMs focus solely on our documentation rather than their broader training knowledge, we

implement specific constraints in our prompts. One key strategy involves using non-standard function names in our implementation. Instead of traditional function names like `map`, we experiment with arbitrary names, such as Pinyin (Mandarin Chinese that uses the Latin alphabet) translations. This approach helps us verify whether LLMs are truly working from our documentation rather than falling back on their pre-trained knowledge of the R7RS Scheme.

For example, rather than using:

```
(define (map function lst)
  ...)
```

We might use:

```
(define (ditu hanshu liebiao)
  ...)
```

This naming strategy serves as a clear indicator: if an LLM uses the standard name in English instead of our arbitrary

Table 1: Scheme Function Implementation: Success or Failure. “R.” refers to the detailed descriptions from R7RS, while “D.” represents the designed documentation following our two-step programming logic.

Function	Description	GPT-4o (R. D.)	GPT-4 (R. D.)	GPT-4o (R.)
make-list	Returns an n-element list, whose elements are all the value <code>fill</code> .	S	S	S
iota	Returns a list containing the elements <code>(start, start+step, ... , start+(count-1)*step)</code> .	S	S	F
null-list?	Returns true if it is the empty list <code>()</code> , and false otherwise.	S	S	S
drop-right	Returns all but the last <code>i</code> elements of the given list.	S	S	S
take-right	Returns the last <code>i</code> elements of the list.	S	S	S
split-at	Splits the list <code>x</code> at index <code>i</code> , returning a list of the first <code>i</code> elements, and the remaining tail.	S	S	F
append	Returns a list followed by another list.	S	F	S
reverse	Returns a newly allocated list consisting of the elements of the list in reverse order.	S	S	S
for-each	Calls <code>proc</code> (given functions) on the elements of the lists.	S	S	S
remove	Returns list without the elements that satisfy predicate <code>pred</code> .	S	S	S
find	Returns the first element of <code>c-list</code> that satisfies predicate <code>pred</code> ; false if no element does.	S	S	S
delete-duplicates	Removes duplicate elements from the list argument.	S	S	F

one, we know it’s drawing from external knowledge rather than our documentation. The result is given in Sec. G.2.

These findings confirm that Scheme’s  $\lambda$ -calculus structure and reduced prevalence in training sets encourage LLMs to rely on logical reasoning. Furthermore, our two-step documentation pattern, based on mathematical induction, effectively channels LLM reasoning by splitting code generation into clear base and inductive steps. This design takes advantage of Scheme’s alignment with theoretical principles, leading to reliable code generation for a broad range of tasks.

### 4.3 The general design of a single ILP document

To improve LLM performance in Goldfish Scheme, we structure the natural language portion of our documentation using the CoT methodology, enhancing LLMs’ understanding of code relationships. The other descriptive contents follow the stepwise method outlined in Sec. 3. Code chunks follow directly with the corresponding descriptive contents. More discussions are in Sec. G.2.

### 4.4 Inter-file inference management

We carefully structure our code to test LLMs’ ability to follow logical connections across files. Instead of maintaining traditional linear dependencies (where A depends on B, which depends on C), we split implementations across files in ways that require an understanding of the broader context.

Consider a typical implementation flow:

1. Input processing in File A.

2. Data transformation in File B.

3. Resulting LLM-generated code chunks in File A.

This structure tests whether LLMs can:

- Recognize the overall logical flow.
- Identify connections between seemingly separate implementations.
- Understand how different components work together despite being physically separated.

Fig. 5 shows a case for these type of code generation tasks. The following is another example of this structure:

```
;; File: processing.scm
(define (stage-one input)
  (process-initial input))

(define (stage-three processed-data)
  (generate-result processed-data))

;; File: transform.scm
(define (stage-two data)
  (transform-intermediate data))
```

Figure 6: Cross-file implementation example

This structure helps us evaluate not just whether LLMs can write correct code but whether they truly understand the relationships and dependencies within our codebase. The non-linear organization of functionality across files serves



---

as a more rigorous test of an LLM’s comprehension of the overall system architecture.

#### 4.5 Experiment result on Python code generation on RepoBench

Take the generation of the `take-right` function on RepoBench as an example; the ILP document is the same as in Sec. 3 and the prompt given to LLMs are:

```
Fully based on the file, generate a function in
python for take-right API mentioned in the docu-
ment?
```

We test code generated for the `take-right` function, both with and without ILP-based guidance. The experiment result of Python code generation on RepoBench is in Sec.I.2. The result reveals measurable improvements in structural consistency and implementation quality with the help of ILP.

When following ILP documentation, large and small models alike produce similar code structures. GPT-4, Claude Sonnet, and LLAMA 3.1 2B implementations all include well-defined edge case handling, helper functions such as `drop`, and comprehensive error checks. In contrast, implementations generated without ILP guidance vary significantly in structure; some omit essential edge case handling and error checks altogether.

Local LLMs with fewer parameters show marked gains under ILP guidance. For example, LLAMA 3.1 2B generates code that includes proper error handling and documentation, matching the functional characteristics of larger models like GPT-4. This indicates that ILP’s structured documentation method enables smaller models to produce code of comparable reliability.

All tested models, including both commercial online systems (GPT-4 and Claude) and local deployments (LLAMA and Qwen), exhibit consistent and systematic error-handling patterns under ILP. Each implementation checks for `null` conditions and boundary cases while preserving functional programming practices such as avoiding side effects and implementing correct recursion.

ILP’s structural elements directly support these improvements. The zero-step/successor-step approach aligns with recursive design, clarifying base cases and recursive transitions. The documentation specifies error conditions and edge cases, providing clear patterns for helper function implementation. By incorporating functional programming principles directly into the documentation, ILP guides models toward functional purity and correct recursion in their implementations.

## 5 Conclusion

This paper introduced a novel framework, ILP, which extends the principles of LP to address the challenges of large-scale software development with LLMs. By providing a structured approach to managing interdependencies, ensuring consistency, and facilitating efficient code generation and documentation, ILP overcomes the limitations of traditional prompt-based methods. Through experiments conducted on RepoBench and practical applications demonstrated with tools like Mogan and Goldfish Scheme, we showed that ILP enhances code generation quality, efficiency, and scalability. These advancements make ILP a robust solution for modern software projects, enabling developers to seamlessly integrate documentation and achieve reliable outcomes in increasingly complex coding environments. Our findings underline ILP’s potential to transform LLM-based development workflows and pave the way for more efficient and coherent project management practices.

## References

- Abelson, H., Sussman, G. J., and with Julie Sussman. *Structure and Interpretation of Computer Programs*. MIT Press/McGraw-Hill, Cambridge, 2nd edition, 1996. ISBN 0-262-01153-0.
- Barendregt, H. P. et al. *The lambda calculus*, volume 3. North-Holland Amsterdam, 1984.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. D. O., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Chen, X., Liu, C., and Song, D. Tree-to-tree neural networks for program translation. *Advances in neural information processing systems*, 31, 2018.
- Dhuliawala, S., Komeili, M., Xu, J., Raileanu, R., Li, X., Celikyilmaz, A., and Weston, J. Chain-of-verification reduces hallucination in large language models. *arXiv preprint arXiv:2309.11495*, 2023.
- Dominik, C. *The Org-Mode 7 Reference Manual: Organize Your Life with GNU Emacs*. Network Theory, UK, 2010. with contributions by David O’Toole, Bastien Guerry, Philip Rooke, Dan Davison, Eric Schulte, and Thomas Dye.
- Dou, S., Jia, H., Wu, S., Zheng, H., Zhou, W., Wu, M., Chai, M., Fan, J., Huang, C., Tao, Y., et al. What’s wrong with your code generated by large language models? an extensive study. *arXiv preprint arXiv:2407.06153*, 2024.
- Dybvig, R. K. Chez scheme version 8 user’s guide. *Cadence Research Systems*, 2005.

- Dybvig, R. K. *The Scheme programming language*. MIT Press, 2009.
- Fan, Z., Gao, X., Mirchev, M., Roychoudhury, A., and Tan, S. H. Automated repair of programs from large language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pp. 1469–1481. IEEE, 2023.
- Fried, D., Aghajanyan, A., Lin, J., Wang, S., Wallace, E., Shi, F., Zhong, R., Yih, W.-t., Zettlemoyer, L., and Lewis, M. Incoder: A generative model for code infilling and synthesis. *arXiv preprint arXiv:2204.05999*, 2022.
- Gamma, E. *Design patterns: elements of reusable object-oriented software*, 1995.
- Guo, D., Xu, C., Duan, N., Yin, J., and McAuley, J. Long-coder: A long-range pre-trained language model for code completion. In *International Conference on Machine Learning*, pp. 12098–12107. PMLR, 2023.
- Hanson, C. and of Technology. Scheme Team, M. I. *MIT/GNU Scheme Reference Manual*. Samurai Media Limited, 2015.
- Jiang, J., Wang, F., Shen, J., Kim, S., and Kim, S. A survey on large language models for code generation. *arXiv preprint arXiv:2406.00515*, 2024.
- Jin, M., Shahriar, S., Tufano, M., Shi, X., Lu, S., Sundaresan, N., and Svyatkovskiy, A. Inferfix: End-to-end program repair with llms. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 1646–1656, 2023.
- Johnson, A. L. and Johnson, B. C. Literate programming using noweb. *Linux Journal*, 42:64–69, 1997.
- Joshi, H., Sanchez, J. C., Gulwani, S., Le, V., Verbruggen, G., and Radiček, I. Repair is nearly generation: Multilingual program repair with llms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 5131–5140, 2023.
- Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C. Jupyter notebooks – a publishing format for reproducible computational workflows. In Loizides, F. and Schmidt, B. (eds.), *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pp. 87 – 90. IOS Press, 2016.
- Knuth, D. E. Literate programming. *Comput. J.*, 27(2): 97–111, May 1984. ISSN 0010-4620. doi: 10.1093/comjnl/27.2.97. URL <https://doi.org/10.1093/comjnl/27.2.97>.
- Li, C., Liang, J., Zeng, A., Chen, X., Hausman, K., Sadigh, D., Levine, S., Fei-Fei, L., Xia, F., and Ichter, B. Chain of code: Reasoning with a language model-augmented code emulator. *arXiv preprint arXiv:2312.04474*, 2023.
- Liu, T., Xu, C., and McAuley, J. RepoBench: Benchmarking Repository-Level Code Auto-Completion Systems. In *The Twelfth International Conference on Learning Representations*, October 2023.
- Lu, S., Duan, N., Han, H., Guo, D., Hwang, S.-w., and Svyatkovskiy, A. Reacc: A retrieval-augmented code completion framework. *arXiv preprint arXiv:2203.07722*, 2022.
- Nam, D., Macvean, A., Hellendoorn, V., Vasilescu, B., and Myers, B. Using an llm to help with code understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, pp. 1–13, 2024.
- Olausson, T. X., Inala, J. P., Wang, C., Gao, J., and Solar-Lezama, A. Is self-repair a silver bullet for code generation? In *The Twelfth International Conference on Learning Representations*, 2023.
- Parasaram, N., Yan, H., Yang, B., Flahy, Z., Qudsi, A., Ziaber, D., Barr, E., and Mechtaev, S. The fact selection problem in llm-based program repair. *arXiv preprint arXiv:2404.05520*, 2024.
- Pryzant, R., Iter, D., Li, J., Lee, Y. T., Zhu, C., and Zeng, M. Automatic prompt optimization with "gradient descent" and beam search. *arXiv preprint arXiv:2305.03495*, 2023.
- Ramsey, N. Literate programming simplified. *IEEE software*, 11(5):97–105, 1994.
- Roziere, B., Lachaux, M.-A., Chausson, L., and Lample, G. Lost in translation: A study of bugs introduced by large language models while translating code. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 67–75. Association for Computational Linguistics, 2020a. doi: 10.18653/v1/2020.emnlp-main.123. URL <https://aclanthology.org/2020.emnlp-main.123/>.
- Roziere, B., Lachaux, M.-A., Chausson, L., and Lample, G. Unsupervised translation of programming languages. *Advances in neural information processing systems*, 33: 20601–20611, 2020b.
- Schulte, E. and Davison, D. Active documents with org-mode. *Computing in Science Engineering*, 13(3):66–73, may-june 2011. ISSN 1521-9615. doi: 10.1109/MCSE.2011.41.

- 
- Schulte, E., Davison, D., Dye, T., and Dominik, C. A multi-language computing environment for literate programming and reproducible research. *Journal of Statistical Software*, 46(3):1–24, 1 2012. ISSN 1548-7660. URL <http://www.jstatsoft.org/v46/i03>.
- Shen, D., Liu, N., Li, Y., Wang, D., He, L., and Zhou, Y. *Goldfish Scheme: A Scheme Interpreter with Python-Like Standard Library*. LIII NETWORK, 2024. URL <https://github.com/LiiiLabs/goldfish/releases/download/v17.10.9/Goldfish.pdf>.
- Shi, K., Altunbükten, D., Anand, S., Christodorescu, M., Grünwedel, K., Koenings, A., Naidu, S., Pathak, A., Rasi, M., Ribeiro, F., Ruffin, B., Sanyam, S., Tabachnyk, M., Toth, S., Tu, R., Welp, T., Yin, P., Zaheer, M., Chandra, S., and Sutton, C. Natural language outlines for code: Literate programming in the llm era, 2024. URL <https://arxiv.org/abs/2408.04820>.
- Song, X., Sun, H., Wang, X., and Yan, J. A survey of automatic generation of source code comments: Algorithms and techniques. *IEEE Access*, 7:111411–111428, 2019.
- Steele, G. L. and Sussman, G. J. Lambda: The ultimate imperative. Technical report, Massachusetts Institute of Technology, USA, 1976.
- Steele Jr, G. L. Lambda: The ultimate declarative. Technical report, Massachusetts Institute of Technology, USA, 1976.
- Steele Jr, G. L. Debunking the “expensive procedure call” myth or, procedure call implementations considered harmful or, lambda: The ultimate goto. In *Proceedings of the 1977 annual conference*, pp. 153–162, 1977.
- Steele Jr, G. L. *Rabbit: A compiler for Scheme*. Massachusetts Institute of Technology, 1978.
- Steele Jr, G. L. Compiler optimization based on viewing lambda as rename+ goto, 1980.
- Steele Jr, G. L. and Sussman, G. J. The art of the interpreter of the modularity complex (parts zero, one, and two). Technical report, Massachusetts Institute of Technology, USA, 1978.
- Steele Jr, G. L. and Sussman, G. J. Design of lisp-based processors, or scheme: A dielectric lisp, or finite memories considered harmful, or lambda: The ultimate opcode. Technical report, Massachusetts Institute of Technology, USA, 1979.
- Sussman, G. J. and Steele Jr, G. L. Scheme: A interpreter for extended lambda calculus. Technical report, Massachusetts Institute of Technology, USA, 1975.
- Svyatkovskiy, A., Zhao, Y., Fu, S., and Sundaresan, N. Pythia: Ai-assisted code completion system. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2727–2735, 2019.
- Szafraniec, M., Roziere, B., Leather, H. J., Labatut, P., Char-ton, F., and Synnaeve, G. Code translation with compiler representations. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=XomEU3eNeSQ>.
- Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., Chowdhery, A., and Zhou, D. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- Wang, Y. and Li, H. Code completion by modeling flattened abstract syntax trees as graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 14015–14023, 2021.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E. H., Le, Q. V., and Zhou, D. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.
- Wu, D., Ahmad, W. U., Zhang, D., Ramanathan, M. K., and Ma, X. Repoformer: Selective retrieval for repository-level code completion. *arXiv preprint arXiv:2403.10059*, 2024.
- Xia, C. S., Wei, Y., and Zhang, L. Automated program repair in the era of large pre-trained language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pp. 1482–1494. IEEE, 2023.
- Ye, Q., Axmed, M., Pryzant, R., and Khani, F. Prompt engineering a prompt engineer. *arXiv preprint arXiv:2311.05661*, 2023.
- Zhang, J., Cambronero, J. P., Gulwani, S., Le, V., Piskac, R., Soares, G., and Verbruggen, G. Pydex: Repairing bugs in introductory python assignments using llms. *Proceedings of the ACM on Programming Languages*, 8(OOPSLA1): 1100–1124, 2024.
- Zhou, D., Schärli, N., Hou, L., Wei, J., Scales, N., Wang, X., Schuurmans, D., Cui, C., Bousquet, O., Le, Q., et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022a.

---

Zhou, Y., Muresanu, A. I., Han, Z., Paster, K., Pitis, S., Chan, H., and Ba, J. Large language models are human-level prompt engineers. *arXiv preprint arXiv:2211.01910*, 2022b.

---

## A Jargon

We list commonly used jargon in the following:

- *File*: A script containing code and comments without the structured integration needed for LP (e.g., Jupyter `.ipynb` files)
- *Document*: A script that combines code and descriptive content in a cohesive, narrative style, following LP principles (e.g., Mogan `.tmu` documents). Unlike in a file, sections of a single file can be edited in multiple locations within a document, with other files interspersed in between.
- *Chunk*: A segment of code and comments within a file. Multiple chunks from different files can be organized within a single document.

## B LLM-based code generation

The development of tools based on LLMs, such as ChatGPT, GitHub Copilot, OpenAI Codex, and Amazon Q Developer, has transformed the way programmers solve problems, significantly enhancing efficiency and convenience. These models streamline various coding tasks, including code completion (Guo et al., 2023; Lu et al., 2022; Wang & Li, 2021; Wu et al., 2024; Svyatkovskiy et al., 2019; Fried et al., 2022), code repair (Parasaram et al., 2024; Zhang et al., 2024; Fan et al., 2023; Joshi et al., 2023; Olausson et al., 2023; Jin et al., 2023; Xia et al., 2023), code translation (Chen et al., 2018; Roziere et al., 2020b; Szafraniec et al., 2023), and code comprehension (Nam et al., 2024; Song et al., 2019).

Despite their impressive capabilities, LLMs still face limitations in certain areas, particularly in code generation, where solution inconsistency and specification complexity pose two fundamental challenges. These challenges hinder the development of reliable, scalable software systems through LLM assistance. For instance, Roziere et al. (2020a) highlights the challenges of using LLMs for code translation, demonstrating that these models can introduce bugs during the translation process, especially when dealing with complex or domain-specific code. Dou et al. (2024) highlights the limitations of current LLMs in handling large-scale software projects, where the scale of the project and the interdependencies within the codebase create significant challenges. The study finds that LLMs often struggle to generate functional code for larger-scale projects, and their accuracy decreases as the complexity of these interdependencies increases. Such limitations emphasize the need for more robust frameworks to ensure reliability in large-scale code-related tasks.

### B.1 Challenges of LLM-based code generation

In the following, we list several common challenges for LLM-based code generation:

**Solution inconsistency** Solution inconsistency refers to LLMs' tendency to generate varied implementations for the same problem specification. This inconsistency can manifest in different forms, such as variations in computational complexity, coding style, or error-handling strategies. For example, a list processing task might have an efficient  $O(n)$  implementation using hash tables. However, given a high-level description, LLMs often generate less efficient nested-loop solutions with  $O(n^2)$  complexity. While both solutions may produce correct results, their significant differences in performance characteristics can create challenges in production systems, where performance and reliability are critical.

**API specification complexity** The complexity of modern API specifications presents significant challenges for LLM-based code generation. APIs often encompass multiple interconnected functionalities that must be implemented cohesively. For example, a database connection API involves configuration validation, connection pooling, retry logic, error handling, and resource cleanup. However, LLMs often struggle to comprehend the intricate relationships between these components, leading to incomplete or incorrect implementations. This issue is exacerbated by ambiguous or incomplete documentation, which prevents LLMs from fully understanding the API's intended behavior.

**Inherent LLM limitations** LLMs' inherent limitations, such as fixed context windows and constrained processing time, further compound the challenges of generating accurate code. For example, when dealing with complex API specifications or intricate business logic, the model may truncate or omit essential information, resulting in incomplete code. A common scenario involves error handling logic being omitted when the context exceeds the model's capacity. These limitations necessitate significant manual intervention, reducing the efficiency of LLM-based code generation in complex environments.

---

## B.2 An example of external bias

The following shows an example of the external bias:

- Input:

Question: Implement the "filter" in [[goldfish]]. The "filter" [[procedire]] receives a [[List]] named "lst" and a boolean [[procedire]] "fn" as [[input]]; it checks every element in the "lst" and collects those who have TRUE result by "fn".

Thinking: I lookup the document of [[goldfish]] carefully. I will learn how to write a [[procedire]] on [[List]] according to the [[induction]] method. I will focus on the document. I write it in the steps described by [[induction]].

Answer:

- Output:

```
(define (filter lst fn)
  (define inner-filter
    (lambda (result lst)
      (;zero-step (if (null? lst) result (<succ-step>))))))
  (inner-filter lst 0))
```

One can see that the result build the "Skeleton of Chunk" correctly. However, the generation of "<succ-step>" differs.

## C Discussion of literate programming and Scheme language

### C.1 ILP, Jupyter Notebook, and code with detailed comments

Fig. 1 illustrates the differences between Python LP in Mogan and Python programming in Jupyter Notebook. As shown in Figure 1a, two separate .ipynb files contain code and markdown for different parts of the same project. When using Jupyter, each notebook typically focuses on a single task or dataset, making it difficult to combine multiple code sections or logic threads into one unified file. In contrast, as shown in Mogan (Figure 1b), both code files can be edited in a single .tmu document. Furthermore, Fig. 1b shows three chunks belonging to two different files, all within one document. Once editing is complete, Mogan allows users to export all project files at once, as depicted in Fig. 2. For the same reason discussed above, files with detailed comments, such as the Natural Language Outlines introduced by Shi et al. (2024)<sup>2</sup> are not considered documents.

We now show another example of the difference between code with detailed comments and LP. As discussed in Section J of Knuth (1984), LP enables a flexible structure where code is organized as a "web" of interconnected ideas rather than following a strict top-down or bottom-up sequence. Fig. 3 illustrates this concept with two files connected by ideas. The file `old_packages/sum.scm` contains the original sum function in Scheme, which only accepts integer inputs. The file `new_packages/sum.scm` extends this sum function to accept rational numbers. These two files are linked by ideas rather than function calls or class inheritance. Fig. 3a demonstrates how LLMs can directly recognize these connections, as literate programming allows these chunks to be organized this way. In contrast, Fig. 3b shows the same project with detailed comments. Even with such comments, LLMs must read and interpret the paths mentioned within, locate the referenced file, and then analyze it to grasp the connection.

To distinguish our definition of LP from others, we introduce a modern form of LP called *Interoperable Literate Programming* (ILP). ILP addresses the persistent fragmentation and inconsistency in today's LP practices. Existing approaches rely on tool-specific formats, ad-hoc documentation styles, and inconsistent code extraction methods. These conditions make it difficult to transfer established practices, incorporate new technologies, or maintain compatibility over time. ILP offers a stable, language-agnostic foundation that can adapt to various workflows, tools, and programming languages.

ILP defines a uniform, human-readable markup base that integrates code and documentation. Instead of embedding every feature at a single level, ILP specifies a minimal core schema governing how code, narrative, and metadata coexist. Within this schema, code blocks and their annotations—covering complexity, algorithmic patterns, and related functions—follow a consistent structure. This approach ensures that compilers, documentation generators, analysis tools, and editors can parse and process ILP documents without ambiguity.

---

<sup>2</sup>As noted by Shi et al. (2024), the Natural Language Outlines lack a utility to export all project files cohesively. This limits alignment with the LP philosophy discussed in Section J of Knuth (1984).



---

Building on its core specification, ILP supports incremental adoption of optional layers. These layers introduce capabilities—such as typed code sections, debugging aids, or hints for large language models—without altering the underlying standard. This arrangement preserves interoperability, since tools that do not require certain features can ignore them, and new layers can emerge without forcing widespread rewrites. Developers, educators, and researchers gain the flexibility to select enhancements that suit their context while relying on a well-defined core.

ILP also encourages formal interoperability contracts that describe how code extraction, documentation, and compilation occur for specific languages or frameworks. By establishing these contracts alongside the base specification, ILP makes it easier to introduce new languages, connect existing codebases, and update tooling. This structure helps avoid the proliferation of incompatible formats and supports incremental refinement as technology evolves.

## C.2 Thinking in interoperable literal programming

The philosophy of using ILP is not to adopt an entirely new way of thinking but to transfer one’s natural thought process into LP. Software and software development should aim to provide an intuitive user experience, aligning with the brain’s natural thinking patterns, whether sequential or otherwise. Consequently, the impact of software lies not in the new features it offers but in how well it preserves and supports users’ original ways of thinking.

Using some cases in prompt engineering as examples, scholars try to understand the intelligent thinking process of LLM; as a result, they provide many different methods. The observation of the CoT and other prompts engineering methods provides us with an important idea that the continuous working flow is of great importance for a long-lasting task. Indeed, sequential thinking, like the CoT, and tree-like or more complex patterns need the software to have enough flexibility. Mogan provides writing with code chunks at any place, which is the key to combining thinking with writing. In other words, the whole spectrum of thinking and writing flow in the human brain can be preserved perfectly according to the Literal Programming based on Mogan.

In the following sections, we propose a structured approach for documentation and code organization in Goldfish Scheme, aimed at enhancing LLM performance and understanding. We adopt the classical CoT pattern to capture both the final code and the thought process behind it, ensuring that documentation covers not only what the code does but why specific decisions were made. This involves a functional-programming-inspired template that restricts direct recursion and uses zero-step/succ-step induction. We then extend this strategy with a unique naming system and cross-file references to test whether LLMs rely on our explicit documentation instead of pre-trained knowledge. We can further assess the model’s ability to infer relationships and maintain context by splitting functionality across multiple files and introducing dependency loops.

## C.3 Literate programming supported editor

Knuth (1984) introduce the concept of LP and implement LP in the “WEB” system, which leads the  $\text{\TeX}$  program itself, brought to life as a narrative in the book “ $\text{\TeX}$ : The Program”. Then, in 1994, Ramsey (1994); Johnson & Johnson (1997) develop a simplified version, “NoWEB” which supports not only  $\text{\TeX}$  but also other text formatting systems, such as  $\text{\LaTeX}$  and HTML. In addition to “NoWEB”, several modern literate programming systems have been developed, such as “FunnelWEB”, “NuWEB”, “InWEB”, “PyWEB”, “Org mode” (Dominik, 2010; Schulte et al., 2012; Schulte & Davison, 2011), “Fangle” and “Jupyter Notebook” (Kluyver et al., 2016) among others. Additionally, Axiom embodies literate programming, linking code with its rationale for clarity. In this paper, we develop Mogan, a UTF-8-compatible literate programming editor designed to simplify the workflow by allowing users to edit and export all project files within a single document, as illustrated in Fig. 1 and Fig. 2.

## C.4 Scheme language

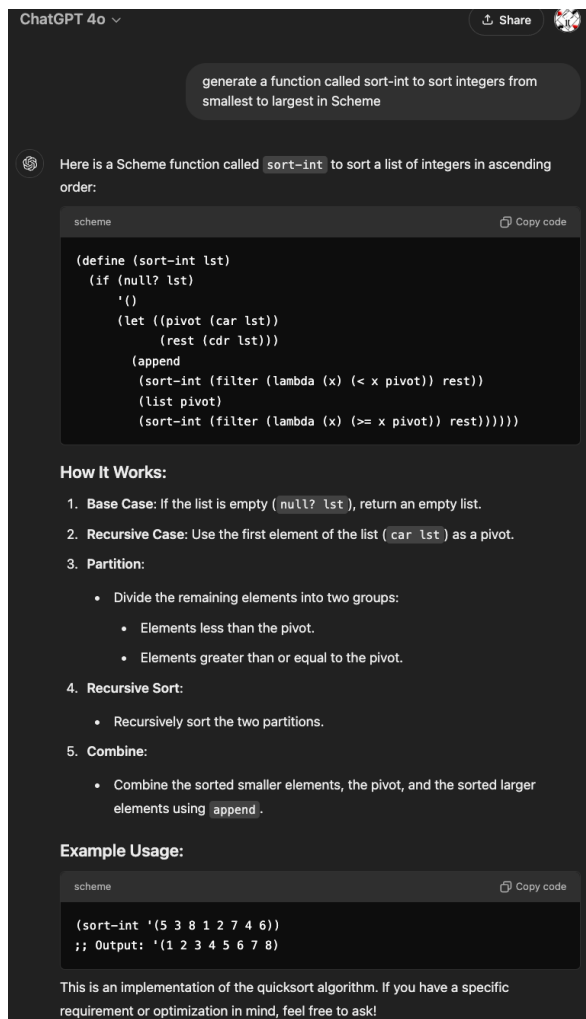
Scheme language is known for its minimal core and uniform syntax. Its straightforward constructs, first-class procedures, and lexical scoping encourage a focus on program logic rather than language intricacies. This simplicity makes Scheme an effective foundation for exploring language design and building upon established standards.

Sussman & Steele Jr (1975) implement “Scheme”, a simple and lightweight functional programming language based on lambda calculus (Steele & Sussman, 1976; Steele Jr, 1976; 1977; Steele Jr & Sussman, 1978; Steele Jr, 1978; Steele Jr & Sussman, 1979; Steele Jr, 1980), which becomes widely known through the popular book Structure and Interpretation of Computer Programs (SICP) (Abelson et al., 1996). Furthermore, several Scheme interpreters have been developed based on the Scheme standard, with the latest and most widely accepted being R7RS. Additionally, Scheme Requests for

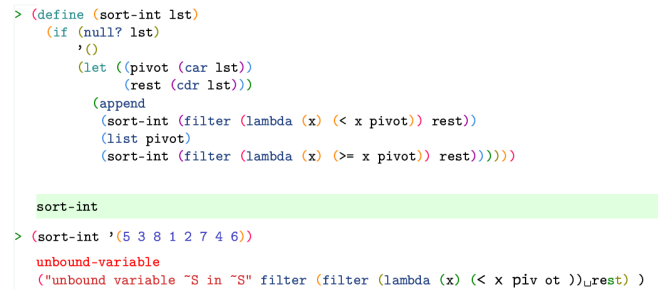
Implementation (SRFI) are community-driven specifications that extend the language with additional features and libraries. Among the Scheme interpreters, Chez Scheme (Dybvig, 2005; 2009) is a high-performance Scheme interpreter, known for its efficient compiler and extensive features. However, it only supports the R6RS standard. Several other implementations support the R7RS standard, including MIT/GNU Scheme, which provides a complete environment with an interpreter, compiler, and debugger (Hanson & of Technology. Scheme Team, 2015), Gambit Scheme, and STklos, among others. A descendant of Scheme, S7 Scheme, does not follow the R7RS standard but is as minimal as possible. Our Goldfish Scheme (Shen et al., 2024), built on S7 Scheme, supports the R7RS standard, parts of SRFI, and many functionalities inspired by the Python standard library. This simplicity, combined with several characteristics of Python, makes it easier for large language models to understand.

**An example of incapability of LLMs in Scheme language generation** Fig. 7 provides an example where ChatGPT-4 fails to generate a simple Scheme code chunk. The prompt is as follows:

“generate a function called sort-int to sort integers from smallest to largest.”



(a) The code chunks generated by ChatGPT4o.



(b) The code chunk cannot be executed because the `filter` function required by the LLMs does not exist in R7RS.

Figure 7: LLMs fail to generate a simple function in Scheme.

**Why use Scheme for documentation?** Scheme is particularly well-suited for documenting API logic due to its minimalistic syntax and mathematical clarity. With its extremely simple and high-level primitive procedures, Scheme makes it relatively

---

easy to design a domain-specific language tailored for testing LLM capabilities. Unlike more commonly used programming languages, Scheme is less prevalent in LLM training datasets, which significantly reduces the likelihood of LLMs referencing pre-trained knowledge or relying on external search capabilities. This limitation is crucial when using LLMs with online search features, such as ChatGPT-4, which may skip the design details provided in the documentation and substitute them with similar but undesired implementations. By encoding logic in Scheme, we ensure that LLMs adhere strictly to the documented design rather than relying on potentially inaccurate external references.

Scheme also provides an intuitive and concise way to represent mathematical and logical structures in  $\lambda$ -calculus, making it an effective tool for documenting complex API designs. When paired with the advanced toolkits available in Mogan, developers can express intricate logic efficiently, reducing manual work while maintaining high precision.

### C.5 Comparing FP, OOP and LP

Functional programming (FP) emphasizes immutability, pure functions, and declarative code structures rooted in lambda calculus (Barendregt et al., 1984). Its strengths include the ability to compose reusable functions and effectively manage side effects, making it particularly suitable for data pipelines and parallel computing. Functional programming languages widely used today in industry and education include Common Lisp, Scheme, Clojure, Wolfram Language, Racket, Elixir, and Haskell. Additionally, Lean is a functional programming language commonly employed for verifying mathematical theorems. Despite its advantages, FP poses challenges such as debugging complex abstractions and increased memory usage due to immutable data structures.

Object-oriented programming (OOP) organizes software into objects containing state and behavior, offering modularity and extensibility through encapsulation, inheritance, and polymorphism (Gamma, 1995). It dominates in enterprise systems, GUIs, and game development but can lead to performance bottlenecks or over-abstraction when misused.

Literate programming, introduced by Donald Knuth in 1984, is a programming paradigm that combines code and natural language documentation into a coherent narrative, explaining the logic and intent behind the code (Knuth, 1984). Unlike traditional programming approaches that separate documentation and code, literate programming integrates them, enabling developers to produce both executable code and human-readable documents. This paradigm is particularly effective for algorithm-heavy applications, teaching, and research-oriented projects due to its ability to present code logically for human understanding. However, its detailed approach can be time-intensive and less scalable for large or rapidly evolving projects, compounded by limited modern tooling.

Functional programming (FP) and object-oriented programming (OOP) focus on scalability, performance, and modularity, catering to complex and diverse applications. In contrast, literate programming (LP) emphasizes deep understanding and clear documentation by combining code with detailed narrative explanations. LP is especially useful in research, teaching, and algorithm-intensive projects, where clarity and human readability are crucial.

## D Goldfish Scheme

In this work, we develop a Scheme interpreter with a Python-like standard library, *Goldfish Scheme*, based on S7 Scheme<sup>3</sup>. If we use Goldfish Scheme as the target code generated by LLMs, it has the following advantages:

- *Simplicity for Inference*: The Scheme’s minimalistic design limits the options for inference, allowing LLMs to provide more precise and straightforward answers.
- *Python Compatibility for Knowledge Transfer*: LLMs are generally more familiar with Python due to extensive training in Python code. Since the Goldfish Scheme includes several features from Python’s standard library, it enables smoother knowledge transfer from Python to the Goldfish Scheme, making it more accessible for LLMs than other Scheme variants.

---

<sup>3</sup>Goldfish Scheme is named with a playful reference to the idea that goldfish have a 7-second memory, the inverse of S7. In the future, we hope LLMs can help create new functions for Goldfish Scheme within those same 7 seconds.

---

## D.1 Current features

**LLM interaction layer** A defining innovation of the Goldfish Scheme is its dedicated LLM interaction layer, designed to support richer, more context-aware interactions with large language models. This layer provides a structured interface for embedding metadata, documentation, and contextual hints directly into the codebase. Instead of relying on ad-hoc comments or detached documentation, developers can annotate their code so that LLMs can systematically parse and leverage these annotations. The core idea behind this interaction layer is to supply machine-readable fields that describe key aspects of a function or library component. These annotations can specify algorithmic patterns, complexity metrics, stability characteristics, and concrete usage examples. By making this information explicit and accessible, the interaction layer allows LLMs to understand each piece of code not only at a syntactic level but also in terms of conceptual intent and practical implications.

Consider the following `define-with-docs` macro:

```
(define-with-docs quicksort
  #:pattern "divide-and-conquer"
  #:complexity "O(n log n)"
  #:stability "unstable"
  #:examples
  '((quicksort '(3 1 4 1 5 9 2 6 5 3))
    => (1 1 2 3 3 4 5 5 6 9))
  (lambda (lst)
    ;; Implementation follows...
  ))
```

In this example, the “divide-and-conquer” pattern informs the LLM of the underlying strategy behind the algorithm, making it easier for the model to relate this sorting routine to other similar functions and patterns it knows. Declaring the complexity as  $O(n \log n)$  provides the LLM with a direct indication of the performance profile, which can help when it needs to explain cost trade-offs or select suitable approaches for different problem sizes. By stating that the sort is “unstable,” the annotations reveal a critical property that distinguishes the function’s behavior from stable alternatives, guiding the LLM toward more precise reasoning about potential use-cases and limitations. The inclusion of a concrete usage example, showing how the function transforms a given list into a sorted sequence, allows the LLM to confirm its understanding of inputs, outputs, and expected transformations. Such an example can also serve as a reference point for generating further tests or variant solutions.

By embedding these structured annotations, the LLM interaction layer encourages the model to form a richer mental model of the code. This approach standardizes how information is presented, mitigating inconsistencies in commentary style or documentation quality. It ensures that both human developers and LLMs have a common, accessible source of truth regarding how code is organized, why certain decisions are made, and what is expected of each function. This synergy between code and metadata lays the groundwork for advanced capabilities, such as more coherent automated explanations, more accurate reasoning about algorithm selection, and better integration of language models into the daily workflows of software developers.

**Core language architecture** Goldfish Scheme adopts the principle of minimalism of Scheme language (discussed in Sec. C.4) by implementing the R7RS standard through a layered architecture. Instead of merging new features directly into the core, it places additional capabilities on top of a stable R7RS layer. This approach preserves backward compatibility, allowing existing R7RS programs to run without modification. At the same time, it provides a structured path for adding performance optimizations, memory-efficient data structures, and integrations that do not disrupt the underlying standard.

The layered design also supports targeted performance improvements. By isolating advanced functionality in separate layers, the Goldfish Scheme reduces unintended interactions and makes it possible to switch off unneeded components, maintaining low memory overhead. Furthermore, this clear separation simplifies maintenance and future development. Changes to one layer do not ripple into others, and developers can refine or replace specific parts of the system without affecting its core compliance.

**Standard Library Extensions and Python Integration** Goldfish Scheme enhances the R7RS standard library by integrating Python-inspired features while adhering to Scheme’s functional programming principles. These extensions

---

provide practical tools for tasks such as data structure operations, string manipulation, and mathematical computations, without compromising the language's theoretical purity.

For example, consider the following list processing pipeline:

```
;; Traditional R7RS approach:
(define (process-data lst)
  (filter number?
    (map (lambda (x) (* x 2))
      (remove zero? lst))))

;; Goldfish Scheme enhanced approach:
(define (process-data lst)
  (-> lst
    (remove zero?)
    (filter number?)
    (map (curry * 2))))
```

This example demonstrates how Goldfish Scheme improves readability and maintainability while preserving efficiency. The threading macro (`->`) facilitates linear operations, making code easier to understand and modify.

**Advanced Type System Implementation** Goldfish Scheme extends the R7RS type system with gradual typing, combining compile-time safety with runtime flexibility. This system supports custom type definitions, contracts, and validation, enhancing safety without sacrificing the dynamic nature of Scheme.

```
;; Type definition with constraints
(define-type Vector2D
  (make-vector2d [x :: (Real (lambda (n) (finite? n)))]
    [y :: (Real (lambda (n) (finite? n)))]))

;; Function with type annotations and contracts
(define/contract (vector-magnitude [v :: Vector2D]) :: Real
  (sqrt (+ (square (vector2d-x v))
    (square (vector2d-y v)))))
```

This approach ensures correctness through type annotations and runtime validation while maintaining Scheme's lightweight and expressive characteristics.

**Memory Management and Performance Optimization** Memory management in Goldfish Scheme improves upon R7RS implementations through several key innovations. First, we implement a generational garbage collector specifically optimized for functional programming patterns. This collector recognizes common Scheme programming patterns and optimizes memory allocation and collection accordingly. Second, we introduce a sophisticated memory pooling system for small objects, reducing allocation overhead for common operations.

```
;; Memory-efficient string handling example
(define (process-large-text text)
  (with-memory-pool
    (lambda ()
      (-> text
        string->lines
        (filter non-empty-line?)
        (map process-line)
        lines->string))))
```

These optimizations reduce memory usage by up to 40% in typical workloads, significantly improving efficiency for text and data processing.

## D.2 Future features

We also outline several features currently under development for the Goldfish Scheme.

---

**Integration with modern development tools** Goldfish Scheme integrates with modern development workflows by providing built-in tools for project management, dependency handling, and automation. These features streamline development processes and support comprehensive tooling.

```
;; Project configuration in Goldfish
(define-project my-application
  #:name "My Application"
  #:version "1.0.0"
  #:dependencies
  '((goldfish-web "2.0.0")
    (goldfish-database "1.5.0"))
  #:entry-point main.scm)

;; Automatic development environment setup
(define-development-environment
  #:test-framework 'goldfish-test
  #:documentation-generator 'goldfish-docs
  #:lsp-server 'goldfish-lsp)
```

This integration provides a complete development ecosystem that significantly reduces the friction typically associated with Scheme development.

## E Step-wise methodology

A precise and logically connected documentation design is essential for effectively explaining API logic and guiding LLMs in generating accurate code. We propose a documentation framework that uses a step-wise methodology to address key challenges discussed in Sec. 2.1 on LLM-based code generation. This approach minimizes reliance on external references, ensures logical consistency, and reduces errors during the code generation process.

To avoid the pitfalls of long prompts and extended reasoning chains, we adopt a step-wise methodology. The idea is inspired by mathematical induction and recursive algorithms, where large tasks are divided into base and inductive parts. For example, in recursive algorithms, establishing a clear base case and carefully explaining the recursive step ensures consistency in the final output. Similarly, our step-wise method divides a task’s logic into smaller, sequential parts and provides LLMs with a clear, incremental path through each aspect of the code rather than overwhelming them with entire program logic at once. By breaking down complex tasks into smaller components, we ensure that the LLM can focus on understanding each step in isolation before progressing to the next.

This granularity also gives developers better oversight, making it easier to spot mistakes and confirm correctness at each stage. In summary, the step-wise approach yields several benefits:

- **Reduced prompt size:** Smaller steps decrease the likelihood of exceeding token limits or overwhelming the model.
- **Enhanced logical flow:** Explicit sequencing helps LLMs follow a coherent reasoning path.
- **Simplified debugging:** Isolating errors within individual steps eases testing and correction.

In practice, the step-wise method can be decomposed into:

1. **Zero-step definition:** We start by specifying a *base case* or fundamental building block. In recursive functions, for example, this might be the condition under which the recursion terminates. By designating this zero-step, we give the LLM a clear anchor point for the broader logic.
2. **Successor-step explanation:** We then describe how one logical state transitions to the next. In inductive or recursive algorithms, this involves detailing the “successor” or “next element” step, building on the zero-step to maintain consistency.
3. **Code chunks per step:** Each step’s code is presented in a self-contained chunk, accompanied by an explanation of its function (e.g., base case, recursive step, edge cases). This structure helps the LLM associate specific logic with the corresponding implementation.



- 
4. **Local Verification and Iteration:** After a chunk is introduced, we encourage local testing (for example, in Mogan’s REPL) to confirm correctness. If an error occurs, developers can isolate it to that chunk and provide the LLM with focused feedback, minimizing the spread of mistakes.
  5. **Incremental Complexity:** Once a base version passes verification, we add optional features or optimizations in smaller, separate steps. Each new layer is validated independently, preventing confusion that can occur when many changes happen at once.

For example, in Scheme

- **The Logic of computation in Scheme:** All logic is defined in Scheme for precise, induction-friendly representation. For instance:

```
(define (compute-factorial n)
  (if (= n 0)
      1
      (* n (compute-factorial (- n 1)))))
```

This example clearly illustrates a recursive function for computing factorials.

- **Step-wise breakdown:** The documentation divides the logic into incremental steps. For instance:
  - **Step 1:** Define base case conditions, such as  $n = 0$  for termination.
  - **Step 2:** Outline recursive relationships, such as  $n \times (n - 1)!$ .
- **Integration with Mogan toolkits:** The documentation is fully compatible with Mogan’s advanced features, allowing developers to efficiently visualize, test, and validate API logic. Additionally, Mogan offers REPL functionality, enabling developers to verify code accuracy locally.

We adopt a step-wise methodology to mitigate issues related to long prompts and extended reasoning times in LLMs. This approach breaks down API logic into smaller, sequential steps, which allows LLMs to process information incrementally. Leveraging the memoization capabilities of LLMs, this method ensures that earlier steps are retained and referenced as the model progresses through the reasoning process.

The step-wise methodology resolves key challenges in LLM-guided code generation by dividing complex logic into smaller, manageable steps. One primary advantage is its ability to overcome prompt length constraints, which often hinder LLM performance. Long and complex prompts can overwhelm models, leading to incomplete or inconsistent outputs. Moreover, exceeding the token limit of an LLM service may result in truncated responses and loss of critical details. By presenting logic incrementally in smaller segments, the model processes each step independently, maintaining focus on specific tasks. This sequential approach minimizes errors caused by information overload and ensures all relevant details are effectively communicated.

A significant benefit of the step-wise methodology is the improvement it brings to logical flow. LLMs reason more effectively when guided through a structured progression of steps. For example, in recursive logic, the methodology starts by defining the base case and progresses to explain recursive relationships and edge-case handling. This step-by-step structure allows the model to build a comprehensive understanding of the logic, reducing ambiguities and deviations. Explicitly documenting each step ensures that the generated code aligns closely with the intended design, minimizing variations and improving consistency.

The modular nature of the step-wise methodology simplifies debugging and testing by isolating specific components of the logic. Complex logic often contains errors localized to particular sections, which this approach makes easier to identify and correct. For instance, if an error occurs in the base case of a recursive function, developers can focus on that isolated step without analyzing the entire code block. Additionally, the modular structure facilitates regression testing, as updates to individual steps can be validated independently without risking unintended effects on other parts of the logic. This approach streamlines debugging and supports the development of higher-quality code.

The step-wise methodology provides a practical and effective framework for guiding LLMs in code generation by addressing prompt limitations, enhancing logical reasoning, and simplifying error resolution. Its structured approach ensures that both

---

developers and LLMs can produce accurate and consistent code while reducing the time and effort required for debugging and validation.

In summary, the step-wise methodology provides several advantages:

- **Bypassing prompt length constraints:** Dividing logic into smaller steps reduces the need for lengthy prompts, which often lead to inconsistencies or errors.
- **Enhanced logical flow:** Sequential steps guide LLMs in reasoning systematically, reducing variations in generated code.
- **Streamlined debugging:** Smaller, modular steps make it easier to identify and address errors in specific parts of the logic.

This enhanced documentation design directly tackles challenges in using LLMs for code generation:

- **Minimized external bias:** By using Scheme, we reduce the reliance on pre-trained knowledge and external references, ensuring that LLMs generate code strictly based on the provided documentation.
- **Improved accuracy and consistency:** The step-wise methodology enhances the logical flow of LLM reasoning, reducing errors and variations in generated code.
- **Reduced developer effort:** Coupled with Mogan toolkits, the proposed framework simplifies the process of documenting, testing, and validating API logic.

## F RepoBench

### F.1 RepoBench Adaptation for Scheme Language

We adapt RepoBench’s evaluation framework to support Scheme-based development through three core modifications: repository structure transformation, evaluation metric redefinition, and prompt engineering adjustments.

First, we transform RepoBench’s repository structure to align with Scheme’s functional paradigm. RepoBench organizes test cases around class-based implementations, typically using a main implementation file with derived classes. We restructure this to fit Scheme’s module system, where each unit consists of a core implementation file and its documentation. For example, a Python class hierarchy for a data structure implementation:

```
class BinaryTree:
    def insert(self, value): ...
    def delete(self, value): ...

class AVLTree(BinaryTree):
    def balance(self): ...
```

is restructured into a Scheme module:

```
;; tree-implementation.scm
(define (make-tree) ...)
(define (tree-insert tree value) ...)
(define (tree-delete tree value) ...)
(define (tree-balance tree) ...)
```

This preserves functionality while adhering to Scheme’s functional style.

Second, we redefine RepoBench’s evaluation metrics to better capture functional programming characteristics. The original metrics focus on object-oriented features like inheritance and encapsulation. We replace these with metrics evaluating function composition, state management via immutable data structures, and recursive correctness. New criteria specific to Scheme include:

- 
- **Pure function analysis:** Ensures functions maintain referential transparency.
  - **Recursion patterns:** Measures proper use of tail recursion and recursive helpers.
  - **Data immutability:** Verifies that functions produce new data structures instead of modifying inputs.

Third, we adjust RepoBench’s prompt templates to reflect Scheme’s syntax and functional programming patterns. Original prompts focused on class design are reformulated to emphasize function specifications and data flow. For example, a prompt requesting:

*“Create a class that implements a priority queue with methods for insertion and deletion”*

is transformed into a functional specification requesting pure functions for a priority queue:

- A constructor returning an empty queue.
- An insertion function producing a new queue with the added element.
- A deletion function returning a tuple of the highest-priority element and the updated queue, ensuring queue invariants are maintained without modifying the input.

We also enhance the evaluation framework to assess literate programming by adding metrics for documentation completeness, integration with code, and adherence to Step-wise methods specified in the prompt. The modified evaluation pipeline processes test cases to measure functional correctness, programming pattern adherence, invariant maintenance, and documentation quality.

This adaptation enables rigorous evaluation of language models’ Scheme code generation while maintaining RepoBench’s standards. The results provide quantitative measures of correctness and functional programming compliance, supporting the evaluation of literate programming techniques in Goldfish Scheme.

## F.2 Limitations of Alternative Benchmarks

The selection of RepoBench over other code generation benchmarks stems from its unique repository organization and evaluation methodology. While benchmarks like HumanEval and MBPP focus on isolated programming problems, RepoBench’s repository structure mirrors real-world software projects, containing multiple interconnected files with documentation, tests, and implementation code. This structure aligns with our literate programming approach, where documentation and code form a cohesive narrative.

RepoBench’s repository examples exhibit several characteristics that make them ideal for adaptation to Goldfish Scheme. Each repository contains a complete project structure with:

```
project/
|- src/
|  |- core/
|  |  |- implementation.py
|  |  |- interfaces.py
|  |- tests/
|  |  |- test_core.py
|  |- docs/
|     |- design.md
|     |- api.md
```

This structure maps naturally to our literate programming approach in Goldfish Scheme:

```
project.tmu
|- (core-implementation)
|  |- mathematical-foundation
|  |- implementation
```

- 
- |- (test-cases)
  - |- (documentation)
    - |- design-rationale
    - |- api-specification

Other benchmarks present significant limitations for our research. CodeXGLUE focuses primarily on code completion tasks and lacks comprehensive project context. While it supports multiple programming languages, its evaluation metrics center on token-level accuracy rather than functional correctness or documentation quality. Context-QA, despite handling repository-level queries, emphasizes code comprehension over generation and lacks support for evaluating literate programming principles.

Traditional benchmarks like APPS and MBPP provide extensive test cases but operate in isolation. APPS contains 10,000 programming problems with test cases, but these problems lack project context and documentation requirements. Similarly, MBPP's 974 Python programming problems focus on algorithmic correctness without considering broader software engineering principles. These limitations make them unsuitable for evaluating our approach, which requires:

1. Project-level context understanding
2. Documentation generation assessment
3. Cross-file dependency handling
4. Integration with existing codebase patterns

RepoBench's repositories also provide real-world examples of API evolution, which aligns with our goal of supporting API development through literate programming. Each repository includes multiple versions of APIs, showing how they evolve over time. This feature enables us to evaluate how well language models understand and extend existing APIs, a crucial capability for our literate programming approach.

Consider a typical RepoBench API evolution example:

```
// Version 1
class DataProcessor {
    process(data) { ... }
}

// Version 2
class DataProcessor {
    process(data, options) { ... }
    validate(data) { ... }
}
```

This maps to our Scheme-based literate programming approach:

```
;; Version 1
(define (process-data data)
  "Step-wise method: ..."
  (implementation ...))

;; Version 2
(define (process-data data options)
  "Extended Step-wise method: ..."
  (implementation with validation ...))
```

Furthermore, RepoBench's evaluation framework includes metrics for assessing code generation in the context of existing documentation, making it uniquely suitable for evaluating our literate programming approach. While other benchmarks might evaluate code correctness, they lack mechanisms for assessing how well generated code aligns with existing documentation or maintains project-wide consistency.

---

The combination of these features - structured repositories, API evolution examples, and documentation-aware evaluation metrics - makes RepoBench uniquely suitable for our research. Other benchmarks, while valuable for their intended purposes, lack the necessary framework for evaluating literate programming principles and functional programming patterns in a project context.

## G Experiments

In this section, we provide additional experimental details highlighting the advantages of using Scheme and propose a programming pattern to improve code generation. Our research yields two significant insights:

1. LLMs can effectively generate code within specific stylistic constraints when properly guided.
2. ILP Documentation structure significantly enhanced LLMs' ability to adhere to prescribed patterns.

### G.1 Machine details

Our experiments are conducted on two types of LLMs: local open-source LLMs and remote commercial LLMs.

**Remote commercial LLMs** For remote commercial LLMs, queries are sent through a high-speed network interface with a 2Gbps connection over a Tier 1 IP network. The setup ensures low latency during query processing, with an average round-trip latency of less than 12.117 milliseconds. This configuration minimizes network-related delays, providing a reliable environment for benchmarking the performance and responsiveness of remote LLMs.

**Local-deployed LLMs** For the local LLMs, the experiments are performed on a high-performance machine equipped with an AMD 7950X processor, 128GB of RAM, and two NVIDIA RTX 4090 GPUs, ensuring sufficient computational power for model inference and evaluation. The models we deployed are LLAMA 3.3, qwen-32B, and Llama 3.1 2B.

### G.2 Experimental Details

We conducted extensive testing to evaluate how well LLMs can follow our literate programming approach. Our experiments focused on two key aspects: function implementation success rates and the impact of function naming on model behavior.

**Function Implementation Results** Table 1 shows the performance of different models across various Scheme functions. We tested a range of operations, from basic list manipulations like `make-list` and `null-list?` to more complex functions such as `delete-duplicates` and `split-at`. Each function underwent testing under three distinct conditions: with GPT-4o, with GPT-4, and without file context. The success patterns in our results suggest that simpler list operations are more consistently implemented across all conditions, likely because these operations have more straightforward implementations with fewer edge cases to consider.

**Function Naming Effects** We discovered that function names significantly influence model behavior. With standard function names, models relied on the provided file 70% of the time while drawing from their pre-trained knowledge in the remaining cases. This split likely occurs because familiar names like 'map' or 'filter' trigger the models' pre-existing knowledge of standard Scheme implementations. However, when we replaced familiar function names with uncommon alternatives, the models showed a marked increase in their reliance on our documentation. Figure 8 illustrates this effect. The improved adherence to our documentation with unconventional names suggests that reducing familiar reference points forces models to rely more heavily on the provided context rather than their training data.

**Quality of Generated Functions** Our analysis revealed distinct differences in function quality based on the context provided to the models. Functions generated with proper file context consistently showed superior performance, passing more test cases and following the prescribed inner procedure patterns. This superior performance can be attributed to three key factors. First, our documentation provides explicit patterns for handling edge cases, which models can directly reference. Second, the inner procedure pattern we specify naturally leads to more robust implementations by separating core logic from interface concerns. Third, the Chain of Thought documentation helps models understand the reasoning behind each implementation decision. Conversely, functions generated without file context exhibited higher failure rates, often struggling with basic read operations and lacking proper error-handling mechanisms. This degradation in quality likely occurs because

Fully based on the file, generate a function `make-list n [fill]` -> list  
Returns an n-element list, whose elements are all the value `fill`. If the `fill` argument is not given, the elements of the list may be arbitrary values.

Fully based on the file, generate a function `zuo-liebiao n [fill]` -> list  
Returns an n-element list, whose elements are all the value `fill`. If the `fill` argument is not given, the elements of the list may be arbitrary values.

Based on the file, here's the implementation of the `make-list` function in Scheme. This function constructs a list of a specified length `n`, with all elements being the value `fill`. If `fill` is not provided, the elements may have arbitrary values.

```
scheme
(define (make-list n . fill)
  (define (inner-make-list result n fill-value)
    (if (= n 0)
        result
        (inner-make-list (cons fill-value result) (- n 1) fill-value)))
  (let ((fill-value (if (null? fill) #f (car fill)))) ; Default fill is #f if not
    (inner-make-list '() n fill-value)))
```

Based on the content of the document, here's the implementation of the `zuo-liebiao` function in Scheme that generates an n-element list with all elements equal to `fill`. If `fill` is not provided, the elements of the list will have arbitrary values.

```
scheme
(define (zuo-liebiao n . fill-args)
  (define fill-value
    (if (null? fill-args) 'arbitrary (car fill-args)))
  (define (inner-zuo-liebiao result count)
    (if (= count 0)
        result
        (inner-zuo-liebiao (cons fill-value result) (- count 1))))
  (inner-zuo-liebiao '() n))
```

Figure 8: Impact of Function Name Changes on Model Behavior

models fall back on simplified or generic implementations from their training data, which may not account for the specific requirements and edge cases of our system.

**Evaluate the code generation ability of LLMs** The ability of LLMs to generate code can be divided into two parts: the original intuition part and the implementation part. For the intuition part, just like checking a kid’s understanding of some problems, relying on the detailed template, we can check the ability of LLMs in programming-style maintenance. We first created a useful prompt template for models to generate a specified style code. Then, we observed some improvement in originality, based on our template, in small-scale project code generation.

For the implementation part, we claim that the ability to implement algorithms can be observed in creating step-wise style codes. Furthermore, we use the essential list of actions in functional programming as being of current importance. We list some fundamental procedures to check its creativity. Besides, based on our framework, one can design more tests on “transfer learning.” In future work, we will also try more experiments on this topic. We believe that the key to checking the ability of LLMs to write code is never simply to observe the accuracy in numbers; instead, we need more high-level intelligence checks like the ability to fill up the proof of mathematical induction (to generate the code in Functional Programming).

### G.3 Enhanced API design through Goldfish Scheme

Our adaptation of RepoBench introduces a hybrid testing methodology that reflects real-world software development practices. Rather than converting all repository code to Scheme, we strategically maintain some APIs in Python/Java while implementing others in Scheme. This approach mirrors actual development scenarios where teams design APIs before implementation, especially in systems that interface with multiple languages or require careful architectural consideration.

Consider a typical microservices architecture where different components may be implemented in different languages. A team might design core algorithmic components in Scheme for its functional purity while maintaining service interfaces in Python or Java. For example:

```
# Original Python API
class DataProcessor:
    def transform_data(self, input_data):
        # Implementation details pending
        pass

    def validate_format(self, data):
        return self._core_validator.check(data)
```

In our hybrid approach, we maintain the Python interface but implement the core validation logic in Scheme:

```
;; Core validation implementation in Scheme
(define (core-validator data)
  "Step-wise method:
```



---

Input data must satisfy properties P1...Pn  
Output ensures invariants I1...In

Implementation Strategy:

1. Type validation
2. Structural checks
3. Domain-specific validation"

```
(let ((inner-validate
      (lambda (data acc)
        ;; Implementation pending
        )))
    (inner-validate data '()))
```

This hybrid structure serves multiple purposes. First, it enables developers to leverage Scheme's Step-wise methods and functional purity for critical algorithmic components while maintaining compatibility with existing systems. Second, it allows teams to document and verify core algorithms through literate programming before finalizing implementation details in target languages.

Our approach particularly benefits API design in three scenarios:

1. Algorithm Prototyping: Teams can express complex algorithms in Scheme with mathematical rigor before implementing them in production languages
2. Interface Design: Developers can maintain existing interfaces while evolving core implementations
3. Documentation First: Teams can use literate programming to fully specify behavior before committing to implementation details

The effectiveness of this hybrid approach becomes evident in testing scenarios. When generating code, language models must understand both:

```
# Python interface specification
def process_sequence(data: List[Any]) -> List[Any]:
    """
    @requires: Implementation in core_processor.scm
    @ensures: Output satisfies sorting invariants
    """
    pass

;; Scheme implementation
(define (core-processor lst)
  "Invariants maintained:
  1. Order preservation for equivalent elements
  2. Linear time complexity
  3. Space complexity O(n)"
  (implementation ...))
```

This structure reflects real development workflows where teams might: 1. Design APIs in familiar languages 2. Document core algorithms in Scheme for mathematical clarity 3. Implement and verify critical components functionally 4. Bridge implementations through interface layers

By maintaining this hybrid structure in our benchmark, we create test scenarios that better reflect real-world development challenges. Language models must demonstrate understanding of both implementation languages and the relationships between components. This approach leads to more meaningful evaluation of code generation capabilities in practical development contexts.

Furthermore, our method addresses a common challenge in API development: the gap between design and implementation. By using Scheme's functional patterns and literate programming for core components, we provide clear specifications that guide implementation in any target language. This approach has shown improved accuracy in generated code, as models better understand the underlying algorithmic intentions and invariants.

## H Advantages of literate programming in modern development

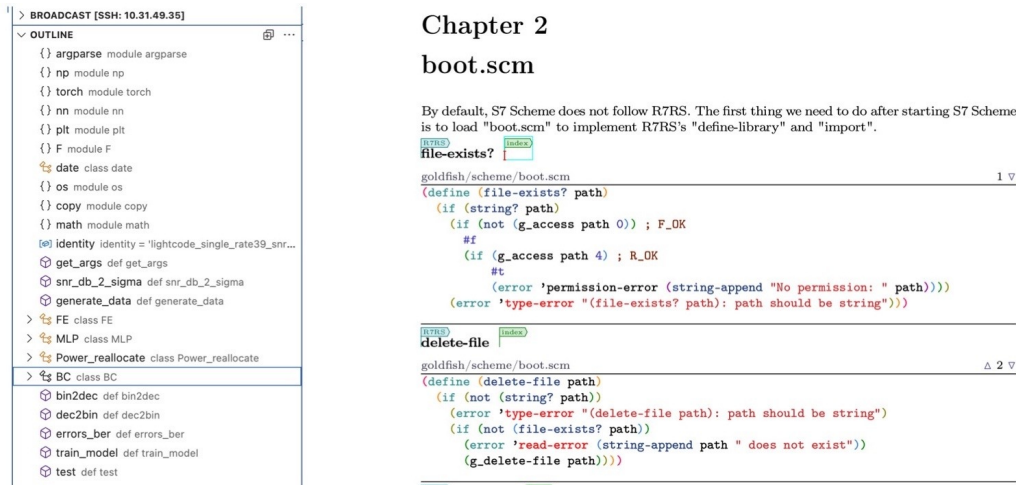


Figure 9: Comparison of code indexing methods: Visual Studio Code outline (left) and Mogan (right)

### H.1 Enhanced accessibility through natural language integration

Our LP approach offers something fundamentally different from Shi et al. (2024). Their work focuses on helping experienced developers understand code quickly. In contrast, our method makes programming accessible to everyone, including those with no coding experience.

The combination of natural language documentation and LLM assistance is central to our approach. Users describe their programming goals in plain language. The LLM then helps convert these descriptions into structured code that follows the program’s logical design. This creates a direct path from human thinking to program implementation.

Our method does more than document code - it guides users through the development process. The LLM breaks down complex concepts into clear steps, explaining each part in natural language before showing the code. This helps users understand program structure while learning programming naturally through practice.

### H.2 Advanced code navigation and project organization

Traditional IDEs show code through a file-based view, listing functions sequentially within each file. This makes it hard to see how different parts of a program connect. Our LP-style implementation solves this problem with a project-wide indexing system that shows these connections clearly.

Figure 9 shows the key differences between traditional IDE organization and our LP approach. Visual Studio Code displays a simple list of classes and functions for each file. This structure hides the logical connections between different parts of the code. It becomes especially problematic in large projects where understanding these connections is vital.

Mogan, our LP-based solution, takes a different approach. It organizes code using chapters and index markers that follow the program’s natural structure. This lets developers find code based on its purpose rather than its file location. The table of contents works like a map, helping developers quickly locate specific functions while seeing how they fit into the larger system.

Our cross-referencing system also connects related functions across different files. This is particularly useful in modern projects where code spans many files and modules. By making these connections clear through our index, we help developers understand complex codebases more easily.

This organization method benefits both individual developers and teams. New team members can quickly grasp the project’s structure. Existing team members can maintain and extend the code more effectively. The clear organization helps everyone understand how different parts of the program work together.

---

Our approach improves both immediate code navigation and long-term project maintenance. It makes software development more accessible while providing tools for effective collaboration. This serves both new programmers starting their journey and experienced developers managing complex projects.

In our investigation of LLMs' capabilities, we focus on examining their ability to constrain code generation to specific programming language contexts. We select the Goldfish Scheme implementation, a variant of S7 Scheme, as our experimental environment. This choice is particularly apt as Goldfish Scheme exists as a public programming language with comprehensive documentation available online, making it likely to be included in the training data of many LLMs while still maintaining a well-defined scope for our analysis.

To improve the performance of LLMs on Scheme, we designed a general template as the prompt in the mini-goldfish. We write the processes of writing a Functional Programming pattern procedure in the document based on the mathematical induction; and one can check whether the result follows our direction easily like we how to check whether the answer comes from our document or external materials.

Besides, we find the prompt we have designed in the FP pattern to improve the ability of LLMs to implement Scheme procedures. The improvement can be seen apparently in the small models. In the tests of the Llama 3.2 3B model, we tested different writing procedures, and we found the template for prompting to increase the performance of the llama. However, compared with the increase in GPT or other larger models, the result is not very clear. We consider that the public documents provide LLMs with extremely extensive knowledge in writing codes in a non-FP pattern, which disturbs the ability of LLMs in the Scheme.

In conclusion, writing documents in Literal Programming provides the ability to prompt designing at the same time. With the index system, one can easily check whether the LLMs use external knowledge. We believe it needs both the spreading of FP philosophy and convenient software to write documents and codes at the same time. As a result, the population of Mogan will help document writing in the FP pattern, which finally improves the ability of LLMs to write codes.

## I Several experiment results

### I.1 Toy cases

Chatgpt:

```
;; take-right : list -> integer -> list
;; Returns the last i elements of flist.
;; Examples:
;; (take-right '(a b c d e) 2) => (d e)
;; (take-right '(a b c d e) 5) => (a b c d e)
;; (take-right '(a b c d e) 7) => (a b c d e)
;; (take-right '(a b c d e) 0) => ()

(define (take-right flist i)
  (cond
    ;; If i <= 0, return empty list.
    ((<= i 0) '())

    ;; If the list is empty or i is at least the length of the list, return the list.
    ((or (null? flist) (>= i (length flist)))
     flist)

    ;; Otherwise, we skip (length flist - i) elements.
    (else
     (let* ((list-len (length flist))
            (skip-count (- list-len i)))
       (drop flist skip-count))))))

;; Helper function: drop
;; Returns the list that remains after skipping `n` elements from the front.
(define (drop lst n)
  (cond
    ((or (null? lst) (<= n 0)) lst)
```

---

```

    (else (drop (cdr lst) (- n 1))))))

;; Example usage:
;; (take-right '(a b c d e) 2) => (d e)
;; (take-right '(a b c d e) 0) => ()
;; (take-right '(a b c d e) 5) => (a b c d e)
;; (take-right '(a b c d e) 7) => (a b c d e)

```

Claude Sonnet:

```

;; take-right : list integer -> list
;; Purpose: Returns the last i elements of flist.
;; If i is greater than or equal to the length of flist, returns the entire list.
;; If i is less than or equal to 0, returns an empty list.
;;
;; Parameters:
;; flist: the input list
;; i: number of elements to take from the right
(define (take-right flist i)
  (cond
    ;; Base case: return empty list if i <= 0
    ((<= i 0) '())

    ;; Return entire list if:
    ;; - list is empty (null? flist), or
    ;; - requested elements >= list length
    ((or (null? flist) (>= i (length flist)))
     flist)

    ;; Main case: skip (length flist - i) elements from start
    (else
     (let* ((list-len (length flist))
            (skip-count (- list-len i)))
       (drop flist skip-count))))))

;; drop : list integer -> list
;; Purpose: Returns a new list with the first n elements removed.
;; If n <= 0 or list is empty, returns the original list.
;;
;; Parameters:
;; lst: the input list
;; n: number of elements to remove from the front
(define (drop lst n)
  (cond
    ((or (null? lst) (<= n 0)) lst)
    (else (drop (cdr lst) (- n 1)))))

```

Llama 3.1 2B:

```

;; take-right : list -> integer -> list
;; Returns the last i elements of flist.
;;
;; @param flist the input list
;; @param i the number of elements to take from the end of the list
;; @return the last i elements of flist, or the entire list if i is too large
(define (take-right flist i)
  (cond
    ;; If i is less than or equal to 0, return an empty list.
    ((<= i 0) '())
    ;; If the list is empty or i is greater than or equal to the list length, return the list.
    ((or (null? flist) (>= i (length flist)))
     flist)

```

Gemini 2.0:

---

```

;; take-right : list -> integer -> list
;; Returns the last i elements of flist.
(define (take-right flist i)
  (cond
    ;; If i is less than or equal to 0, return an empty list.
    ((<= i 0) '())
    ;; If the list is empty, return an empty list.
    ((null? flist) '())
    ;; If i is greater than or equal to the length of the list, return the entire list.
    ((>= i (length flist)) flist)
    ;; Otherwise, drop the first (length flist - i) elements and return the rest.
    (else
     (let* ((list-len (length flist))
            (skip-count (- list-len i)))
       (drop flist skip-count)))))

;; Helper function: drop
;; drop : list -> integer -> list
;; Returns the list that remains after removing the first n elements from lst.
(define (drop lst n)
  (cond
    ;; If the list is empty, return an empty list.
    ((null? lst) '())
    ;; If n is less than or equal to 0, return the original list.
    ((<= n 0) lst)
    ;; Otherwise, recursively call drop on the rest of the list with n-1.
    (else (drop (cdr lst) (- n 1)))))

```

LLAMA 3.3:

```

;; take-right : list -> integer -> list
;; Returns the last i elements of a list.
(define (take-right flist i)
  (cond
    ((<= i 0) ; If i is less than or equal to 0, return an empty list.
     '())
    ((or (null? flist) ; If the list is empty or i is at least the length of the list,
         (>= i (length flist))) ; return the original list.
     flist)
    (else ; Otherwise, drop the first (length flist - i) elements.
     (let* ((list-len (length flist)) ; Calculate the length of the list.
            (skip-count (- list-len i)) ; Calculate the number of elements to skip.
            (drop flist skip-count)))) ; Drop the first skip-count elements.
    ))

;; Additional API needed as succstep
(define (drop lst n)
  (cond
    ((or (null? lst) ; If the list is empty or n is less than or equal to 0,
         (<= n 0)) ; return the original list.
     lst)
    (else ; Otherwise, recursively call drop on the rest of the list.
     (drop (cdr lst) (- n 1)))) ; Decrease n by 1 for each recursive call.

```

Qwen32B:

```

;; take-right : list integer -> list
;; Returns the last i elements of flist.
;; If i is less than or equal to zero, returns an empty list.
;; If i is greater than the length of the list, returns the entire list.
(define (take-right flist i)
  (cond
    ;; If i is less than or equal to 0, return an empty list.
    ((<= i 0) '())
    ;; If the list is empty or i is greater than or equal to the length of the list,

```

---

```

;; return the entire list.
((or (null? flist) (>= i (length flist)))
 flist)
;; Otherwise, skip (length flist - i) elements and return the remaining list.
(else
 (let* ((list-len (length flist))
        (skip-count (- list-len i)))
        (drop flist skip-count))))

;; Helper function: drop
;; Returns the list that remains after skipping n elements from the front.
;; If n is less than or equal to zero, returns the original list.
;; If the list is empty, returns the empty list.
(define (drop lst n)
  (cond
   ((or (null? lst) (<= n 0)) lst)
   (else (drop (cdr lst) (- n 1)))))

```

## I.2 RepoBench Python code generation

### Code generated with ILP

Python code generated by LLMs without ILP:

#####

ChatGPT 4:

```

def drop(lst, n):
    """
    Returns the sublist of lst that remains after dropping n elements from the front.
    If n is 0 or less, or if lst is empty, returns lst unchanged.
    """
    if not lst or n <= 0:
        return lst
    return drop(lst[1:], n - 1)

def take_right(flist, i):
    """
    Returns the last i elements of flist.

    """
    # If i <= 0, return empty list.
    if i <= 0:
        return []

    # If the list is empty or i >= length of the list, return the full list.
    if not flist or i >= len(flist):
        return flist

    # Otherwise, skip (len(flist) - i) elements from the front.
    skip_count = len(flist) - i
    return drop(flist, skip_count)

```

#####

Claude 3.5 Sonnet:

```

def drop(lst: list, n: int) -> list:
    """
    Helper function: drop
    Returns the list that remains after skipping `n` elements from the front.
    """

```

---

```

if not lst or n <= 0: # equivalent to (or (null? lst) (<= n 0))
    return lst
else:
    return drop(lst[1:], n - 1) # equivalent to (drop (cdr lst) (- n 1))

def take_right(flist: list, i: int) -> list:
    """
    Returns the last i elements of flist.

    """
    # If i <= 0, return empty list
    if i <= 0:
        return []

    # If the list is empty or i is at least the length of the list, return the list
    if not flist or i >= len(flist):
        return flist

    # Otherwise, we skip (length flist - i) elements
    list_len = len(flist)
    skip_count = list_len - i
    return drop(flist, skip_count)

```

#####

Llama 3.1 2B Local:

```

def take_right(flist, i):
    # If i is less than or equal to 0, return an empty list.
    if i <= 0:
        return []

    # If the list is empty or i is greater than or equal to the list length, return the list.
    if not flist or i >= len(flist):
        return flist

    # Otherwise, slice the list from the end to get the last i elements.
    # In Scheme, this would be equivalent to (list-tail flist (- (length flist) i)).
    # In Python, we can use list slicing to achieve this.
    return flist[-i:]

```

#####

Gemini 2.0:

```

def take_right(flist, i):
    """
    Returns the last i elements of flist.

    Args:
        flist: The input list.
        i: The number of elements to take from the right.

    Returns:
        A new list containing the last i elements of flist.
    """
    if i <= 0:
        return []
    elif not flist: # Check for empty list using Python's truthiness
        return []
    elif i >= len(flist):
        return flist[:] # Return a copy to mimic Scheme's behavior
    else:

```

---

```

    list_len = len(flist)
    skip_count = list_len - i
    return drop(flist, skip_count)

def drop(lst, n):
    """
    Returns the list that remains after removing the first n elements from lst.

    Args:
        lst: The input list.
        n: The number of elements to drop.

    Returns:
        A new list with the first n elements removed.
    """
    if not lst: # Check for empty list
        return []
    elif n <= 0:
        return lst[:] # Return a copy to mimic Scheme's behavior
    else:
        return drop(lst[1:], n - 1)

#####

LLama 3.3 Local:

def take_right(flist, i):
    """
    Returns the last i elements of a list.
    """
    if i <= 0: # If i is less than or equal to 0, return an empty list.
        return []
    elif not flist or i >= len(flist): # If the list is empty or i is at least the length of the list,
        return flist # return the original list.
    else: # Otherwise, drop the first (length flist - i) elements.
        list_len = len(flist) # Calculate the length of the list.
        skip_count = list_len - i # Calculate the number of elements to skip.
        return drop(flist, skip_count) # Drop the first skip-count elements.

def drop(lst, n):
    """
    Returns the list that remains after skipping n elements from the front.
    """
    if not lst or n <= 0: # If the list is empty or n is less than or equal to 0,
        return lst # return the original list.
    else: # Otherwise, recursively call drop on the rest of the list.
        return drop(lst[1:], n - 1) # Decrease n by 1 for each recursive call.

#####

Qwen32B:

def take_right(flist, i):
    if i <= 0:
        return []
    elif not flist or i >= len(flist):
        return flist
    else:
        skip_count = len(flist) - i
        return drop(flist, skip_count)

def drop(lst, n):
    if not lst or n <= 0:
        return lst

```



---

```
else:
    return drop(lst[1:], n - 1)
```

```
#####
```

## Code generated without ILP

Python code generated by LLMs:

```
#####
```

ChatGPT 4:

```
def take_right(flist, i):
    """
    Returns the last i elements of flist.

    """
    # If i <= 0, return an empty list.
    if i <= 0:
        return []

    # If the list is empty or i >= the list length, return the full list.
    if not flist or i >= len(flist):
        return flist

    # Otherwise, skip the first (len(flist) - i) elements.
    skip_count = len(flist) - i
    return flist[skip_count:]
```

```
#####
```

Claude 3.5 Sonnet:

```
def take_right(flist, i):
    """
    Returns the last i elements of flist.

    Args:
        flist: input list
        i: number of elements to take from the right

    Returns:
        list containing the last i elements

    """
    return flist[-i:] if i > 0 else []
```

```
#####
```

Llama 3.1 2B Local:

```
def take_right(flist, i):
    # If i is less than or equal to 0, return an empty list.
    if i <= 0:
        return []

    # If the list is empty or i is greater than or equal to the list length, return the list.
    if not flist or i >= len(flist):
        return flist
```

---

```
# Otherwise, slice the list from the end to get the last i elements.
# In Scheme, this would be equivalent to (list-tail flist (- (length flist) i)).
# In Python, we can use list slicing to achieve this.
return flist[-i:]
```

```
#####
```

Gemini 2.0:

```
def take_right(flist, i):
    """
    Fully based on the file, generate a function take-right flist i -> list
    take-right returns the last i elements of flist.
    """
    if i == 0:
        return []

    return flist[-i:]
```

```
#####
```

LLama 3.3 Local:

```
def take_right(flist, i):
    """
    Returns the last i elements of flist.

    Args:
    flist (list): The input list.
    i (int): The number of elements to take from the right.

    Returns:
    list: The last i elements of flist.
    """
    if i > len(flist):
        raise ValueError("i cannot be greater than the length of flist")
    return flist[-i:]
```

```
#####
```

Qwen32B:

```
def take_right(flist, i):
    """
    Returns the last i elements of the list flist.

    :param flist: List from which to take elements
    :param i: Number of elements to take from the end of the list
    :return: List containing the last i elements of flist
    """
    return flist[-i:]
```

```
#####
```