

LLM4Decompile: Decompiling Binary Code with Large Language Models

Hanzhuo Tan, Qi Luo, Jing Li, Yuqun Zhang
Southern University of Science and Technology
The Hong Kong Polytechnic University

Abstract

Decompilation aims to restore compiled code to human-readable source code, but struggles with details like names and structure. Large language models (LLMs) show promise for programming tasks, motivating their application to decompilation. However, there does not exist any open-source LLM for decompilation. Moreover, existing decompilation evaluation systems mainly consider token-level accuracy and largely ignore code executability, which is the most important feature of any program. Therefore, we release the first open-access decompilation LLMs ranging from 1B to 33B pre-trained on 4 billion tokens of C source code and the corresponding assembly code. The open-source LLMs can serve as baselines for further development in the field. To ensure practical program evaluation, we introduce Decompile-Eval, the first dataset that considers re-compilability and re-executability for decompilation. The benchmark emphasizes the importance of evaluating the decompilation model from the perspective of program semantics. Experiments indicate that our LLM4Decompile has demonstrated the capability to accurately decompile 21% of the assembly code, which achieves a 50% improvement over GPT-4. Our code, dataset, and models are released at <https://github.com/albertan017/LLM4Decompile>.

1 Introduction

Decompilation is the process of converting compiled machine code or bytecode back into a high-level programming languages. This is often done to analyze the workings of software when its source code is not accessible (Brumley et al., 2013; Katz et al., 2018; Hosseini and Dolan-Gavitt, 2022; Xu et al., 2023; Armengol-Estapé et al., 2023; Jiang et al., 2023; Wong et al., 2023). There have been numerous tools developed for decompilation, such as Ghidra (Ghidra, 2024) and IDA Pro (Hex-Rays,

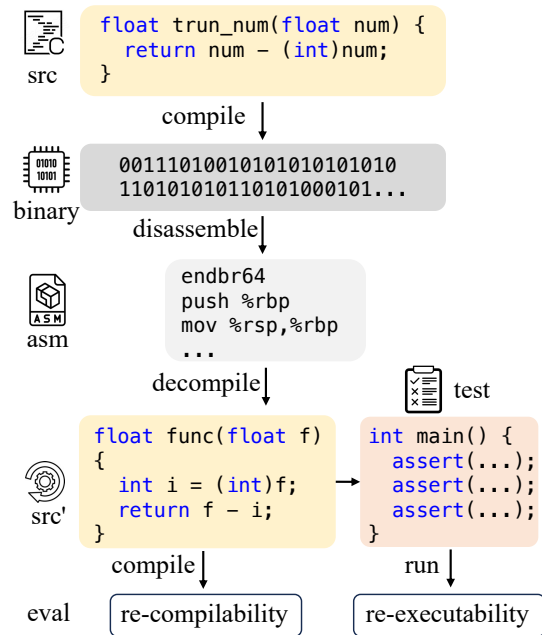


Figure 1: Pipeline to evaluate the decompilation.

2024). Although these tools have the capability to revert binary code to source code in specific scenarios, they often fall short in producing code that is easily readable by humans. The inherent difficulty of decompilation lies in the inability to fully recreate the source, especially finer details like variable names (Lacomis et al., 2019) and the primary structure (Wei et al., 2007), such as loops and conditional statements, which are often lost during the compilation process.

Recent advances in large language models (LLMs) have led researchers to approach programming languages as distinct linguistic systems, using pre-trained code LLMs for various coding tasks (Lippincott, 2020; Rozière et al., 2023; Guo et al., 2024). These models have shown impressive performance improvements over traditional techniques (Zeng et al., 2022; Xu et al., 2022), which leads us to the possibility of apply-

ing LLMs to cope with the decompilation challenge. To illustrate, Transformer-based models such as Slade (Armengol-Estapé et al., 2023) and Nova (Jiang et al., 2023) have showcased the potential of using language models to turn binary code back into source code that is much closer to the source code in readability and structure. However, the scope of their models is somewhat constrained at 200M and 1B parameters for Slade and Nova which could result in a reduced capacity for learning and generalization, whereas larger models typically exhibit a marked improvement in these areas by leveraging their extensive parameters to process and integrate a broader range of information (Rozière et al., 2023; OpenAI, 2023). Moreover, their lack of public availability limits their contribution to promoting further progress in this domain. Furthermore, to the best of our knowledge, no standardized benchmark dataset exists for evaluating and comparing decompilation techniques. Researchers tend to employ different datasets (da Silva et al., 2021; Collie et al., 2020; Tan et al., 2017) to evaluate their results, making direct comparison difficult. Therefore, there is a strong need to establish a benchmark to evaluate the decompilation performance, which can significantly facilitate the formulation of coherent and standard evaluation criteria for the decompilation domain.

Thus, our objective is to create and release the first open-source LLM dedicated to decompilation, and to assess its capabilities by constructing the first decompilation benchmark focused on re-compilability and re-executability. We start by compiling a million C code samples from AnghaBench (da Silva et al., 2021) into assembly code using GCC (Stallman et al., 2003) with different configurations, forming a dataset of assembly-source pairs in 4 billion tokens. We then finetune the DeepSeek-Coder model (Guo et al., 2024), a leading-edge code LLM, using this dataset. Followed by constructing the evaluation benchmark, Decompile-Eval, based on HumanEval (Chen et al., 2021) questions and test samples. Specifically, we formulate the evaluation from two perspectives: whether the decompiled code can recompile successfully, and whether it passes all assertions in the test cases. Figure 1 presents the steps involved in our decompilation evaluation. First, the source code (denoted as *src*) is compiled by the GCC compiler with specific parameters, such as optimization levels, to produce the executable binary. This binary is then disassembled into assembly language

(denoted as *asm*) using the `objdump` tool. The assembly instructions are subsequently decompiled to reconstruct the source code in a format which is readable to humans (denoted as *src'*). To assess the quality of the decompiled code (*src'*), it is tested for its ability to be recompiled with the original GCC compiler (re-compilability) and for its functionality through test assertions (re-executability).

On Decompile-Eval, the `llm4decompile` models demonstrated promising results in their ability to decompile binaries, with an impressive 90% of the decompiled code being recompileable using the same settings in the GCC compiler, signifying a solid understanding of code structure and syntax. As for the ability to execute the code, 21% of the decompiled codes from the 6B version successfully capture the semantics of a program and pass all the test cases.

In conclusion, our contributions are twofold:

- We provide the first open-source LLM ranging from 1B to 33B tailored for decompilation, which also facilitates compilation and diverse binary tasks.
- We construct the first decompilation benchmark which targets re-compilation and re-execution rate, which indicate syntax recovery and semantic preservation—both essential for usable and robust decompilation.

2 Related Work

2.1 Decompilation

The practice of reversing executable binaries to their source code form, known as decompilation, has been researched for decades. Traditional decompilation relies on analyzing the control and data flows of program (Brumley et al., 2013), and employing pattern matching, as seen in tools like Hex-Rays Ida pro (Hex-Rays, 2024) and Ghidra (Ghidra, 2024). These systems attempt to identify patterns within a program’s control-flow graph (CFG) that corresponds to standard programming constructs such as conditional statements or loops. However, crafting these rule-based systems can be challenging and prone to mistakes as the rules are complex to create, often only partially cover the CFG, and take extensive time to develop (Armengol-Estapé et al., 2023; Jiang et al., 2023). They are particularly weak when facing code that has been optimized, which is a common practice for commercially compiled software. The code output from such decompilation processes tends to be a source-

code-like representation of assembly code, including direct translations of variables to registers, use of `gotos`, and other low-level operations instead of the original high-level language constructs. This output, while often functionally similar to the original code, is difficult to understand and may not be efficient enough to allow for recompilation (Liu and Wang, 2020).

Drawing inspiration from neural machine translation, researchers have reformulated decompilation as a translation exercise, converting machine-level instructions into readable source code. Initial attempts in this area utilized recurrent neural networks (RNNs) (Katz et al., 2018) for decompilation, complemented by error-correction techniques to enhance the outcomes. Nonetheless, these efforts were constrained in their effectiveness.

The latest advancements in natural language processing (NLP) have enabled the use of pre-trained language models (LMs) for coding-related tasks (Rozière et al., 2023; Lippincott, 2020; Guo et al., 2024). These models generally incorporate the Transformer architecture (Vaswani et al., 2017), using self-attention mechanisms, and are pre-trained on extensive text datasets. This approach allows LMs to capture contextual nuances and aids in the acquisition of general language understanding. In the realm of binary decompilation, BTC (Hosseini and Dolan-Gavitt, 2022) was one of the first to fine-tune an LM for this purpose. Following this, Slade (Armengol-Estapé et al., 2023) utilized the BART model and trained an LM-based decompiler with 200 million parameters, while Nova (Jiang et al., 2023) developed a binary LM with 1 billion parameters starting from the StarCoder checkpoint and fine-tuned it for decompilation. Although these models show potential in decompilation, they are limited in size, e.g., the Code Llama model (Rozière et al., 2023), launched in 2023, has at least 7 billion parameters.

2.2 Evaluation

There is a notable gap in the field of decompilation: a lack of a unified, accepted benchmark for measuring the quality of decompilation tools. Various sources are used for evaluation purposes. BTC (Hosseini and Dolan-Gavitt, 2022), for example, leverages web data including interview-style coding problems and the extensive Debian Linux repository to evaluate decompilation accuracy. Meanwhile, Slade (Armengol-Estapé et al., 2023) is tested using both Synth, a synthetic

code generation framework, and the subset of ExeBench (da Silva et al., 2021; Collie et al., 2020), a benchmark consisting of executable C programs. Nova (Jiang et al., 2023) evaluates its decompilation capabilities on a synthetic dataset that it created, in addition to the CodeFlaws (Tan et al., 2017) dataset, which is designed to identify common coding errors.

The metrics employed for these evaluations predominantly focus on N-gram similarity, with the use of BLEU or Token Accuracy, as well as Edit Similarity (ES). Slade (Armengol-Estapé et al., 2023) goes a step further by incorporating Input Output (IO) Accuracy (Le et al., 2014; Liu and Wang, 2020) into its evaluation framework. This metric assesses semantic equivalence through behavioral equality, meaning it checks whether the decompiled code and the original code produce the same outputs when given the same inputs. However, IO Accuracy relies on external processes to generate input and output samples for comparison. The generation of input and output samples often involves randomness, leading to non-deterministic results and making it difficult to consistently assess the performance of a decompiler.

Consequently, our goal is to develop and make the first open-source large language model (LLM) tailored for decompilation. We also aim to establish the first benchmark for re-compilability and re-executability to set a standard for performance evaluation in the field of decompilation.

3 LLM4Decompile

In this section, we describe the pre-training data, present different model configurations, and discuss the pre-training objectives involved in pre-training the LLM.

3.1 Pre-training Data

We construct `asm-source` pairs based on Ang-habench (da Silva et al., 2021), which is a public collection of one million compilable C files. Following the practice of previous works (Armengol-Estapé et al., 2023), we first compile the source code into an binary object file, disassemble the object file into assembly code, and pair it with the source code, where we only consider the x86 Linux platform. In the real deployment, programmers will choose different compiler optimization flags to optimize the execution performance. Compiler optimization refers to the process of tweaking and trans-

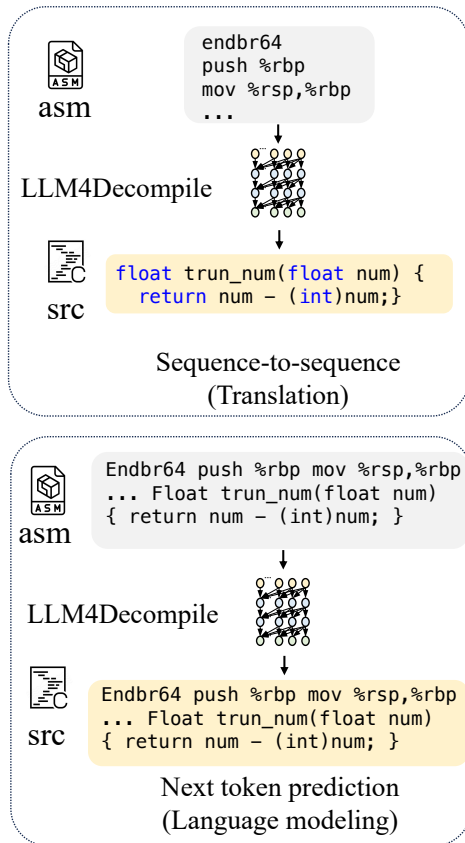


Figure 2: Training objectives

forming source code to generate faster and more efficient machine code. It involves techniques like eliminating redundant instructions, better register allocation, loop transformations, etc. The different optimization levels trade off compile time against execution time and debugging ability. The key optimization levels are: **O0** (default, No optimization) to **O3** (Aggressive optimizations, compilation time consuming). We compile the source code into all four stages, i.e., O0, O1, O2, O3, and pair each of them with the source code. To inform the model with the optimization stage, we use the following prompt: # This is the assembly code with [optimization state] optimization: [asm code] # What is the source code?.

3.2 Model Configurations

Our LLM4Decompile uses the same architecture as DeepSeek-Coder and we initialize our model with the corresponding DeepSeek-Coder checkpoints. As a neural translator, the training objectives can be categorized into two categories, as shown in Figure 2.

1) Next token prediction (NTP) or language modeling, which is the pre-training objective of most

LLM models targeting predicting the next token given previous inputs. As depicted in Equation 1, it minimizes the negative log probability for the ground truth token x_i :

$$\mathcal{L} = - \sum_i \log P_i(x_i|x_1, x_2, \dots, x_{i-1}; \theta) \quad (1)$$

where the conditional probability P is modeled using LLM4Decompile model with parameters θ . These parameters are optimized by applying the gradient descent algorithms (Ruder, 2016) with respect to the input sequence x_1, x_2, \dots, x_{i-1} preceding the given token x_i .

2) Sequence-to-sequence prediction (S2S), which is the training objective adopted in most neural machine translation models that aims to predict the output given the input sequence. As depicted in Equation 2, it minimizes the negative log probability for the C code tokens x_i, \dots, x_j :

$$\mathcal{L} = - \sum_i \log P_i(x_i, \dots, x_j|x_1, \dots, x_{i-1}; \theta) \quad (2)$$

where the loss is calculated only for the output sequence $x_i \dots x_j$, or the C code.

The main difference lies in whether the input sequence or the input assembly code is included in the training loss or not, whereas in language modeling, all the inputs are included for loss calculation. We conduct different ablation studies on these two training objectives to explore their effectiveness in decompilation.

4 Experiment Setups

4.1 Decompile-Eval Benchmark

Currently, there appears to be no benchmark for decompilation evaluation that considers whether code can be recompiled or executed correctly. When assessing decompilation model performance, researchers rely on metrics that measure N-gram similarity (such as BLEU or CodeBLEU) or edit similarity (ES). However, these metrics, commonly used in machine translation and text generation, fail to adapt to the evaluation of programming languages.

Programming languages are highly structured and logical, insensitive to the naming of functions and variables, yet very sensitive to the flow of data and logic. Changing variable or function names does not affect the meaning of a program, but a single logical error can alter its entire function and purpose. As illustrated in Figure 3, the use of BLEU

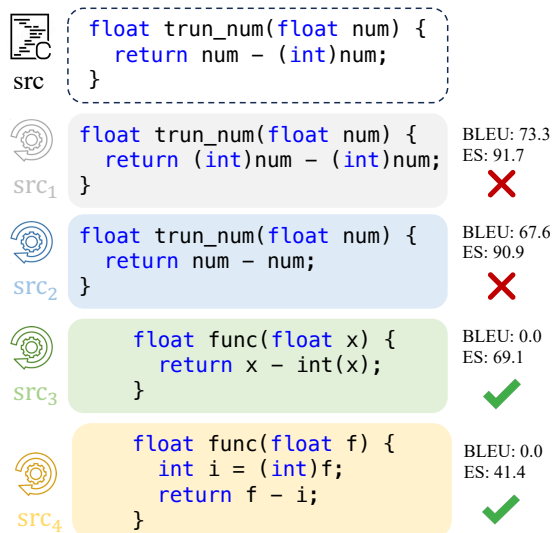


Figure 3: Limitation to use BLEU and Edit Similarity for evaluating decompilation results.

and ES in evaluating code similarity is problematic. For *src₁*, the variation from the original *src* is confined to type conversion of variable *num*, which leads to high BLEU and ES scores. However, this alteration completely changes the intent of the code. Similarly, *src₂* achieves high BLEU and ES scores, yet the semantics of the function are lost. Conversely, *src₃* undergoes normalization of function and variable names, causing no semantic shift yet scoring zero in BLEU against the original code. The example of *src₄* is more extreme: if the program logic is broken down into multiple lines, the ES drops to 41.4%, falsely indicating a low similarity. However, during compilation, names are typically standardized by the compiler, and source code is often broken down into basic operations depending on optimization. For this reason, the ability to recompile and execute the code is far more indicative than N-gram or edit similarity for evaluating decompilation efficacy.

To address the gap in decompilation assessment, we introduce Decompile-Eval, the first benchmark to evaluate the re-compilability and re-executability of decompilation systems. This benchmark is derived from HumanEval (Chen et al., 2021), which is the leading benchmark for code generation assessment and includes 164 programming challenges with accompanying Python solutions and assertions. We converted these Python solutions and assertions into C, making sure that they compile with the GCC compiler using standard C libraries and pass all the original assertions. In our eval-

uation process (Figure 1), the C source code is first compiled into a binary, then disassembled into assembly code, and finally fed into the decompilation system to be reconstructed back into C source code. This regenerated C code is compiled with GCC to test re-compilability and combined with the original assertions to check if it can successfully execute and pass those assertions. Re-compilability and re-executability serve as critical indicators in validating the effectiveness of a decompilation process. When decompiled code can be recompiled, it provides strong evidence of syntactic integrity. It ensures that the decompiled code is not just readable, but also adheres to the structural and syntactical standards expected by the compiler. However, syntax alone does not guarantee semantic equivalence to the original pre-compiled program. Re-executability provides this critical measure of semantic correctness. By re-compiling the decompiled output and running the test cases, we assess if the decompilation preserved the program logic and behavior. Together, re-compilability and re-executability indicate syntax recovery and semantic preservation—both essential for usable and robust decompilation.

In alignment with established evaluation practices, following Slade (Armengol-Estapé et al., 2023), we partition 1000 samples from the AngHaBench into a test set. We then utilize BLEU and ES as our primary metrics for assessment.

4.2 Baselines

To benchmark against SOTA decompilers, we selected two key baselines. First, **GPT-4** represents the most capable LLMs, providing an upper bound on LLM performance. As one of the largest language models, GPT-4 significantly surpasses previous LLMs across modalities. Second, **DeepSeek-Coder** is selected as the current SOTA open-source Code LLM. It represents the forefront of publicly available models specifically tailored for coding tasks. While recent academic works like BTC (Hosseini and Dolan-Gavitt, 2022) and Slade (Armengol-Estapé et al., 2023) showcase LLMs for decompilation, these models present significant integration challenges, such as complex pre-processing settings, non-standardized tokenizer/model loading, and necessitating significant effort to modify and adapt them. We thus selected GPT-4 and DeepSeek-Coder as representative cutting-edge and open-source baselines accessible for evaluation.

Table 1: Evaluation Results on Decompile-Eval

Model	Re-compilability					Re-executability				
	O0	O1	O2	O3	Avg.	O0	O1	O2	O3	Avg.
GPT4	0.92	0.94	0.88	0.84	0.895	0.1341	0.1890	0.1524	0.0854	0.1402
DeepSeek-Coder-33B	0.0659	0.0866	0.1500	0.1463	0.1122	0.0000	0.0000	0.0000	0.0000	0.0000
LLM4Decompile-1b	0.8780	0.8732	0.8683	0.8378	0.8643	0.1573	0.0768	0.1000	0.0878	0.1055
LLM4Decompile-6b	0.8817	0.8951	0.8671	0.8476	0.8729	0.3000	0.1732	0.1988	0.1841	0.2140
LLM4Decompile-33b	0.8134	0.8195	0.8183	0.8305	0.8204	0.3049	0.1902	0.1817	0.1817	0.2146

Table 2: Evaluation Results on Anghabench.

model	BLEU					Edit Similarity				
	O0	O1	O2	O3	Avg.	O0	O1	O2	O3	Avg.
DeepSeek-Coder-33B	0.0362	0.0367	0.0306	0.0313	0.0337	0.1186	0.1196	0.1124	0.1133	0.116
LLM4Decompile-1b	0.5099	0.493	0.487	0.4835	0.4934	0.6223	0.5946	0.5825	0.5822	0.5954
LLM4Decompile-6b	0.8219	0.8246	0.8143	0.8148	0.8189	0.8562	0.8551	0.8422	0.8453	0.8497
LLM4Decompile-33b	0.7724	0.7477	0.7485	0.7514	0.755	0.8252	0.7974	0.7993	0.8056	0.8069

4.3 Implementation

We use the Python implementation of the DeepSeek-Coder models (1.3B, 6.7B, and 33B) obtained on Hugging Face (Wolf et al., 2019). We set a global *batch size* = 2048 and *learning rate* = $2e-5$ and train the models with the AdamW optimizer (Loshchilov and Hutter, 2019) for 2 epochs. During the evaluation, we set the *max_new_tokens* to 512. To ensure fairness in the analysis of time and space complexity, all experiments are performed on a cluster equipped with 8 NVIDIA A100-80GB GPUs.

5 Experiment Results

5.1 Main Results

Table 1 presents the primary findings of our study. Initially, the base version of DeepSeek-Coder was unable to accurately decompile binaries. It could generate code that seemed correct and was sometimes compilable but failed to retain the original program semantics. After fine-tuning, the LLM4Decompile models demonstrated a significant improvement in their ability to decompile binaries, with an impressive around 90% of the code being compilable, signifying a solid understanding of code structure and syntax. As for the ability to execute the code, the 6B version of LLM4Decompile shows a remarkable advantage over the 1B version, 21% of the decompiled codes from the 6B version successfully capture the semantics of a program and pass all the test cases, while for the 1B version only 10% can be re-executed. The improvement highlights the benefits of larger model sizes in capturing the semantics of a program. Nonetheless, the increase in model size to 33B yields only a

marginal improvement, with an average increase of less than one percentage point in re-executability. This plateau may be due to the challenge of tuning the 33B model.

Table 2 summarizes the results on AnghaBench, where LLM4Decompile shows notably high BLEU and ES scores, e.g., the 6B model achieves 0.82 BLEU scores, almost identical to the source code. This outstanding performance suggests a significant data leakage issue within the test set. Decompiled code, with its variables normalized, should not realistically allow for such high N-gram/ES scores. This anomaly underscores the importance of establishing an independent, reliable benchmark for decompilation evaluation, as similarly high BLEU and ES scores have been reported in prior research.

5.2 Ablations

As discussed in Section 3.2, our LLM4Decompile model adopts a sequence-to-sequence (S2S) forecasting approach, which outperforms other training techniques for several reasons. In this training methodology, the input—specifically the assembly code is not included in the calculation of the loss function. This allows the model to focus solely on generating accurate output source code, enabling it to better understand the underlying patterns and structures of the decompiled code. In contrast, integrating the assembly code into the training process, as in the next token prediction (NTP) task, encompasses both the input assembly code and the output source code, which can decrease performance by around 4 points, as shown in Table 3. The complexity of assembly code is another factor; being inherently complex and low-level, it is harder for the

Table 3: Ablation study on training methodology.

Model	Re-compilability					Re-executability				
	O0	O1	O2	O3	Avg.	O0	O1	O2	O3	Avg.
S2S	0.8817	0.8951	0.8671	0.8476	0.8729	0.3000	0.1732	0.1988	0.1841	0.2140
NTP	0.8329	0.8598	0.8317	0.8329	0.8393	0.2805	0.1390	0.1573	0.1341	0.1777
NTP+S2S	0.8963	0.8598	0.8963	0.8720	0.8811	0.3232	0.1463	0.1951	0.1707	0.2088

model to learn meaningful patterns when assembly code is included in the training process. By excluding the assembly code from the loss calculation, the S2S approach enables the model to avoid this complexity and concentrate on high-level source code patterns. Although an alternative strategy involves an initial training step with both assembly and C code followed by fine-tuning focused on the translation task (NTP+S2S), this approach still doesn't perform as well as the S2S.

6 Conclusions

We presented the first open-source decompilation-focused LLM and standardized re-compilability/re-executability benchmark. Analyses on this diverse compiled C code dataset revealed promising capabilities—our 6B LLM4Decompile achieved 87% re-compilability, indicating syntactic understanding, and 21% re-executability, suggesting semantic preservation. As an initial exploration into data-driven decompilation, our work establishes an open benchmark to motivate future efforts. The public dataset, model, and analyses represent encouraging first steps toward enhancing decompilation through novel techniques.

7 Limitations

The scope of this research is limited to the compilation and decompilation of C language targeting the x86 platform. While we are confident that the methodologies developed here could be adapted to other programming languages and platforms with relative ease, these potential extensions have been reserved for future investigation. Additionally, our current study is constrained to the decompilation of single functions, without taking into account factors such as cross-references and external type definitions. This presents a simplified view of the decompilation process, omitting the complexities introduced by these elements. Addressing these aspects would provide a more comprehensive understanding of decompilation across a broader spectrum of scenarios and is an important avenue for subsequent research.

References

- Jordi Armengol-Estapé, Jackson Woodruff, Chris Cummins, and Michael F. P. O'Boyle. 2023. [Slade: A portable small language model decompiler for optimized assembler](#). *CoRR*, abs/2305.12520.
- David Brumley, JongHyup Lee, Edward J. Schwartz, and Maverick Woo. 2013. [Native x86 decompilation using semantics-preserving structural analysis and iterative control-flow structuring](#). In *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*, pages 353–368. USENIX Association.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *CoRR*, abs/2107.03374.
- Bruce Collie, Jackson Woodruff, and Michael F. P. O'Boyle. 2020. [Modeling black-box components with probabilistic synthesis](#). In *GPCE '20: Proceedings of the 19th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences, Virtual Event, USA, November 16-17, 2020*, pages 1–14. ACM.
- Anderson Faustino da Silva, Bruno Conde Kind, José Wesley de Souza Magalhães, Jerônimo Nunes Rocha, Breno Campos Ferreira Guimarães, and Fernando Magno Quintão Pereira. 2021. [ANG-HABENCH: A suite with one million compilable C benchmarks for code-size reduction](#). In *IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2021, Seoul, South Korea, February 27 - March 3, 2021*, pages 378–390. IEEE.
- Ghidra. 2024. [Ghidra software reverse engineering framework](#).

- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y Wu, YK Li, et al. 2024. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv preprint arXiv:2401.14196*.
- Hex-Rays. 2024. [Ida pro: a cross-platform multi-processor disassembler and debugger](#).
- Iman Hosseini and Brendan Dolan-Gavitt. 2022. [Beyond the C: retargetable decompilation using neural machine translation](#). *CoRR*, abs/2212.08950.
- Nan Jiang, Chengxiao Wang, Kevin Liu, Xiangzhe Xu, Lin Tan, and Xiangyu Zhang. 2023. [Nova⁺: Generative language models for binaries](#). *CoRR*, abs/2311.13721.
- Deborah S. Katz, Jason Ruchti, and Eric M. Schulte. 2018. [Using recurrent neural networks for decompilation](#). In *25th International Conference on Software Analysis, Evolution and Reengineering, SANER 2018, Campobasso, Italy, March 20-23, 2018*, pages 346–356. IEEE Computer Society.
- Jeremy Lacomis, Pengcheng Yin, Edward J. Schwartz, Miltiadis Allamanis, Claire Le Goues, Graham Neubig, and Bogdan Vasilescu. 2019. [DIRE: A neural approach to decompiled identifier naming](#). In *34th IEEE/ACM International Conference on Automated Software Engineering, ASE 2019, San Diego, CA, USA, November 11-15, 2019*, pages 628–639. IEEE.
- Vu Le, Mehrdad Afshari, and Zhendong Su. 2014. [Compiler validation via equivalence modulo inputs](#). In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, pages 216–226. ACM.
- Thomas Lippincott. 2020. [Starcoder: A general neural ensemble technique to support traditional scholarship, illustrated with a study of the post-atlantic slave trade](#). In *15th Annual International Conference of the Alliance of Digital Humanities Organizations, DH 2020, Ottawa, Canada, July 20-25, 2020, Conference Abstracts*.
- Zhibo Liu and Shuai Wang. 2020. [How far we have come: testing decompilation correctness of C decompilers](#). In *ISSTA '20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, USA, July 18-22, 2020*, pages 475–487. ACM.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. [Code llama: Open foundation models for code](#). *CoRR*, abs/2308.12950.
- Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.
- Richard M Stallman et al. 2003. Using the gnu compiler collection. *Free Software Foundation*, 4(02).
- Shin Hwei Tan, Jooyong Yi, Yulis, Sergey Mechtaev, and Abhik Roychoudhury. 2017. [Codeflaws: a programming competition benchmark for evaluating automated program repair tools](#). In *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, Argentina, May 20-28, 2017 - Companion Volume*, pages 180–182. IEEE Computer Society.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Tao Wei, Jian Mao, Wei Zou, and Yu Chen. 2007. [A new algorithm for identifying loops in decompilation](#). In *Static Analysis, 14th International Symposium, SAS 2007, Kongens Lyngby, Denmark, August 22-24, 2007, Proceedings*, volume 4634 of *Lecture Notes in Computer Science*, pages 170–183. Springer.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.
- Wai Kin Wong, Huaijin Wang, Zongjie Li, Zhibo Liu, Shuai Wang, Qiyi Tang, Sen Nie, and Shi Wu. 2023. [Refining decompiled C code with large language models](#). *CoRR*, abs/2310.06530.
- Frank F. Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. [A systematic evaluation of large language models of code](#). In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming, MAPS 2022*, page 1–10, New York, NY, USA. Association for Computing Machinery.
- Xiangzhe Xu, Zhuo Zhang, Shiwei Feng, Yapeng Ye, Zian Su, Nan Jiang, Siyuan Cheng, Lin Tan, and Xiangyu Zhang. 2023. [Lmpa: Improving decompilation by synergy of large language model and program analysis](#). *CoRR*, abs/2306.02546.

Zhengran Zeng, Hanzhuo Tan, Haotian Zhang, Jing Li, Yuqun Zhang, and Lingming Zhang. 2022. [An extensive study on pre-trained models for program understanding and generation](#). In *ISSTA '22: 31st ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, South Korea, July 18 - 22, 2022*, pages 39–51. ACM.