



Università
di Genova

DIBRIS DIPARTIMENTO
DI INFORMATICA, BIOINGEGNERIA,
ROBOTICA E INGEGNERIA DEI SISTEMI

Using LLMs to Rank Decompiled Code Variants

Luigi Timossi

Master's Thesis

Università di Genova, DIBRIS
Via Dodecaneso, 35 16146 Genova, Italy
<https://www.dibris.unige.it/>



**Università
di Genova**

Computer Science MSc
Computer Security and Engineering Curriculum

Using LLMs to Rank Decompiled Code Variants

Luigi Timossi

Advisors: Matteo Dell'Amico, Giovanni Lagorio
Examiner: Marina Ribaudò

March, 2026

Abstract

We use large language models (LLMs) to evaluate and rank decompiled code variants to address the inherently subjective challenge of assessing decompiler output quality and “human-likeness”. We compare two main approaches: using statistical metrics (perplexity and Total Loss) and employing LLMs as qualitative judges (LLM-as-a-Judge). We evaluate two open-weight models, *qwen-3* and *deepseek-r1*, on a dataset of real-world C projects decompiled with Ghidra (comparing different commits and Pull Requests) and other decompilers.

Our findings reveal that perplexity is a poor proxy for human-likeness, as human code naturally exhibits higher entropy compared to the rigid, highly repetitive boilerplate generated by decompilers. We demonstrate that Total Loss is a superior quantitative metric, because it accumulates prediction errors across all tokens, penalizing the highly predictive verbosity of decompilers. We show that abstracting code into Abstract Syntax Trees (ASTs) mitigates decompiler-specific lexical biases. In parallel, the qualitative LLM-as-a-Judge approach shows significant promise, with *deepseek-r1* achieving an alignment of $\sim 74\%$ with human consensus. However, we also identify critical vulnerabilities in LLMs, such as lexical hyper-fixation and hallucinations.

This work establishes a baseline for automated decompiler evaluation and highlights both the potential and the limitations of using LLMs in reverse engineering workflows.

Table of Contents

| | |
|--|-----------|
| Chapter 1 Introduction | 8 |
| Chapter 2 Related Work | 10 |
| 2.1 Generative Refinement of Decompiler Output | 10 |
| 2.2 LLM-based Benchmarking | 11 |
| 2.3 LLM-as-a-Judge | 12 |
| 2.4 Why Perplexity | 13 |
| Chapter 3 Background | 14 |
| 3.1 Ghidra | 14 |
| 3.2 SLEIGH and P-code | 15 |
| 3.2.1 P-code Semantics and Varnodes | 15 |
| 3.3 The Decompilation Pipeline | 16 |
| 3.3.1 Actions and Rules | 16 |
| 3.3.2 DefaultGroups | 17 |
| 3.4 Logic of Control Flow Structuring | 18 |
| 3.4.1 Basic Block Formulation | 18 |
| 3.4.2 The Structuring Algorithm | 22 |
| 3.4.3 The <i>for</i> Special Case | 25 |
| 3.4.4 The <i>Goto</i> Problem | 27 |
| 3.5 Code Emission | 27 |

| | | |
|------------------------------|---|-----------|
| 3.6 | Large Language Models | 28 |
| 3.6.1 | Transformer | 28 |
| 3.6.2 | Tokens | 30 |
| 3.6.3 | Softmax | 31 |
| 3.6.4 | Quantization | 31 |
| 3.7 | Decoding | 32 |
| 3.7.1 | Temperature | 32 |
| 3.7.2 | Top-p and Top-k | 33 |
| 3.8 | Perplexity | 34 |
| 3.9 | Human-likeness | 34 |
| Chapter 4 Methodology | | 36 |
| 4.1 | Dataset Maker | 36 |
| 4.1.1 | Dataset Collection | 38 |
| 4.2 | LLM Server | 38 |
| 4.2.1 | Models | 39 |
| 4.2.2 | Configuration | 44 |
| 4.2.3 | Decoding Strategy (temperature and top-p) | 44 |
| 4.2.4 | Routes | 45 |
| 4.2.5 | Metrics | 46 |
| 4.3 | Client | 46 |
| 4.3.1 | Building Ghidra | 47 |
| 4.3.2 | Ghidra Headless | 48 |
| 4.3.3 | Evaluation | 48 |
| 4.3.4 | Abstraction and Anonymization | 50 |
| 4.3.5 | Prompting | 52 |
| 4.4 | Pull requests | 54 |
| 4.5 | Reporting | 55 |

| | | |
|--------------------------|--|-----------|
| 4.6 | Dogbolt | 56 |
| Chapter 5 Results | | 57 |
| 5.1 | LLM performance | 57 |
| 5.1.1 | Generation | 57 |
| 5.1.2 | Score | 60 |
| 5.2 | Perplexity as a Metric for “Humanness” | 62 |
| 5.2.1 | Other Decompilers | 69 |
| 5.3 | LLM-as-a-Judge Evaluation | 76 |
| 5.3.1 | Handling Negative Constants – PR #8628 | 79 |
| 5.3.2 | Indexing Pattern Recognition – PR #8587 | 83 |
| 5.3.3 | Final Transform for While – PR #8161 | 86 |
| 5.3.4 | Sorting Switch Cases – PR #7253 | 91 |
| 5.3.5 | Improving Pointer-Expression Recovery – PR #6722 | 94 |
| 5.4 | Dogbolt Evaluation | 98 |
| 5.4.1 | Length Bias | 100 |
| 5.4.2 | Dogbolt Results | 100 |
| 5.5 | Correlation Between Perplexity & LLM judge | 104 |
| 5.5.1 | Ghidra | 104 |
| 5.5.2 | Dogbolt | 105 |
| 5.6 | Judge vs. Human Evaluation | 106 |
| 5.6.1 | Quantitative Alignment | 107 |
| 5.6.2 | Perplexity Correlation | 108 |
| 5.7 | Total Loss as a Quality Metric | 109 |
| 5.7.1 | Ghidra | 109 |
| 5.7.2 | Dogbolt | 111 |
| 5.8 | Discussion | 113 |

| | |
|---|------------|
| Chapter 6 Conclusion | 115 |
| 6.1 Limitations | 116 |
| Appendix A Quiz | 123 |
| A.1 Participant instructions | 123 |
| A.2 Quiz questions | 124 |
| A.3 Quiz results | 134 |
| Appendix B Future Ghidra work | 135 |
| B.1 Alternative decompilation flows | 135 |
| B.2 Implementing in GitHub Actions | 137 |
| Appendix C Prompts | 138 |
| C.1 LLM-as-a-Judge prompt templates | 138 |

Chapter 1

Introduction

Reverse engineering is a critical process in software security, enabling analysts to understand, debug, and modify software without access to its source code [CC90]. Decompilation tools like *Ghidra* and *Hex-Rays* have long been the backbone of this process, translating binary executables back into high-level code [Eag11; Nat19]. However, the output from these tools often suffers from issues such as poor readability, non-idiomatic constructs, and a lack of meaningful variable names, which can significantly hinder the ability of the analyst to comprehend and work with the decompiled code.

The advent of LLMs has opened new avenues for enhancing reverse engineering workflows. LLMs have demonstrated remarkable capabilities in understanding and generating code, making them promising candidates for improving the quality of decompiled output. Recent research has explored using LLMs to refine decompiler output, generate comments, and even act as judges to evaluate code quality. However, much of this work has focused on either generative refinement or broad benchmarking of decompilers, often relying on proprietary models and tools [Tan+24; HLC24]. In this master’s thesis, we take a different approach by leveraging LLMs to evaluate the “humanness” of decompiled code using intrinsic model metrics like *perplexity* [Hin+12]. We investigate how well local LLMs can distinguish between different versions of the same codebase, such as *pull requests*, without modifying the code itself. This fine-grained analysis is crucial for assessing incremental changes in code quality and readability, which is often more relevant in real-world reverse engineering tasks than wholesale comparisons of different decompilers. Our work also addresses the practical constraints of reverse engineering, such as privacy and cost, by exploring the feasibility of running these evaluations on local hardware, rather than relying on cloud-based Application Programming Interfaces (APIs). This makes our approach more accessible and applicable in security-sensitive contexts where data privacy is paramount [Sta+24; Car+21].

This study explores the application of LLMs to automate the evaluation and ranking of decompiled code variants. Specifically, we investigate whether LLMs can effectively proxy human judgment in determining which decompiled variant is more readable, idiomatic, and structurally sound. For limitations of time and resources, we focus on two open-weight models, *qwen-3* and *deepseek-r1*, which are representative of the current state-of-the-art in local LLM capabilities in the range of 14B of parameters. We explore two distinct methodologies: using the statistical predictability of the code (perplexity) as a quantitative metric, and prompting the models to act as qualitative evaluators (LLM-as-a-Judge). By testing on real-world datasets across different decompilers and specific pull requests within the Ghidra ecosystem, we aim to uncover the biases, strengths, and limitations of this approach, creating a framework with anonymizations and bias limitations.

Chapter 2

Related Work

The intersection of LLMs and reverse engineering has rapidly evolved, transforming how analysts interact with decompiled code. While recent literature has indeed explored utilizing LLMs to assist in reverse engineering (using framework like Model Context Protocol servers), the vast majority of this work focuses on two distinct areas:

- Generative refinement of decompiler output
- LLM-based evaluation and benchmarking of decompilation tools

Unlike generative approaches that aim to produce better code, our work focuses on measuring the “humanness” of existing code using intrinsic model metrics like perplexity and using those LLMs as Judges (Section 2.3). Furthermore, unlike broad benchmarking frameworks that rely on proprietary APIs to rank tools, our research investigates the granular utility of local LLMs in distinguishing between specific versions (or pull requests) of the same codebase.

2.1 Generative Refinement of Decompiler Output

The most prominent use of LLMs in this field is the attempt to improve the readability of the raw output produced by traditional decompilers. This body of work is complementary to ours; while we do not attempt to modify the code, understanding the deficits of raw decompiler output explains why our metrics are necessary to quantify “humanness”.

One example of this approach is LLM4Decompile [Tan+24] which is an LLM model that was trained to decompile binary code into high-level language, acting as a decompiler itself (LLM4Decompile-End is the model to decompile, LLM4Decompile-Ref is the model to

refine another decompiler output). Another example is DeGPT [HLC24], which introduces an end-to-end framework designed to optimize decompiler output directly employing a “three-role mechanism” (Referee, Advisor, and Operator) to guide an LLM in renaming variables, appending comments, and simplifying structure.

These works demonstrate that LLMs can significantly reduce the cognitive burden on analysts by rewriting code to be more idiomatic.

2.2 LLM-based Benchmarking

Closer to our specific problem domain is the emerging field of using LLMs to evaluate code quality, using LLMs to scale and automatize human evaluation, a technique which is often referred to as “LLM-as-a-Judge.”

DecompileBench [Gao+25] is the state-of-the-art in this area, It presents a comprehensive framework for evaluating decompilers, introducing the concept of using “LLM-as-a-Judge” to rate code understandability across 12 specific dimensions (e.g., Variable Naming, Control Flow Clarity). Their work validates that LLMs can align well with human experts in ranking different decompilers (e.g., comparing Ghidra vs. Hex-Rays vs. LLM-based decompilers). While DecompileBench is similar to our work in its use of LLMs for assessment, it differs from ours on the type of validation used: DecompileBench relies only on prompting the model to output a score based on Function Source Code, Decompiler A’s Pseudo Code, and Decompiler B’s Pseudo Code to calculate an ELO rating (a rating value that describes the score). In contrast, our work leverages also with intrinsic model metrics like *perplexity* to quantify the “surprise” of a specific model. Perplexity provides a more objective, quantifiable measure of how natural the code appears to the model, rather than relying only with subjective ratings, see Section 2.4.

Our work also diverges from DecompileBench in its focus on evaluating code variants within the same codebase (e.g., different pull requests). This fine-grained analysis is crucial on cases where is required to assess incremental changes rather than wholesale tool comparisons (es. GitHub Actions [Git]).

The last difference is that DecompileBench heavily utilizes proprietary, closed-source models (like GPT-4 or Claude-3.5) and licensed decompilers (like Hex-Rays or Binary Ninja). Our work specifically explores the feasibility of running these evaluations on local hardware. This addresses the privacy and cost constraints often present in security-sensitive reverse engineering tasks, which large-scale benchmarks often overlook. We also evaluate some other decompilers using Dogbolt, see Section 4.6, but since it does not provide APIs to automate the evaluation, we had to do it manually, and we only used it for a subset of our dataset.

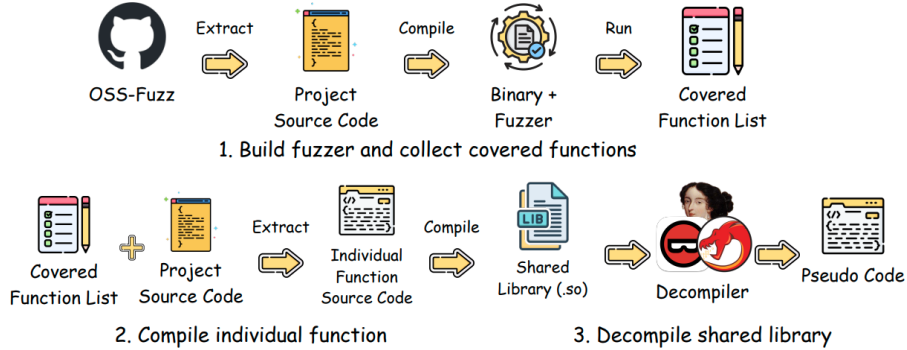


Figure 2.1: Creation of the dataset used in DecompileBench [Gao+25].

The dataset that we will use is a subset of that used in DecompileBench, since the entire dataset was overly large (more than 1 Tb) for our local hardware, and we needed to focus on a smaller subset of code variants to test the utility of perplexity and “LLM-as-a-Judge” in a more controlled setting, we decided to use a subset of that dataset. To build the dataset we adopted the technique proposed by the DecompileBench team: extracting source code from OSS-Fuzz projects [Goo], identifying functions covered during execution using Clang’s coverage sanitizer, and then extracting individual function implementations with their dependencies using `clang-extract`. These functions are compiled into standalone binaries (.so) with everything (binary position, function name, and source code) saved in a dataset object [pyt]. We can see the process in Figure 2.1.

2.3 LLM-as-a-Judge

The “LLM-as-a-Judge” [Li+25] is a technique that leverages the reasoning capabilities of LLMs to substitute human evaluators in tasks that require subjective judgment. These aspects can be various (e.g., overall quality, logic, readability, etc.) and are often difficult to quantify with traditional metrics. In the context of reverse engineering, using an LLM as a judge allows us to evaluate the “humanness” of decompiled code. Some problems with this approach are the potential for *bias* in the model’s judgments, such as position bias (prefer always the first option), or length bias (prefer the longer option). These problems must be carefully managed in a study that relies on LLMs for evaluation.

2.4 Why Perplexity

If we accept that human source code is “natural” and predictable, we can model it stochastically using neural (Transformer) language models. From this perspective, the decompiler acts as a noisy channel that introduces distortions into the original signal. The goal of LLM-based evaluation is to quantify how much the output signal (the decompiled code) deviates from the expected statistical distribution of “natural” human code [Hin+12]. A language model trained on human source code learns a probability distribution P over token sequences. When this model observes a sequence of decompiled code $S = t_1, t_2, \dots, t_N$, it assigns a probability to each token based on the preceding context, as we will discuss in Section 3.8. If the decompiled code uses alien or “unnatural” constructs, the model, expecting human patterns, will assign these tokens a very low probability. This statistical “surprise” is the foundation of perplexity.

Chapter 3

Background

This chapter provides the necessary background knowledge to understand how Ghidra works, in particular the decompilation process, with its architecture, main components, and the decompilation pipeline. The decompiler of Ghidra is an enormous and complex software, (over 200k lines of C++) so we focus only on the parts that are relevant to this study. Then we will introduce the main concepts behind LLMs, their architecture, sampling, and metrics for evaluation.

3.1 Ghidra

Ghidra [Nat19], released by the National Security Agency (NSA) in 2019, employs a bifurcated design that separates the user-facing interaction layer from the core analysis engine. This separation is not merely an implementation detail but a fundamental architectural constraint that dictates how data flows during the reverse engineering process.

The framework operates across two distinct memory spaces: a frontend implemented in Java and a backend analysis engine written in C++. The Java frontend is responsible for the Graphical User Interface (GUI), project database management, and plugin orchestration. It provides the high-level API exposed to users and scripts (e.g., Python or Java scripts via the GhidraScript framework). However, the computationally intensive tasks of data-flow analysis, variable inference, and control flow structuring are offloaded to a native C++ executable, typically named `decomp` or `decomp_dbg` (for debugging). These executables and the code are located at `Ghidra/Features/Decompiler/src/decompile/cpp`.

Communication is mediated by the `ghidra.app.decompiler.DecompInterface`. This interface manages a dedicated input/output stream to the native process, utilizing an XML-based protocol to exchange data. When a function decompilation is requested, the Java

client does not simply invoke a library function; it serializes the request into an XML command (e.g., `<decompile_at>`) and transmits it to the backend. The C++ process, holding its own representation of the function’s data flow in `Funcdata` objects, performs the analysis and returns the results as a serialized XML stream describing the high-level code structure and syntax tokens.

3.2 SLEIGH and P-code

As written in the documentation created by running `<make doc>` [NSA], the decompiler provides its own Register Transfer Language (RTL), referred internally as P-code; some examples are shown in Table 3.1. The disassembly of processor specific machine-code languages, and subsequent translation into p-code, forms a major sub-system of the decompiler. There is a processor specification language, referred to as SLEIGH, which is dedicated to this translation task; this piece of the code can be built as a standalone binary translation library, for use by other applications.

3.2.1 P-code Semantics and Varnodes

Unlike intermediate languages in compilers, P-code is designed specifically for reverse engineering, prioritizing the explicit representation of memory and register modifications.

The fundamental unit of data in P-code is the *Varnode*. A Varnode is defined by the triple (*Space, Offset, Size*), representing a contiguous sequence of bytes in a specific address space.

Table 3.1: Some P-code Operations and Semantics `opcodes.hh` [NSA; Pco]

| Opcode | Operands | Semantics |
|--------------|------------------------------|---|
| CPUI_COPY | $in_0 \rightarrow out$ | Copy one operand to another. |
| CPUI_LOAD | $space, ptr \rightarrow out$ | Load from a pointer into a specific address. |
| CPUI_STORE | $space, ptr, val$ | Store at a pointer into a specified address space. |
| CPUI_INT_ADD | $in_0, in_1 \rightarrow out$ | Integer addition, signed or unsigned. |
| CPUI_CBRANCH | $dest, cond$ | Conditional jump to <i>dest</i> if <i>cond</i> is non-zero. |

We must distinguish between two forms of P-code used during analysis:

1. **Raw P-code:** The direct, unoptimized output of the SLEIGH translation. It is represented by the class *PcodeOpRaw* (or by unprocessed *PcodeOp*), and contains the bare essentials: an opcode, a sequence number (address), and the input/output Varnodes.
2. **High P-code:** The result of the analysis pipeline. In this form, the code has been converted to Static Single Assignment (SSA) form, where every Varnode is defined exactly once for each function, if a variable is assigned multiple times, each assignment is given a new instance called low-level variable, dead code has been eliminated, and high-level concepts like function calls (replacing jump-and-link semantics) have been recovered. It is represented by the class *HighVariable*; this is an abstraction that groups multiple low-level Varnodes (which may reside in different registers or stack locations during execution) into a single logical variable, similar to a variable in C code.

The transformation from Raw to High P-code is where the majority of the decompilation logic resides. It is an inference process that attempts to raise the abstraction level of the code, often relying on heuristics that may fail in the presence of obfuscation or aggressive compiler optimizations.

3.3 The Decompilation Pipeline

The C++ decompiler engine processes a function at a time through a series of iterative passes. The architecture organizes these passes into *Actions* and *Rules*, managed by the *ActionDatabase*. Inside the `ActionDatabase::universalAction` we have two main types of objects:

- **ActionGroup:** Represents a list of Actions that are applied sequentially. The group's properties (eg., `rule_repeatapply`) influence how the contained actions are executed.
- **ActionPool:** It is a pool of Rules that are applied simultaneously to every *PcodeOp*. Each Rule triggers on a specific localized data-flow configuration. The Rules are applied repeatedly until no Rule can make any additional transformations.

3.3.1 Actions and Rules

Actions represent large-scale transformations applied to the graph of Varnodes and operations. They are the base class for objects that make modifications to a function's

(Funcdata) syntax tree. Their purpose is to manage complex stages of the workflow, such as recovering the control-flow structure or generating SSA form.

Rules, on the other hand, are a class designed to perform a single specific transformation on a PcodeOp or a Varnode. A Rule triggers when it recognizes a particular local configuration in the data flow and specifies a sequence of modification operations to transform it.

3.3.2 DefaultGroups

Actions and Rules are selected and activated according to the type of *DefaultGroup* they belong to. These groups represent standardized workflows for different analysis phases and are built by the method `ActionDatabase::buildDefaultGroups`. The main groups are:

- **decompile**: the standard workflow for full decompilation, composed of all of the phases.
- **jumpable**: optimized for analyzing jump tables.
- **normalize**: used for code normalization.
- **paramid**: for parameter identification.
- **register**: for register analysis.
- **firstpass**: a first fast analysis pass.

Each DefaultGroup is a list of names that refer to specific `ActionGroup`, `ActionPool` or individual `Action` to execute in that configuration. These lists define subsets of all the Actions.

The decompiler can be customized by selecting different DefaultGroups in Java with the method `setSimplificationStyle` of the decompiler interface but only the group named *decompile* returns C code to Ghidra, since in `ghidra_process.cc` we have:

Listing 3.1: `ghidra_process.cc`

```
1  [...]
2  fd->encode(encoder,0,ghidra->getSendSyntaxTree());
3  if (ghidra->getSendCCode() &&
4  (ghidra->allacts.getCurrentName() == "decompile")) //HERE WE HAVE THE CHECK
5  ghidra->print->docFunction(fd);
6  [...]
```

3.4 Logic of Control Flow Structuring

Recovering high-level control structures (loops, conditionals) from the unstructured Control Flow Graph (CFG) is arguably the most challenging phase of decompilation. It is effectively a pattern-matching problem on a directed graph, aimed at finding subgraphs that correspond to structured programming constructs.

3.4.1 Basic Block Formulation

The decompiler first aggregates P-code operations into *BasicBlocks*, that is, sequences of instructions with a single entry point and a single exit point (excluding internal calls). The CFG is formed by the edges representing jumps and branches between these blocks. Ghidra normalizes this graph to ensure a unique entry block, often inserting empty placeholder blocks to handle re-entrant loops or complex function entries.




Figure 3.1: Control Flow Graph of a function [BEP25]

In Figure 3.1 we can see an example of a CFG, we used it to create a C function and then

we have extracted its corresponding P-code representation in the table below¹.

| Source Code (C) | P-Code / Basic Blocks |
|---|--|
| <pre> 1 int a2_local; 2 int a1_local; 3 putchar(L'1'); 4 if ((a1 == 1) </pre> | <p>Basic Block 0</p> <pre> 0x0010118d:1: RSP(0x0010118d:1) = RSP(i) + #0xffffffffffffff8 0x0010118d:2: *(ram,RSP(0x0010118d:1)) = RBP(i) 0x00101195:d: u0x00004780(0x00101195:d) = RSP(i) + #0xffffffffffffff4 0x00101195:f: *(ram,u0x00004780(0x00101195:d)) = EDI(i) 0x00101198:10: u0x00004780(0x00101198:10) = RSP(i) + #0xffffffffffffff0 0x00101198:12: *(ram,u0x00004780(0x00101198:10)) = ESI(i) 0x001011a0:14: RSP(0x001011a0:14) = RSP(i) + #0xffffffffffffffe0 0x001011a0:15: *(ram,RSP(0x001011a0:14)) = #0x1011a5 0x001011a0:67: u0x10000008:1(0x001011a0:67) = *(ram,RSP(0x001011a0:14)) 0x001011a0:16: call jputchar(free)(#0x31:4,u0x10000008:1(0x001011a0:67)) 0x001011a5:17: u0x00004780(0x001011a5:17) = RSP(i) + #0xffffffffffffff4 0x001011a5:18: u0x00011e80:4(0x001011a5:18) = *(ram,u0x00004780(0x001011a5:17)) 0x001011a5:1e: ZF(0x001011a5:1e) = u0x00011e80:4(0x001011a5:18) == #0x1:4 0x001011a9:23: goto Block_2:0x001011bd if (ZF(0x001011a5:1e) != 0) else Block_1:0x001011ab </pre> |
| <pre> 1 (a2 != 2)){ </pre> | <p>Basic Block 1</p> <pre> 0x001011ab:24: u0x00004780(0x001011ab:24) = RSP(i) + #0xffffffffffffff0 0x001011ab:25: u0x00011e80:4(0x001011ab:25) = *(ram,u0x00004780(0x001011ab:24)) 0x001011ab:2b: ZF(0x001011ab:2b) = u0x00011e80:4(0x001011ab:25) != #0x2:4 0x001011af:31: goto Block_2:0x001011bd if (ZF(0x001011ab:2b) != 0) else Block_4:0x001011b1 </pre> |

¹P-codes varies during all phases of the decompilation process; due to optimization rules, dead code elimination, and other transformations. The P-code shown here are taken from the `collapseInternal` method using `printRaw` of the `FlowBlock` class. The `BasicBlock` order may not correspond directly to the original source code order

| Source Code (C) | P-Code / Basic Blocks |
|--|--|
| <pre> 1 putchar(L'2'); 2 if (a1 != a2) { </pre> | <p>Basic Block 2</p> <pre> 0x001011c2:46: RSP(0x001011c2:46) = RSP(i) + #0xffffffffffffe0 0x001011c2:47: *(ram,RSP(0x001011c2:46)) = #0x1011c7 0x001011c2:69: u0x10000011:1(0x001011c2:69) = *(ram,RSP(0x001011c2:46)) 0x001011c2:48: call jputchar(free)(#0x32:4,u0x10000011:1(0x001011c2:69)) 0x001011c7:49: u0x00004780(0x001011c7:49) = RSP(i) + #0xfffffffffffff4 0x001011c7:4a: u0x00011e80:4(0x001011c7:4a) = *(ram,u0x00004780(0x001011c7:49)) 0x001011ca:4d: u0x00004780(0x001011ca:4d) = RSP(i) + #0xfffffffffffff0 0x001011ca:4e: u0x00006a00:4(0x001011ca:4e) = *(ram,u0x00004780(0x001011ca:4d)) 0x001011ca:54: ZF(0x001011ca:54) = u0x00011e80:4(0x001011c7:4a) == u0x00006a00:4(0x001011ca:4e) 0x001011cd:59: goto Block_3:0x001011cf if (ZF(0x001011ca:54) == 0) else Block_5:0x001011db </pre> |
| <pre> 1 putchar(L'4'); 2 goto LAB_001011e5; </pre> | <p>Basic Block 3</p> <pre> 0x001011d4:5b: RSP(0x001011d4:5b) = RSP(i) + #0xffffffffffffe0 0x001011d4:5c: *(ram,RSP(0x001011d4:5b)) = #0x1011d9 0x001011d4:6b: u0x1000001a:1(0x001011d4:6b) = *(ram,RSP(0x001011d4:5b)) 0x001011d4:5d: call jputchar(free)(#0x34:4,u0x1000001a:1(0x001011d4:6b)) 0x001011d9:5e: goto Block_6:0x001011e5 </pre> |
| <pre> 1 } else { 2 putchar(L'3'); 3 } </pre> | <p>Basic Block 4</p> <pre> 0x001011b6:33: RSP(0x001011b6:33) = RSP(i) + #0xffffffffffffe0 0x001011b6:34: *(ram,RSP(0x001011b6:33)) = #0x1011bb 0x001011b6:6d: u0x10000023:1(0x001011b6:6d) = *(ram,RSP(0x001011b6:33)) 0x001011b6:35: call jputchar(free)(#0x33:4,u0x10000023:1(0x001011b6:6d)) 0x001011bb:36: goto Block_5:0x001011db </pre> |
| <pre> 1 putchar(L'5'); 2 } </pre> | <p>Basic Block 5</p> <pre> 0x001011e0:38: RSP(0x001011e0:38) = RSP(i) + #0xffffffffffffe0 0x001011e0:39: *(ram,RSP(0x001011e0:38)) = #0x1011e5 0x001011e0:6f: u0x1000002c:1(0x001011e0:6f) = *(ram,RSP(0x001011e0:38)) 0x001011e0:3a: call jputchar(free)(#0x35:4,u0x1000002c:1(0x001011e0:6f)) </pre> |

| Source Code (C) | P-Code / Basic Blocks |
|---|--|
| <pre> 1 LAB_001011e5: 2 putchar(L'6'); 3 return; 4 } </pre> | <p>Basic Block 6</p> <pre> 0x001011ea:3c: RSP(0x001011ea:3c) = RSP(i) + #0xffffffffffffe0 0x001011ea:3d: *(ram,RSP(0x001011ea:3c)) = #0x1011ef 0x001011ea:71: u0x10000035:1(0x001011ea:71) = *(ram,RSP(0x001011ea:3c)) 0x001011ea:3e: call jputchar(free)(#0x36:4,u0x10000035:1(0x001011ea:71)) 0x001011f1:44: return(#0x0) </pre> |

BasicBlocks are created in `flow.cc` by the method `FlowInfo::splitBasic`. The routine partitions the P-code instruction stream at control-flow boundaries: conditional and unconditional jumps, call sites that alter control flow, and return instructions. Each such instruction ends the current block and/or starts a new one.

The CFG with the BasicBlocks can also be seen in Ghidra by entering the Display Function Graph window and enabling the P-code field in the layout of Instruction/Data (Figure 3.2)

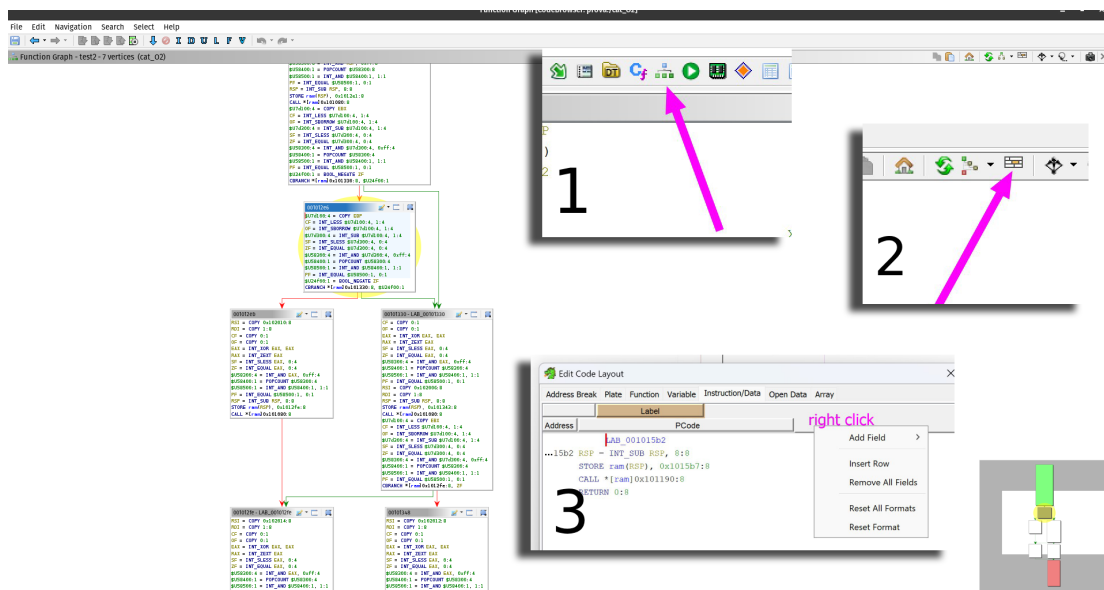


Figure 3.2: Control Flow Graph with P-code in Ghidra

3.4.2 The Structuring Algorithm

To transform the CFG into C statements, Ghidra employs a structuring algorithm implemented in the `ActionBlockStructure` class (an `Action`, see Section 3.3.1). The process involves identifying regions of the graph that match known schemas (or patterns) of control flow: inside the `apply` method of `ActionBlockStructure` we have a call to `collapseAll` that is the main loop of the algorithm:

```
1 void CollapseStructure::collapseAll(void)
2 {
3     int4 isolated_count;
4
5     finaltrace = false;
6     graph.clearVisitCount();
7     orderLoopBodies();
8
9     collapseConditions();
10
11    isolated_count = collapseInternal((FlowBlock *)0);
12    while(isolated_count < graph.getSize()) {
13        FlowBlock *targetbl = selectGoto();
14        isolated_count = collapseInternal(targetbl);
15    }
16 }
```

The method implements a deterministic sequence of passes that progressively transform the `BasicBlocks` into structured `FlowBlocks` and performs the following steps:

1. Preparation

The algorithm first clears previous visitation state (`graph.clearVisitCount`, line 6) and invokes `orderLoopBodies` (line 7). This pass discovers loop headers and back-edges, establishing a partial ordering among loop bodies. Detecting loops early is essential to prevent later structuring passes from erroneously breaking loop semantics.

2. Conditional simplification

Next, `collapseConditions` (line 9) attempts to simplify complex boolean logic and fold adjacent blocks that form logical AND/OR patterns (for example, transforming sequences that represent `if (A && B)` or `if (A || B)` into single conditional constructs). This phase applies local rules such as `ruleBlockOr` to reduce predicate complexity before higher-level structuring.

3. Initial collapse

The engine then calls `collapseInternal ((FlowBlock *)0)` (line 11), which scans

the graph and applies standard structuring rules (e.g. `ruleBlockIfElse`, `ruleBlockWhileDo`, `ruleBlockSwitch`) to collapse perfectly structured regions. The routine returns an `isolated_count` indicating how many blocks have been fully resolved without introducing `gotos`.

4. Unstructured flow handling

If the graph is not fully collapsed (`isolated_count < graph.getSize`, line 12–14), the method iterates: it selects a problematic edge with `selectGoto` and marks that edge as unstructured (to be emitted as a `goto/break/continue` in the final code). The selection is driven by heuristics to minimize disruption to surrounding structure. After marking the edge, `collapseInternal (targetbl)` is invoked again (often passing the target block of the newly created `goto`) so the structuring engine can resume collapsing other regions. This loop repeats until every block is resolved.

In the `collapseInternal` method we have the main pattern recognition method, some patterns have precedence over others, since it may occur that a region matches multiple schemas. For example, a `switch` may also match an `if-else` pattern.

These are the preferred patterns, in order:

- `goto`
- `cat` (block concatenation)
- `proper if` (if without else)
- `if-else`
- `while-do`
- `do-while`
- `infinite loop`
- `switch`

These “rules” are implemented inside a loop that tries every pattern till no more changes are possible.

In the `ruleBlockWhileDo` method we can see how the pattern matching is done:

```
1 bool CollapseStructure::ruleBlockWhileDo(FlowBlock *bl)
2 {
3   FlowBlock *clauseblock;
```

```

4   int4 i;
5
6   if (bl->sizeOut() != 2) return false; // Must be binary condition
7   if (bl->isSwitchOut()) return false;
8   if (bl->getOut(0) == bl) return false; // No loops at this point
9   if (bl->getOut(1) == bl) return false;
10  if (bl->isInteriorGotoTarget()) return false;
11  if (bl->isGotoOut(0)) return false;
12  if (bl->isGotoOut(1)) return false;
13  for(i=0;i<2;++i) {
14      clauseblock = bl->getOut(i);
15      if (clauseblock->sizeIn() != 1) continue; // Nothing else must hit clause
16      if (clauseblock->sizeOut() != 1) continue; // Only one way out of clause
17      if (clauseblock->isSwitchOut()) continue;
18      if (clauseblock->getOut(0) != bl) continue; // Clause must loop back to bl
19
20      bool overflow = bl->isComplex(); // Check if we need to use overflow syntax
21      if ((i==0)!=overflow) { // clause must be true out of bl unless we use
22          overflow syntax
23          if (bl->negateCondition(true))
24              dataflow_changecount += 1;
25          }
26          BlockWhileDo *newbl = graph.newBlockWhileDo(bl,clauseblock);
27          if (overflow)
28              newbl->setOverflowSyntax();
29          return true;
30      }
31  }

```

Firstly it checked that the block has exactly two outgoing edges (line 6) and it is not already part of a switch or a loop. Then, for each outgoing edge, it checks if the clauseblock (the potential loop body) has exactly one incoming edge (from the condition block) and one outgoing edge (back to the condition block). If these conditions are met, it confirms the presence of a while-do loop structure.

A condition is considered complex when the basic block that computes it contains too many instructions to be cleanly represented within a single conditional expression. The method `BlockBasic::isComplex` performs this check according to the following criteria.

The algorithm counts the number of *statements* in the block and:

- A conditional jump (branch) counts as 1 statement.

- `CALL` instructions count as 1.
- Operations that produce outputs used only inside the block or marker instructions do not count, but if a variable is used many times or is tied to memory, it contributes to the count.

If the total number of statements in the block exceeds 2, the block is considered complex.

The overflow syntax (`f_whiledo_overflow`) is a specific state assigned to a `BlockWhileDo` when its loop control condition is determined to be complex. It indicates that, although a logical `while` structure exists, the conditional block is too long or complicated to be emitted as a single boolean expression `while(condition){}`. Instead of printing `while(<complex condition>){}`, the decompiler emits an alternative form, typically an infinite loop with an internal `break` to preserve semantics.

After identifying a structure in the next iteration of the main loop in `collapseInternal`, a single `FlowBlock` representing the high-level construct (e.g., a `BlockWhileDo` for a `while` loop) is created. This new block encapsulates the original matched block, maintaining their internal P-code operations while providing a structured interface for further processing and eventual emission.

3.4.3 The *for* Special Case

As can be seen in Section 3.4.2, the Ghidra decompiler does not have an explicit rule to recognize `for` loops. Indeed, `for` loops in Ghidra are treated as special cases of `while-do` loops:² The check is performed in the method `BlockWhileDo::finalTransform`, this method proceeds only if the block is not marked with overflow syntax.

1. **Loop variable identification:** `findLoopVariable` is called to search for a variable controlling the iteration (e.g., `i` in `i < 10`). This variable must appear in the exit condition and be modified within the loop body.
2. **Initializer identification:** `findInitializer` searches for the instruction that initializes the variable (e.g., `i = 0`) in the block immediately preceding the loop.
3. **Relocation:** If both an iterator (`iterateOp`) and an initializer (`initializeOp`) are found, the decompiler physically moves the P-code operations (using `opUninsert` / `opInsertAfter`) so they lie adjacent to the loop boundaries, preparing them for syntactic emission.

²The transformation is triggered only if the architecture option `analyze_for_loops` is enabled.

4. **Non-printing marking:** In `finalizePrinting` these operations are marked with `opMarkNonPrinting`. This instructs the emitter not to print them as separate statements inside the body or before the loop, but to include them in the `for (...)` header.

```

1 void BlockWhileDo::finalTransform(Funcdata &data)
2 {
3     // Simplification style
4     BlockGraph::finalTransform(data);
5     if (!data.getArch()->analyze_for_loops) return;
6     if (hasOverflowSyntax())
7         return; // Still too complex
8     FlowBlock *copyBl = getFrontLeaf();
9     if (copyBl == (FlowBlock *)0) return;
10    BlockBasic *head = (BlockBasic *)copyBl->subBlock(0);
11    if (head->getType() != t_basic) return;
12    PcodeOp *lastOp = getBlock(1)->lastOp(); // There must be a last op in body,
        // for there to be an iterator statement
13    if (lastOp == (PcodeOp *)0) return;
14    BlockBasic *tail = lastOp->getParent();
15    if (tail->sizeOut() != 1) return;
16    if (tail->getOut(0) != head) return;
17    PcodeOp *cbranch = getBlock(0)->lastOp();
18    if (cbranch == (PcodeOp *)0 || cbranch->code() != CPU_CBRANCH) return;
19    if (lastOp->isBranch()) { // Convert lastOp to -point- iterateOp must
        // appear after
20        lastOp = lastOp->previousOp();
21        if (lastOp == (PcodeOp *)0) return;
22    }
23
24    findLoopVariable(cbranch, head, tail, lastOp);
25    if (iterateOp == (PcodeOp *)0) return;
26
27    if (iterateOp != lastOp) {
28        data.opUninsert(iterateOp);
29        data.opInsertAfter(iterateOp, lastOp);
30    }
31
32    // Try to set up initializer statement
33    lastOp = findInitializer(head, tail->getOutRevIndex(0));
34    if (lastOp == (PcodeOp *)0) return;
35    if (!initializeOp->isMoveable(lastOp)) {
36        initializeOp = (PcodeOp *)0; // Turn it off

```

```

37     return;
38 }
39 if (initializeOp != lastOp) {
40     data.opUninsert(initializeOp);
41     data.opInsertAfter(initializeOp, lastOp);
42 }
43 }

```

If all conditions are met, the decompiler effectively transforms the `while-do` structure into a `for` loop by relocating and marking the relevant P-code operations.

3.4.4 The *Goto* Problem

A significant limitation of this approach arises when the CFG contains irreducible control flow that does not match any predefined schema. This is common in binaries optimized with aggressive compiler techniques or those containing manual assembly optimizations.

When `ActionBlockStructure` fails to find a matching pattern, the jump inside the Flow-Block remains and it will be represented as a *goto*³ statement to preserve semantic correctness, this phenomenon significantly degrades the readability of the output.

3.5 Code Emission

The final phase of the pipeline is the translation of the structured High level P-code into C syntax. This is not a simple text dump but a structured generation of an AST represented by `ClangToken` objects.

Before emission, the *ActionNameVars* pass attempts to assign meaningful names to the recovered `HighVariable` objects. If debug symbols (DWARF, PDB) are available, they are utilized. In their absence, Ghidra relies on heuristics based on variable usage (e.g., loop counters named `i`, `j`) or storage location (e.g., `iVar1`, `uVar2`). This process is highly complex and often results in generic, non-descriptive identifiers.

The C++ backend generates a stream of `ClangToken` objects representing the code structure. This tokenized representation is sent to the Java frontend via a protocol based on XML messages. This structured data allows the Ghidra GUI to provide interactive features, such as cross-referencing and dynamic renaming, since the UI elements remain linked to the underlying `Varnode` and `HighVariable` objects.

³Or a `break/continue` if it jumps out of/into a loop structure.

3.6 Large Language Models

The advent of LLMs marks a fundamental discontinuity in the history of artificial intelligence and Natural Language Processing. It is not merely an increase in computational capacity, but an ontological redefinition of how machines process, represent, and generate semantic information. At the heart of this revolution lies the Transformer architecture, introduced in 2017 by Vaswani [Vas+17], which enabled overcoming the sequential limitations of previous Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) architectures.

3.6.1 Transformer

The shift from recurrent architectures to the Transformer was motivated by the need for parallelization and the handling of long-range dependencies. Whereas RNNs processed tokens sequentially (t_1, t_2, \dots, t_n) , accumulating error and dispersing the gradient over long sequences, the Transformer processes the entire sequence simultaneously, relying entirely on the attention mechanism.

As shown in Figure 3.3, the Transformer architecture consists of an encoder-decoder structure, where both components are composed of multiple layers of self-attention mechanisms and feed-forward neural networks.

The **encoder** (left) ingests a token sequence $x = (x_1, \dots, x_n)$ and produces continuous representations $z = (z_1, \dots, z_n)$. The **decoder** (right) conditions on z and generates the output tokens $y = (y_1, \dots, y_m)$ autoregressively, emitting one token per step. Both sides are built by stacking identical blocks composed of Multi-Head Attention and position-wise Feed-Forward layers, typically wrapped with residual connections and normalization. Inputs and targets are first embedded into a n -dimensional space, and a positional encoding is added to each embedding to encode token order [Vas+17].

At the core there is Multi-Head Attention, which allows the model to jointly attend to information from different representation subspaces at different positions. Each attention head computes scaled dot-product attention, enabling the model to focus on relevant parts of the input sequence when generating each output token.

Most of the state of art LLMs use an architecture with only the decoder part, omitting the encoder entirely [Bai]. This design choice is preferred for its simplicity, its good zero-shot generalization, and cheaper training cost to attain a reasonable performance.

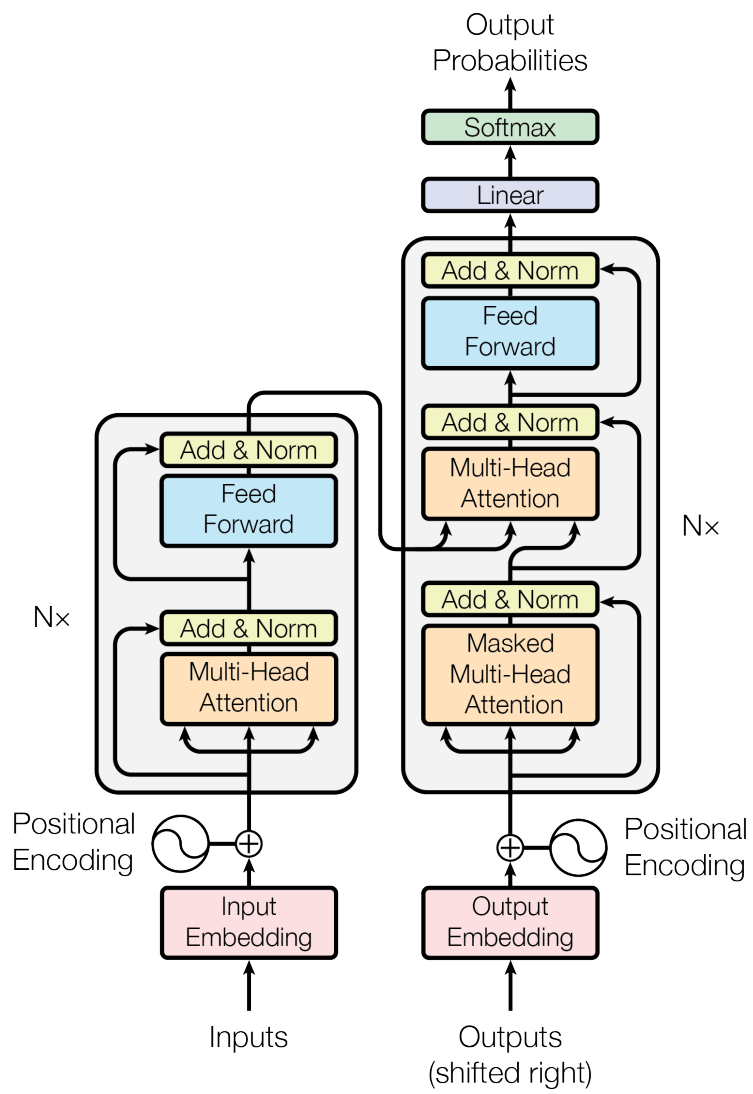


Figure 3.3: The Transformer model architecture [Vas+17]

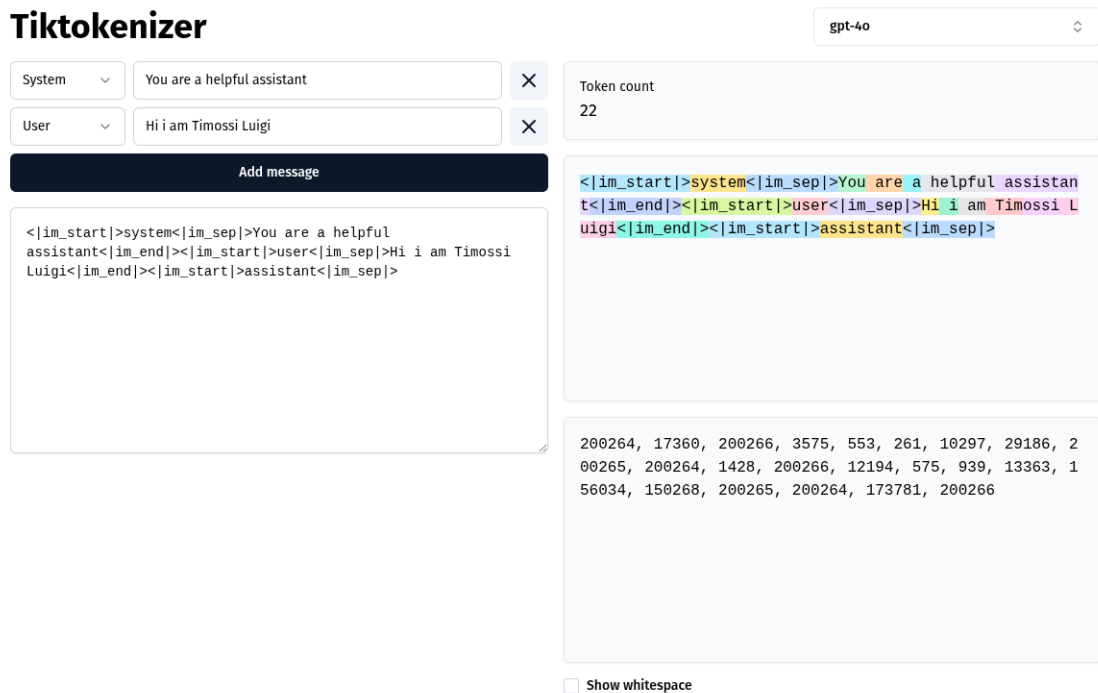


Figure 3.4: Tokenization example using tiktokenizer

3.6.2 Tokens

In LLMs, text is processed in chunks called **tokens**. A token can represent a word, a subword, or even a single character, depending on the tokenization scheme used. The choice of tokenization method significantly impacts the model’s performance, as it affects how the model interprets and generates text [Mul+18].

We can see an example using tiktokenizer,⁴ a webtool for visualizing tokenization for different models, to tokenize a sentence.

As shown in Figure 3.4, the sentence “Hi i am Timossi Luigi” is tokenized into a sequence of tokens. Each token corresponds to a specific integer ID in the model’s vocabulary.

For this example, for the model gpt-4o the token “i” corresponds to the ID 575. The tokenization process is crucial for LLMs, as it transforms raw text into a format that the model can process. Different models use different tokenization values.

In the tokens we count also special tokens like < |im_start| > that indicates boundaries

⁴<https://tiktokenizer.vercel.app/>

for roles messages like system, assistant, or user. In this case the token $\langle |im_start| \rangle$ corresponds to “input message start” after that we have the role then the token “input message separator” the message and, finally, the token “input message end”. Different models use different special tokens or do not use them at all.

3.6.3 Softmax

The **Softmax** function is used as an output function in the last layer of a neural networks to transform a vector of real values into a probability distribution. This function, for each component of the vector, computes the exponential normalized by the sum of the exponentials of all components, producing in output a probability vector of the same dimension with values in the interval $[0,1]$ [Zvo].

$$\text{softmax}(y_i) = \frac{e^{y_i}}{\sum_{j=1}^n e^{y_j}}$$

Where $y = (y_1, y_2, \dots, y_n)$ is an input vector whose values are in range from $-\infty$ to $+\infty$.

3.6.4 Quantization

Quantization is a technique used to reduce the memory footprint and computational requirements of neural networks by representing weights and activations with lower precision. In the context of LLMs, quantization can be applied to the model’s parameters (weights) and activations (intermediate outputs) to enable faster inference and reduce memory usage, especially on resource-constrained devices [Dev; Facb]. There are different quantization schemes, such as:

- **Post-training quantization:** This method quantizes a pre-trained model without requiring additional training. It can be applied to both weights and activations, but it may lead to a drop in model accuracy if not done carefully.
- **Quantization-aware training:** This method incorporates quantization into the training process, allowing the model to learn to compensate for the reduced precision. This approach typically results in better accuracy compared to post-training quantization.

Quantization can significantly reduce the computational requirements of LLMs, enabling faster inference and making it feasible to deploy large models on edge devices or in real-time applications. However, it is important to carefully evaluate the impact of quantization on model performance, as aggressive quantization can lead to a significant drop in accuracy.

3.7 Decoding

When the model generates text, it produces a vector of raw scores, called *logits*, for each token in the vocabulary at each timestep. These logits represent the unnormalized likelihood of each token being the next token in the sequence. By applying the `softmax` function to these logits, the model obtains a probability distribution over the vocabulary. Once the probability distribution is computed, the model must select the next token. This process, called decoding [IBM], can be influenced by different strategies such as:

- **Greedy decoding:** Always select the token with the highest probability, produces output that closely matches the most common language in the model’s pretraining data and in the prompt text, which is desirable in less creative or fact-based use cases. This can cause the model to produce repetitive or generic output.
- **Sampling decoding:** The model chooses a subset of tokens, and then one token is chosen randomly from this subset to be added to the output text. Sampling adds variability and randomness to the decoding process, which can be desirable in creative use cases. This can cause the model to produce unexpected or incorrect output.

3.7.1 Temperature

Temperature (T) is a hyperparameter that acts directly on the `softmax` function, which can be rewritten as:

$$\text{softmax}(y_i) = \frac{e^{\frac{y_i}{T}}}{\sum_{j=1}^n e^{\frac{y_j}{T}}}$$

Where $T > 0$ is the temperature parameter. The temperature modifies the distribution of probabilities over the tokens.

- $T < 1$ (Cooling): Differences between logits are amplified. The token with the highest logit receives a probability close to 1. The distribution becomes “peaked”, reducing variety and increasing determinism. Useful for logical or mathematical tasks.
- $T > 1$ (Heating): Differences are flattened. The distribution tends toward uniformity. Tokens with lower logits gain probability mass, increasing “creativity” but also the risk of incoherence (hallucinations).
- $T \rightarrow 0$: Equivalent to the **Greedy decoding**, where always the single most probable token is chosen.

Miklos and Rebeka [Mik], in their study on the impact of temperature on text generation have tested different temperature settings using OpenAI GPT-4.1 with the prompt “Why do researchers use control groups in experiments?.” They send two times the same prompt with different temperature settings and analyzed the outputs:

- At $T = 0.1$, the output was identical for both runs and both answers offered a clear, textbook-style explanation
- At $T = 1.4$, the outputs varied between runs, providing more expressive, creative answers with illustrative examples.
- At $T = 2$, the outputs became incoherent and nonsensical with grammatical errors, illogical statements, and gibberish characters.

These results highlight how much temperature settings influence the balance between coherence and creativity in LLM outputs.

Even with an optimal temperature, the long tail of the distribution (thousands of tokens with infinitesimal but nonzero probability) can introduce errors if sampled. To mitigate this, techniques like Top-k and Top-p sampling are employed.

3.7.2 Top-p and Top-k

Top-p (nucleus) sampling is a stochastic decoding strategy that at each generation step restricts sampling to the smallest subset of tokens whose cumulative probability is at least a threshold p [wikib].

Formally, given a vocabulary V and a context $x_{<t}$, the nucleus $V^{(p)}$ is the minimal subset satisfying

$$V^{(p)} \subseteq V, \quad \sum_{x \in V^{(p)}} P(x \mid x_{<t}) \geq p$$

Tokens outside $V^{(p)}$ are assigned zero probability; probabilities inside the nucleus are renormalized

Top-k sampling limits the candidate tokens to the k most probable ones at each step, and nucleus sampling dynamically adjusts the candidate set based on the cumulative probability threshold p .

The combined use of Temperature (to model the shape of the curve) and Top-p (to intelligently truncate the tail) represents the current industry standard for high-quality text generation.

3.8 Perplexity

Evaluating the quality of an LLM is intrinsically challenging because language judgments are often subjective. Nonetheless, there are rigorous quantitative metrics. *Perplexity* is the standard measure used during LLM pre-training; it stems from information theory and quantifies the model’s uncertainty when predicting the next token [wika].

Mathematically, perplexity is the exponential of the average negative log-likelihood (i.e., the exponentiated cross-entropy) of the predicted tokens. Exponentiating the cross-entropy restores the measure to probability-like units, yielding an intuitive “effective branching factor” the average number of plausible next-token choices the model considers [Mor].

Let the per-token negative log-likelihood (token loss) be defined as [Rua+25]

$$\ell_i = -\log P(x_i | x_{<i})$$

Then the perplexity of a sequence $X = (x_1, \dots, x_N)$ is

$$\text{PPL}(X) = \exp\left(\frac{1}{N} \sum_{i=1}^N \ell_i\right) = \exp\left(-\frac{1}{N} \sum_{i=1}^N \log P(x_i | x_{<i})\right)$$

Thus, perplexity can be interpreted as the exponential of the average token loss across the sequence.

A perplexity of K indicates that, on average, the model behaves as if it were choosing among K equally likely alternatives at each prediction step [Sax].

- Relation to entropy: $\text{PPL} = 2^{H(P)}$, where $H(P)$ denotes the Shannon entropy of the distribution.
- Interpretation: Lower perplexity means the model assigns higher probability mass to the true tokens from the test set.

Note that a low perplexity reflects statistical predictability relative to the training corpus and does not guarantee factual accuracy or correctness.

3.9 Human-likeness

LLMs are designed to generate text that closely mimics human language, capturing nuances, context, and stylistic elements. This capability is achieved through extensive training on vast corpora of text, enabling the models to learn patterns and structures inherent

in human communication. This mimic of human-like text generation has profound implications across various domains, including customer service, content creation, and education. Since the generated text is often indistinguishable from that written by humans, we can use LLM to measure how “human-like” is a piece of text. If we have two C functions that perform the same task but generated with different decompilers (or different settings/version of the same decompiler), we can probably assume that using an LLM to measure how “human-like” is the generated code can be a good proxy for code quality/readability.

Here we have two main paths to measure the human-likeness of a piece of code:

1. Ask the LLM directly via prompt to rate the human-likeness of the code. This approach can lead to subjective and inconsistent results, as the model’s responses may vary based on the prompt phrasing, the context, and hallucinations.
2. Use perplexity as a quantitative metric to evaluate the human-likeness of the code. This approach leverages the statistical properties of the language model to assess how well the generated code aligns with patterns learned from human-written code.

As shown in Section 3.8, perplexity measures how well a language model predicts a sequence of tokens. A lower perplexity indicates that the model finds the sequence more predictable, which often correlates with human-like text. By calculating the perplexity of code snippets generated by different decompilers, we can objectively compare their human-likeness. For this reason, perplexity is used only to evaluate the human-like quality of the decompiled code, it **does not** evaluate its functional correctness.

Chapter 4

Methodology

The framework of our work is based on a client-server architecture illustrated in Figure 4.1, where the server hosts the LLM and provides an API for interacting with it, while the client is responsible for building specific Ghidra versions, preparing the code samples and prompts, invoking the server, and collecting results in JavaScript Object Notation (JSON) format. Every service is designed to be modular, allowing for the integration of different LLM models and evaluation metrics. For reproducibility we used *Docker Compose* to containerize both the server and client components, ensuring consistent environments across different machines and operating systems. The dataset creation is also a containerized process, and the result are mounted as volumes to the client container, allowing for easy access and manipulation of the data without the need for complex data transfer mechanisms.

4.1 Dataset Maker

As written in Section 2.2 and shown in Figure 4.2, we decided to use a subset of the complete OSS Fuzz dataset used by DecompileBench [Gao+25] for our evaluation, specifically the four Open Source projects: `file`, `libxls`, `readstat`, `xz`, which are written in C and have a rich history of commits and pull requests on GitHub. This choice was motivated by the need to have a manageable dataset size for local evaluation, while still covering a set of real world code and functions to evaluate our approach. For every project the dataset maker extracts all the functions and recompiles them into standalone binaries; this process creates different optimization levels of the binary, specifically `-O0` and `-O2`, and `-O3` which are the most common optimization levels used in real world scenarios, and which can have a significant impact on the decompilation output and its readability.

For creating our dataset we forked the original dataset maker script and modified it to

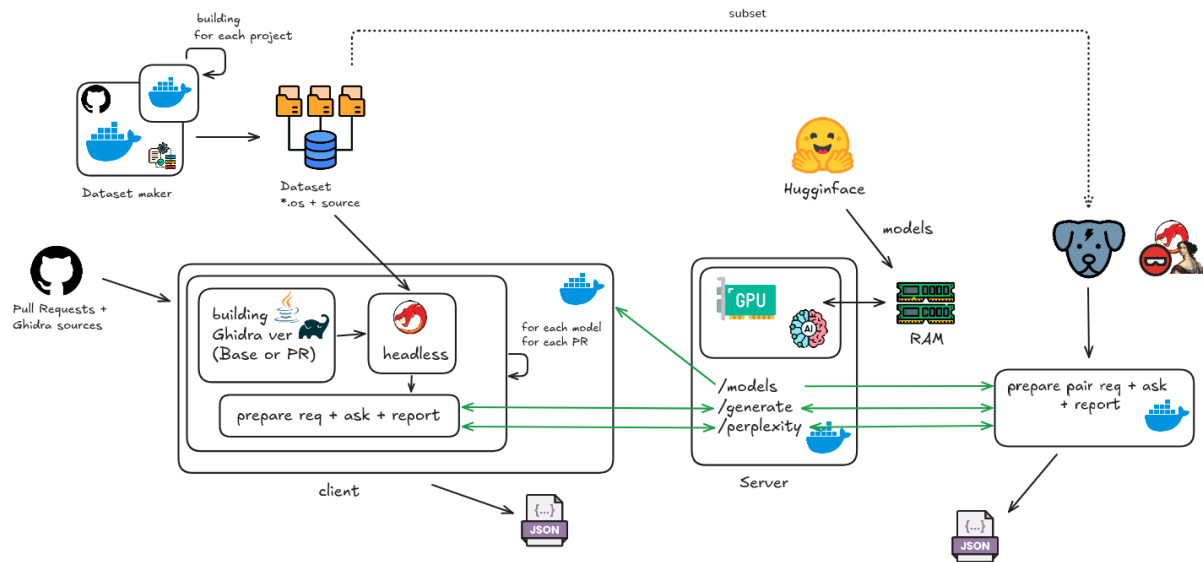


Figure 4.1: Overview of the client-server architecture used for evaluating decompiled code with LLMs.

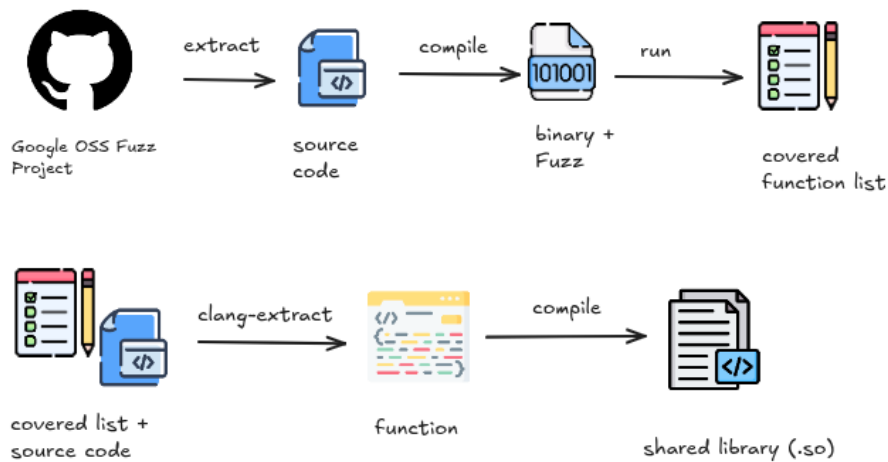


Figure 4.2: The dataset maker process, which extracts functions from selected open source projects, compiles them with different optimization levels, and organizes the resulting binaries and metadata into a structured format for evaluation.

fit our edits; we fixed specific optimization levels (in the original repository they were using all the optimization levels) and some bug fixes pointed out by one pull request on the original repo [edm]. We cloned our fork into a container (which will also run docker inside for building the projects), edited with the patches, as shown in the README file of DecompileBench, selected just our four projects and then run the dataset maker script.

4.1.1 Dataset Collection

The result of the dataset maker is a folder named `Dataset` which contains three other subfolders:

- **binary**: contains the compiled binary of the functions; every file is named with the `task-project_functionName-OX.so` format, and can be used for decompilation and evaluation.
- **compiled_ds**: contains a file structure of the dataset format used by the `Datasets` library [pyt], which is a Python library for handling large datasets efficiently, and which we use for loading the dataset in our client code. This structure has a “.arrow” file storing data and two JSON files for the metadata such as field names and types. In our case, we are interested only in three fields: `file`, which contain the name of the function; `path` which contains the path “binary/namefile”; and `func` which contains the source code of the function.
- **eval**: contains also a dataset structure, but we will not use it since is used for recompile success and other metrics we are not interested in, since we want to focus on the evaluation of the decompilation output rather than the compilation process.

4.2 LLM Server

The server is responsible for hosting the LLM and providing an API for interacting with it, specifically for receiving code samples and prompts from the client, processing them with the LLM, and returning the results. The server is designed to be modular, allowing for the integration of different LLM models and evaluation metrics, and it is containerized using Docker for reproducibility and ease of deployment.

It uses Gunicorn [Che] as the WSGI HTTP server for handling incoming requests, and it is built on top of a Python web framework (Flask [Ron]) to define the API endpoints and handle the logic for processing requests and interacting with the LLM.

4.2.1 Models

The heavy part of the framework is without doubt the server, and the models that run on it. In our case the local environment is a single Graphics Processing Unit (GPU) machine with 16 GB of Video Random Access Memory (VRAM), so we had to select models that can run on this hardware, and that can provide a good performance for our evaluation. We also used the VRAM Calculator¹ to estimate the memory requirements of different models and ensure they fit within our hardware constraints. Inside the VRAM different areas have to coexist:

- **Base Model Weights:** The trained parameters of the model, the “weights” with their precision (could be quantized for reduced memory usage).
- **Activations:** Intermediate computation results during forward passes through the layers. This grows with batch size and input length, and is critical for stability during inference.
- **KV Cache:** Key-Value cache used to avoid recomputing attention for previously processed tokens. Given the lengthy decompilation prompts containing source code, this cache grows proportionally with input length.
- **Framework Overhead:** Fixed memory cost from PyTorch, CUDA drivers, and buffer management. This overhead exists regardless of model size.

Initially we wanted to use models in the range of 8B parameters. The candidates were Meta Llama 3.1 (8B), Qwen2.5-Coder-Instruct (7B), DeepSeek-R1-Distill-Qwen (7B) and Google Gemma 2 (9B) for reasons, such as the availability of the model, the performance on code-related tasks, and the memory requirements. After extensive testing and variation of prompts, we decided to use models in range of 14B parameters since the first were not able to provide consistent and meaningful scores for our evaluation, falling into the Position Bias problem, which is a common issue in LLM-based evaluation where the model’s scoring is heavily influenced by the position of the input text rather than its content, this led to unreliable assessments even when we tried to mitigate it with different prompt engineering techniques, such as randomizing the order of the inputs, or using different templates for the prompts.

As example of this problem in Figure 4.3 when the majority of the scores are in the “TIE” category, meaning that the model is always selecting the same choice as the winner (either first or second), even after swapping the input order, which indicates a strong position bias and a lack of meaningful differentiation between the two inputs based on their content. Even after checking the non TIE cases, we found that the motivation for the choice was

¹<https://apxml.com/tools/VRAM-calculator>




Figure 4.3: Distribution of winner bias, the TIE is when a model always select one choice (BASE, PR) even after swapping the input.

```

lVar8 = 0;
do {
  switch(*(undefined1*)(lVar15 + 0x102361 + lVar8)) {
  case 0:
    iVar16 = (int)*pcVar12 + iVar16 * 10 + -0x30;
    break;
  }
  lVar8 = 0;
  do {
    switch(_dta_timestamp_parse_actions[lVar8 + lVar15 + 1]) {
    case 0:
      iVar16 = (int)*pcVar12 + iVar16 * 10 + -0x30;
      break;
    }
  }
}

```

Figure 4.4: Example of a function with only one difference between the BASE and the PR.

| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION |
|-------------|-----------------------------------|--------|--|
| qwen-coder | Code: Humanity & Readability | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| | Code: Fidelity & Cleanliness (GT) | PR | Both candidates accurately capture the control flow and logical structure of the original source code, making them equally faithful to the human intent. However, Candidate B follows a more structured format typical of modern C code, which might make it slightly easier to read and maintain for developers familiar with this style. |
| | AST: Humanity & Readability | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| | AST: Fidelity & Cleanliness (GT) | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| deepseek-r1 | Code: Humanity & Readability | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| | Code: Fidelity & Cleanliness (GT) | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| | AST: Humanity & Readability | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| | AST: Fidelity & Cleanliness (GT) | BASE | Preserving `goto` ensures that the decompiled code retains the Source's lower-level control flow, making A structurally closer to the Source's intent. |
| | Code: Humanity & Readability | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |
| | Code: Fidelity & Cleanliness (GT) | TIE | Detected potential bias in LLM response (Position Bias); declaring TIE. the LLM gave PR in both original and swapped prompts. |

Figure 4.5: Example of the motivation for the choice of the model, where the motivation is not based on the content of the decompilation and the difference between the two inputs, but rather on some superficial aspect of the input.

not based on the content of the decompilation and the difference between the two inputs. In this case where we have a function with only one difference between the BASE and the PR, that can be seen in Figure 4.4, and the corresponding responses in Figure 4.5 where, among the TIEs; we have two cases where the model consistently selected a winner, in the case of PR the motivation was consistent with the content of the decompilation and the difference between the two inputs, while in the case of BASE the motivation was on some superficial aspect of the input ignoring the context.

For these reasons we decided to use bigger models, since with these results we were not able to draw any meaningful conclusion from the evaluation, and we wanted to have a more reliable and consistent evaluation for our study.

4.2.1.1 Model Selection

Based on these considerations, we wanted to select 4 models for our evaluation, the range of 14B parameters is the last range that can be run on our hardware with 4-bit quantization,

and that can provide a good performance for our evaluation. Unfortunately, due to the scarcity of models in this range, and the tendency of releasing only very big or very small models, we did not have much choice in the selection (we cannot find any models by Google, OpenAI or Meta in this range), we found these state-of-the-art models:

- Qwen 3 (14B)
- DeepSeek-R1 Distill Qwen (14B)
- Phi 4 (14B)
- Mistral Nemo instruct 2407 (14B)

Unfortunately after some testing we found out that Phi 4 and Mistral Nemo were not able to provide consistent and meaningful scores for our evaluation, falling into the Position Bias problem as can be seen in Figure 4.6. We also tried to replace them with Qwen2.5 Coder and Starcoder2 (both 15B) but even with those models we were not able to get consistent results, so we decided to use only Qwen3 and DeepSeek-R1 Distill Qwen, since they were able to provide more reliable and consistent scores.

- **Qwen 3 (14B)**: Is the latest generation in the Qwen family and comes as a suite of both dense and mixture-of-experts (MoE) models. A distinctive feature of the series is the ability to switch within the same model between a *thinking* mode (aimed at complex reasoning, mathematics, and coding) and a *non-thinking* mode (optimized for efficient, general-purpose dialogue), allowing the same backbone to adapt to different interaction styles and difficulty levels [Ten].
- **DeepSeek-R1 Distill Qwen (14B/15B²)**: Part of the DeepSeek-R1 family, which introduces a reasoning-focused post-training pipeline combining cold-start supervised data and large-scale Reinforcement Learning (RL) to improve reasoning quality and readability over purely RL-trained variants (e.g., DeepSeek-R1-Zero). The *Distill* checkpoints transfer (distill) the reasoning patterns learned by the larger DeepSeek-R1 model into smaller dense backbones (including Qwen-based models), providing strong math/code/reasoning capability in a size that is practical to run locally on our hardware [Dee25; Gup].

We wanted to include different types of models but in the end we had to select models that are all based on the same architecture (Qwen), since they were the only ones that

²In the name and the description it is shown to have 14B parameters, but in the huggingface website it is tagged as 15B, if we calculate via python the number of parameters we found out it has 14.77B parameters

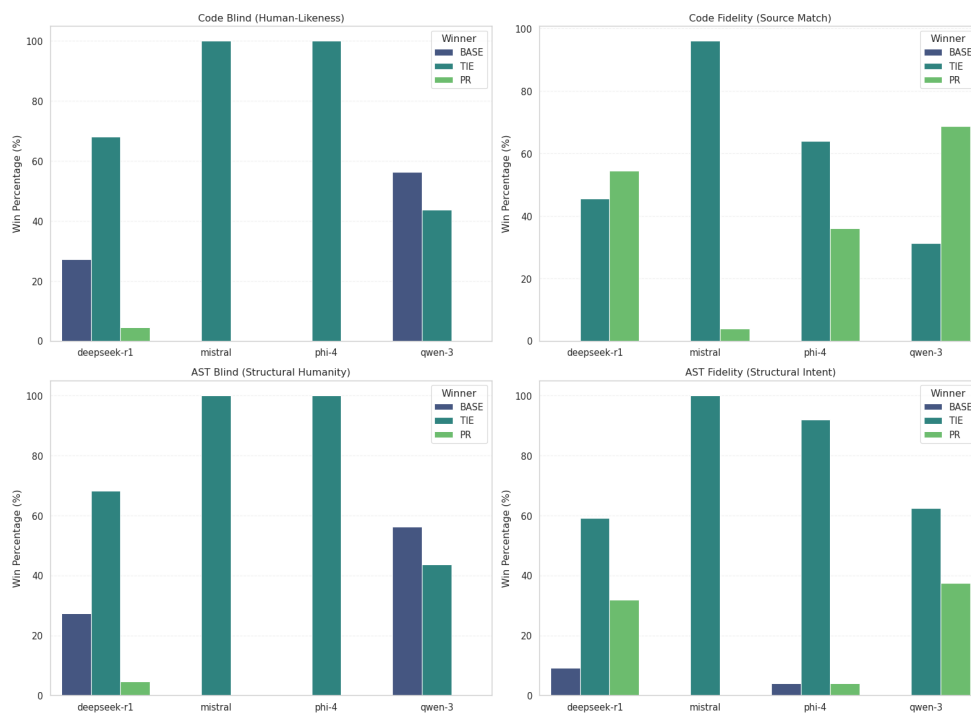


Figure 4.6: Distribution of winner bias for the new models, where we can see that Phi 4 and Mistral Nemo are still falling into the Position Bias problem, while Qwen3 and DeepSeek-R1 Distill Qwen are providing more reliable and consistent scores for our evaluation.

can provide consistent and meaningful scores for our evaluation, and that can fit in our hardware constraints, but we acknowledge that this is a limitation of our study and that it would be interesting to include models with different architectures and training paradigms in future work, to assess the generality of our findings across a wider range of LLM designs.

4.2.2 Configuration

The server supports multiple local LLMs through a simple configuration layer that maps a short, client-facing identifier to the corresponding Hugging Face repository name. Concretely, a dictionary (`MODELS_CONFIG`) defines the available models and is the single source of truth for both the `/models` endpoint and for request-time model switching.

All models are loaded using 4-bit quantization via `bitsandbytes` [Faca].

```
bnb_config = BitsAndBytesConfig(
    load_in_4bit=True,
    bnb_4bit_quant_type="nf4",
    bnb_4bit_use_double_quant=True,
    bnb_4bit_compute_dtype=torch.bfloat16
)
```

At startup, the server checks the availability of the GPU, CUDA and, to reduce cold-start delays, checks and optionally downloads all model snapshots before accepting requests, ensuring that evaluation runs are not affected by network variability. The greatest bottleneck in our setup is the model loading time and context switch (changing from one model to another), which can take various minutes per model due to the size and quantization overhead. We have organized the request pipeline in the Client to minimize the number of context switches but to mitigate further, the server implements a simple caching mechanism. At the start, when checking if the weights are in Huggingface cache folder we load the model into the GPU for quantization, after that we move the quantized weights and its tokenizer to CPU memory (RAM) and unload the model from VRAM. When every model have been loaded, the server opens up requests maintaining the quantized weights in RAM, so that every context switch can be fulfilled by quickly moving the quantized weights from RAM to VRAM, which is much faster than loading from disk and quantizing on the fly.

4.2.3 Decoding Strategy (temperature and top-p)

For the `/generate` endpoint, we configured the decoding parameters through a dedicated function that returns a dictionary of `transformers` generation arguments. In our experi-

ments we rely on *nucleus sampling* (`top_p`) with a low `temperature`, to balance determinism (useful for fair comparisons across decompilers) and the ability to escape repetitive or low-quality completions.

Concretely, the default configuration is:

- `temperature=0.4`: reduces randomness by sharpening the token distribution. Lower values make outputs more stable across runs, which is desirable for evaluation.
- `top_p=0.9`: nucleus sampling, i.e., tokens are sampled only from the smallest set whose cumulative probability mass is p . This prevents the model from selecting very unlikely tokens while still allowing variation.
- `max_new_tokens=4096`: upper bound on completion length, used as a safety and latency-control measure. The models are expected to generate many output tokens due to their “reasoning” style, so this limit is set high to avoid truncation while still preventing wasting time because of repetition.

4.2.4 Routes

The server exposes a minimal Representational State Transfer (REST) API. All endpoints exchange JSON payloads and are intentionally kept coarse-grained (a lot of work in a single request) to decouple the client implementation from model-specific details. The available routes are:

- **GET** `/`: health check endpoint. It returns the server readiness status, whether CUDA is available, and the currently loaded model identifier (if any). This is used by docker compose to ensure healthcheck status for required services.
- **GET** `/models`: returns the list of supported model keys (the abstract identifiers used by the client), mapped server-side to Hugging Face repository IDs.
- **POST** `/generate`: main inference endpoint. The request body includes `model_id` and a `prompt`. The server loads (or switches to) the requested model, wraps the prompt into a chat-style template via the tokenizer, runs text generation, and returns the generated completion.
- **POST** `/score`: scoring endpoint used to compute a language-model based score for a given text. The request body includes `model_id` and `text`. The server computes the token-level negative log-likelihood and returns the derived perplexity.
- **POST** `/free`: explicit cleanup endpoint to unload the currently resident model and aggressively release GPU memory.

Since different models cannot fit simultaneously in GPU memory, model switching is handled server-side: each request triggers a check on the currently loaded model and, if needed, a full unload/load cycle. To avoid concurrent access to GPU state, all inference and scoring operations are protected by a global lock, enforcing sequential execution.

4.2.5 Metrics

To make the evaluation reproducible and to quantify server-side overhead, the server logs per-request performance metrics to a CSV file (`llm_metrics.csv`). Each entry includes:

- **Model and operation:** `model_id` and `operation` (`generate` or `score`).
- **Latency:** wall-clock duration (seconds) measured around the full operation, including tokenization and GPU synchronization.
- **Peak GPU memory:** peak VRAM allocated during the operation, obtained via CUDA peak memory statistics.
- **Tokens:** number of prompt/input tokens and number of generated output tokens; these are also used to derive an approximate throughput (tokens per second).

Metric collection is implemented via a dedicated monitoring context manager that resets CUDA peak counters before execution and synchronizes the device before reading final statistics. This design provides a uniform measurement procedure across both generation and perplexity scoring, and enables later analysis of the impact of prompt length, time, and memory usage.

4.3 Client

The client is responsible for orchestrating the entire evaluation workflow, including building specific Ghidra versions, preparing code samples and prompts, invoking the server for decompilation and scoring, and collecting results in JSON format for analysis. It is designed to be modular and flexible, allowing for easy integration of different evaluation strategies and metrics, and it is containerized using Docker for reproducibility and ease of deployment.

The evaluation stage is an end-to-end pipeline over a set of target Ghidra Pull Requests (PRs). At a high level, the client:

1. Ensures that the *base* version is built and has produced decompilations for all dataset binaries

2. Iterates over the selected PR, building each corresponding Ghidra revision and extracting the related decompilations
3. Selects a limited subset of binaries that actually exhibit decompilation differences, to focus the evaluation on meaningful cases and reduce noise
4. For each model, queries the LLM server to score and compare the outputs, producing per-PR and aggregate JSON reports for later analysis.

4.3.1 Building Ghidra

The build process is automated via Python scripts that interact with Git and Gradle, an open source build system of choice for Java, Android, and Kotlin (we use an Ubuntu image for the container) [Gra]. Firstly we clone and build the Ghidra repository from GitHub; this version is used as the base for all our evaluations. Then we build the base version and extract the functions from binary and we get the PR number that we want to evaluate against base from a function that calls the GitHub API and returns the list of all Ghidra’s PR. Then for each PR we checkout the specific version of Ghidra, with a script that takes as input the PR number, fetches the specific head reference from the GitHub repository (`pull/ID/head:pr-ID`) and checks it out.

For building Ghidra two prerequisites are necessary: **Java 11** and **Gradle** (optionally). Java is required to run the build scripts and Gradle is used for managing dependencies and building the project. With newer versions of Ghidra that include a wrapper for Gradle (`gradlew`), it is not necessary to install Gradle globally.

One problem is that every Ghidra version needs a specific Java version, so we have to check the `application.properties` file inside the repository for the required minimal Java version, and then install it in the container before building Ghidra. So, inside the container we manage more than one Java version, and we switch between them based on the requirements of the Ghidra version we are building. Another problem is that some PRs are based on older versions of Ghidra which do not have the Gradle wrapper, so for building those versions we have to install more than one version of Gradle and switch between them based on the requirements of the Ghidra version we are building. This is the main reason for using an Ubuntu image for the container, since it allows us to easily manage multiple versions of Java and Gradle using the package manager and environment variables.

After building a specific version of Ghidra (Base or PR), for every binary found in the dataset folder, we check if it is not already decompiled by that specific version of Ghidra (i.e., if the corresponding JSON file with the decompilation output does not exist); after creating the list of the files not yet decompiled, we start the decompilation process.

This incremental strategy prevents re-running expensive steps and makes the workflow resumable.

4.3.2 Ghidra Headless

For decompilation we use the headless mode of Ghidra, which allows us to run Ghidra in a command-line environment without the need for a GUI [Ghi]. This is particularly useful for automating the decompilation process and integrating it into our evaluation workflow. The headless mode is invoked via a command-line script; the entry point is `support/pyghidraRun` executed in headless mode (`--headless`), creating a temporary per-binary Ghidra project, importing the binary, and finally running a post-script (`extract.py`) that performs the actual decompilation and exports the results to JSON. We use parallel execution to speed up the decompilation of multiple binaries; To avoid race conditions and project-file locks during parallel execution, we create an isolated project directory for each binary and delete it at the end of the run.

The headless invocation follows this template:

```
$GHIDRA_HOME/support/pyghidraRun --headless <proj_dir> <proj_name> \  
-deleteProject \  
-import <binary_path> -overwrite \  
-scriptPath <scripts_dir> \  
-postScript extract.py
```

For older Ghidra versions where `pyghidraRun` is not present, we fall back to the standard headless launcher `support/analyzeHeadless` with the same arguments. The client passes configuration to the post-script via environment variables: the output directory, the evaluated Ghidra version tag, and the comma-separated list of target function names to decompile. If not specified, all functions are decompiled, in our case we pass only the function names present in the database for that binary because in the compilation process some other functions are added in the binary for making it able to compile. This allows the same `extract.py` script to be reused across runs and versions without hardcoding paths or dataset-specific information.

4.3.3 Evaluation

Not every decompilation difference is relevant for our study: superficial variations (e.g., whitespace, renaming, minor formatting) would introduce noise. For this reason, before invoking the LLM, the client performs a structural comparison between the baseline and PR outputs:

- it loads the JSON produced by the decompiler for both `base` and `pr_#ID`,
- it compares the abstracted representation (via our `tree-sitter`-based abstraction, see Section 4.3.4) of the two outputs,
- it retains only binaries for which the abstracted forms differ, i.e., where control-flow or structure plausibly changed.

For each retained binary we compute a complexity proxy (*cyclomatic complexity*) based on the control-flow graph of the decompilation, which is defined as:

$$M = E - N + 2P,$$

where:

- E is the number of edges in the control-flow graph,
- N is the number of nodes in the control-flow graph,
- P is the number of connected components (one function is considered a single connected component, since we decompile at function level this is always 1 in our case).

The cyclomatic complexity quantifies code complexity based on decision points (e.g., if, while, case), where higher scores indicate harder-to-test, maintain, and error-prone code [Son]. Then for sampling reasons, we sort candidates by decreasing complexity and keep only the top `MAX_SAMPLES`. This strategy ensures that our evaluation focuses on the most challenging and relevant cases, where improvements in decompilation quality are most impactful.

4.3.3.1 LLM-based scoring

For each selected binary and for each model exposed by the server, we score:

- the baseline decompilation,
- the PR decompilation,
- the original source function (dataset ground truth),

using the server `/score` endpoint, which returns token-level negative log-likelihood aggregated into a *perplexity* value. In addition, we compute the perplexity of the abstracted AST forms (base/PR/source), enabling an evaluation that is less sensitive to naming and formatting. To reduce redundant server calls across repeated evaluations, the client maintains a cache keyed by (`func_name`, `model_id`, `test_binary_name`, `ppl_type`) for perplexity of base, source and their abstracted forms.

4.3.3.2 Comparison metric

For each sample we derive a simple *relative* metric:

$$\Delta\text{PPL} = \text{PPL}_{PR} - \text{PPL}_{base}$$

(and analogously for the abstracted representation). A negative ΔPPL indicates that the PR output is more *likely* under the model than the baseline output, which we interpret as a proxy for improved fluency/regularity. Conversely, a positive ΔPPL suggests a regression in perceived likelihood.

4.3.4 Abstraction and Anonymization

To evaluate the structural quality of the decompilation independently of variable naming and formatting, we implemented an abstraction mechanism using `tree-sitter` a Parser used by ATOM [Wik], specifically the language for C with `tree-sitter-c`. The Python client parses the decompiled C code into an AST and traverses it to generate a “skeletal” representation of the code.

In this representation, specific identifiers, literals, and types are replaced with generic placeholders (e.g., `id`, `num`, `type`), while control flow keywords (`if`, `while`, `for`, `switch`, `goto`) and block structures are preserved. This process effectively anonymizes the code and create a filter/standard for indentation and formatting, cleaning out the possible noise and forcing the LLM to focus purely on the control flow logic and structural complexity (e.g., the presence of “goto” statements vs structured loops) rather than being biased by variable names, comments or formatting.

We can see from this example how the original code (left) is transformed into the abstracted version (right), where all identifiers and literals are replaced with placeholders, while the control flow and structure are maintained.

| Source Code | Abstracted Code |
|--|---|
| <pre> 1 //This function does random stuff dont try to understand it 2 void complex(int a, char *b) { 3 long *f; 4 int h[10]; 5 if (a > 0) { 6 while (a < 10) { 7 printf("Value: %d\n", a); 8 a++; 9 } 10 } else { 11 goto end;//random comment 12 } 13 h[0] = 42; 14 end: 15 char c = b[0]; 16 f->g(h[i]); 17 (*(int *) (p + 4)) = 5; 18 } </pre> | <pre> 1 type id(type id, type *id){ 2 type *id; 3 type id[10]; 4 if(id > 0){ 5 while(id < 10){ 6 call(str, id); 7 id++; 8 } 9 }else{ 10 goto lbl; 11 } 12 id[0] = 42; 13 lbl: 14 type id = id[0]; 15 call(id[id]); 16 (*(type)(id + 4)) = 5; 17 } </pre> |

This allows us to evaluate the “humanness” of the decompilation output based on its logical structure and flow, rather than being influenced by specific naming choices or formatting styles that may vary widely across different decompilers and Ghidra versions.

This anonymization also allows us to compute a separate perplexity score for the abstracted AST representation, which can be compared against the perplexity of the original source code decompilation. By analyzing the correlation between these two scores, we can gain insights into how well the LLM is capturing the structural and logical aspects of the code, as opposed to just surface-level token patterns.

As shown in Figure 4.7, we observed a strong positive correlation (Pearson’s $r > 0.7$, $p < 0.001$) between the perplexity of the fully abstracted AST and that of the original source code across all evaluated models. Within each subplot, the black dashed line represents the linear regression model (line of best fit), while the narrow shaded area denotes the 95% confidence interval, indicating a high degree of statistical certainty in this trend.

In other words, samples that were rated as more “surprising” or “incoherent” in their original form also tended to have higher perplexity in their abstracted AST representation, and vice versa. This mathematically confirms that while the AST abstraction systematically shifts absolute perplexity values by removing lexical entropy, it perfectly preserves




Figure 4.7: Pearson correlation between the abstracted AST perplexity and source code perplexity for the same decompilation output, across all evaluated samples.

the relative predictability of the underlying logic.

4.3.5 Prompting

One key element of our evaluation framework is the design of the prompts used to query the LLM for scoring the decompilation outputs. We divide our analysis in two main types of prompts:

- **Human evaluation prompt:** This type of prompt is designed to elicit a qualitative assessment from the LLM about which version of the decompilation output is better in terms of human readability.
- **Closer to source prompt:** This type of prompt is designed to elicit a qualitative assessment from the LLM about which version of the decompilation output is closer to the original source code.

For each type of prompt we have two types, the one with the code in the original form and the one with the code in the abstracted form, to check if the model prefers one version over the other and if the abstraction process is effective in making the model focus on the structural aspects of the code rather than being influenced by specific naming choices or formatting styles.

4.3.5.1 Biases and prompt design

We had to design the prompts carefully to avoid introducing bias and help the model to focus on the relevant aspects of the decompilation outputs, the prompts are available in the appendix (Appendix C). To avoiding “update version” bias, we used a template that does not explicitly mention the concept of “BASE” and “PR”, but rather presents the two decompilation outputs as “Version A” and “Version B”. We then ask the model to compare them based on specific criteria. To help the model from possible hallucinations, it is important to make it generate the motivation for the choice *before* revealing the correct answer, so we ask the model to first provide an explanation of which version it thinks is better and why, since it will use the token generated to fill its context window.

To avoid “position bias” we perform a consistency check by swapping the inputs (Version A and Version B) and re-querying the model to see if it still prefers the same version, which can help us to identify and mitigate cases where the model’s preference is influenced by the position of the input rather than its content. If the model’s choice changes after swapping; it suggests that the original preference may have been due to position bias rather than a genuine assessment of the decompilation quality, so is flagged as *TIE*. To optimize the context window for the input, we designed the prompt to put version A and version B in a Diff format, where we include all the code of the first version but only the lines that differ in the second version. In this way we can reduce the amount of tokens in the input and make it easier for the model to focus on the relevant differences between the two versions.

To avoiding the bias of the model where it always selects the version with ++ since it is a common pattern where the PR version or the “update” version is better than the one with --, we decided to use different symbols for the versions, instead of using ++ and -- we used % and &, this way we can avoid the bias of the model and make it focus on the content of the decompilation rather than being influenced by the symbols used to represent the versions (informing the model inside the prompt).

Here we can see an example of the Diff design for the prompt, where version A and version B are represented with % and & respectively, and the prompt includes only the lines that differ between the two versions, while the rest of the code is included in full for both versions (the function is random, do not try to understand it):

| ver. A | ver. B | ver. prompt |
|---|---|---|
| <pre> 1 void complex(int a, char *b) { 2 long *f; 3 int h[10]; 4 if (a > 0) { 5 while (a < 10) { 6 printf("Value: %d\n", a); 7 a++; 8 } 9 } else { 10 goto end; 11 } 12 h[0] = 42; 13 end: 14 char c = b[id]; 15 f->g(h[i]); 16 (*(int *) (p + -4)) = 5; 17 }</pre> | <pre> 1 void complex(int a, char *b) { 2 long *f; 3 int h[10]; 4 if (a > 0) { 5 while (a < 10) { 6 printf("Value: %d\n", a); 7 a++; 8 } 9 } else { 10 goto end; 11 } 12 h[0] = '*'; 13 end: 14 char c = b[id]; 15 f->g(h[i]); 16 (*(int *) (p - 4)) = 5; 17 }</pre> | <pre> 1 void complex(int a, char *b) { 2 long *f; 3 int h[10]; 4 if (a > 0) { 5 while (a < 10) { 6 printf("Value: %d\n", a); 7 a++; 8 } 9 } else { 10 goto end; 11 } 12 % h[0] = 42; 13 & h[0] = '*'; 14 end: 15 char c = b[id]; 16 f->g(h[i]); 17 % (*(int *) (p + -4)) = 5; 18 & (*(int *) (p - 4)) = 5; 19 }</pre> |

4.4 Pull requests

Ghidra’s development is organized around PRs on GitHub, where contributors propose changes to the codebase that can include bug fixes, new features, or improvements to existing functionality. Each PR represents a specific set of changes that can affect the decompilation output in different ways. Fortunately in the GitHub repository the PRs have tags (in our case with the “Feature: Decompiler” label), since a lot of PR do not modify decompiler logic, we could easily select only those that were relevant for our evaluation, and that could provide meaningful insights into the impact of specific changes on decompilation quality.

For our evaluation, we selected a subset of PRs:

- **PR #8628**: improves the cleanup rule Rule2Comp2Sub so that it also handles *constant* subtractions that are rewritten during decompilation as additions with nega-

tive constants (e.g., `x + -0x1a` \rightarrow `x - 0x1a`), improving readability and bringing the output closer to typical C source style.

- **PR #8587**: extends the `constantptr` rule to automatically detect and correct one indexing patterns for spacebase constants (e.g., `*(undefined *) ((long)i * 0x30 + 0xaddr + (long)j * 4)` \rightarrow `globalArray[(long)i + -1].field[j]`).
- **PR #8161**: fixes an issue in `BlockWhileDo` where the decision to use overflow syntax could be made using unoptimized P-code and then treated as permanent: after P-code optimization the exit-condition block may no longer be “complex,” but the stale overflow flag would still block for-loop recovery; the patch re-checks that overflow-marked `BlockWhileDo` blocks are still complex at the time of for-loop recovery and clears the overflow decision when it is no longer justified.
- **PR #7253**: sorts `switch case` entries by their target address (i.e., the destination basic-block address) rather than leaving them in an arbitrary/disassembler-dependent order.
- **PR #6722**: fixes cases where the decompiler fails to recover submember array access due to `RulePtrArith` distributing an addition through an `INT_MULT` (e.g., `(idx + -0x30) * 4 + 4` \rightarrow `idx * 4 + (-0x30 * 4 + 4)`), which prevents matching the struct member; the rule now undoes this distribution and retries, allowing clean output such as `PTR_0041a1b8->ar[local_18]` instead of `*(undefined4 *)((int)PTR_0041a1b8 + local_18 * 4 + 4)`.

Ghidra is a large and complex codebase, PRs are modifications that can affect some specific aspect of the decompilation output, it will never be the case that a PR will improve all the decompilations or change the decompilation core aspects in a drastic way, but rather it will improve some specific cases and make them more readable, while in other cases it can make them worse, so we want to evaluate the impact of these specific changes on the decompilation quality and readability, and see if the LLM-based evaluation can capture these improvements or regressions in a meaningful way.

4.5 Reporting

For each PR the client writes a dedicated report file (`reports/#PR.json`) containing aggregate statistics such as mean Δ PPL across samples, mean baseline/PR/source perplexities, and their abstracted counterparts, and then a list of results for each model: (i) the list of evaluated samples with their individual perplexities and Δ PPL values, and (ii) results of qualitative analysis from the LLM. At the end of the full run, a `final_report.json`




Figure 4.8: Example of the Dogbolt interface, where we can see the decompilation outputs from three different decompilers

aggregates all per-PR summaries. If a report already exists, the client loads it and skips re-evaluation, enabling robust recovery after interruptions.

Finally, to avoid leaving a large model resident in GPU memory, the client calls the server cleanup endpoint at the end of the run, explicitly unloading the active model.

4.6 Dogbolt

Our analysis is focused on evaluating the impact of specific PRs on decompilation readability and humanness for Ghidra, but it is also interesting to see how the decompilation quality has evolved over time across different decompilers, and to have a more general overview of the trends and patterns in the decompilation output horizontally (not vertical on only Ghidra). For this reason, we also created a framework of three different decompilers (Ghidra 12.0.1, Binary Ninja 5.2, and Hex-Rays 9.2) using Dogbolt,³ an interactive online *Decompiler Explorer* that displays outputs from multiple decompilers for the same input binary, as can be seen in Figure 4.8. Dogbolt is community-maintained and open source but intentionally does not provide a public automation API. We used it in a restricted version of our dataset (only the project “file”), while all large-scale and reproducible measurements in this study were executed in our own local pipeline.

Other differences with our main evaluation framework are the fact that the outputs of decompilers are not directly comparable using Diff since they generate different code, names, and formatting, so we rely, in this case, on giving all the code to the LLM. For the evaluation we use the same prompting strategy described in Section 4.3.5, but with a pipeline where for every function we create three different prompts, one for each pair of decompilers (Binary Ninja vs Ghidra, Ghidra vs Hex-Rays, Binary Ninja vs Hex-Rays), and then we aggregate the results to have a more general overview of the trends and patterns in the decompilation output across different decompilers. We also had to parse the results for extracting the functions from the Dogbolt results since this is external to the dataset maker framework discussed in Section 4.1.

³<https://dogbolt.org/>

Chapter 5

Results

5.1 LLM performance

An empirical evaluation of the *qwen-3* and *deepseek-r1* models was conducted across two distinct tasks: evaluating the “humanity” of decompiled code via generation (LLM Judge) and calculating code perplexity (Score). The performance was measured in terms of execution time and peak VRAM usage.

As shown in Figure 5.1, there is a dichotomy in execution times between the two operations. The *score* operation is inherently fast, with both models completing most passes in under 50 seconds. However, the *generate* operation exhibits significantly higher and more erratic execution times, particularly for *qwen-3*, which frequently exceeds 200 seconds and reaches up to nearly 500 seconds.

Figure 5.2 presents a counter-intuitive finding: calculating perplexity (*score*) requires a higher mean and maximum Peak VRAM than generating text (*generate*).

5.1.1 Generation

A closer inspection of Figure 5.3, that is the data of only generation, reveals the root cause of the discrepancy between the two models: The scatter plot for “Generated Tokens vs Execution Time” shows a linear correlation, confirming that generation time is bottlenecked heavily by the length of the model’s output rather than the input context. Critically, *qwen-3* routinely generates between 2,000 and 4,000+ tokens during the evaluation task and often reaches our limit of 4096 new tokens.

Given that the task is an LLM Judge producing an evaluation of code “humanity,” a




Figure 5.1: Execution time for each evaluated model, score vs generate.




Figure 5.2: Peak VRAM for each evaluated model.

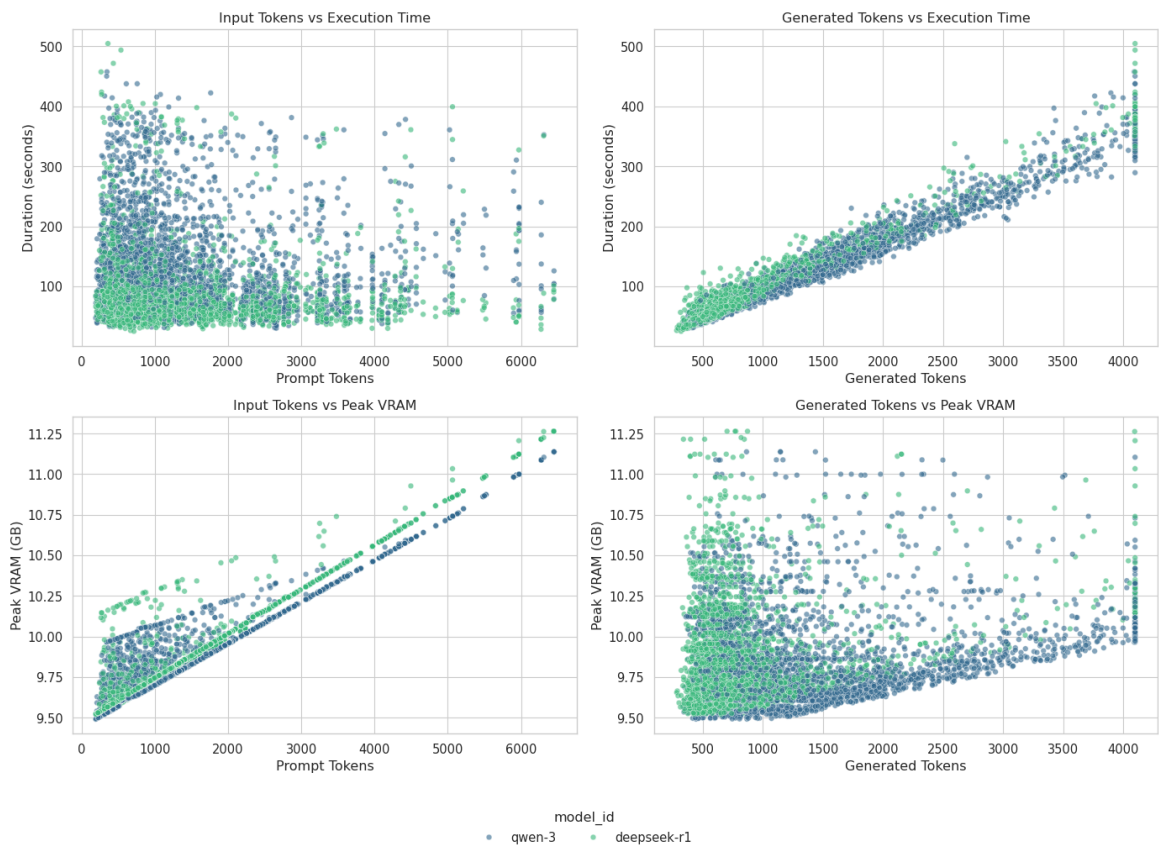


Figure 5.3: Generated tokens vs execution time and peak VRAM for each evaluated model.

rationale exceeding 4,000 tokens is highly suspicious. This suggests that *qwen-3* may be suffering from severe verbosity, repeating the input code, or failing to trigger stop tokens appropriately. In contrast, *deepseek-r1* is significantly more concise, rarely exceeding 2,500 tokens; it can still reach the 4096-token limit, but much less frequently, which translates directly to more predictable and efficient execution times.

In Figure 5.4 we can see an example of the output of *qwen-3* for a specific function, we can see that the model is repeating the same pattern of thoughts, which is a common symptom of verbosity and lack of proper stop token triggering. This behavior not only leads to unnecessarily long outputs but also significantly increases execution time and VRAM usage, as the model continues to generate tokens without producing meaningful content. This verbosity issue is a critical flaw for the LLM Judge task, as it undermines the model’s ability to provide concise and relevant evaluations of code “humanity,” and it also creates significant computational inefficiencies that could be prohibitive in larger-scale evaluations or with longer input contexts. Maybe this could be mitigated by implementing stricter stop conditions or by fine-tuning the model to better understand the task requirements and avoid unnecessary verbosity, but this will need further investigation.

Furthermore, analyzing the memory consumption in the bottom row of Figure 5.3 reveals a distinct behavior regarding how VRAM scales during the generation process. The “Input Tokens vs Peak VRAM” subplot demonstrates a strange phenomenon: Under 2500 tokens we can observe a cloud of data points but exceeding that number of tokens we observe an uninterrupted linear relationship, this line is also the minimum boundary of the data. In the line, *deepseek-r1* exhibits a slightly higher memory scaling compared to *qwen-3*, consuming marginally more VRAM for the same input length.

5.1.2 Score

The linear scaling of VRAM during the *score* operation (Figure 5.5, right panel) is a major scalability risk. At roughly 3,800 input tokens, both models push past 12 GB of VRAM. If the dataset contains an enormous larger input (e.g., 8k to 32k tokens), the current perplexity scoring methodology will inevitably result in Out of Memory failures on standard consumer GPUs.

Fortunately, the *generate* operation does not exhibit this issue, as it processes tokens sequentially and can leverage the KV cache to manage memory more efficiently.

The duration is mostly flat (0–20 seconds), but spikes occur for both models. This implies either batch-processing artifacts or that the evaluation dataset contains many decompiled functions of the exact same token length that trigger specific internal computational bottlenecks. A possible explanation is that runtime is influenced not only by input length, but also by the distribution of per-token loss values. This hypothesis requires additional

So, the structure of the switch in the Diff is different from the Source. Therefore, neither A nor B is structurally isomorphic to the Source.

But the question is which of A or B is closer. Since the Source's switch is based on m->type, which are integer constants, while the Diff's switch is based on characters, which are different.

But perhaps the code in the Diff is a different function, and the question is about whether the changes in the Diff make it closer to the Source structure.

Alternatively, perhaps the code in the Diff is a different function, but the question is about whether the changes in the Diff (from A to B) make it closer to the Source.

Wait, perhaps the function in the Diff is a different function, but the question is about whether the changes in the Diff make it closer to the Source structure.

Alternatively, perhaps the code in the Diff is a different function, but the question is about whether the changes in the Diff (from A to B) make it closer to the Source.

Alternatively, perhaps the function in the Diff is a different function, but the question is about whether the changes in the Diff make it closer to the Source structure.

Alternatively, perhaps the function in the Diff is a different function, but the question is about whether the changes in the Diff (from A to B) make it closer to the Source.

Alternatively, perhaps the function in the Diff is a different function, but the question is about whether the changes in the Diff make it closer to the Source structure.

Wait, perhaps I'm stuck. Let me try to think differently.

The evaluation criteria are:

1. Structural Isomorphism: Does the version have the same control flow as the Source?

Figure 5.4: Example of repetitions in *qwen-3* output

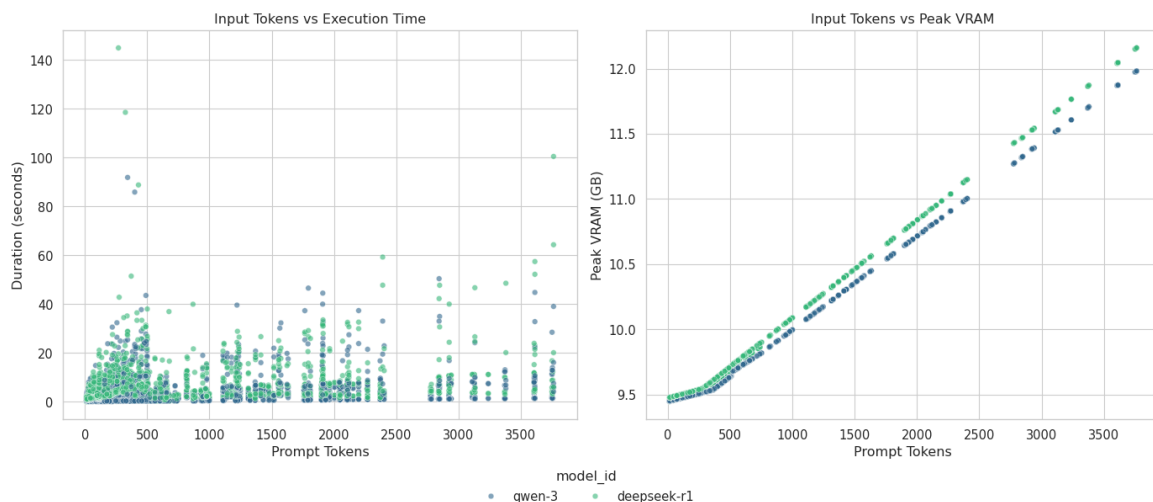


Figure 5.5: Score vs execution time and peak VRAM for each evaluated model.

targeted experiments, and further studies will be necessary to confirm it.

While both models are capable of performing the required tasks, *deepseek-r1* is objectively better suited for the LLM Judge (*generate*) role due to its restraint and conciseness, avoiding the computationally expensive verbosity traps that plague *qwen-3*. However, for the perplexity (*score*) task, the architecture of the operation itself poses a severe hardware bottleneck that scales poorly with larger decompiler outputs.

5.2 Perplexity as a Metric for “Humanness”

In Figure 5.6 we can see the distribution of perplexity values for the original source code, the base code, and the PR code, as well as their abstracted representations. We can observe that the original source code has generally higher perplexity values compared to the decompiled versions, which is unexpected since the original source code should be more “natural” and predictable than the decompiled output. This suggests that the decompilation process may introduce certain patterns or structures that are more familiar to the language model, leading to lower perplexity scores, while the original source code may contain more variability and less predictable constructs that result in higher perplexity. We can observe that the perplexity distribution calculated with *deepseek-r1* is generally higher than the one calculated with *qwen-3*.

Another observation is that the abstracted representations of the code tend to have higher perplexity values compared to their original counterparts. This is likely because the abstraction process removes specific identifiers and literals, which can make the code less

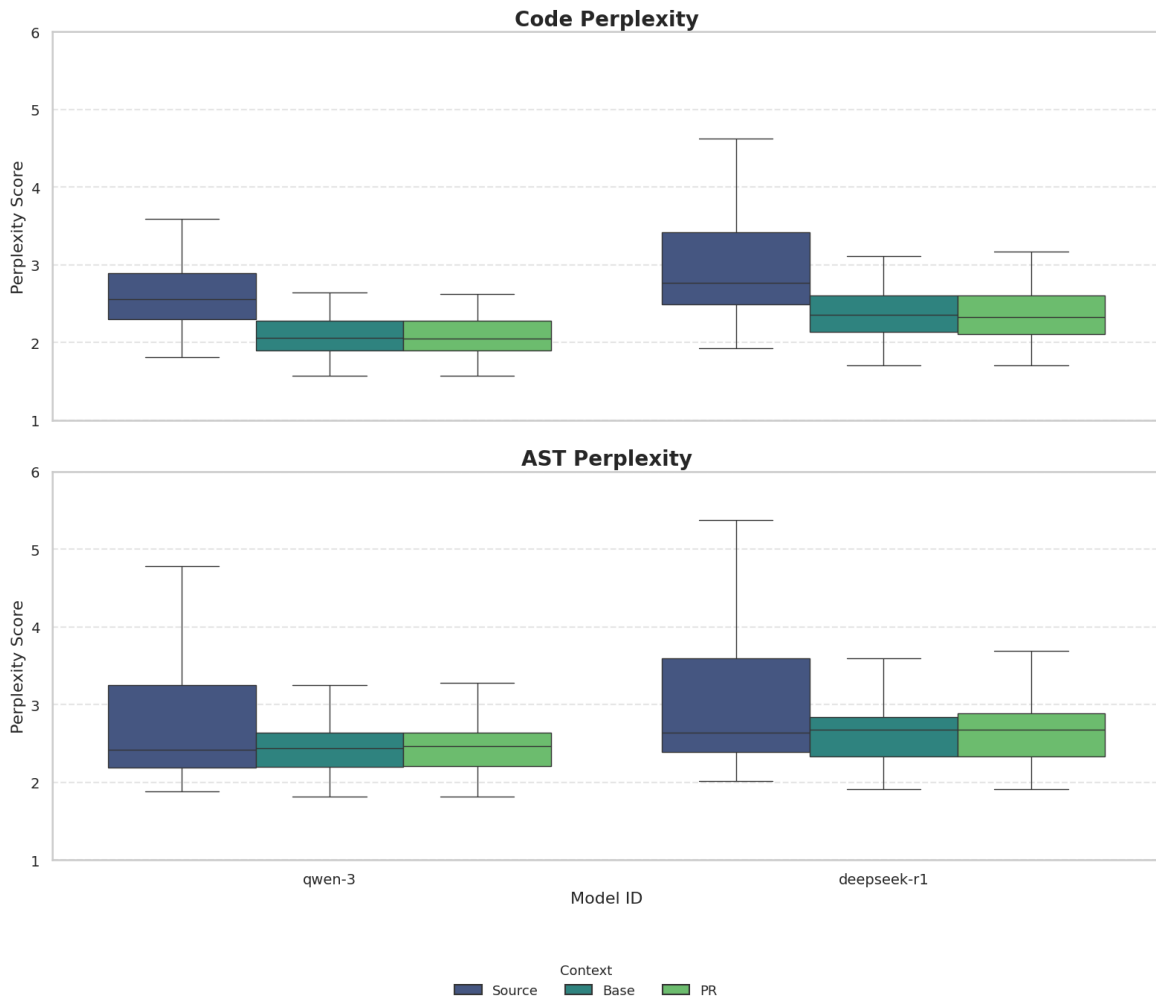


Figure 5.6: On the top, the perplexity values across original source code, base code, and PR code. On the bottom, perplexity values across abstracted representations of the same functions.

predictable and more “surprising” to the language model, but even in this case, the original source code still has higher perplexity than the decompiled versions, reinforcing the idea that the decompilation process may be introducing more predictable patterns into the code.

In Figure 5.7 we can observe that the original code exhibits a wider, slightly flatter distribution with a higher mean loss compared to the decompiled output, which is characterized by a sharper peak shifted towards zero. This visualization highlights a counter-intuitive phenomenon: despite the original source code being the “human ground truth”, the language model finds the decompiled code significantly more predictable.

We attribute this behavior to Token Inflation and Loss Dilution: As indicated by the token counts in the figure (e.g., $\sim 168k$ tokens for decompiled vs $\sim 120k$ for original source), the decompilation process introduces a substantial *token inflation*. Ghidra generate verbose, explicit code full of boilerplate structures (e.g., redundant casts, standard control flow patterns, explicit initializations, and restricted vocabulary). These pattern tokens are syntactically rigid and in a context where are used, they are easy to predict for the model, leading to many tokens with very low loss values (close to zero). Their sheer volume effectively dilutes the mean loss, artificially lowering the overall perplexity score compared to the denser, more information-rich human code. In contrast, the original source code reflects human authorship, which includes domain-specific naming conventions, creative syntactic choices, and stylistic variability. This “human entropy” flattens the density curve and shifts the mean loss to the right, as the model is more frequently “surprised” by the programmer’s unique choices compared to the machine’s standardized output.

In Figure 5.8 we can see the AST version of the previous analysis, where we abstracted away variable names and literals to focus on the structural aspects of the code. Comparing these results with the previous analysis on raw code tokens, we observe two critical phenomena:

1. **Persistence of structural inflation:** Even in the anonymized form, the *token inflation* remains significant. The decompiled AST contains $\sim 107k$ tokens compared to $\sim 70k$ for the original source (+52%). This confirms that the verbosity of the decompiled code is not merely lexical (e.g., long variable names) but syntactical. The decompiler introduces explicit casts, redundant blocks, and verbose control flow structures that persist even after anonymization, continuing to dilute the mean loss with predictable tokens.
2. **The closer entropy:** Unlike the raw code analysis, where the gap between the distributions was pronounced, the AST distributions for Source and Decompiled code are more similar in shape. The difference in Mean Loss has narrowed (e.g., for *deepseek-r1*, the gap represents only ~ 0.047 , compared to larger margins in the raw code).

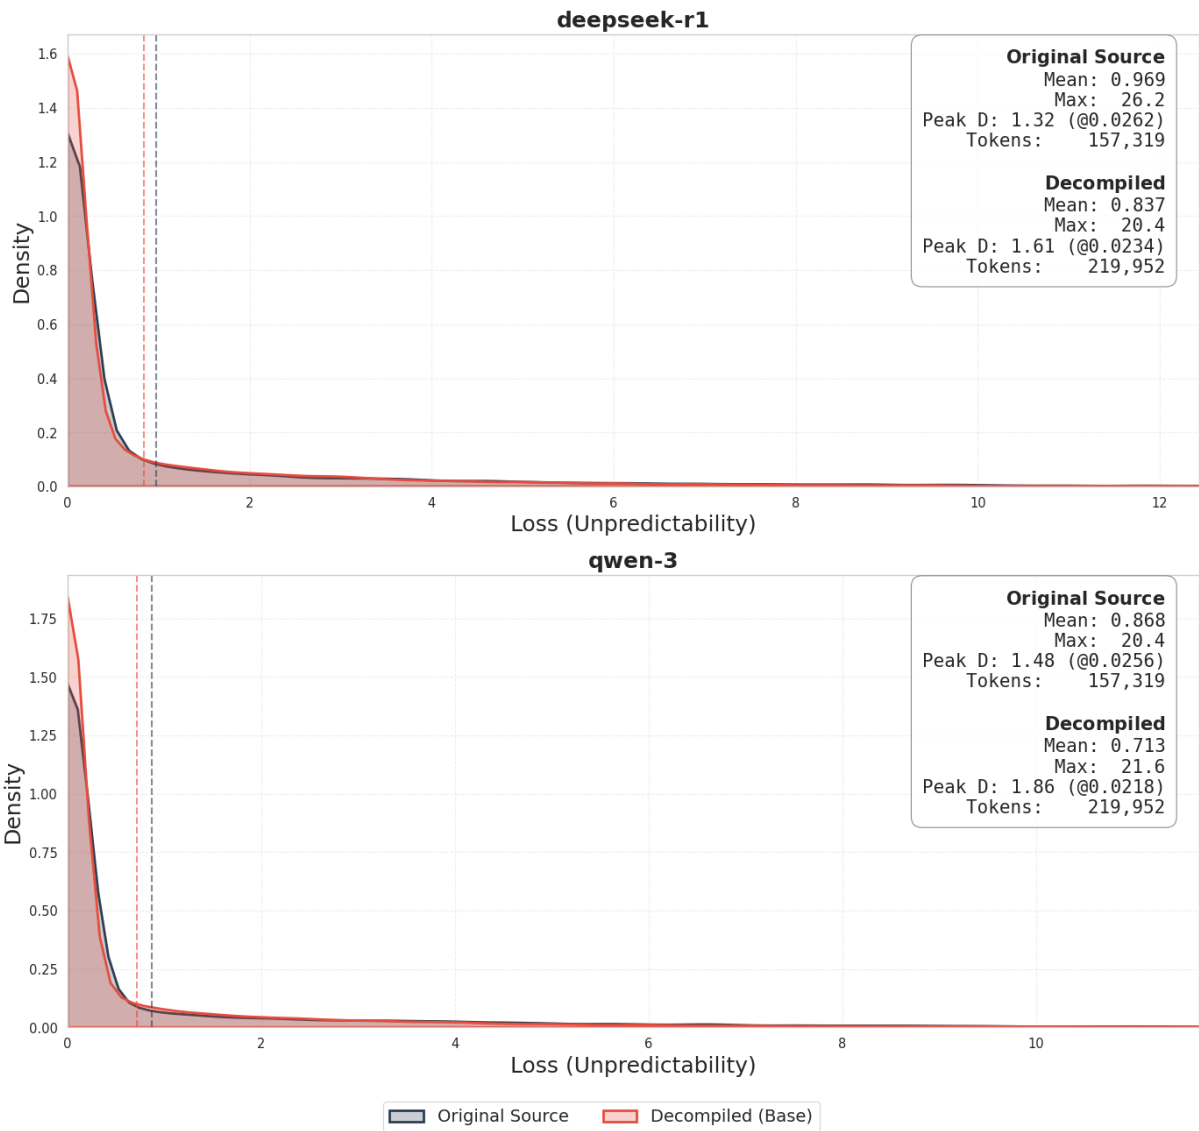


Figure 5.7: Density distribution of cross-entropy loss values for all tokens in the original source code and the decompiled versions.

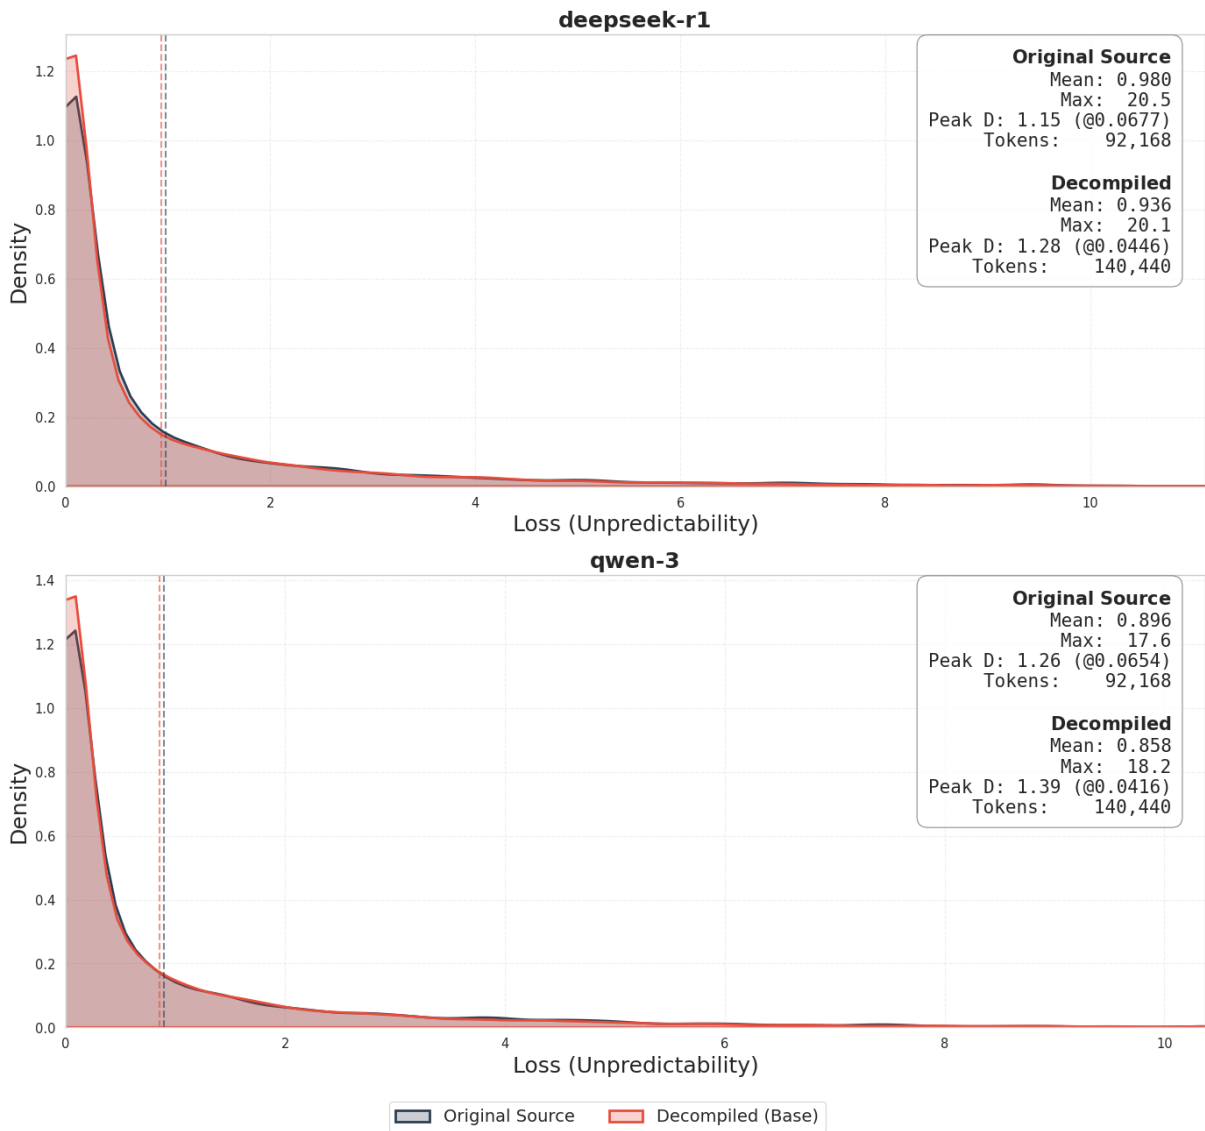


Figure 5.8: Density distribution of cross-entropy loss values for all tokens in the original source code and the decompiled versions.

This convergence suggests that the *lexical entropy* was a discriminator in the previous analysis.

- In the raw code, human-written names provided high variability (surprisal), while decompiled names were generic.
- In the AST version, the anonymization process effectively “standardizes” the two codes; Consequently, the source code becomes more predictable by the model.

The token-level analysis allows us to verify the tokens that contribute the most to the loss, and consequently to the perplexity, in a specific function. We can observe that the token itself is not the only factor that contributes to the loss, but especially the context in which it is used. For example, in Figure 5.9 we can see the loss values for the token “LAB”, which is a common label used in the decompiled code to indicate jump targets. This token has a low loss value (1.664) when it appears after the `dowhile` loop, but it has a much higher loss value (11.812) when it appears inside the `if` allowing a flow branch to ignore the condition and enter the scope without checks.

```
    } while (uStack_38._4_2_ != 10);
LAB_00101b6a:
    if (__ptr != (ushort *)0x0) {
LAB_00101bb4:
        free(__ptr);
    }
}
```

Figure 5.9: Loss values for `xls_parseWorkbook` in the decompiled base code

Obviously every function has a different distribution of loss values, and some functions may be more “natural” than others. In Figure 5.10 we can see an example of the loss values for `xls_parseWorkbook` function from the DeepSeek analysis, remembering that the perplexity is calculated as the exponential of the mean loss (Section 3.8), we can see that the original source code has a perplexity of ~ 3.06 (mean loss ~ 1.12), while the decompiled base version has a perplexity of ~ 2.4 (mean loss ~ 0.89). The max values for the loss are higher for the decompiled version, in contrary to the global distribution where the decompiled code had a sharper peak towards zero, but then when we look at the anonymized version of the same function, the original source code became the one with a higher max loss value than the decompiled version. However, the anonymization process manages to bring the two versions significantly closer, both in terms of loss values and, consequently, perplexity.

In Figure 5.11 we investigate the *Entropy Variance* of the code by analyzing the standard deviation of the token loss. While the mean loss indicates the average predictability, the standard deviation reveals the *dynamic range* of the code complexity.

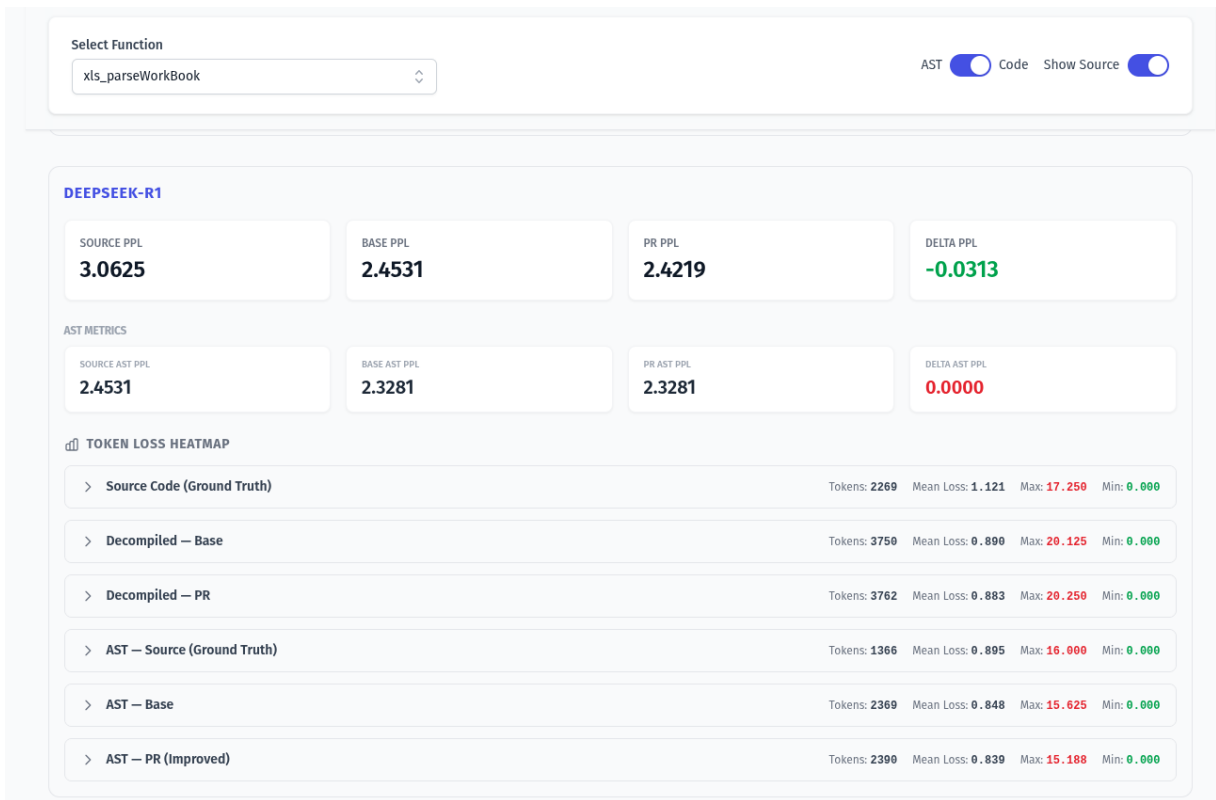


Figure 5.10: Example of loss values for a function in the decompiled code.

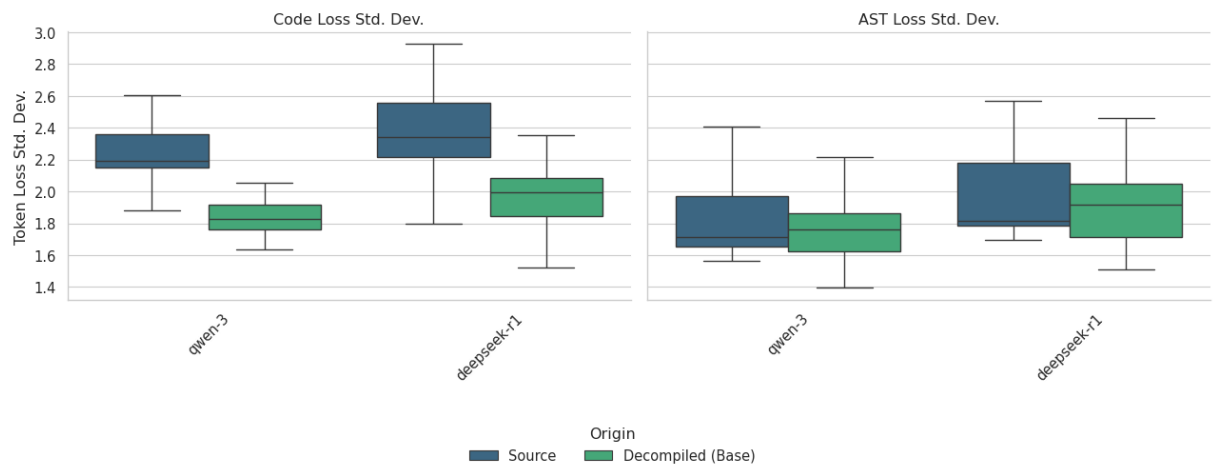


Figure 5.11: Standard deviation of loss values across all functions in the normal version and anonymized version.

We can observe a clear trend across both models:

- **Difference in Variance:** The original source code consistently exhibits higher standard deviation compared to the decompiled versions. This confirms that human-written code is characterized by *burstiness*: it alternates between low-entropy boilerplate and high-entropy domain-specific logic. The language model struggles to predict this rhythm, leading to fluctuating loss values. The decompiled code shows significantly lower variance. This reflects the *monotonicity* of machine-generated code. In our case Ghidra applies consistent transformation rules throughout the binary, resulting in more predictable results.
- **The Lexical Factor:** Comparing the code panel with the AST one, we notice that the gap between Source and Decompiled shrinks significantly in the AST representation. This implies that a substantial portion of the entropy variance in human code is driven by *lexical choices* (variable naming and literals) as we predicted. Once these are removed, the structural variability of human code is only marginally higher than that of the decompiled code.

This result reinforces our conclusion: in our case (Ghidra vs Source) Human-Likeness is not defined by raw predictability (where the machine wins), but by the **variance of unpredictability**. A “natural” code signature is one that surprises the model in inconsistent, context-dependent bursts, rather than being uniformly predictable. Meanwhile, for the anonymized code, the gap in variance is smaller, suggesting that even at the structural level, human code retains a degree of unpredictability that machine-generated code lacks.

5.2.1 Other Decompilers

The Dogbolt analysis allows us to compare the perplexity values of the decompiled versions from other decompilers (Hex-Rays and Binary Ninja) with the original source code and the Ghidra decompilation, with a subset of the original database (we cannot compare in an absolute way the plot in Figure 5.6 with these but only the relative ordering of the distributions). We can see in Figure 5.12 that the perplexity values for the decompiled versions resulted by the Dogbolt analysis, are generally in line with the Ghidra ones. We can see that the original source code still has higher perplexity values compared to the decompiled versions, and the abstracted representations of the code still tend to have higher perplexity values compared to their original counterparts. Another key observation is that the binary ninja has a perplexity distribution that is more similar to the original source code compared to the Ghidra decompilation, meanwhile Hex-Rays has a distribution in the middle between the other two decompilers, then we can observe that every distribution with *deepseek-r1* has a higher mean perplexity than the same distribution with *qwen-3* like we have seen previously.




Figure 5.12: Perplexity values for the original source code and the decompiled versions from other decompilers.

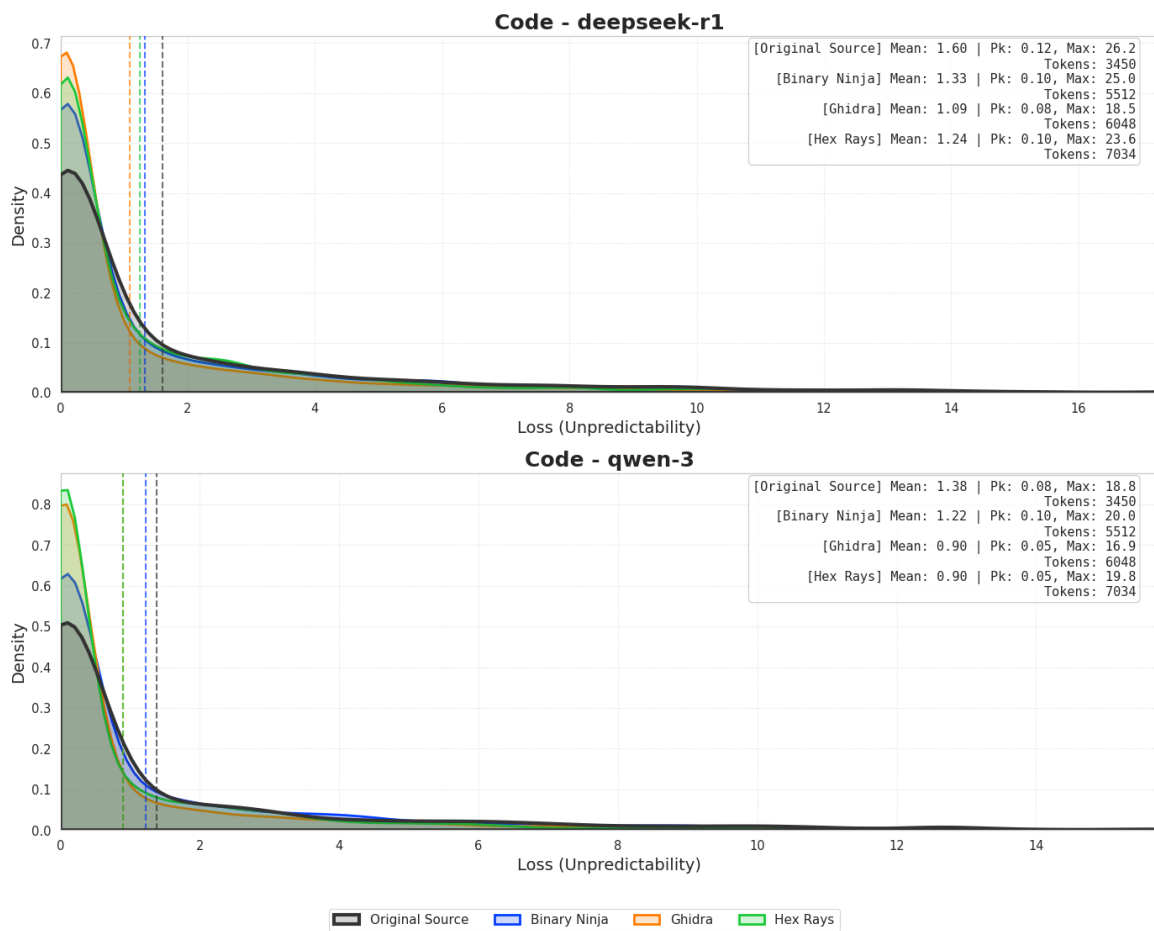


Figure 5.13: Loss values for the original source code and the decompiled versions from other decompilers. Code version

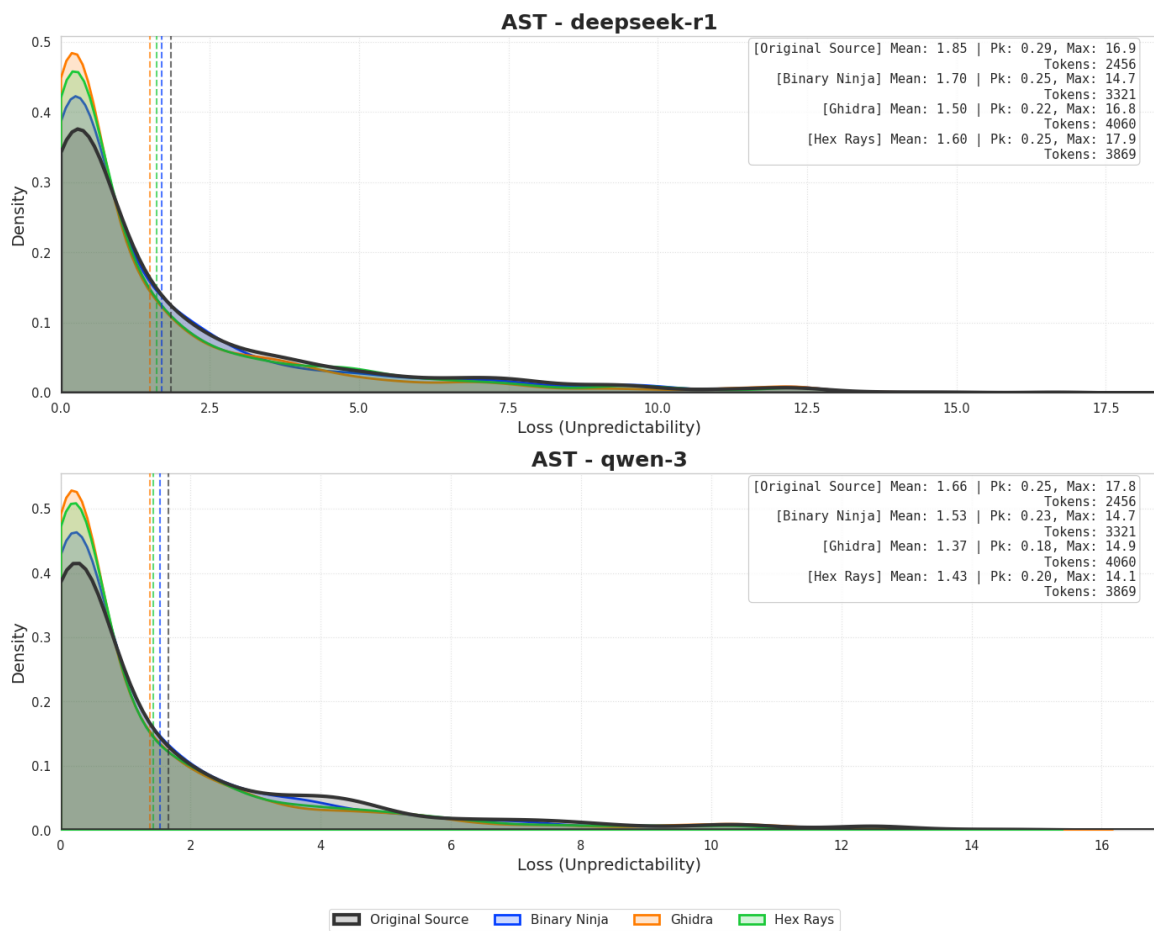


Figure 5.14: Loss values for the original source code and the decompiled versions from other decompilers. Abstract version

The density plots in Figure 5.13 confirm the phenomena observed in the previous sections across both LLM judges. Focusing on the raw *Code Tokens* (top row), we can observe a distinct hierarchy in predictability:

- **Ghidra** exhibits the sharpest peak near zero and the lowest mean loss (~ 1.09 for DeepSeek-R1 and ~ 0.90 for Qwen-3). This reaffirms its tendency to generate highly rigid and predictable code.
- **Hex-Rays** follows closely, showing a similar sharp peak but with a slightly higher mean loss and a much larger token count (over 7,000 tokens in this sample), in particular with the *deepseek-r1* while with *qwen-3* we can see that the mean is the same as Ghidra but with a density at peak higher than the others (as the Max loss is much higher than Ghidra). The massive token inflation in Hex-Rays suggests extremely verbose syntactic and lexical choices that artificially dilute the cross-entropy loss.
- **Binary Ninja** stands out as the decompiler that most closely approximates the original source code. Its loss distribution is flatter, and its mean loss (~ 1.33 for DeepSeek-R1) is significantly closer to the original source (~ 1.60) than the other tools. It also exhibits less token inflation compared to Ghidra and Hex-Rays with a Max loss closer to the original source.

Meanwhile, in the *AST* plots in Figure 5.14, where identifiers and literals are abstracted, we observe the expected flattening of all distributions. By stripping away lexical entropy, the gap between human-written code and machine-generated code narrows. Notably, the massive token inflation seen in Hex-Rays raw code drops significantly during AST abstraction (from 7,034 down to 3,869 tokens), indicating that a vast majority of its predictability stems from highly repetitive lexical tokens, types, and literal declarations rather than purely structural blocks. Even in the abstracted form, Binary Ninja maintains the distribution shape most similar to the original source code and with Hex-Rays having less tokens, the density at peak now is less than the Ghidra one for *qwen-3*. Finally, comparing the two models across all subplots, we note that while *qwen-3* consistently produces lower absolute loss values (higher confidence/predictability) than *deepseek-r1*, both models almost perfectly agree (except for the code tokens for *qwen-3* where Hex-Rays have a higher peak density than Ghidra) on the relative ordering of the distributions. Original source code always retains the highest mean loss and widest variance, followed by Binary Ninja, Hex-Rays, and finally Ghidra.

A token-level analysis of the Hex-Rays decompilation (e.g., Figure 5.15) for the model *qwen-3*, reveals that Hex-Rays generates comments and other repetitive tokens that are highly predictable. The most predictable tokens are often those associated with those artifacts of the decompilation process. These tokens create a dense cluster of low-loss



Figure 5.15: Heatmap of loss values for Hex-Rays decompilation with model *qwen-3*



Figure 5.16: Heatmap of loss values for Hex-Rays decompilation with model *deepseek-r1*




Figure 5.17: Distribution of winners across all evaluated models.

values that significantly lower the overall perplexity score, despite the presence of more complex and less predictable tokens elsewhere in the code.

When we look at the same decompilation with the model *deepseek-r1* (Figure 5.16), we can see that the loss values are generally higher, but the distribution of low-loss tokens is still present, confirming that the anonymization process is a key factor in the analysis of the perplexity, as it standardizes the code through the removal of personalized identifiers, making the predictability of the code more dependent on its structural patterns rather than on specific lexical choices.

5.3 LLM-as-a-Judge Evaluation

In Figure 5.17 we can see the distribution of winners across all evaluated models, based on the qualitative judgments of the LLM-as-a-Judge. The two models, Qwen3 and DeepSeek,

have a significant number of Ties, but they also show a balanced distribution of wins between the base code and the PR code. We can see a bin for “Error” winners as well, this happens primarily when the model exceeded the token limit of 4096 tokens (reasonable limit set by us), since they are reasoning models (they create a context with the generated tokens inside `<think>` tags) sometimes the context becomes too large (often because they repeat thoughts entering a loop) and they finish the limit without giving inside the response a clear winner (e.g., `'Winner': 'X'`). Sometimes the response does not give a clear winner with the pattern “Winner: X” but it is still able to give a clear preference in the rationale, in this case we consider it as an error since the model did not follow the instructions correctly, but we can still extract some information from the rationale.

We previously said that a “Tie” is when the model always judges the same result regardless the switch of the base and PR code. We have a significant number of Ties for all models, watching the result we can observe a ratio of $\sim \frac{2}{3}$ (204/308) of times the model prefers the PR version (204) over the BASE version (104). This suggests a position bias, even using our Diff format explained in Section 4.3.5, this bias could be caused by the model’s tendency to favor the second option presented in a comparison, or it could be due to some subtle cues in the prompt formatting that inadvertently signal a preference for the PR version.

We used as BASE the version in commit `14c2495cb06ccd09324050c787ad0fc756d35dae`, which is at the moment of writing this master’s thesis, the latest commit in the Ghidra repository, some of these PRs are derived from a BASE version that is older than the one we used, so the differences between the BASE and the PR version may be more significant than what we expect from only the description of the PR, we decided to include these results to evaluate the history of version changes.¹

Using the BASE version as a reference point, we can analyze the impact of specific PRs on the model’s judgments of code “humanity” and fidelity.

Unfortunately, when evaluating fidelity to the original source code (“Code Fidelity” and “AST Fidelity”), the level of uncertainty is too high to draw meaningful conclusions. We observed a significant number of Ties and often no clear winner emerges, which suggests that the models struggle to discern fidelity differences between the BASE and PR versions. This could be due to the fact that both versions are relatively close in terms of their structural and lexical features, making it difficult for the models to detect subtle improvements or regressions in fidelity. We have decided to focus our analysis on the “humanity” judgments, where we observed more decisive outcomes and clearer preferences from the models, but we also decided to include the results of fidelity in the figures.

Figure 5.18 quantifies the extent to which the LLM-as-a-Judge decisions are driven by *snip-*

¹We ran a simulation also with the BASE version where each PR is derived, unfortunately in our dataset of 600+ functions there were no differences between the two decompilations for the older PRs that suffered this problem (8161 and 6722)




Figure 5.18: Length Bias Analysis

pet length rather than by perceived structural quality. For each comparison, we compute a signed length delta between the winning and losing candidates ($\Delta L = L_{\text{winner}} - L_{\text{loser}}$). Therefore, negative values indicate that the model tends to select the *shorter* decomposition as more human-like, while positive values indicate a preference for more verbose code.

Aggregating the results by judge on the AST prompts, we observe a consistent but mild conciseness preference: *deepseek-r1* has a mean bias of -2.95 characters, while *qwen-3* shows a stronger mean bias of -9.08 characters. In practical terms, both judges are slightly more likely to reward outputs that are marginally shorter once lexical noise is removed, with *qwen-3* being noticeably more sensitive to this factor.

However, the magnitude of these mean shifts is small relative to typical function sizes and to the scale of changes introduced by the evaluated PRs. This suggests that, in the Ghidra PR setting, length is not the primary driver of the models' preferences; rather, it acts as a weak prior that favors conciseness when two candidates are otherwise similar in structure and readability.

We include a consistency plot for each PR to investigate whether the models' judgments are consistent across the different prompts. These data are calculated as the percentage of times the model's preference remains the same when evaluating raw code versus evaluating the abstracted AST representation, we *do not* take into account the errors and Ties in this




Figure 5.19: Distribution of winners across the four prompts for PR #8628

evaluation, but only when the model have given a clear preference for both versions.

5.3.1 Handling Negative Constants – PR #8628

This PR improves handling of constant subtractions that are rewritten during decompilation as additions with negative constants (e.g., $x + -0x1a \rightarrow x - 0x1a$).

As shown in Figure 5.19 (the four prompts results explained in Section 4.3.5), the qualitative evaluation strongly favors the PR branch. In the “Code Blind (Human-Likeness)” and “AST Blind (Structural Humanity)” assessments, the PR version overwhelmingly wins against the BASE branch across both *deepseek-r1* and *qwen-3* models, often exceeding a 60% win rate. This indicates that rewriting expressions to use explicit subtraction rather than adding a negative immediate is universally recognized by the models as a more natural, human-readable idiom.

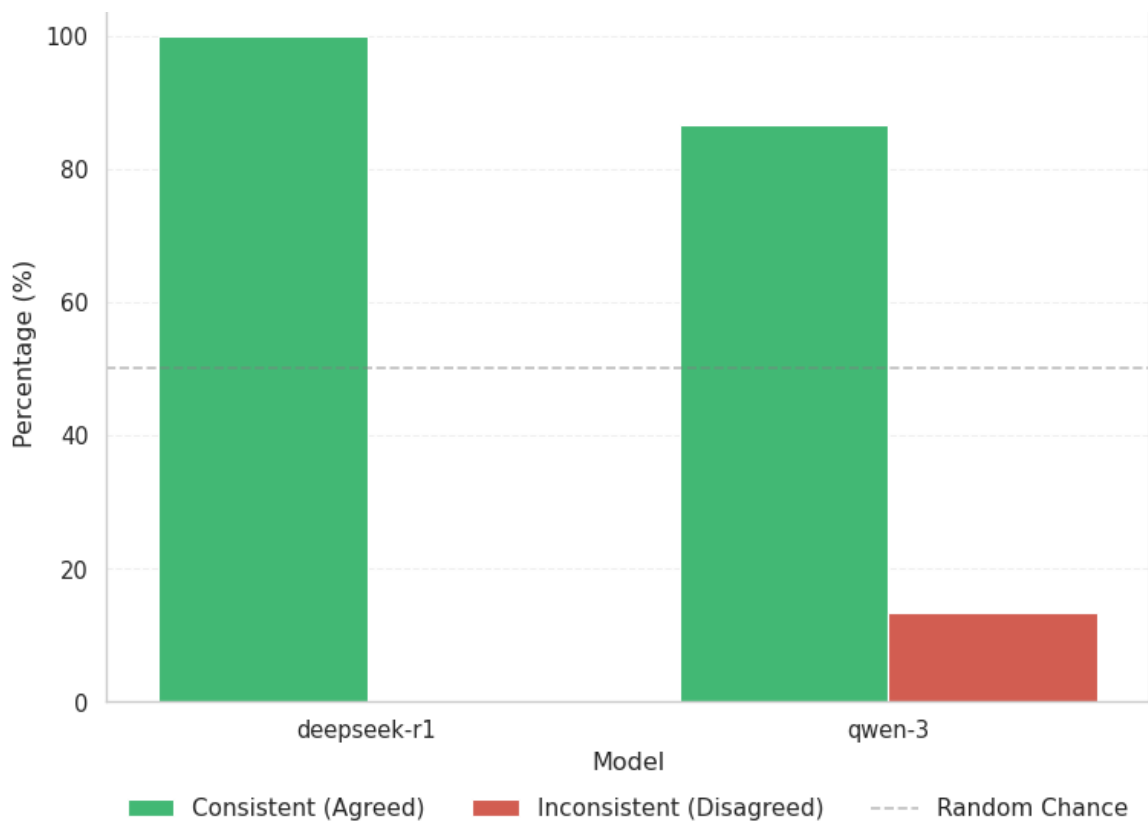


Figure 5.20: Consistency across the prompts for evaluating “humanity” for PR #8628

| | | | |
|--------|------------------------------|------|--|
| qwen-3 | Code: Humanity & Readability | PR | Candidate B uses parentheses to group the arithmetic operation, making the calculation of <code>pcVar16</code> more readable and idiomatic. This grouping clarifies the intent of the expression as a single logical unit, which aligns with human engineering practices that prioritize clarity and natural expression over mechanical, step-by-step decomposition. |
| | AST: Humanity & Readability | BASE | Candidate A's expression <code>id * (type)id * 2 + -2</code> avoids redundant parentheses around the arithmetic operation, aligning with natural C-style idioms where operator precedence is trusted. Candidate B's <code>((type)id * 2 - 2)</code> introduces an extra layer of parentheses that feels mechanically generated, prioritizing explicit grouping over human-readable brevity. Since the logic is identical, Candidate A's version has less artificial nesting, making it more idiomatic. |

Figure 5.21: Inconsistency in file `task-readstat_sav_parse_date-O3`

| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION |
|--------|------------------------------|--------|--|
| qwen-3 | Code: Humanity & Readability | BASE | Candidate A uses a direct and idiomatic comparison with <code>-1</code> , which is more readable and natural for a human programmer. Candidate B introduces an unsigned cast (<code>-1U</code>), which complicates the logic and may obscure the original intent. The original version avoids unnecessary type casting and maintains a cleaner, more straightforward expression. |
| | AST: Humanity & Readability | PR | Candidate B uses the unsigned suffix <code>-1U</code> in the condition, which explicitly indicates the comparison is between unsigned values. This aligns with human-like engineering by avoiding potential signed/unsigned type ambiguity, improving clarity and correctness. The <code>U</code> suffix reflects a deliberate choice to handle unsigned arithmetic, which is more idiomatic in C-like languages than the raw <code>-1</code> used in Candidate A. |

Figure 5.22: Inconsistency in file `task-readstat_readstat_parse_por-O0`

To further investigate the robustness of these judgments, we analyzed the consistency of the models' decisions between evaluating raw code and evaluating its abstracted AST representation (Figure 5.20). *Deepseek-r1* shows 100% consistency, meaning it never changed its preference between the raw code and the AST. In contrast, *qwen-3* exhibits an inconsistency rate of approximately 14% (two times). These two inconsistencies occur in the files `task-readstat_sav_parse_date` and `task-readstat_readstat_parse_por`.

In the first inconsistent case (Figure 5.21), the *qwen-3* model shifts its preference based on the presence of parentheses. When evaluating the raw code, it praises the PR for using parentheses to group arithmetic operations (`((type)id * 2 - 2)`), arguing it makes the intent clearer and aligns with human engineering practices. It does not talk about the sum of negative number but rely his judgment on parenthesis. However, when evaluating the AST, the model flips its vote to the BASE version. It suddenly argues that the extra parentheses introduced in the PR feel “mechanically generated” and that the BASE version is more idiomatic because it trusts operator precedence, ignoring again the sum of negative number. We can see that in this case, the model put in priority the parentheses over the sum of negative number. But the fact that the model is not consistent in its judgment when the parentheses are removed in the AST evaluation, suggests that it may be overfitting to superficial cues rather than truly understanding the underlying code structure and semantics.

The second inconsistency (Figure 5.22) highlights a conflict regarding type casting and

| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION |
|--------|------------------------------|--------|---|
| qwen-3 | Code: Humanity & Readability | PR | Candidate B uses more idiomatic C syntax for arithmetic operations. The expression <code>((ulong)uVar9 * 2 - 2)</code> is visually clearer than adding a negative value, and <code>iVar3 = iVar3 - 1</code> is the standard decrement pattern in C. These choices align with natural human programming patterns rather than mechanical transformations. |
| | AST: Humanity & Readability | PR | Candidate B's changes use more natural arithmetic expressions (e.g., <code>id - 1</code> instead of <code>id + -1</code>) and clearer operator precedence grouping (parentheses around compound expressions). These choices align with human programmer conventions of prioritizing readability through direct mathematical notation rather than literal arithmetic phrasing. The structural depth remains identical, but B's syntax better reflects idiomatic C-style programming patterns. |

Figure 5.23: Consistency in evaluating arithmetic expressions

literal suffixes. In the raw code evaluation, *qwen-3* prefers the **BASE** version because it uses a direct comparison with `-1`, calling the PR's use of an unsigned cast (`-1U`) an unnecessary complication that obscures intent. Yet, when evaluating the AST, the model flips to favor the PR. In this abstracted view, it argues that the explicit unsigned suffix (`U`) improves clarity and correctness, avoiding signed/unsigned ambiguity. This flip suggests that the model's judgment is heavily influenced by surface-level syntax rather than a deeper understanding of the code's semantics. The presence of the unsigned suffix in the PR version may have triggered a bias in the model when evaluating the raw code, leading it to view the **BASE** version as more straightforward. However, when the syntax is abstracted away in the AST, the model's preference shifts, indicating that it may be overfitting to specific tokens or patterns rather than consistently evaluating the underlying code quality and readability.

Despite these edge cases of inconsistency, the models generally demonstrate a strong ability to recognize the PR's structural improvements regardless of the lexical context. Figure 5.23 provides a clear example of this robust evaluation.

In this instance, *qwen-3* correctly identifies the core structural improvement in both the raw code and the AST representation. When evaluating the raw code, it notes that subtracting a positive integer (e.g., `iVar3 - 1`) is a standard decrement pattern and significantly more idiomatic than adding a negative value. When transitioning to the AST evaluation, the model maintains this exact reasoning, explicitly favoring the abstracted `id - 1` over `id + -1`.

This demonstrates that when a structural transformation is sufficiently distinct and maps clearly to standard human programming conventions (like direct mathematical notation), the model's judgment remains stable. The removal of specific variable names and literal values during the AST anonymization does not disrupt its ability to identify the more human-like syntactic pattern.

When the models have preferred **BASE** over the PR, we can see that is simply hallucinating reasons to justify, not recognizing the context of the change like in Figure 5.24.



Figure 5.24: Hallucination of reasons to prefer BASE to PR

5.3.2 Indexing Pattern Recognition – PR #8587

This PR focuses on automatically detecting and correcting indexing patterns for space-base constants (e.g., `*(undefined *)((long)i * 0x30 + 0xaddr + (long)j * 4) → globalArray[(long)i + -1].field[j]`).

As seen in Figure 5.25, the results for the raw code blind evaluation are exceptionally decisive: both *deepseek-r1* and *qwen-3* unanimously prefer the PR output 100% of the time. Transforming obscure, hardcoded pointer arithmetic into clean array indexing is universally recognized by the models as a massive improvement in human readability.

We can observe in our study that the *Deepseek-r1* model suffers from a light position bias in contrast to the *qwen-3* model, which always shows preference to the PR version.

This PR in particular only modifies 7 functions from the base branch.

Because of this small sample size, Figure 5.26 shows binary extremes. *Deepseek-r1* and *qwen-3* are 100% consistent in their evaluations across the 7 functions.

In Figures 5.27 and 5.28 we can observe two examples where the majority of cases correctly identify the modifications of the PR. In this case the two models identify that the PR’s output avoids raw pointer arithmetic in the raw code, and when evaluating the AST, it continues to recognize the underlying structural pattern of “direct array access” as the superior, more human-like engineering choice, regardless of whether the specific variable names are visible or not.

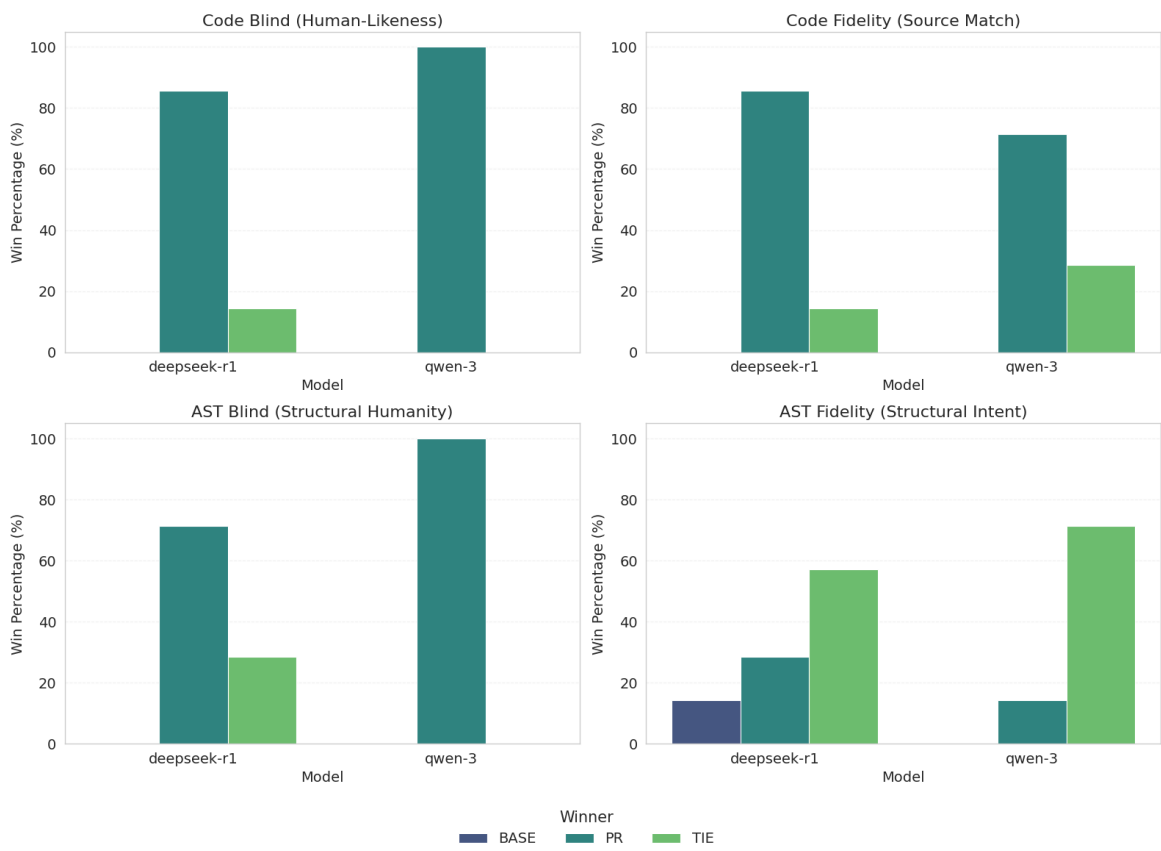


Figure 5.25: Distribution of winners across the four prompts for PR #8587

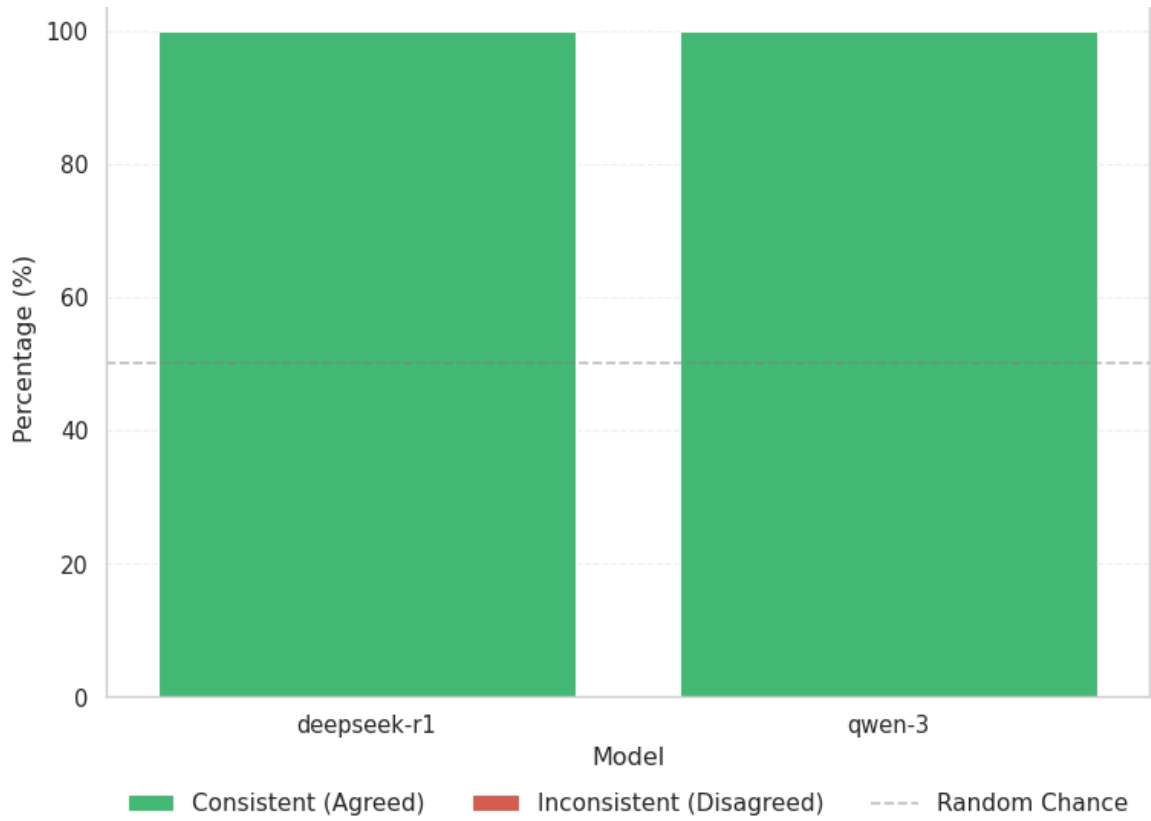


Figure 5.26: Consistency across the prompts for evaluating “humanity” for PR #8587

| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION |
|--------|------------------------------|--------|---|
| qwen-3 | Code: Humanity & Readability | PR | Candidate B uses a named array reference (<code>_dta_timestamp_parse_actions</code>) with a calculated index, which is more readable and maintainable than Candidate A's hardcoded offset (<code>0x102361</code>). This approach avoids magic numbers and aligns with human engineering practices by explicitly referencing a named data structure, making the code's intent clearer and more maintainable. ► Raw Response |
| | AST: Humanity & Readability | PR | Candidate B's switch condition uses a more abstract expression (<code>id[id + id + 1]</code>) compared to Candidate A's hardcoded offset (<code>0x102361 + id</code>). Avoiding magic numbers and using dynamic expressions aligns with human-like code design, which prioritizes readability and abstraction over low-level, machine-like constructs. The structural change reduces reliance on opaque constants, improving idiomatic clarity. ► Raw Response |

Figure 5.27: Example of consistency in evaluation by qwen-3

| | | | |
|-------------|------------------------------|----|--|
| deepseek-r1 | Code: Humanity & Readability | PR | Candidate B uses a cleaner and more natural approach by accessing an array with an index that includes variable offsets, which is a common human-readable practice. The code is more modular and less dependent on raw pointer arithmetic, making it easier to follow and maintain. |
| | AST: Humanity & Readability | PR | Candidate B uses a cleaner and more readable switch statement with direct array access, which is a natural choice for an experienced human programmer. This approach avoids complex pointer arithmetic and is more idiomatic, aligning with human-like code structure and readability. |

Figure 5.28: Example of consistency in evaluation by deepseek-r1

5.3.3 Final Transform for While – PR #8161

This PR addresses a possible issue with the handling of the `final_transform` method of the `while` BasicBlocks, which could lead to not recover a `for` loop where it could be possible instead.

Unfortunately, the difference between the BASE and PR versions in our dataset is not inherent to the PR itself, but rather a consequence of the fact that the BASE branch was generated with a newer version of Ghidra (post-PR) while the PR branch was generated with an older version of Ghidra. We can see that in the fork repo of the PR branch, where it contains the message “This branch is 2 commits ahead of and 1948 commits behind NationalSecurityAgency:master”, and consequently the evaluation is not really about the PR but about the differences between two versions of Ghidra, which makes it difficult to draw conclusions about the specific impact of the PR’s changes.

Looking at the aggregate data in Figure 5.29, the impact of the AST abstraction on this specific dataset is profound. In the “Code Blind” evaluation, the models are largely undecided (high “TIE” rate) or slightly favor the older PR branch due to localized quirks like the one seen in Figure 5.33. However, in the anonymized evaluation, there is a massive shift in preference toward the BASE branch. This suggests that the structural improvements accumulated in the 1948 commits of the BASE branch become much more apparent to the models once distracting lexical artifacts are removed.

Consequently, as shown in Figure 5.30, the inconsistency rate for this dataset is notable with more than 25% for *deepseek-r1* and over 30% for *qwen-3*.

Nevertheless, we can still observe a few interesting examples:

In Figure 5.31, we can clearly see a paradox: although this PR was explicitly designed to improve `for`-loop recovery, the evaluated BASE output still appears more idiomatic. In our dataset, this is explained by the branch mismatch discussed above: the BASE branch was generated from a much newer Ghidra state, while the PR branch is based on an older snapshot. As a result, the older PR output falls back to a more convoluted `do-while` structure with manual counter management. Encouragingly, *deepseek-r1* consistently identifies the BASE version as more human-readable in both raw-code and AST evaluations, explicitly

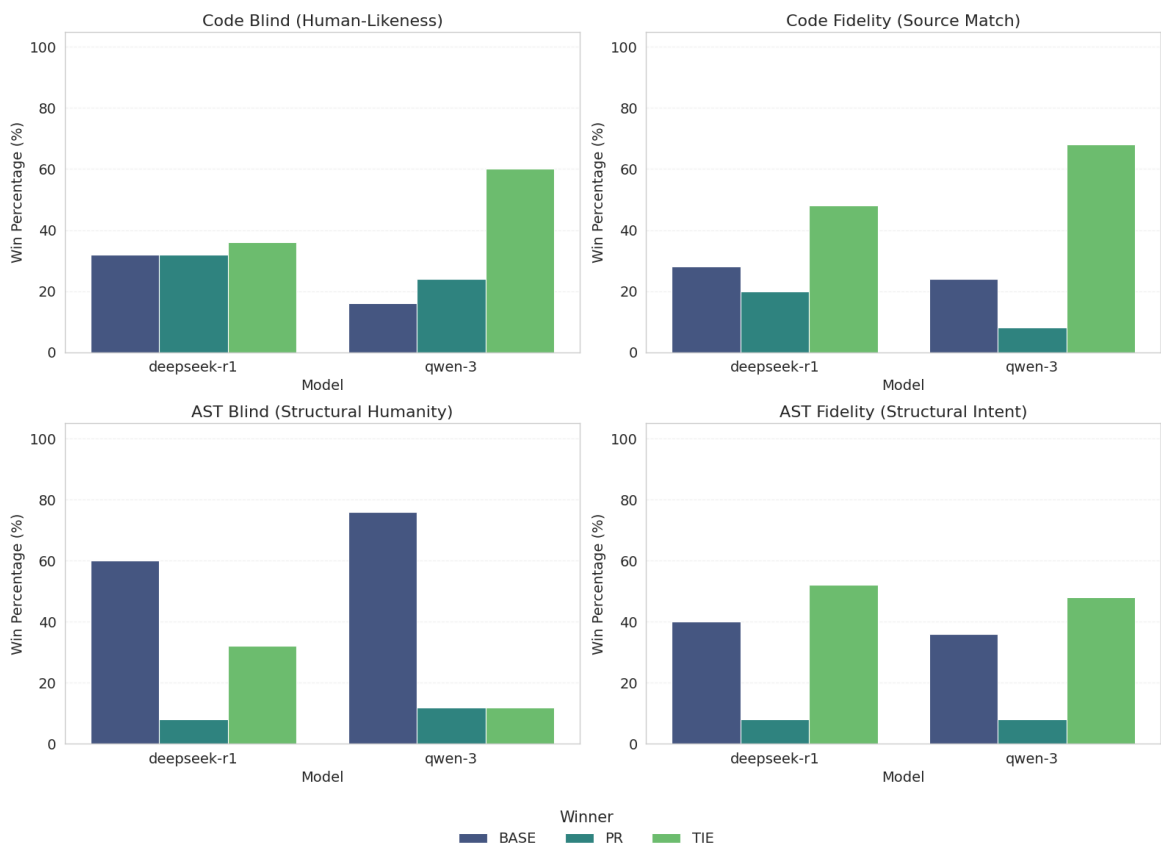


Figure 5.29: Distribution of winners across the four prompts for PR #8161

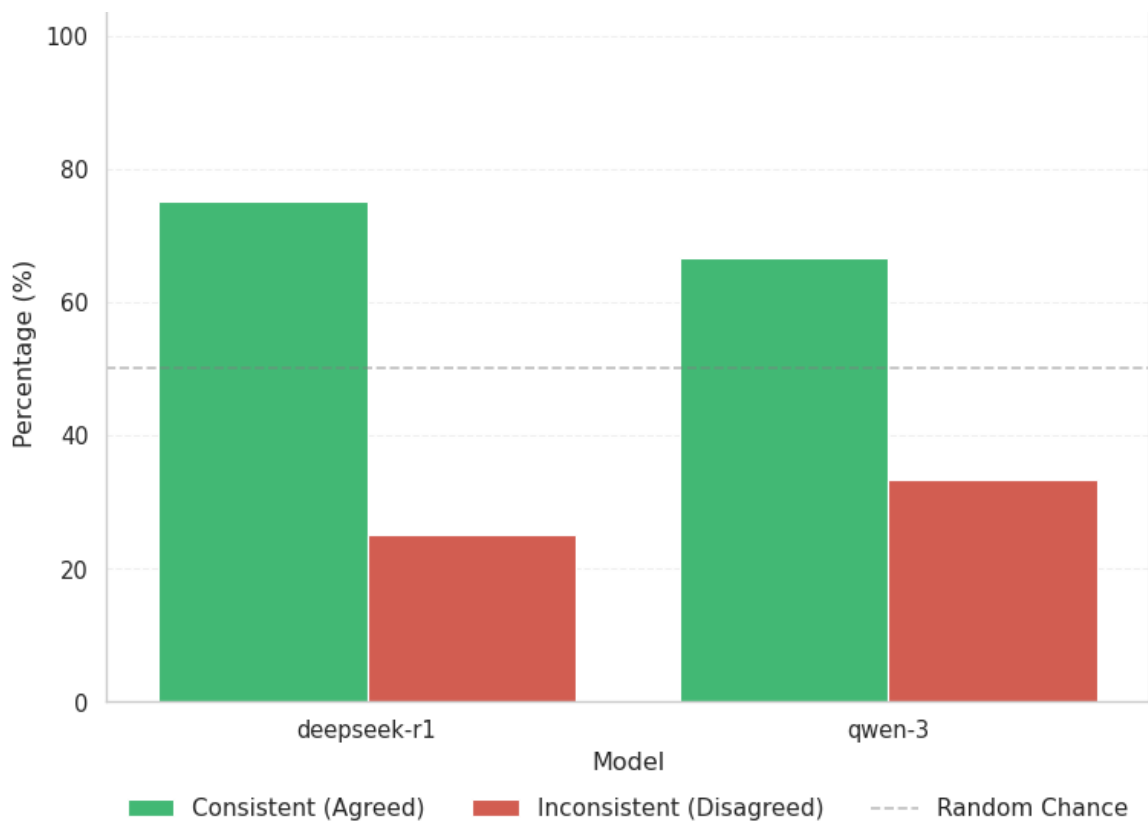


Figure 5.30: Consistency across the prompts for evaluating “humanity” for PR #8161

| | |
|---|--|
| <pre> for (; param_4 != 0; param_4 = param_4 - 1) { if ((param_3[param_4 - 1] & 0xdfU) != 0) goto LAB_00101160; param_4 = 0; LAB_00101160: uVar2 = 0x13; if (param_2 != 0) { local_20 = param_4; local_18 = param_3; if (param_5 == (iconv_t)0x0) { if (param_2 < param_4 + 1) { return 0x13; } memcpy(param_1, param_3, param_4); param_1[param_4] = '\0'; } else { local_30 = param_2 - 1; local_28 = param_1; sVar3 = iconv(param_5, &local_18, &local_20, &local_28, &local_30); if (sVar3 == 0xffffffff) { piVar4 = _errno_location(); iVar1 = *piVar4; if (iVar1 == 7) { return 0x13; } if (iVar1 != 0x16) { if (iVar1 != 0x54) { return 0x10; } return 0x11; } } param_1[param_2 + ~local_30] = '\0'; } } } uVar2 = 0; </pre> | <pre> if (param_4 != 0) { do { if ((param_3[param_4 - 1] & 0xdfU) != 0) goto LAB_00101160; param_4 = param_4 - 1; } while (param_4 != 0); param_4 = 0; LAB_00101160: uVar2 = 0x13; if (param_2 != 0) { local_20 = param_4; local_18 = param_3; if (param_5 == (iconv_t)0x0) { if (param_2 < param_4 + 1) { return 0x13; } memcpy(param_1, param_3, param_4); param_1[param_4] = '\0'; } else { local_30 = param_2 - 1; local_28 = param_1; sVar3 = iconv(param_5, &local_18, &local_20, &local_28, &local_30); if (sVar3 == 0xffffffff) { piVar4 = _errno_location(); iVar1 = *piVar4; if (iVar1 == 7) { return 0x13; } if (iVar1 != 0x16) { if (iVar1 == 0x54) { return 0x11; } return 0x10; } } } } } uVar2 = 0; </pre> |
| <p>Code: Humanity & Readability</p> <p>AST: Humanity & Readability</p> <p>deepseek-r1</p> | <p>BASE</p> <p>Candidate A uses a clean for loop and straightforward conditionals, which are more idiomatic and readable. Candidate B employs a do-while loop and nested conditionals, which are less natural and more machine-like.</p> <p>BASE</p> <p>Candidate A uses a clean for loop, a natural choice for decrementing counters, which is more readable and idiomatic. Candidate B's do-while loop adds unnecessary complexity and nesting, making it less human-like.</p> |

Figure 5.31: The BASE version is better at for loop recovery

| | | | |
|-------------|---|------|---|
| deepseek-r1 | Code: <i>Humanity & Readability</i> | PR | Candidate B includes an explicit cast to (ushort *) when reallocating the pointer, which is more readable and avoids potential pointer type issues. This explicitness is a human-like trait, showing attention to detail and clear intent. |
| | AST: <i>Humanity & Readability</i> | BASE | Candidate A avoids unnecessary type casting, which simplifies the code and makes it more idiomatic. The explicit cast in Candidate B adds redundancy, making it appear more like a machine translation than a human's clean, abstract code. |

Figure 5.32: Inconsistency in file task-libxls_xls_parseWorkbook-O3

| <pre>----- uVar3 = -uVar1; if ((int)-uVar1 < 0) { uVar3 = uVar1; }</pre> | <pre>----- uVar3 = -uVar1; if (0 < (int)uVar1) { uVar3 = uVar1; }</pre> | | |
|---|--|--------|--|
| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION |
| qwen-3 | Code: <i>Humanity & Readability</i> | PR | Candidate B's condition 'if (0 < (int)uVar1)' is a direct, idiomatic check for non-zero unsigned values, which is more readable and natural for a human programmer. Candidate A's 'if ((int)-uVar1 < 0)' is a convoluted mechanical translation that obscures the intent (checking non-zero) with unnecessary negation. B's approach aligns with human engineering practices by simplifying logic and avoiding artificial complexity. |
| | AST: <i>Humanity & Readability</i> | PR | Candidate B simplifies the conditional logic by directly checking if 'id' is positive ('0 < (type)id'), which is a natural, idiomatic way to express this intent. Candidate A's original condition '((type)-id < 0)' is mechanically equivalent but obfuscates the intent, resembling a machine-generated translation rather than a human's direct expression of the logic. The simplified form in B aligns with human-readable code patterns. |
| deepseek-r1 | Code: <i>Humanity & Readability</i> | PR | Candidate B's condition is more straightforward and readable. It directly checks if uVar1 is positive without unnecessary negation, which is a cleaner and more idiomatic approach. |
| | AST: <i>Humanity & Readability</i> | PR | Candidate B's condition 'if(0 < (type)id)' is more direct and readable, avoiding unnecessary operations and better aligns with idiomatic code by using a straightforward comparison. This choice is cleaner and more natural for a human programmer, whereas Candidate A's condition is redundant and less efficient. |

Figure 5.33: Older is better in file task-file_file_zmagic-O3

favoring the recovered for loop as the more natural control-flow pattern.

However, the differences in Ghidra versions across thousands of commits introduce other artifacts that cause evaluation inconsistencies. In Figure 5.32, we observe an inconsistency from *deepseek-r1* regarding explicit type casting. When evaluating the raw code, the model praises the PR branch for including an explicit cast, interpreting it as a “human-like trait, showing attention to detail.” Yet, when the lexical details are abstracted away in the AST evaluation, the model’s judgment flips completely. It penalizes the exact same cast, now viewing it as “redundancy” that makes the code appear “more like a machine translation”.

Interestingly, despite the BASE branch generally representing a much newer and more capable decompiler state, there are isolated instances where the older PR branch generated cleaner localized logic. Figure 5.33 illustrates a scenario where both models *consistently* prefer the PR branch across both raw code and AST evaluations.

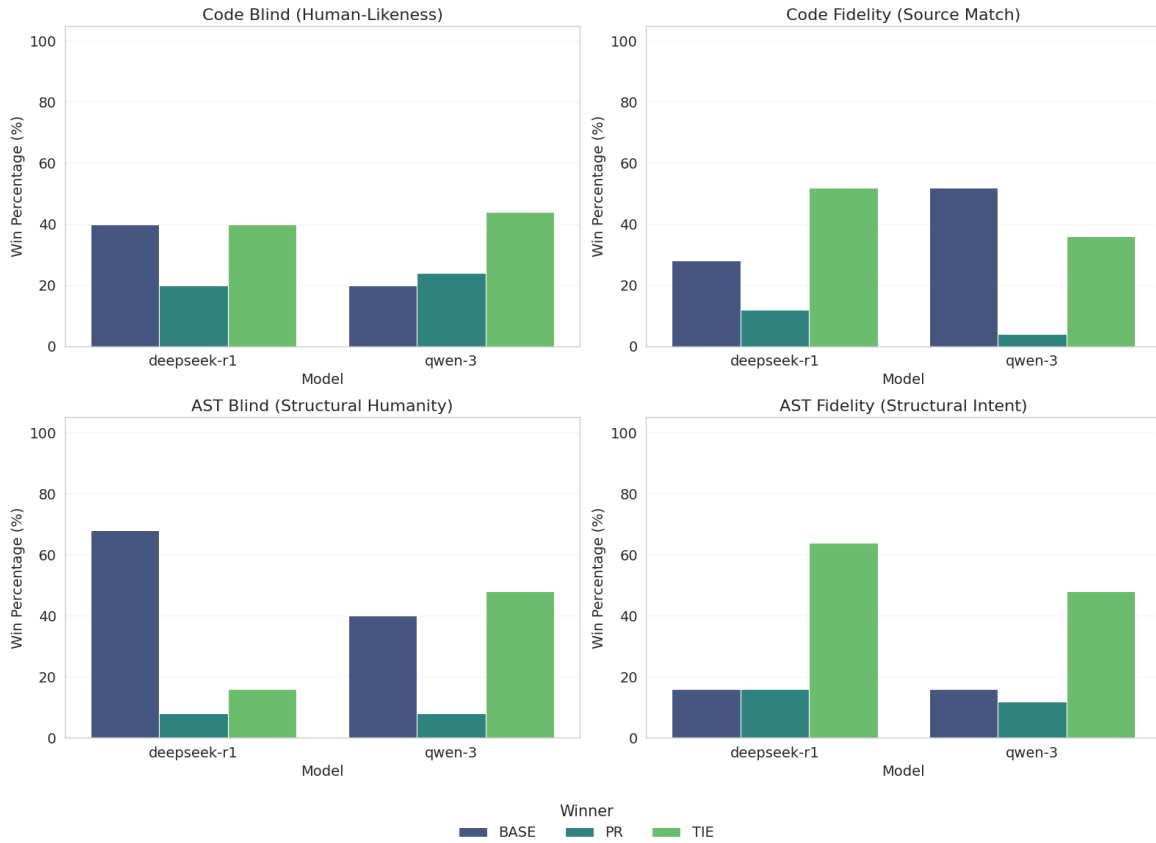


Figure 5.34: Distribution of winners across the four prompts for PR #7253

5.3.4 Sorting Switch Cases – PR #7253

This PR sorts `switch` case entries by their target address

As shown in Figure 5.34, we can observe that in the “AST Blind” evaluation, the BASE branch is heavily preferred, with *deepseek-r1* strongly favoring the BASE version (nearly 70% win rate) and *qwen-3* predominantly defaulting to TIE or BASE., while we can see in the other plots a prominent indecision with a lot of TIEs and a slight preference for the BASE branch. This behavior can be explained by two main factors: (i) As the PR #8161 was generated with an older version of Ghidra, the BASE branch already contains improvements (ii) We will see that in this case of reordering switch cases, the models struggle to recognize the transformation as an improvement, and they may even hallucinate reasons to justify their bias

This is clearly reflected in Figure 5.35, where *deepseek-r1* exhibits a substantial inconsis-




Figure 5.35: Consistency across the prompts for evaluating “humanity” for PR #7253

| | | |
|------------------------------|----|---|
| Code: Humanity & Readability | PR | Version B introduces explicit 'break;' statements for the case blocks that were missing them in Version A. This prevents unintended fall-through in the switch statement, which is a common source of bugs. The addition of 'break;' improves readability and aligns with human engineering practices by making the code's intent clearer and more maintainable. ▶ Raw Response |
| AST: Humanity & Readability | PR | Candidate B reduces redundant case entries in the switch statement, leading to a cleaner structure. The elimination of repetitive 'case str:' blocks aligns with human engineering principles by avoiding mechanical, verbose code generation. This simplification improves readability and reduces visual noise compared to Candidate A's excessive case listings. ▶ Raw Response |

Figure 5.36: Not recognized the modification in task-file_file_signextend-O0

| | | |
|------------------------------|------|--|
| Code: Humanity & Readability | BASE | Candidate A groups similar cases together, making the code more concise and readable. It uses natural, idiomatic grouping which is characteristic of experienced human programmers, avoiding the machine-like redundancy of handling each case separately. ▶ Raw Response |
| AST: Humanity & Readability | PR | Candidate B reorganizes the switch cases to group similar return conditions together, improving readability and clarity. This structure is more natural and idiomatic, as it groups related cases, making the code easier to understand and maintain. ▶ Raw Response |

deepseek-r1

Figure 5.37: Inconsistency in evaluation in task-file_file_signextend-O2

tency rate of roughly 35%, and *qwen-3* proves wildly unstable, disagreeing with its own raw-code evaluations in over 60% of the cases when presented with the AST.

In the next two figures, the only modification is the order of the cases in the switch statement (No other changes are introduced).

In Figure 5.36, we observe a severe instance of model hallucination by *qwen-3*. Unable to find a structural or logical improvement in the simple reordering, the model fabricates reasons to justify its preference for the PR. During the raw code evaluation, it falsely claims that the PR “introduces explicit break statements” to prevent fall-through. During the AST evaluation, it similarly hallucinates that the PR “reduces redundant case entries”. This exposes a significant blind spot in using LLMs as judges: when forced to compare logically identical blocks that differ only in arrangement, the model may invent non-existent structural diffs rather than correctly identifying the transformation.

Figure 5.37 demonstrates a different vulnerability: contradictory reasoning driven by lexical abstraction. In this instance, *deepseek-r1* evaluates the raw code and selects the BASE version, arguing that it “groups similar cases together”, avoiding machine-like redundancy. However, when the specific `case` values and internal variables are abstracted into an AST, the model flips its decision to the PR version, ironically using the exact same claiming.

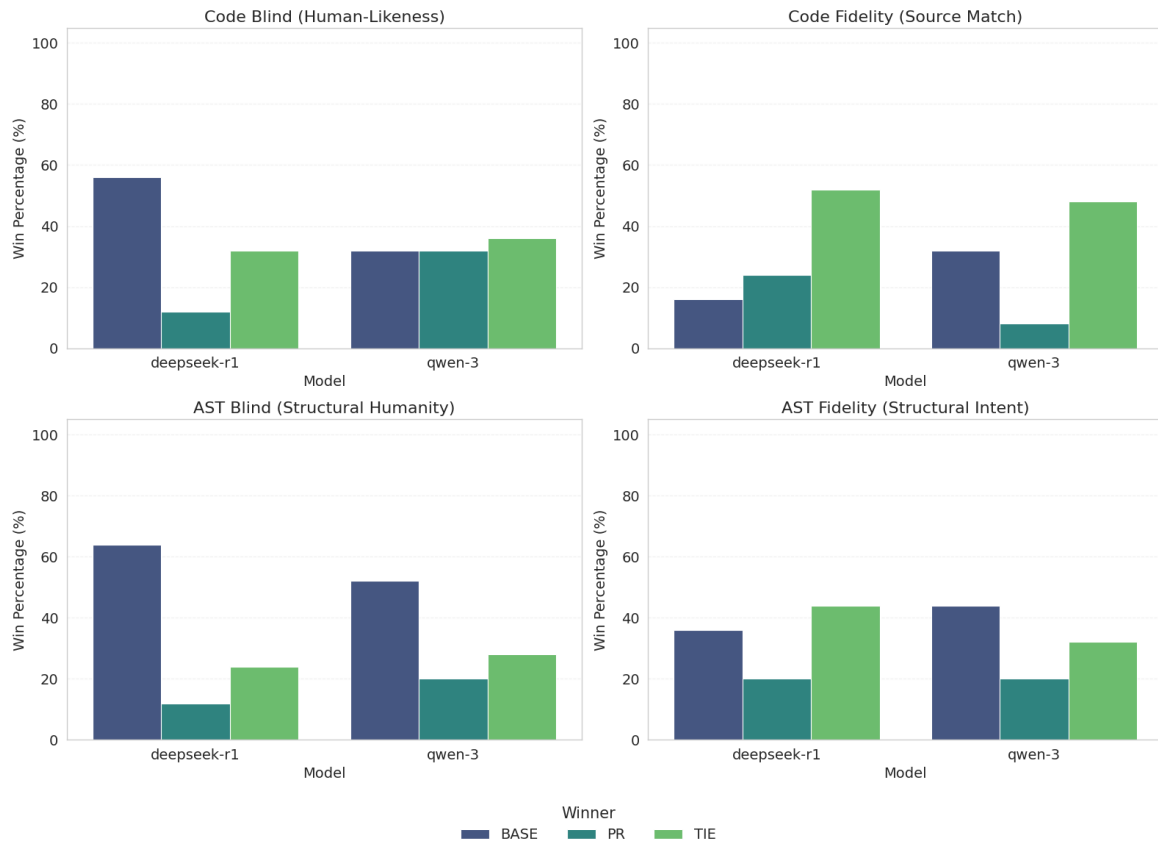


Figure 5.38: Distribution of winners across the four prompts for PR #6722

5.3.5 Improving Pointer-Expression Recovery – PR #6722

This PR focuses on improving pointer-expression recovery, allowing clean output such as `PTR_0041a1b8->ar[local_18]` instead of raw pointer casts and offset arithmetic.

This PR suffers from the same branch mismatch issue as PR #8161, where the BASE branch is generated with a much newer Ghidra version than the PR branch. Because of this significant version mismatch, the overall results in Figure 5.38 are heavily skewed toward the BASE branch. While the PR itself improves pointer-expression recovery, the thousands of newer commits in the BASE branch include numerous other control-flow and structural enhancements. Consequently, the LLM judges, particularly in the “AST Blind” evaluation where structural integrity takes precedence over lexical details, strongly favor the BASE version. *Deepseek-r1* chooses BASE in over 60% of the AST cases, recognizing the overall maturity of the newer decompiler output over the isolated pointer fixes in the older PR branch.




Figure 5.39: Consistency across the prompts for evaluating “humanity” for PR #6722

| | | | |
|-------------|------------------------------|------|--|
| deepseek-r1 | Code: Humanity & Readability | PR | Candidate B explicitly casts the result of <code>bsearch</code> to <code>(char *)</code> , enhancing clarity and readability. This explicitness is a human-like trait, as it demonstrates clear understanding of pointer types, making the code more maintainable and easier to comprehend for other developers. |
| | AST: Humanity & Readability | BASE | Candidate A omits unnecessary explicit casting, which is cleaner and more in line with idiomatic practices where implicit type conversions are trusted when appropriate. This reduces redundancy and enhances readability. |

Figure 5.40: Inconsistency in evaluation in `sav_parse_long_variable_names_record-00`

Despite this imbalance, the consistency rates for PR #6722 (Figure 5.39) are relatively high compared to other PRs, with *deepseek-r1* achieving over 85% consistency and *qwen-3* nearing 80%. However, the inconsistencies that do occur highlight recurring vulnerabilities in the models’ evaluation logic, particularly regarding type casting and complex control flow.

Figure 5.40 illustrates a classic contradiction in how LLMs perceive type casting, mirroring issues seen in previous PRs. In the raw code evaluation, *deepseek-r1* prefers the PR branch, praising the explicit cast of the `bsearch` function result to `(char *)` as a “human-like trait” that enhances clarity and maintainability. Yet, when evaluating the abstracted AST where specific types are hidden, the model completely reverses its logic. It awards the win to the BASE branch, penalizing the explicit cast as an “unnecessary” and “redundant” feature, suddenly arguing that trusting implicit conversions is the true idiomatic human practice.

More alarmingly, Figure 5.41 exposes a severe case of double hallucination and confusion caused by the AST abstraction process. In this complex function, both the BASE and PR versions utilize `goto` statements to handle edge-case returns.

- During the raw code evaluation, *qwen-3* favors the PR, falsely claiming that it “avoids using a `goto` statement” to justify its preference for the PR’s `do-while` loop structure.
- When presented with the AST, the model flips its preference to BASE and hallucinates again, this time falsely claiming that BASE is the one that “avoids complex nested conditions and `goto` statements.”

Despite the before mentioned edge cases and hallucinations, the vast majority of evaluations for this dataset follow a highly consistent path that strongly favors the BASE branch. Figure 5.42 serves as the quintessential example of this dominant trend, perfectly illustrating why the newer Ghidra version consistently wins both the raw code and AST evaluations.

In this snippet, the older PR version introduces a convoluted and non-idiomatic pattern for array access. Both *deepseek-r1* and *qwen-3* correctly and consistently penalize the PR branch for this regression. During the raw code evaluation, they easily identify the PR’s approach as verbose, mechanical, and overly complex.

| <pre> uVar3 = 0xb; if (param_1 != (long *)0x0) { lVar7 = *param_1; if (lVar7 != -1){ uVar4 = 0; uVar5 = 0; uVar3 = 8; bVar1 = true; do { param_1 = param_1 + 2; bVar0 = 1; bVar2 = false; switch(lVar7) { case 3: </pre> | <pre> uVar3 = 0xb; if ((param_1 != (long *)0x0) && (lVar7 = *param_1, lVar7 != -1)) { param_1 = param_1 + 2; uVar4 = 0; uVar5 = 0; uVar3 = 8; bVar1 = true; do { bVar0 = 1; bVar2 = false; switch(lVar7) { case 3: case 4: case 5: </pre> | | | | | | | | | | |
|---|---|---------------------|--|------------|--------|------------------------------|----|--|-----------------------------|------|--|
| [...] | [...] | | | | | | | | | | |
| <pre> case 0x1c: case 0x1d: case 0x1e: case 0x1f: case 0x20: goto switchD_00101150_case0_c; case 0x21: break; default: bVar0 = 1; bVar2 = false; if (1 < lVar7 + 0xbfffffffffffffffU) goto switchD_00101150_case0_c; } if (!bVar1) goto switchD_00101150_case0_c; uVar4 = uVar4 + bVar0; uVar5 = uVar5 + 1; lVar7 = *param_1; bVar1 = bVar2; } while (lVar7 != -1); if ((bVar0 == 1 && uVar5 < 5) && (uVar4 < 4)) { *param_2 = uVar5; uVar3 = 0; } } switchD_00101150_case0_c: } return uVar3; } </pre> | <pre> case 0x1e: case 0x1f: case 0x20: goto LAB_001011a8; case 0x21: break; default: bVar0 = 1; bVar2 = false; if (1 < lVar7 + 0xbfffffffffffffffU) { return 8; } } if (!bVar1) { return 8; } uVar4 = uVar4 + bVar0; uVar5 = uVar5 + 1; lVar7 = *param_1; param_1 = param_1 + 2; bVar1 = bVar2; } while (lVar7 != -1); if ((bVar0 == 1 && uVar5 < 5) && (uVar4 < 4)) { *param_2 = uVar5; uVar3 = 0; } } LAB_001011a8: return uVar3; } </pre> | | | | | | | | | | |
| <table border="1"> <thead> <tr> <th>MODEL</th> <th>EVALUATION CRITERIA</th> <th>WINNER</th> <th>MOTIVATION</th> </tr> </thead> <tbody> <tr> <td rowspan="2">qwen-3</td> <td>Code: Humanity & Readability</td> <td style="text-align: center;">PR</td> <td>Candidate B avoids using a "goto" statement, which is generally considered a code smell in human-authored code. It also structures the loop and conditionals in a more linear, readable way without artificial nesting. The use of a "do-while" loop with clear variable updates and error checks (like "if (!bVar1) return 8;") improves readability compared to Candidate A's use of labels and jumps. The logic is more streamlined and avoids mechanical, step-by-step constructs that a human would abstract.</td> </tr> <tr> <td>AST: Humanity & Readability</td> <td style="text-align: center;">BASE</td> <td>Candidate A avoids complex nested conditions and goto statements, which are often indicative of machine-generated code. It uses straightforward conditional checks and a structured loop, aligning with human-like readability. Candidate B introduces a combined condition with potential typos (e.g., "idid"), redundant assignments, and a "goto" statement, which disrupts linear flow and increases cognitive load. The tie-breaker also favors A due to its less artificial nesting.</td> </tr> </tbody> </table> | MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION | qwen-3 | Code: Humanity & Readability | PR | Candidate B avoids using a "goto" statement, which is generally considered a code smell in human-authored code. It also structures the loop and conditionals in a more linear, readable way without artificial nesting. The use of a "do-while" loop with clear variable updates and error checks (like "if (!bVar1) return 8;") improves readability compared to Candidate A's use of labels and jumps. The logic is more streamlined and avoids mechanical, step-by-step constructs that a human would abstract. | AST: Humanity & Readability | BASE | Candidate A avoids complex nested conditions and goto statements, which are often indicative of machine-generated code. It uses straightforward conditional checks and a structured loop, aligning with human-like readability. Candidate B introduces a combined condition with potential typos (e.g., "idid"), redundant assignments, and a "goto" statement, which disrupts linear flow and increases cognitive load. The tie-breaker also favors A due to its less artificial nesting. |
| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION | | | | | | | | |
| qwen-3 | Code: Humanity & Readability | PR | Candidate B avoids using a "goto" statement, which is generally considered a code smell in human-authored code. It also structures the loop and conditionals in a more linear, readable way without artificial nesting. The use of a "do-while" loop with clear variable updates and error checks (like "if (!bVar1) return 8;") improves readability compared to Candidate A's use of labels and jumps. The logic is more streamlined and avoids mechanical, step-by-step constructs that a human would abstract. | | | | | | | | |
| | AST: Humanity & Readability | BASE | Candidate A avoids complex nested conditions and goto statements, which are often indicative of machine-generated code. It uses straightforward conditional checks and a structured loop, aligning with human-like readability. Candidate B introduces a combined condition with potential typos (e.g., "idid"), redundant assignments, and a "goto" statement, which disrupts linear flow and increases cognitive load. The tie-breaker also favors A due to its less artificial nesting. | | | | | | | | |

Figure 5.41: Inconsistency in evaluation in task-xz_lzma.validate_chain-O3

| MODEL | EVALUATION CRITERIA | WINNER | MOTIVATION |
|-------------|------------------------------|--------|--|
| qwen-3 | Code: Humanity & Readability | BASE | Version A uses more precise and idiomatic C expressions, such as direct array indexing with <code>undefined*</code> for byte access, which aligns with human-readable conventions. Version B's use of <code>undefined*</code> and explicit pointer arithmetic introduces unnecessary verbosity and potential ambiguity, making the code feel more mechanical. The original version (A) maintains clarity and precision in type handling, which is characteristic of human-engineered code. |
| | AST: Humanity & Readability | BASE | Candidate A uses straightforward pointer arithmetic (<code>id[id]</code>) which aligns with idiomatic C practices for array access. Candidate B's <code>&id[id]</code> introduces unnecessary complexity by dereferencing the address of <code>id</code> itself, which is non-standard and obscures the intent. Human engineers typically avoid such convoluted pointer manipulations, favoring clarity and directness. |
| deepseek-r1 | Code: Humanity & Readability | BASE | Candidate A uses direct array indexing without unnecessary address-taking, which is cleaner and more idiomatic. The code is concise and avoids redundant operations, making it more readable and natural for human programmers. Candidate B's approach is more verbose and less elegant, resembling a machine-like translation of the logic. |
| | AST: Humanity & Readability | BASE | Candidate A's code is more human-like as it avoids unnecessary pointer dereferencing, making it cleaner and more readable. The changes in Candidate B introduce redundant pointer operations, which are more machine-like and less idiomatic. |

Figure 5.42: Example of majority consistency in evaluation across models

Because the redundancy is purely structural and does not rely on specific variable names to appear unnatural, the models easily maintain their judgment without falling into the trap of lexical bias. This demonstrates that when a decompiler generates an objectively inferior syntactic construct, such as redundant address-taking or dereferencing operations, the LLM judges can reliably and consistently detect the “machine-like” quality regardless of whether the code is raw or anonymized.

5.4 Dogbolt Evaluation

Based on the data presented in Figure 5.43 and detailed in Table 5.1, several critical insights emerge regarding how the LLM judges perceive the outputs of the different decompilers. There is a stark contrast between the “Blind” evaluations (where the LLM judges purely on abstract readability and perceived “humanity”) and the “+ Source” evaluations (where the LLM judges fidelity to the original code) for *binary-ninja*. In the blind tests (*Code Blind* and *AST Blind*), *binary-ninja* frequently emerges as winner, particularly according to *deepseek-r1*, securing 16 and 21 wins respectively. The models evidently perceive its output as highly readable and natural. However, the moment the original source code is introduced




Figure 5.43: Distribution of winners across the four prompts for the Dogbolt evaluation

as a baseline (*Code + Source* and *AST + Source*), its win rate collapses drastically. This heavily implies that while *binary-ninja* applies aggressive transformations to make the code look clean and human-readable, these transformations structurally distance the output from the original human-written logic.

Another key indicator of the massive differences between the versions is the exceptionally high number of “Ties”. When the source code is provided, the tie rate spikes, peaking at 30 ties out of the dataset for *deepseek-r1* in the *Code + Source* test. This high rate of indecision suggests that none of the decompilers successfully reconstruct the original source in a recognizable way. When the LLM attempts to match the decompiler outputs to the human ground truth, it finds that all candidates have hallucinated distinct structural artifacts, applied different control-flow recovery heuristics, or inflated the code with boilerplate to such a degree that choosing a “closest match” becomes arbitrary. It could be position Bias, (since the ratio of wins after switch is 99/47 for Version B) but it could also be that the LLM judges are overwhelmed by the sheer unpredictability of two different decompiled code, leading to a default tie decision when no clear winner emerges.

Hex-Rays manages to maintain the most stable performance across both blind and source-provided prompts, winning almost consistently across all prompts, suggesting a more balanced approach that preserves some of the original code’s structure while still applying transformations for readability.

Ghidra, on the other hand, consistently underperforms for *deepseek-r1*, rarely winning and often being overshadowed by the other two decompilers, but it also performs slightly better when the raw code is provided. Furthermore, the inconsistencies in the Ghidra results across the two models suggest that its output may be more sensitive to the specific evaluation criteria emphasized by each LLM, which further complicates the notion of a single “best” decompiler.

Table 5.1: Dogbolt evaluation winners by judge and prompt type.

| Judge | Test Type | Tie | Binary-Ninja | Ghidra | Hex-Rays |
|-------------|---------------|-----|--------------|--------|----------|
| deepseek-r1 | AST + Source | 25 | 11 | 6 | 14 |
| | AST Blind | 19 | 21 | 6 | 14 |
| | Code + Source | 30 | 3 | 6 | 18 |
| | Code Blind | 18 | 16 | 9 | 17 |
| qwen-3 | AST + Source | 17 | 11 | 11 | 16 |
| | AST Blind | 14 | 18 | 11 | 17 |
| | Code + Source | 23 | 5 | 13 | 15 |
| | Code Blind | 16 | 12 | 17 | 15 |

5.4.1 Length Bias

Figure 5.44 further quantifies the drastic variations among the decompilers. The boxplots illustrate the delta length between the winning snippet and the losing alternatives. The presence of massive outliers, some exceeding a difference of ± 1000 tokens, proves that the tools generate code of vastly different lengths for the exact same binary function.

While the medians hover near zero, we can observe a slight negative bias in the blind tests (*winner* and *winner_ast*), meaning the LLM judges generally exhibit a slight preference for more concise code when evaluating purely for readability. However, in the source-provided tests (*winner_s* and *winner_ast_s*), this distribution flattens or shifts slightly positive. When forced to prioritize fidelity, the models occasionally reward verbosity if the longer decompiler output happens to capture explicit logic or variable declarations that were present in the original source but optimized away by the more concise decompilers.

5.4.2 Dogbolt Results

The evaluation of the Dogbolt dataset corroborates the core findings from the Ghidra PR analysis, albeit on a much larger scale. While the Ghidra dataset compared minor, localized modifications between two commits of the same engine, the Dogbolt dataset tasks the LLM judges with comparing entirely different decompilers (Binary Ninja, Hex-Rays, and Ghidra). Consequently, the judges have a significantly broader range of structural, lexical, and control-flow modifications to choose from.

Despite this increased variance, the fundamental behaviors, biases, and inconsistencies of both *qwen-3* and *deepseek-r1* remain perfectly in line with our previous observations.

In Figure 5.45 we can see a case of inconsistency driven by lexical abstraction, mirroring




Figure 5.44: Distribution of winners across the four prompts for the Dogbolt evaluation, showing length bias

| | |
|--|--|
| <p>Decompiler A: binary-ninja</p> <pre> type id(type id, type id){ if(!id)return id; type id = id; type id = 0; type id = id; type id = 0; do{ type id = call(&id, id); id = id; if(id == 0xe)return id; if(id == 0xf)return id; id = call(&id, id); }while(id == id); return id; } </pre> | <p>Decompiler B: hex-rays</p> <pre> type *id(type *id, type *id){ type id; type id; type *id; type *id; type *id; type *id; type *id; id = id; id = id; if(!id)return id; id = id; id = 0; id = id; id = 0; while(1){ id = call(&id, id); if(id == 14)break; if(id == 15)return id; id = call(&id, id); if(id != id)return id; } return id; } </pre> |
| <p>QUALITATIVE COMPARISON (BLIND)</p> <p>CODE WINNER: A</p> <p>Candidate A uses a "do-while" loop, which is a semantic structure that aligns with high-level human thinking. Its logic is structured with clear loop conditions and minimal artificial nesting. Candidate B uses a "while(1)" loop with "break", which is a common but less idiomatic pattern in human-readable code. Additionally, Candidate B's variable names (e.g., 'v3', 'v5') are less descriptive, and its logic involves more return points within the loop, increasing complexity. Candidate A's structure is cleaner and avoids deep nesting, making it more maintainable.</p> <p>AST WINNER: B</p> <p>Candidate B uses a natural infinite loop ("while(1)") with explicit "break" and "return" statements, which aligns with idiomatic C patterns. It avoids redundant loop conditions (like "id == id" in A) and organizes control flow with clearer exit points. Although both have questionable logic, B's structure reflects standard practices for infinite loops with conditional exits, whereas A's "do-while" with a tautological condition feels less intentional and more assembly-like.</p> | <p>QUALITATIVE COMPARISON (GROUND TRUTH)</p> <p>GT CODE WINNER: TIE</p> <p>Detected potential bias in LLM response (Position Bias); the model have chosen always version B. declaring TIE.</p> <p>GT AST WINNER: B</p> <p>Candidate B matches the Source AST's node types (e.g., function parameters with pointers "type *id") present in the Ground Truth but absent in Candidate A. While both candidates exhibit loops and conditional structures, Candidate B preserves the pointer type in function parameters and aligns more closely with the Source's variable declaration hierarchy and nesting depth.</p> |

Figure 5.45: Incoherent results from Dogbolt with Qwen

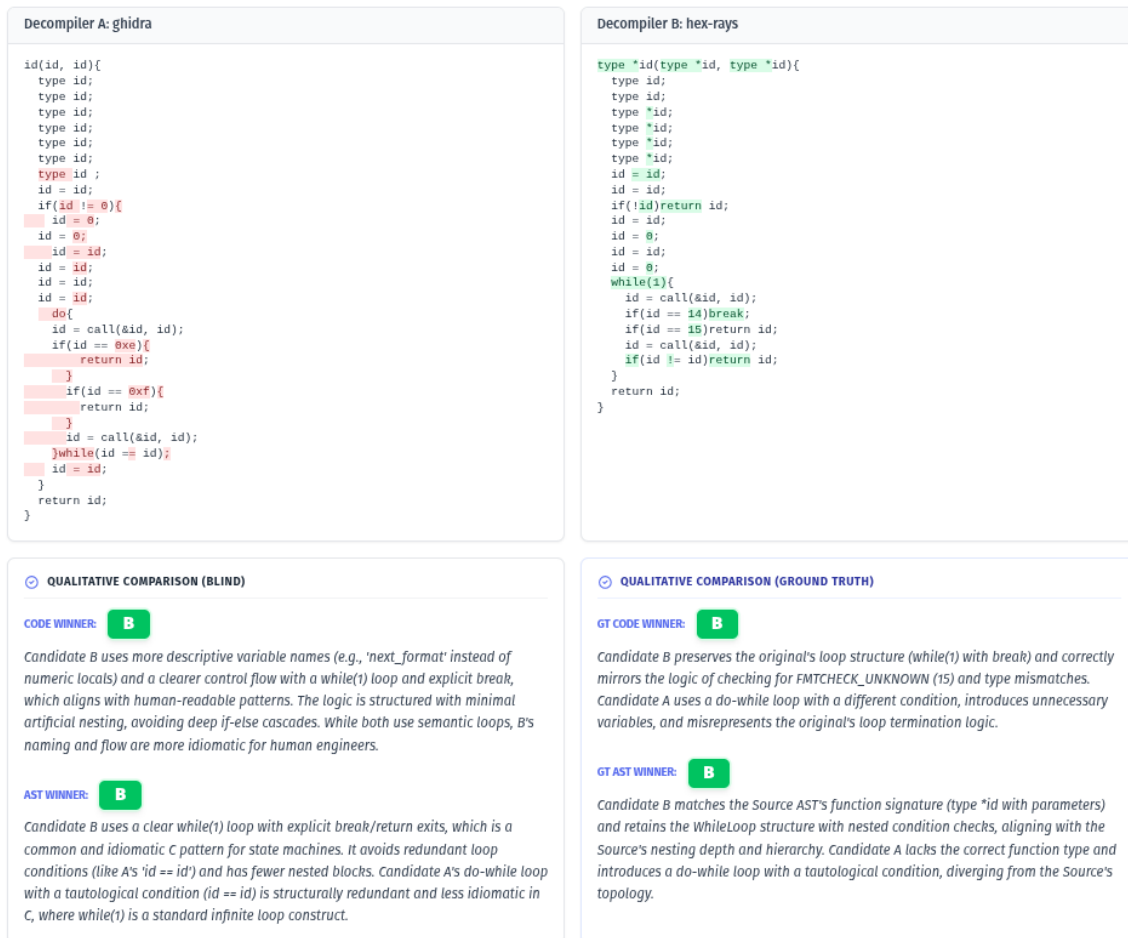


Figure 5.46: Coherent results from Dogbolt with Qwen

the vulnerabilities seen in the Ghidra evaluations. When *qwen-3* evaluates the raw code (comparing Binary Ninja and Hex-Rays), it prefers Candidate A (Binary Ninja), praising its **do-while** loop as a “semantic structure that aligns with high-level human thinking” and criticizing Candidate B’s **while(1)** with a **break** as less idiomatic. However, when the exact same logic is abstracted into an AST, the model flips its decision to Candidate B. It suddenly contradicts its previous motivation, now claiming that the **while(1)** loop is a “natural infinite loop” that aligns with idiomatic C patterns, while dismissing Candidate A’s **do-while** as “tautological” and “assembly-like”.

Conversely, in Figure 5.46 the LLM compares Ghidra against Hex-Rays. In this scenario, Hex-Rays generates a **while(1)** loop, whereas Ghidra produces a structure with a **do-while** loop nested inside multiple conditional checks and variable updates. *Qwen-3* consistently rewards Hex-Rays across all four evaluation prompts (Code Blind, AST

Table 5.2: Point-Biserial Correlation between LLM Judge Preference and Δ Perplexity

| Model | Code r | Code p -value | N | AST r | AST p -value | N |
|-------------|----------|-----------------|----|---------|----------------|----|
| qwen-3 | -0.0805 | 0.5443 | 59 | -0.2263 | 0.0509 | 75 |
| deepseek-r1 | 0.0128 | 0.9204 | 63 | -0.1036 | 0.3866 | 72 |

Blind, Code Fidelity, and AST Fidelity). It recognizes that avoiding artificial nesting and deep `if-else` cascades is a universal hallmark of human engineering, a judgment it can confidently make even when variable names are stripped away in the AST.

Interestingly, the LLM evaluates the same structural pattern as a negative in Figure 5.45 (where the `do-while` is penalized) but as a positive in Figure 5.46 (where the `do-while` is praised). This highlights the extreme sensitivity of the models to contextual cues.

5.5 Correlation Between Perplexity & LLM judge

5.5.1 Ghidra

To determine whether the quantitative predictability of the decompiled code (measured via Perplexity) aligns with the qualitative “human-likeness” evaluations provided by our LLM judges, we conducted a Point-Biserial correlation analysis.² We also removed the results with clear hallucinations in the motivation field (e.g., where the model referenced code elements or modifications that were not present), as these cases do not reflect genuine judgments about code quality (12 in raw code and 11 in AST, with 17 *deepseek-r1* and 6 *qwen-3*). In this setup, the qualitative winner is treated as a binary variable ($\text{PR} = 1$, $\text{BASE} = 0$), and it is correlated against the Δ Perplexity ($\text{PPL}_{\text{PR}} - \text{PPL}_{\text{BASE}}$).

Consequently, a *negative* correlation coefficient ($r < 0$) would indicate the expected behavior: when the LLM judge prefers the `PR` branch (1), the Δ PPL is negative (meaning the `PR` code had a lower, “better” perplexity).

Looking at the raw code correlations (*Code r*) in Table 5.2, we observe values extremely close to zero for both models (-0.0805 and 0.0128). Furthermore, the associated p -values (> 0.5) clearly indicate that these results are not statistically significant. This demonstrates a complete lack of correlation between raw code perplexity and qualitative LLM judgment.

²The point-biserial correlation (r_{pb}) measures the strength and direction of the association between a continuous interval/ratio variable and a binary (dichotomous) variable, ranging from -1 to +1. As a special case of the Pearson correlation, it is commonly used to determine if a naturally occurring dichotomy (e.g., pass/fail, yes/no) relates to a continuous score [Vol]

This disconnect highlights a crucial insight: the features that make code “predictable” to an LLM’s next-token-prediction engine (lower perplexity) are fundamentally different from the features that an LLM judge evaluates when prompted to look for “human-like” readability. As discussed in previous sections, a decompiler might generate highly repetitive, boilerplate-heavy code that drives perplexity down through token inflation, but an LLM judge will correctly penalize this exact same code for being verbose and mechanical.

When evaluating the abstracted AST representations (*AST r*), the negative correlation slightly strengthens for both *qwen-3* (-0.2263) and *deepseek-r1* (-0.1036). Additionally, the sample size (N) increases (as the number of TIEs decreases).

While the p -values (0.0509 and 0.3866) still fall short of the threshold for statistical significance ($p < 0.05$ to confirm if the relationship exists in the population), the shift is notable. By removing the “lexical noise” and standardizing the decompiler output artificially, the structural predictability of the AST aligns slightly better with the structural judgment of the LLM. However, the correlation remains weak overall even if for the *qwen-3* model the p -value is approaching the threshold.

Ultimately, this data confirms that **Perplexity is not a reliable proxy for human-likeness in decompiled code**. While perplexity effectively measures how well the code aligns with the statistical distribution of the model’s training data, qualitative human-likeness is defined by higher-level engineering concepts, such as idiomatic control flow, logical grouping, and abstraction, which often introduce necessary structural “surprises” that inherently increase perplexity.

5.5.2 Dogbolt

We extended our correlation analysis to the Dogbolt dataset, where the LLM judges compared outputs from different decompilers (e.g., Binary Ninja vs. Hex-Rays vs. Ghidra) in an A/B testing format. We binarized the qualitative winner (Candidate A = 1, Candidate B = 0) and correlated it against the Δ Perplexity ($PPL_A - PPL_B$). In this case, cleaning the data from hallucinations was not feasible: the differences between candidates were often substantial, and the models tended to provide high-level, generic narratives of “improvements”, sometimes citing a few true examples, but also overlooking (or implicitly collapsing) many other simultaneous changes. Therefore, we included all the data points in the correlation analysis, acknowledging that some motivations may not reflect genuine judgments about code quality but rather post-hoc rationalizations.

As established previously, if the LLM judge preferred the code that is “most predictable” to a language model, we would expect a *negative* correlation (i.e., when A wins, A’s perplexity is lower, making Δ PPL negative).

Table 5.3: Point-Biserial Correlation between LLM Judge Preference and Δ Perplexity

| Model | Code r | Code p -value | N | AST r | AST p -value | N |
|-------------|----------|-----------------|----|---------|----------------|----|
| qwen-3 | 0.1953 | 0.2040 | 44 | 0.2292 | 0.1255 | 46 |
| deepseek-r1 | 0.0352 | 0.8250 | 42 | 0.2283 | 0.1511 | 41 |

The data in Table 5.3 reveals a striking phenomenon: the correlation coefficients (**Code r**) are *positive*. A positive r value (0.1953 for *qwen-3* and 0.0352 for *deepseek-r1*) dictates that when the LLM judge selects a candidate as the “human-like” winner, that candidate actually tends to have a *higher* perplexity score than the loser.

While the high p -values (0.2040 and 0.8250) confirm that this is not a statistically significant linear rule, the inversion of the sign perfectly corroborates our findings in Sections 5.2.1 and 5.4. Decompilers that score high in qualitative human-likeness achieve this by utilizing diverse, domain-specific, and sometimes complex idioms that carry higher entropy. Conversely, decompilers with artificially low perplexity (by inflating the number of tokens) are penalized by the judges for being overly verbose and machine-like.

When evaluating the AST versions (AST r), the positive correlation strengthens noticeably to ~ 0.229 for both models, and the p -values improve (0.1255 and 0.1511), though they still remain above the strict 0.05 threshold for significance.

This parallel strengthening across both *qwen-3* and *deepseek-r1* indicates that the preference for “higher perplexity” structures is not merely a lexical artifact. Even when the code is anonymized, the LLM judges prefer abstracted syntactic structures that contain more variance and unpredictability over the repetitive, low-loss boilerplate generated by the most predictable decompilers.

The Dogbolt correlation data acts as a definitive counter-proof to the assumption that language models inherently prefer code that perfectly matches their own predictive distributions. By analyzing the data, and the correlations, we saw that perplexity is not a sufficient measure of human-likeness; Therefore, it should not be used as a standalone metric to measure the readability or human-likeness of a decompiler’s output.

5.6 Judge vs. Human Evaluation

To validate the results, the biases, and inconsistencies observed in the LLM judges, we conducted a human-subject experiment. We designed a blind survey using Google Forms, presenting real software developers with pairs of decompiled C code (labeled simply as “Snippet A” and “Snippet B”) and asking them to identify which version appeared more

Table 5.4: Mean Alignment Score between LLM Judges and Human Consensus

| Model | Raw Code Alignment | AST Alignment |
|-------------|--------------------|---------------|
| qwen-3 | ~64% | ~69% |
| deepseek-r1 | ~74% | ~74% |

natural, idiomatic, and “human-like”.

The dataset for this quiz consisted of 21 carefully selected questions: 15 questions comparing Ghidra *BASE* vs. *PR* branches, and 6 questions comparing different decompilers from the Dogbolt dataset (Binary Ninja, Hex-Rays, and Ghidra), these questions are available in the appendix (Appendix A). The snippets were strategically chosen to include a mix of scenarios: cases where the LLM judges showed strong consensus, cases of severe inconsistency (flipping decisions between Raw Code and AST), and cases where the models hallucinated structural differences (e.g., PR #7253).

Responses were collected from 11 participants. We aggregated their votes to establish a human “Ground Truth” consensus for each question and compared it against the predictions made by *qwen-3* and *deepseek-r1*.

It’s interesting to note that the human consensus was not always unanimous. In some cases, the votes were split, reflecting genuine ambiguity in the code’s readability. However, in the majority of cases, a clear preference emerged, giving a mean alignment score of 78% between the human choices (min 55%, max 100%).

5.6.1 Quantitative Alignment

To measure how well the LLM judges proxy human intuition, we calculated an alignment score. A score of 1.0 indicates a perfect match between the model’s winner and the human consensus, 0.0 indicates a complete mismatch, and 0.5 is awarded for partial matches (e.g., the model declared a TIE).

The results in Table 5.4 provide two critical insights into the capabilities of LLMs as code judges:

1. **Deepseek-R1 is superior for human judgment:** With a consistent ~73.8% alignment, *deepseek-r1* demonstrates a much more robust understanding of human programming idioms than *qwen-3*. It is important to note that the same results for raw code and AST code are identical just in the final number as the intermediate results are different.

2. **AST Abstraction improves Qwen-3:** Strikingly, *qwen-3*'s alignment with human developers increases from $\sim 64.28\%$ to $\sim 69.04\%$ when evaluating the anonymized AST. This means that probably *qwen-3* suffers from a light lexical bias.

Unlike the LLM pipeline, the human study was conducted exclusively on *raw code* snippets and did not include an AST-anonymized condition. This was a deliberate design choice: we assumed human reviewers would be less affected by the lexical anchoring effects observed in the models, and we wanted to preserve a realistic reading experience close to practical reverse-engineering workflows. In addition, several diffs were extremely subtle and already challenging in raw form. In some questions, the only effective change was minimal (e.g., the addition of a single explicit cast on a variable), making any further abstraction likely to remove the remaining discriminative signal and turn the comparison into near-random guessing.

We intentionally included snippets from PR #7253 (where only the vertical order of `switch` cases was altered). Previously, we noted that the LLM judges, especially *qwen-3*, often hallucinated non-existent structural improvements to justify picking a winner in these scenarios. Human developers, conversely, quickly recognized the logical equivalence. Human votes were often split, or respondents correctly noted that the differences were purely stylistic, heavily contrasting with the models' tendency to confidently invent false technical justifications.

The final 6 questions of the quiz pitted the outputs of Binary Ninja, Hex-Rays, and Ghidra against each other. The human consensus overwhelmingly favored Binary Ninja and Hex-Rays over Ghidra. This definitively validates the LLM evaluations from the Dogbolt dataset. It is important to note that the human votes were not unanimous, and in some cases, Ghidra's output was preferred by a 40+% of respondents. This reflects the inherent subjectivity in code readability and the fact that different developers may have varying preferences for certain coding styles or patterns. However, the overall trend clearly supports the LLM judges' conclusions about the relative quality of the decompilers' outputs.

5.6.2 Perplexity Correlation

We also calculated the Point-Biserial correlation between the human consensus winners and the Δ Perplexity for the raw code snippets used in the quiz, for both the raw code and the AST conditions. This analysis aimed to determine whether the human preferences align with the same perplexity trends observed in the LLM judges.

As shown in Table 5.5, the correlation with human consensus is positive in both settings, and stronger for the AST condition ($r = 0.4609$) than for raw code ($r = 0.3278$). This direction is coherent with the previous analyses: snippets judged as more human-like are

Table 5.5: Point-Biserial correlation between Human Consensus and Δ Perplexity (quiz subset).

| Evaluator | Code r | Code p -value | N | AST r | AST p -value | N |
|-----------------|----------|-----------------|----|---------|----------------|----|
| Human Consensus | 0.3278 | 0.2330 | 15 | 0.4609 | 0.0838 | 15 |

not associated with lower perplexity, but tend to show equal or higher perplexity. However, with $N = 15$, neither result is statistically significant at $\alpha = 0.05$ ($p = 0.2330$ for code, $p = 0.0838$ for AST), so these values should be interpreted as an exploratory trend rather than conclusive evidence.

5.7 Total Loss as a Quality Metric

Given the failure of average perplexity to proxy human-likeness (due to its vulnerability to token inflation and boilerplate verbosity), we shifted our focus to the *Total Loss* (\sum Loss). Because Total Loss accumulates the prediction error for every single token, it inherently penalizes verbose, artificially inflated decompiler outputs.

5.7.1 Ghidra

As shown in Figure 5.47, examining the distribution of the Total Loss provides a much clearer and more intuitive separation between human-written source code and decompiled outputs. Unlike average perplexity, the Total Loss accurately reflects the cumulative predictive burden of verbose code. Across both *qwen-3* and *deepseek-r1*, the original source code consistently exhibits a significantly lower and tighter Total Loss distribution compared to both the **BASE** and **PR** decompilation branches. This confirms that summing the token prediction errors effectively penalizes the structural bloat and unnatural verbosity introduced by the decompilation process.

To evaluate its effectiveness, we repeated the Point-Biserial correlation analysis, this time correlating the judges binary preferences (**PR** = 1, **BASE** = 0) against the Δ Total Loss ($\text{Loss}_{\text{PR}} - \text{Loss}_{\text{BASE}}$). A negative correlation ($r < 0$) indicates that the preferred version consistently exhibits a lower (better) Total Loss.




Figure 5.47: Distribution of the Total Loss (Sum of token losses) across Source, BASE, and PR branches for both evaluated LLM judges.

Table 5.6: Point-Biserial Correlation between LLM Judge Preference and Δ Total Loss

| Model | Code r | Code p -value | N | AST r | AST p -value | N |
|-------------|----------|-----------------|----|---------|----------------|----|
| qwen-3 | -0.5396 | < 0.0001 | 59 | -0.4233 | 0.0002 | 75 |
| deepseek-r1 | -0.2477 | 0.0503 | 63 | -0.4577 | < 0.0001 | 72 |

The results in Table 5.6 present a high contrast to the perplexity baseline. *Qwen-3* demonstrates a strong, highly significant negative correlation in raw code ($r = -0.5396, p < 0.0001$), meaning it systematically votes for the snippet that minimizes the Total Loss. *Deepseek-r1* shows a moderate correlation in raw code, but when transitioning to the AST representation, its correlation strengthens dramatically to -0.4577 ($p < 0.0001$).

Unlike average perplexity, Total Loss, in our case, seems to be a statistically viable predictor of LLM preferences, *especially* when computed on the anonymized AST representation, where most of the decompiler-specific lexical boilerplate is removed and the metric is driven primarily by structural differences in the recovered logic.




Figure 5.48: Distribution of the Total Loss (Sum of token losses) across the three decompilers for both evaluated LLM judges in the Dogbolt dataset.

5.7.2 Dogbolt

As illustrated in Figure 5.48, the distribution of Total Loss across the three distinct decompilation engines (Binary Ninja, Ghidra, and Hex-Rays) provides further insights into the metric’s behavior. Interestingly, when compared to our previous observations in Figure 5.47, Ghidra’s Total Loss distribution in this chart aligns much more closely with the original Source code.

This behavior can be directly attributed to the dataset composition. As detailed in Section 4.6, the Dogbolt evaluation relies on a smaller subset of binaries (restricted solely to the `file` project), whereas the PR versus BASE study utilized a much larger and more diverse dataset. Consequently, this restricted subset likely excludes many of the massively bloated and highly complex functions that previously inflated Ghidra’s overall token loss. Despite this dataset-specific variance, the visual separation between the decompilers confirms that Total Loss actively captures the varying verbosity levels of each specific decompiler.

To ensure that the effectiveness of the Total Loss metric is not merely an artifact of the Ghidra ecosystem, we replicated the correlation analysis on the Dogbolt dataset.

In this setup, the Point-Biserial correlation measures the association between the LLM’s chosen winner (Candidate A = 1, Candidate B = 0) and the Δ Total Loss ($Loss_A - Loss_B$).

Table 5.7: Point-Biserial Correlation between LLM Preference and Δ Total Loss (Dogbolt Dataset)

| Model | Code r | Code p -value | N | AST r | AST p -value | N |
|-------------|----------|-----------------|----|---------|----------------|----|
| qwen-3 | -0.4666 | 0.0014 | 44 | -0.5014 | 0.0004 | 46 |
| deepseek-r1 | -0.2359 | 0.1326 | 42 | -0.5268 | 0.0004 | 41 |

The results in Table 5.7 definitively corroborate our previous findings and highlight the critical importance of the AST pipeline.

When evaluating the raw code, *qwen-3* maintains a significant negative correlation ($r = -0.4666$, $p = 0.0014$). However, *deepseek-r1*'s correlation drops to a non-significant -0.2359 ($p = 0.1326$).

When the decompiled outputs are abstracted into ASTs, the correlation for both models converges to a highly significant, strong negative relationship (*qwen-3*: $r = -0.5014$, *deepseek-r1*: $r = -0.5268$, both with $p = 0.0004$).

This phenomenon provides a conclusive insight: different decompilers inject high amounts of arbitrary lexical entropy that skew the raw prediction loss. By abstracting the code into an AST, we completely neutralize this decompiler-specific noise. The resulting AST Total Loss seems to act as a universal quantitative metric that reliably mirrors the LLM's judgments.

5.7.2.1 Human Validation

To determine if this metric also aligns with genuine human intuition, we ran the exact same correlation analysis using the Ground Truth votes collected from our survey of 15 selected comparisons.

Table 5.8: Point-Biserial Correlation between Human Consensus and Δ Total Loss

| Evaluator Metric | Code r | Code p | N | AST r | AST p | N |
|------------------------|----------|----------|----|---------|---------|----|
| Human vs qwen-3 Loss | -0.0505 | 0.8581 | 15 | -0.5633 | 0.0288 | 15 |
| Human vs deepseek Loss | -0.1493 | 0.5955 | 15 | -0.4732 | 0.0748 | 15 |

The data in Table 5.8 provides one of the most compelling insights of this study. When evaluating raw code (*Code r*), human preference exhibits virtually zero correlation with the models' Total Loss.

However, when we correlate the human raw code preferences against the **AST Total Loss**, the correlation surges. For *qwen-3*'s AST Loss, the correlation reaches a remarkable -0.5633 with statistical significance ($p = 0.0288$), despite the small sample size ($N = 15$). For *deepseek-r1*, it reaches -0.4732 but with a p-value that is not statistically significant ($p = 0.0748$).

This quiz has a limited number of samples, it cannot provide a statistically significant result, but it does suggest a trend. Consequently, a lower AST Total Loss proves to be the most accurate quantitative proxy for human-perceived code readability discovered in this research.

5.8 Discussion

The results presented in this study offer a critical perspective on the use of LLMs for evaluating the quality and readability of decompiled code.

First, the quantitative analysis clearly shows that perplexity is not a reliable proxy for code human-likeness. Contrary to the initial intuition, the original source code tends to have systematically higher perplexity values than decompiled versions. This phenomenon is rooted in the nature of human-written code, which is characterized by varied stylistic choices, domain-specific naming conventions, and general high-entropy. In contrast, decompilers such as Ghidra tend to produce rigid syntactic structures, repetitive patterns, and verbose constructs. This structural and lexical token inflation dilutes the cross-entropy mean, artificially lowering overall perplexity and making mechanical code appear more predictable to the model than human code.

To overcome the limitations of average perplexity, this study introduced *Total Loss* (\sum Loss) as an alternative quantitative metric. By accumulating the prediction error across all tokens, Total Loss inherently penalizes the structural bloat and verbosity typical of decompiled outputs. Our findings demonstrate that Total Loss, unlike average perplexity, acts as a statistically viable predictor of qualitative code readability, correctly reflecting the models' penalization of overly complex code.

Crucially, the effectiveness of this metric is maximized when combined with the AST abstraction pipeline. Different decompilers inject a massive amount of arbitrary lexical entropy (e.g., varying variable naming conventions and explicit type casts) that severely skews the raw prediction loss. While human developers intuitively ignore this non-functional noise, LLMs accumulate loss for every character. By removing this decompiler-specific lexical noise through AST anonymization, the resulting AST Total Loss strongly correlated with the LLM judges preferences across both the Ghidra and Dogbolt datasets. Consequently, AST Total Loss emerges as the most accurate quantitative proxy for human-

perceived structural code quality identified in this master’s thesis.

Regarding the use of LLMs as qualitative judges, the two evaluated architectures showed distinct capabilities and limitations. *Deepseek-r1* proved to be clearly superior in approximating human judgment, reaching an alignment rate with human evaluators of about 74%. *Qwen-3*, by contrast, showed significant weaknesses: a marked tendency toward verbosity that harms time and computational efficiency, and strong susceptibility to lexical bias, which improves only when code is abstracted into AST. Moreover, both models showed the risk of analytical hallucinations. When faced with logically equivalent code variants with many lines of differences, or variants differing only in irrelevant stylistic details (such as switch-case ordering), LLMs sometimes invented nonexistent technical justifications to support a preferred choice.

The human-developer experiment, while limited in sample size and therefore not suitable as an absolute ground truth, still provided a useful signal about the overall direction of the findings and helped validate the biases and failures observed in the LLM judges.

Finally, it is important to emphasize that the models evaluated in this study are relatively small in terms of parameter count. This work provides a starting point and a baseline for developing future automated evaluation frameworks for reverse engineering. Although computational bottlenecks emerged and some architectural biases remain, the LLM-as-a-Judge approach is promising. Future developments based on larger foundation models, specialized fine-tuning, or more advanced abstraction pipelines may reduce hallucinations and make these tools essential standards for the continuous improvement of decompilers.

Chapter 6

Conclusion

Evaluating the readability and “human-likeness” of decompiled code has long been a subjective and labor-intensive challenge in reverse engineering. This master’s thesis explored the viability of using LLMs to automate this evaluation, investigating both statistical metrics (Perplexity) and qualitative prompting strategies (LLM-as-a-Judge).

Our quantitative analysis definitively demonstrated that average perplexity is not a reliable proxy for human-likeness in decompiled code. Contrary to initial assumptions, human-written source code consistently exhibits higher perplexity than decompiled output. This is because human code contains domain-specific idioms, creative structural choices, and stylistic variance (high entropy), whereas decompilers generate rigid, repetitive, and verbose boilerplate that artificially drives perplexity down. Thus, the correlation between Δ perplexity and human preference was weak and often inversely related, confirming that perplexity alone cannot capture the qualities that make code more readable or natural to human developers.

To address the shortcomings of average perplexity, we investigated the use of Total Loss as an alternative quantitative metric. By accumulating prediction errors across all tokens, Total Loss inherently penalizes structural bloat and verbosity. Our preliminary findings suggest that Total Loss, particularly when computed on the anonymized AST representation to remove decompiler-specific lexical noise, correlates much more strongly with both LLM preferences and human judgments. While these initial results are promising and indicate that AST Total Loss could serve as a better proxy for structural code quality, using less resources and time respect the LLM-as-a-Judge approach, it should not be viewed as a definitive solution yet; further in-depth studies are required to fully validate its robustness and generalizability.

In exploring the LLM-as-a-Judge paradigm, our results highlighted distinct capabilities and vulnerabilities between the evaluated models. *Deepseek-r1* emerged as a highly ca-

pable evaluator, achieving an alignment of approximately 74% with human developers. We demonstrated that abstracting the decompiled code into an AST successfully mitigated some lexical distractions, forcing the models to evaluate purely structural logic and improving their alignment with human reasoning.

Despite these successes, we identified significant limitations. LLMs remain susceptible to analytical hallucinations, occasionally inventing non-existent technical justifications when evaluating logically equivalent snippets or minor stylistic permutations (such as reordering `switch` cases). Furthermore, the computational cost of processing long contexts for perplexity and loss calculations poses a scalability bottleneck for current hardware.

6.1 Limitations

This study has several limitations that should be acknowledged:

- **Model size and architecture:** The LLMs evaluated in this study are relatively small compared to state-of-the-art models. Larger models with more parameters may have different capabilities and biases, which could affect the generalizability of our findings.
- **Dataset size and diversity:** The dataset of decompiled code snippets used for evaluation is limited in size and may not fully represent the diversity of real-world decompilation scenarios. A larger and more varied dataset could provide a more comprehensive assessment of LLM performance.
- **Human validation sample size:** The human survey conducted to establish a ground truth for readability preferences involved a small number of participants evaluating a limited subset of 21 questions. Consequently, while the survey provides valuable preliminary insights, particularly regarding the correlation with AST Total Loss, it lacks the statistical power required for definitive conclusions.
- **Prompt engineering:** The design of prompts can significantly influence LLM responses. While we attempted to create balanced and informative prompts, different prompt formulations could yield different results, introducing variability in the evaluation.
- **Human evaluation subjectivity:** The human judgments used as a benchmark for evaluating LLM performance are inherently subjective. Different evaluators may have varying opinions on what constitutes more readable or human-like code, which can affect the consistency of the evaluation.

- **Fixed base:** In our evaluation, we compared PRs against a fixed base version. This approach does not isolate the specific changes made in some older PRs but rather evaluates different versions of the same codebase. Future work could focus on more controlled experiments that isolate specific transformations to better understand their impact on readability.

In conclusion, while LLMs are not yet flawless, objective arbiters of code quality, they represent a powerful and promising tool for the reverse engineering community. The findings of this study establish a solid baseline for automated decompiler evaluation. Future work should focus on utilizing larger foundation models, using much bigger datasets, and engaging a significantly larger pool of developer judges to rigorously validate the Total Loss metric and refine the LLM-as-a-Judge approach, while also exploring techniques to mitigate hallucinations and reduce computational overhead.

Acknowledgements

I would like to express my heartfelt thanks to all those who have supported and encouraged me throughout my academic career. First and foremost, I would like to thank my parents and my family for their love, patience, and unconditional support. I would also like to thank my groups of friends, the *ThaSuS* group, the *Office* group, the University Group of bachelor's and master's degree students, and all my other friends scattered here and there whom I would like to see more often, for the laughter, the discussions, the good times, and the mutual support we have given each other. Finally, I would like to thank my advisors and all the professors who have guided me throughout my university career for their expertise, availability, and valuable advice. Without the support of all these people, I would not have made it this far. Thank you all from the bottom of my heart!

Bibliography

- [Bai] Yumo Bai. *Why are most LLMs decoder-only?* — yumo-bai. <https://medium.com/@yumo-bai/why-are-most-llms-decoder-only-590c903e4789>.
- [BEP25] Eva-Maria C. Behner, Steffen Enders, and Elmar Padilla. “SoK: No Goto, No Cry? The Fairy Tale of Flawless Control-Flow Structuring”. In: (2025), pp. 411–431. DOI: [10.1109/EuroSP63326.2025.00032](https://doi.org/10.1109/EuroSP63326.2025.00032).
- [Car+21] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, and Colin Raffel. “Extracting Training Data from Large Language Models”. In: *30th USENIX Security Symposium (USENIX Security 21)*. 2021, pp. 2633–2650.
- [Che] Benoit Chesneau. *Gunicorn - Python WSGI HTTP Server for UNIX* — gunicorn.org. <https://gunicorn.org/>.
- [CC90] Elliot J Chikofsky and James H Cross. “Reverse engineering and design recovery: A taxonomy”. In: *IEEE Software* 7.1 (1990), pp. 13–17. DOI: [10.1109/52.43044](https://doi.org/10.1109/52.43044).
- [Dee25] DeepSeek-AI. *DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning*. 2025. arXiv: [2501.12948](https://arxiv.org/abs/2501.12948) [cs.CL]. URL: <https://arxiv.org/abs/2501.12948>.
- [Dev] Nithin Devanand. *What is Quantization in LLM* — techresearchspace. <https://medium.com/@techresearchspace/what-is-quantization-in-llm-01ba61968a51>.
- [Eag11] Chris Eagle. *The IDA Pro Book: The Unofficial Guide to the World’s Most Popular Disassembler*. 2nd. No Starch Press, 2011.
- [edm] edmcman. *Fix a variety of problems by edmcman · vul337/DecompileBench* — github.com. <https://github.com/vul337/DecompileBench/pull/4>.
- [Faca] Hugging Face. *GitHub - bitsandbytes-foundation/bitsandbytes: Accessible large language models via k-bit quantization for PyTorch.* — github.com. <https://github.com/bitsandbytes-foundation/bitsandbytes>.

- [Facb] Hugging Face. *Quantization* — *huggingface.co*. https://huggingface.co/docs/optimum/concept_guides/quantization.
- [Gao+25] Zeyu Gao, Yuxin Cui, Hao Wang, Siliang Qin, Yuanda Wang, Zhang Bolun, and Chao Zhang. “DecompileBench: A Comprehensive Benchmark for Evaluating Decompilers in Real-World Scenarios”. In: *Findings of the Association for Computational Linguistics: ACL 2025*. Ed. by Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar. Vienna, Austria: Association for Computational Linguistics, July 2025, pp. 23250–23267. ISBN: 979-8-89176-256-5. DOI: [10.18653/v1/2025.findings-acl.1194](https://doi.org/10.18653/v1/2025.findings-acl.1194). URL: <https://aclanthology.org/2025.findings-acl.1194/>.
- [Ghi] Ghidra. *./RuntimeScripts/Common/support/analyzeHeadlessREADME.md at master · NationalSecurityAgency/ghidra* — *github.com*. <https://github.com/NationalSecurityAgency/ghidra/blob/master/Ghidra/RuntimeScripts/Common/support/analyzeHeadlessREADME.md>.
- [Git] GitHub. *GitHub Actions* — *github.com*. <https://github.com/features/actions>.
- [Goo] Google. *GitHub - google/oss-fuzz: OSS-Fuzz - continuous fuzzing for open source software.* — *github.com*. <https://github.com/google/oss-fuzz>.
- [Gra] Gradle. *Gradle Build Tool* — *gradle.org*. <https://gradle.org/>.
- [Gup] Mehul Gupta. *What are DeepSeek-R1 distilled models?* — *medium.com*. <https://medium.com/data-science-in-your-pocket/what-are-deepseek-r1-distilled-models-329629968d5d>.
- [Hin+12] Abram Hindle, Earl T Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu. “On the naturalness of software”. In: *2012 34th International Conference on Software Engineering (ICSE)*. IEEE. 2012, pp. 837–847.
- [HLC24] Peiwei Hu, Ruigang Liang, and Kai Chen. “DeGPT: Optimizing Decompiler Output with LLM”. In: *Proceedings 2024 Network and Distributed System Security Symposium (2024)*. URL: <https://api.semanticscholar.org/CorpusID:267622140>.
- [IBM] IBM. *Foundation model parameters: decoding and stopping criteria* — *ibm.com*. <https://www.ibm.com/docs/en/watsonx/saas?topic=prompts-model-parameters-prompting>.
- [Li+25] Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, Kai Shu, Lu Cheng, and Huan Liu. “From Generation to Judgment: Opportunities and Challenges of LLM-as-a-judge”. In: *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*.

- Ed. by Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng. Suzhou, China: Association for Computational Linguistics, Nov. 2025, pp. 2757–2791. ISBN: 979-8-89176-332-6. DOI: [10.18653/v1/2025.emnlp-main.138](https://doi.org/10.18653/v1/2025.emnlp-main.138). URL: <https://aclanthology.org/2025.emnlp-main.138/>.
- [Mik] Rebeka Kiss Miklós Sebök. *LLM Parameters Explained: A Practical, Research-Oriented Guide with Examples* — *promptrevolution.poltextlab.com*. <https://promptrevolution.poltextlab.com/llm-parameters-explained-a-practical-research-oriented-guide-with-examples/>.
- [Mor] Abby Morgan. *Perplexity for LLM Evaluation* — *comet.com*. <https://www.comet.com/site/blog/perplexity-for-llm-evaluation/>.
- [Mul+18] Lincoln A. Mullen, Kenneth Benoit, Os Keyes, Dmitry Selivanov, and Jeffrey Arnold. “Fast, Consistent Tokenization of Natural Language Text”. In: *Journal of Open Source Software* 3.23 (2018), p. 655. DOI: [10.21105/joss.00655](https://doi.org/10.21105/joss.00655). URL: <https://doi.org/10.21105/joss.00655>.
- [Nat19] National Security Agency. *Ghidra Software Reverse Engineering Framework*. <https://ghidra-sre.org/>. 2019.
- [NSA] NSA. *Ghidra Decompiler Analysis Engine: Decompiler Analysis Engine* — *share.google*. <https://share.google/Q9gHjuuTY3ZlFlm4n>.
- [Pco] Pcode-doc. *P-Code Reference Manual* — *spinsel.dev*. https://spinsel.dev/assets/2020-06-17-ghidra-brainfuck-processor-1/ghidra_docs/language_spec/html/pcoderef.html.
- [pyt] Hugging Face python. *datasets* — *pypi.org*. <https://pypi.org/project/datasets/>.
- [Ron] Armin Ronacher. *Welcome to Flask & Flask Documentation (3.1.x)* — *flask.palletsprojects.com*. <https://flask.palletsprojects.com/en/stable/>.
- [Rua+25] Zhiwen Ruan, Yixia Li, He Zhu, Yun Chen, Peng Li, Yang Liu, and Guanhua Chen. *Enhancing Large Language Model Reasoning via Selective Critical Token Fine-Tuning*. 2025. arXiv: [2510.10974](https://arxiv.org/abs/2510.10974) [cs.CL]. URL: <https://arxiv.org/abs/2510.10974>.
- [Sax] Shubham Saxena. *Understanding Perplexity in Language Models: A Detailed Exploration* — *shubhamsd100*. <https://medium.com/@shubhamsd100/understanding-perplexity-in-language-models-a-detailed-exploration-2108b6ab85af>.
- [Son] SonarSource. *Cyclomatic Complexity Guide — How To Calculate & Test* — *sonarsource.com*. <https://www.sonarsource.com/resources/library/cyclomatic-complexity/>.

- [Sta+24] Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. “Beyond Memorization: Violating Privacy via Inference with Large Language Models”. In: *The Twelfth International Conference on Learning Representations (ICLR)*. 2024.
- [Tan+24] Hanzhuo Tan, Qi Luo, Jing Li, and Yuqun Zhang. “LLM4Decompile: Decompiling Binary Code with Large Language Models”. In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2024, pp. 3473–3487. DOI: [10.18653/v1/2024.emnlp-main.203](https://doi.org/10.18653/v1/2024.emnlp-main.203). URL: <http://dx.doi.org/10.18653/v1/2024.emnlp-main.203>.
- [Ten] Tencent. *Qwen/Qwen3-14B · Hugging Face* — [huggingface.co](https://huggingface.co/Qwen/Qwen3-14B). <https://huggingface.co/Qwen/Qwen3-14B>.
- [Vas+17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. “Attention is All you Need”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- [Vol] Dr. Hannah Volk-Jesussek. *Point-Biserial Correlation: A Beginner’s Guide* — numiqo.com. <https://numiqo.com/tutorial/point-biserial-correlation>.
- [Wik] Wikipedia. *Atom (text editor) - Wikipedia* — en.wikipedia.org. [https://en.wikipedia.org/wiki/Atom_\(text_editor\)#:~:text=Atom%20uses%20Tree%2Dsitter%20to%20provide%20syntax%20highlighting%20for%20multiple%20programming%20languages%20and%20file%20formats.%5B17%5D](https://en.wikipedia.org/wiki/Atom_(text_editor)#:~:text=Atom%20uses%20Tree%2Dsitter%20to%20provide%20syntax%20highlighting%20for%20multiple%20programming%20languages%20and%20file%20formats.%5B17%5D).
- [wika] wikipedia. *Perplexity - Wikipedia* — en.wikipedia.org. <https://en.wikipedia.org/wiki/Perplexity>.
- [wikb] wikipedia. *Top-p sampling - Wikipedia* — en.wikipedia.org. https://en.wikipedia.org/wiki/Top-p_sampling.
- [Zvo] Enes Zvornicanin. *What Is and Why Use Temperature in Softmax?* — *Baeldung on Computer Science* — [baeldung.com](https://www.baeldung.com/cs/softmax-temperature). <https://www.baeldung.com/cs/softmax-temperature>.

Appendix A

Quiz

These are the quiz questions that were used in the human evaluation of the decompiled code snippets.

A.1 Participant instructions

In this short quiz (about 5–10 minutes), you will be presented with pairs of C code snippets (Snippet A and Snippet B). Each pair implements the same logical operation, but with different writing styles.

Your task is to choose which snippet appears more *humanly written*, idiomatic, and maintainable. There is no right or wrong answer; we are interested in your developer intuition. If you cannot decide, leave the answer blank.

All responses are collected anonymously and used only for academic research.

A.2 Quiz questions

```

char *local_28;
size_t local_20;
char *local_18;

for (; param_4 != 0; param_4 = param_4 - 1) {
    if ((param_3[param_4 - 1] & 0xdfU) != 0) goto LAB_00101160;
    param_4 = 0;
LAB_00101160:
    if (param_2 == 0) {
        return 0x13;
    }
    local_20 = param_4;
    local_18 = param_3;
    if (param_5 == (iconv_t)0x0) {

```

```

char *local_28;
size_t local_20;
char *local_18;

if (param_4 != 0) {
    do {
        if ((param_3[param_4 - 1] & 0xdfU) != 0) goto LAB_00101160;
        param_4 = param_4 - 1;
    } while (param_4 != 0);
}
param_4 = 0;
LAB_00101160:
if (param_2 == 0) {
    return 0x13;
}

```

Figure A.1: PR #8161 — task-readstat_readstat_convert-02 (Reason for inclusion: consistency)

```

case 3:
    pcVar21 = bsearch(local_58, __base,
(long)iVar5, 0x2B, compare_key_varlookup);
    pcVar8 = pcVar21;
    if (pcVar21 != (char *)0x0) {

```

```

case 3:
    pcVar21 = (char *)bsearch(local_58, __base,
(long)iVar5, 0x2B, compare_key_varlookup);
    pcVar8 = pcVar21;
    if (pcVar21 != (char *)0x0) {

```

Figure A.2: PR #8161 — task-readstat_sav_parse_very_long_string_record-03 (Reason for inclusion: consistency between DeepSeek; Qwen inconsistency)

```

}
else {
    local_28 = param_3;
    local_20 = param_2;
    local_18 = param_1;
    for (local_38 = 0; local_38 < ncompr; local_38 = local_38 + 1) {
        iVar3 = *(int *)compr + local_38 * 0x20 + 8;
        iVar1 = *(int *)compr + local_38 * 0x20 + 8;
        if (-1 < iVar3) {
            iVar1 = iVar3;
        }
        if ((ulong)(long)iVar1 <= local_78) {
            if (*(int *)compr + local_38 * 0x20 + 8 < 0) {
                local_11c = (**code **)(compr + local_38 * 0x20)(local_78);

```

```

else {
    local_28 = param_3;
    local_20 = param_2;
    local_18 = param_1;
    for (local_38 = 0; local_38 < ncompr; local_38 = local_38 + 1) {
        iVar1 = *(int *)compr + local_38 * 0x20 + 8;
        if (iVar1 < 1) {
            iVar1 = -iVar1;
        }
        if ((ulong)(long)iVar1 <= local_78) {
            if (*(int *)compr + local_38 * 0x20 + 8 < 0) {
                local_11c = (**code **)(compr + local_38 * 0x20)(local_78);
            }
            else {

```

Figure A.3: PR #8161 — task-file_file_zmagic-00 (Reason for inclusion: older PR but better output, possibly due to conciseness)

```

if ('\0' < (char)_sav_date_parse_range_lengths[lVar11]) {
    pcVar16 = pcVar13 + (ulong)uVar12 * 2 - 2;
    if (pcVar13 <= pcVar16) {
        pcVar15 = pcVar13;
        do {

```

```

if ('\0' < (char)_sav_date_parse_range_lengths[lVar11]) {
    pcVar16 = pcVar13 + ((ulong)uVar12 * 2 - 2);
    if (pcVar13 <= pcVar16) {
        pcVar15 = pcVar13;
        do {

```

Figure A.4: PR #8628 — task-readstat_sav_parse_date-03 (Reason for inclusion: inconsistency between () and +-)

| | |
|---|--|
| <pre> if (pbVar3 < pbVar7 + *(long *) (command_lengths + (ulong)(uint)bVar5 * 8)) { return 0xffffffffffffffff; } if (0xf < bVar5) { switch0_001011a4_case0_3: uVar4 = uVar6; goto LAB_00101153; } uVar1 = *param_3 & 0xf; switch((uint)bVar5) { case 0: sVar2 = (ulong)param_3[1] + (ulong)(uVar1 << 8) + 0x40; goto LAB_00101240; case 1: sVar2 = (ulong)param_3[1] + (ulong)(uVar1 << 8) + 0x1040; LAB_00101240: param_3 = param_3 + 2; goto joined_r0x0010130a; case 2: uVar1 = uVar1 + 0x60; goto LAB_00101300; case 3: goto switch0_001011a4_case0_3; case 4: sVar2 = (ulong)param_3[1] + (ulong)(uVar1 << 8) + 0x12; bVar5 = param_3[2]; pbVar7 = param_3 + 3; } </pre> | <pre> if (pbVar3 < pbVar7 + *(long *) (command_lengths + (ulong)(uint)bVar5 * 8)) { return 0xffffffffffffffff; } uVar4 = uVar6; if (0xf < bVar5) goto LAB_00101153; uVar1 = *param_3 & 0xf; switch((uint)bVar5) { case 0: sVar2 = (ulong)param_3[1] + (ulong)(uVar1 << 8) + 0x40; goto LAB_00101240; case 1: sVar2 = (ulong)param_3[1] + (ulong)(uVar1 << 8) + 0x1040; LAB_00101240: param_3 = param_3 + 2; goto joined_r0x0010130a; case 2: uVar1 = uVar1 + 0x60; goto LAB_00101300; case 3: goto LAB_00101153; case 4: sVar2 = (ulong)param_3[1] + (ulong)(uVar1 << 8) + 0x12; bVar5 = param_3[2]; pbVar7 = param_3 + 3; break; case 5: pbVar7 = param_3 + 2; } </pre> |
|---|--|

Figure A.7: PR #6722 — task-readstat_sas_rle_decompress-03 (Reason for inclusion: PR consistency and tie outcome)

| | |
|---|--|
| <pre> local_168 = param_3; do { uVar1 = *puVar12; uVar3 = -uVar1; if ((int)-uVar1 < 0) { uVar3 = uVar1; } } </pre> | <pre> local_168 = param_3; do { uVar1 = *puVar12; uVar3 = -uVar1; if (0 < (int)uVar1) { uVar3 = uVar1; } } </pre> |
|---|--|

Figure A.8: PR #6722 — task-file_file_zmagic-03 (Reason for inclusion: unusual if structure; PR wins)

```

    local_28 = (long)(int)param_3;
    break;
    case '\x18':
    case '\x19':
    case '\x1a':
    case '\x1b':
    case '\x1c':
    case '\x1d':
    case '\x1e':
    case '\x1f':
    case '!':
    case '$':
    case '%':
    case '&':
    case '*':
    case '+':
    case ',':
    case '2':
    case '3':
    case '4':
    break;
    default:
    if (((uint*)(param_1 + 0x44) & 0x40) != 0) {
        file_magwarn(param_1, "cannot happen: m->type=%d\n", (char)param_2[3]);
    }
    return -1;
}
return local_28;
}

```

```

    local_28 = (long)(int)param_3;
    break;
    case 0x18:
    case 0x19:
    case 0x1a:
    case 0x1b:
    case 0x1c:
    case 0x1d:
    case 0x1e:
    case 0x1f:
    case 0x20:
    case 0x24:
    case 0x25:
    case 0x26:
    case 0x2a:
    case 0x2b:
    case 0x2c:
    case 0x32:
    case 0x33:
    case 0x34:
    break;
    default:
    if (((uint*)(param_1 + 0x44) & 0x40) != 0) {
        file_magwarn(param_1, "cannot happen: n->type=%d\n", *(undefined *)
{param_2 + 3});
    }
    return -1;
}
return local_28;
}

```

Figure A.9: PR #6722 — task-file_file_signextend-00 (Reason for inclusion: inconsistency)

| | |
|---|---|
| <pre> uvar0 = read_tag(puvar7); switch(uVar5) { case 0x31: case 0x32: case 0x33: iVar6 = read_string(puVar7, local_538, 0x100); break; case 0x34: iVar6 = read_variable_count_record(puVar7); break; case 0x35: iVar6 = read_precision_record(puVar7); break; case 0x36: iVar6 = read_case_weight_record(puVar7); break; case 0x37: iVar6 = read_variable_record(puVar7); break; case 0x38: iVar6 = read_missing_value_record(puVar7); break; case 0x39: iVar6 = read_missing_value_lo_range_record(puVar7); break; default: goto switchD_00101979_caseD_3a; case 0x41: iVar6 = read_missing_value_hi_range_record(puVar7); break; case 0x42: iVar6 = read_missing_value_range_record(puVar7); break; case 0x43: iVar6 = read_variable_label_record(puVar7); break; case 0x44: iVar6 = read_value_label_record(puVar7); break; case 0x45: iVar6 = read_document_record(puVar7); break; case 0x46: iVar6 = 8; if (((int) *1/1000000000 + 0x2ee) == *(int) *1/1000000000 + </pre> | <pre> uvar0 = read_tag(puvar7); switch(uVar5) { default: goto switchD_00101979_caseD_3a; case 0x41: iVar6 = read_missing_value_hi_range_record(puVar7); break; case 0x31: case 0x32: case 0x33: iVar6 = read_string(puVar7, local_538, 0x100); break; case 0x39: iVar6 = read_missing_value_lo_range_record(puVar7); break; case 0x42: iVar6 = read_missing_value_range_record(puVar7); break; case 0x35: iVar6 = read_precision_record(puVar7); break; case 0x37: iVar6 = read_variable_record(puVar7); break; case 0x34: iVar6 = read_variable_count_record(puVar7); break; case 0x36: iVar6 = read_case_weight_record(puVar7); break; case 0x43: iVar6 = read_variable_label_record(puVar7); break; case 0x45: iVar6 = read_document_record(puVar7); break; case 0x38: iVar6 = read_missing_value_record(puVar7); break; case 0x44: iVar6 = read_value_label_record(puVar7); break; case 0x46: iVar6 = 8; if (((int) *1/1000000000 + 0x2ee) == *(int) *1/1000000000 + </pre> |
|---|---|

Figure A.10: PR #7253 — task-readstat_readstat_parse_por-03 (Reason for inclusion: difficult to assess)

| | |
|---|--|
| <pre> iVar8 = 0; do { switch(*(undefined1 *) (lVar15 + 0x102361 + lVar8)) { case 0: iVar16 = (int) *pcVar12 + iVar16 * 10 + -0x30; </pre> | <pre> iVar8 = 0; do { switch(_dta_timestamp_parse_actions[lVar8 + lVar15 + 1]) { case 0: iVar16 = (int) *pcVar12 + iVar16 * 10 + -0x30; </pre> |
|---|--|

Figure A.11: PR #8587 — task-readstat_dta_parse_timestamp-02 (Reason for inclusion: PR consistency)

| | |
|--|--|
| <pre> do { iVar5 = iVar14; switch(*(undefined1 *) (lVar11 + 0x102311 + lVar8)) { case 0: iVar5 = (int) (char) (*pcVar10 - 0x30U) + iVar14 * 10; </pre> | <pre> do { iVar5 = iVar14; switch(_sav_date_parse_actions[lVar8 + lVar11 + 1]) { case 0: iVar5 = (int) (char) (*pcVar10 - 0x30U) + iVar14 * 10; </pre> |
|--|--|

Figure A.12: PR #8587 — task-readstat_sav_parse_date-02 (Reason for inclusion: the only inconsistency in the PR)

```

uVar3 = 0xb;
if (param_1 != (long *)0x0) {
  lVar7 = *param_1;
  if (lVar7 != -1) {
    uVar4 = 0;
    uVar5 = 0;
    uVar3 = 0;
    bVar1 = true;
    do {
      param_1 = param_1 + 2;
      bVar6 = 1;
      bVar2 = false;
      switch(lVar7) {
        case 3:
        case 4:
        case 5:
        case 6:
        case 7:
        case 8:
        case 9:
        case 10:
        case 0xb:
          bVar2 = true;
          bVar6 = 0;
          break;
        case 0xc:
        case 0xd:
        case 0xe:
        case 0xf:
        case 0x10:
        case 0x11:
        case 0x12:
        case 0x13:
        case 0x14:
        case 0x15:
        case 0x16:
        case 0x17:
        case 0x18:
        case 0x19:
        case 0x1a:
        case 0x1b:
        case 0x1c:
        case 0x1d:
        case 0x1e:
        case 0x1f:
        case 0x20:
          goto switchD_00101156_caseD_c;
        case 0x21:
          break;
        default:
          bVar6 = 1;
          bVar2 = false;
          if (1 < lVar7 + 0xbffffffffffffffu) goto switchD_00101156_caseD_c;
      }
      if (!bVar1) goto switchD_00101156_caseD_c;
      uVar4 = uVar4 + bVar6;
      uVar5 = uVar5 + 1;
      lVar7 = *param_1;
      bVar1 = bVar2;
    } while (lVar7 != -1);
    if ((bVar6 == 1 && uVar5 < 5) && (uVar4 < 4)) {
      *param_2 = uVar5;
      uVar3 = 0;
    }
  }
}

```

```

uVar3 = 0xb;
if (param_1 != (long *)0x0) {
  lVar7 = *param_1;
  if (lVar7 != -1) {
    param_1 = param_1 + 2;
    uVar4 = 0;
    uVar5 = 0;
    uVar3 = 0;
    bVar1 = true;
    do {
      bVar6 = 1;
      bVar2 = false;
      switch(lVar7) {
        case 3:
        case 4:
        case 5:
        case 6:
        case 7:
        case 8:
        case 9:
        case 10:
        case 0xb:
          bVar2 = true;
          bVar6 = 0;
          break;
        case 0x21:
          break;
        default:
          bVar6 = 1;
          bVar2 = false;
          if (1 < lVar7 + 0xbffffffffffffffu) goto switchD_00101156_caseD_c;
          break;
        case 0xc:
        case 0xd:
        case 0xe:
        case 0xf:
        case 0x10:
        case 0x11:
        case 0x12:
        case 0x13:
        case 0x14:
        case 0x15:
        case 0x16:
        case 0x17:
        case 0x18:
        case 0x19:
        case 0x1a:
        case 0x1b:
        case 0x1c:
        case 0x1d:
        case 0x1e:
        case 0x1f:
        case 0x20:
          goto switchD_00101156_caseD_c;
      }
      if (!bVar1) goto switchD_00101156_caseD_c;
      uVar4 = uVar4 + bVar6;
      uVar5 = uVar5 + 1;
      lVar7 = *param_1;
      param_1 = param_1 + 2;
      bVar1 = bVar2;
    } while (lVar7 != -1);
    if ((bVar6 == 1 && uVar5 < 5) && (uVar4 < 4)) {
      *param_2 = uVar5;
      uVar3 = 0;
    }
  }
}

```

Figure A.13: PR #7252 — task-xz_lzma_validate_chain-02 (Reason for inclusion: DeepSeek inconsistency and Qwen error)

```

*{undefined *}(*(long *)(local_10 + 0x20) + (long)local_70) = 0;
if (local_8c == 0) {
  local_90 = &DAT_0010201c;
}
else {
  local_90 = {undefined *}(*(long *)(local_10 + 0x20) + (long)local_6c);
}

```

```

*{undefined *}(*(long *)(local_10 + 0x20) + (long)local_70) = 0;
if (local_8c == 0) {
  local_90 = "";
}
else {
  local_90 = {char *}(*(long *)(local_10 + 0x20) + (long)local_6c);
}

```

Figure A.14: PR #8587 — task-file_file_replace-00 (Reason for inclusion: non-standard modification; PR wins)

| | |
|---|---|
| <pre> xls_showB0F(&local_1c); } if (local_1c._2.2. != 0) { local_20 = realloc(local_20, (ulong)local_1c._2.2.); if (local_20 == (ushort *)0x0) { if (*(int *)PTR_xls_debug_00104fd8 != 0) { fprintf((FILE **)PTR_stderr_00104ff8, "Error: failed to allocate buffer of size %d\n", </pre> | <pre> xls_showB0F(&local_1c); } if (local_1c._2.2. != 0) { local_20 = (ushort *)realloc(local_20, (ulong)local_1c._2.2.); if (local_20 == (ushort *)0x0) { if (*(int *)PTR_xls_debug_00104fd8 != 0) { fprintf((FILE **)PTR_stderr_00104ff8, "Error: failed to allocate buffer of size %d\n", </pre> |
|---|---|

Figure A.15: PR #7252 — task-libxls_xls_parseWorkBook-00 (Reason for inclusion: single-cast modification, but with inconsistency)

| | |
|---|--|
| <pre> undefined0 file_magicfind(long param_1, char *param_2, long *param_3) { long *p1Var1; ulong uVar2; long lVar3; uint uVar4; int iVar5; uint uVar6; ulong uVar7; long *p1Var8; long lVar9; ulong uVar10; } p1Var1 = *(long **)(param_1 + 8); p1Var8 = (long *)p1Var1[4]; if (p1Var8 == p1Var1) { return 0xffffffff; } do { uVar2 = p1Var8[2]; if (uVar2 != 0) { lVar3 = *p1Var8; uVar10 = 0; do { lVar9 = uVar10 * 0x1b0; if ((*(char *) (lVar3 + 6 + lVar9) == '-') && iVar5 = strcmp((char *) (lVar3 + lVar9 + 0x20), param_2), iVar5 == 0) { *param_3 = lVar9 + lVar3; param_3[1] = uVar10 * 8 + p1Var8[1]; uVar4 = 1; goto LAB_001011c0; } uVar10 = (ulong)((int)uVar10 + 1); } while (uVar10 < uVar2); } p1Var8 = (long *)p1Var8[4]; if (p1Var8 == p1Var1) { return 0xffffffff; } } while (true); while (uVar4 = uVar6 + 1, *(char *) (lVar3 + 2 + uVar7 * 0x1b0) != '\0') { LAB_001011c0: uVar6 = uVar4; uVar7 = (ulong)(uVar6 + (int)uVar10); if (uVar2 <= uVar7) break; } param_3[2] = (ulong)uVar6; return 0; } </pre> | <pre> //----- (0000000000001110) ----- int64 __fastcall file_magicfind(__int64 a1, const char *a2, __QWORD *a3) { __int64 v3; // rcx __int64 v4; // rbx __int64 result; // rax unsigned __int64 v6; // r15 __int64 v7; // r12 unsigned __int64 v8; // r13 __int64 v9; // rbp int v10; // ecx unsigned int v11; // eax unsigned __int64 v12; // rcx __int64 v13; // rdx __int64 v15; // [rsp+10h] [rbp-38h] v3 = *(__int64 **)(a1 + 8); v4 = (__int64 *)v3[4]; result = 0xffffffff; v15 = v3; if (v4 != v3) { while (1) { v6 = v4[2]; if (v6) break; LABEL_3: v4 = (__int64 *)v4[4]; if (v4 == v15) return 0xffffffff; } v7 = *v4; v8 = 0; while (1) { v9 = 432 * v8; if (*(_BYTE *) (v7 + 432 * v8 + 6) == 45 && !strcmp((const char *) (v7 + v8 + 32), a2)) break; v8 = (unsigned int)(v8 + 1); if (v6 <= v8) goto LABEL_3; } *a3 = v7 + v9; a3[1] = v4[1] + 8 * v8; v10 = 1; do { v11 = v10; v12 = (unsigned int)(v10 + v8); if (v6 <= v12) break; break; v13 = 432 * v12; v10 = v11 + 1; } while (*(_BYTE *) (v7 + v13 + 2)); a3[2] = v11; return 0; } return result; } </pre> |
|---|--|

Figure A.16: Dogbolt — task-file_file_magicfind-02 (Reason for inclusion: Ghidra vs Hex-Rays comparison)

```

uint8_t file_is_simh(void* arg1, void* arg2)
{
    int64_t rax_1 = *(arg2 + 0x08);
    int32_t rax_6 = *(arg1 + 0x44) & 0x410;
    int32_t var_c;
    if ( !(* (arg1 + 0x44) & 0x1000800) )
    {
        if ( !simh_parse(rax_1, rax_1 + *(arg2 + 0x00)) )
            var_c = 0;
        else if (rax_6 == 0x400)
            var_c = 1;
        else
        {
            int32_t rax_10;
            if (!rax_6)
            {
                rax_10 = 0;
                if ( file_printf(arg1, "SIMHtape\data") != 0xffffffff)
                    var_c = 1 ;
                else
                    var_c = 0xffffffff;
            }
            else
            {
                rax_10 = 0;
                if ( file_printf(arg1, "application/SIMH \tape \data") != 0xffffffff)
                    var_c = 1 ;
                else
                    var_c = 0xffffffff;
            }
        }
        else
            var_c = 0;
    }
    return var_c;
}

//----- (0000000000001110) -----
int64_t __fastcall file_is_simh(__int64 a1, __int64 a2)
{
    int v3; // [rsp+4h] [rbp-2Ch]
    v3 = (_DWORD *)(a1 + 08) & 0x410;
    if ( !(_DWORD *)(a1 + 08) & 0x1000800) != 0 )
    {
        return 0;
    }
    else if ( simh_parse(*(int **)(a2 + 152), *(_QWORD **)(a2 + 160) + *
(_QWORD **)(a2 + 152)) )
    {
        if (v3 == 1024)
        {
            return 1;
        }
        else if (v3)
        {
            if ( (unsigned int)file_printf(a1, "application/SIMH-tape\data") ==
-1 )
                return (unsigned int)-1;
            else
                return 1;
        }
        else if ( (unsigned int)file_printf(a1, "SIMH tape data") == -1 )
        {
            return (unsigned int)-1;
        }
        else
        {
            return 1;
        }
    }
    else
    {
        return 0;
    }
}

```

Figure A.17: Dogbolt — task-file_file_is_simh-00 (Reason for inclusion: Binary Ninja vs Hex-Rays comparison)

```
uint64_t file_is_simh(void* arg1, void* arg2)
```

```
{  
    int32_t rbp = *(arg1 + 0x44);  
    int32_t rbx = 0;  
    if (!(rbp & 0x1000800) ) {  
        int64_t rdi = *(arg2 + 0x98);  
        if (simh_parse(rdi,  
            *(arg2 + 0xa8) + rdi)){  
            int32_t rbp_1 = rbp & 0x410;  
            if (rbp_1 == 0x400)  
                rbx = 1;  
            else {  
                char* const rsi_2;  
                rsi_2 = rbp_1 ? "application/SIMH-tape-data" : "SIMH tape  
data";  
                int64_t rcx_1;  
                rcx_1 = file_printf(arg1, rsi_2) != 0xffffffff;  
                rbx = (rcx_1 << 1) + -fffffffffffffff;  
            }  
        }  
    }  
    return rbx;  
}
```

```
int file_is_simh(long param_1, long param_2)
```

```
{  
    int iVar1;  
    int iVar2;  
    uint uVar3;  
    char *pcVar4;  
    uVar3 = *(uint*)(param_1 + 0x44);  
    iVar2 = 0;  
    if ((uVar3 & 0x1000800) == 0) {  
        iVar1 = simh_parse(*(long*)(param_2 + 0x98),  
            *(long*)(param_2 + 0xa8) + *(long*)(param_2 +  
0x98));  
        if (iVar1 != 0) {  
            uVar3 = uVar3 & 0x410;  
            if (uVar3 == 0x400) {  
                iVar2 = 1;  
            }  
            else {  
                if (uVar3 == 0) {  
                    pcVar4 = "SIMH tape data";  
                }  
                else {  
                    pcVar4 = "application/SIMH-tape-data";  
                }  
                iVar2 = file_printf(param_1, pcVar4);  
                iVar2 = (uint)(iVar2 != -1) * 2 + -1;  
            }  
        }  
    }  
    return iVar2;  
}
```

Figure A.18: Dogbolt — task-file_file_is_simh-03 (Reason for inclusion: Binary Ninja vs Ghidra comparison)

| | |
|--|---|
| <pre> int file_default(long param_1, long param_2) { uint uVar1; int iVar2; char *pcVar3; uVar1 = *(uint *)(param_1 + 0x44); if ((uVar1 & 0x410) == 0) { if ((uVar1 & 0x000) == 0) { if ((uVar1 & 0x1000000) == 0) { return 0; } pcVar3 = "???" } else { pcVar3 = "UNKNUNKN"; } iVar2 = file_printf(param_1, pcVar3); return (uint)(iVar2 & -1) * 2 + -1; } if ((uVar1 & 0x10) != 0) { pcVar3 = "octet-stream"; if (param_2 == 0) { pcVar3 = "x-empty"; } iVar2 = file_printf(param_1, "application/%s", pcVar3); if (iVar2 == -1) { return -1; } } return 1; } </pre> | <pre> //----- (0000000000001580) ----- ---- __int64 __fastcall file_default(__int64 a1, __int64 a2, __int64 a3, __int64 a4, __int64 a5, __int64 a6) { char v6; // a1 __int64 v7; // rcx const char *v8; // rdx __int64 result; // rax const char *v10; // rsi char v11; // [rsp-8h] [rbp-8h] v11 = v6; v7 = *(unsigned int *)(a1 + 68); if ((v7 & 0x410) != 0) { if ((v7 & 0x10) == 0) return 1; v8 = "octet-stream"; if (!a2) v8 = "x-empty"; if ((unsigned int)file_printf(a1, (__int64)"application/%s", (__int64)v6, v7, a5, a6, v6) == -1) return 0xFFFFFFFFLL; else return 1; } else { if ((v7 & 0x800) != 0) { v10 = "UNKNUNKN"; } else { result = 0; if ((v7 & 0x1000000) == 0) return result; v10 = "???" } return 2 * (unsigned int)((unsigned int)file_printf(a1, (__int64)v10, a3, v7, a5, a6, v11) != -1) - 1; } } </pre> |
|--|---|

Figure A.19: Code — task-file_file_default-02 (Reason for inclusion: Ghidra vs Hex-Rays comparison)

| | |
|---|--|
| <pre> int64_t magic_buffer(int64_t arg1, int64_t arg2, int64_t arg3) { if (!arg1) return 0; if (file_reset(arg1, 1) == 0xffffffff) return 0; if (file_buffer(arg1, 0xffffffff, 0, 0, arg2, arg3, arg3, arg2, arg1) != 0xffffffff) return file_getbuffer(arg1); return 0; } </pre> | <pre> //----- (0000000000001130) ----- ---- __int64 __fastcall magic_buffer(__int64 a1, __int64 a2, __int64 a3) { if (!a1) return 0; if ((unsigned int)file_reset(a1, 1) == -1) return 0; if ((unsigned int)file_buffer(a1, 0xFFFFFFFFLL, 0, 0, a2, a3) == -1) return 0; return file_getbuffer(a1); } </pre> |
|---|--|

Figure A.20: Code — task-file_magic_buffer-00 (Reason for inclusion: Binary Ninja vs Hex-Rays comparison)

```

int64_t fmtcheck(int64_t arg1, int64_t arg2)
{
    if (!arg1)
        return arg2;
    int64_t var_20 = arg1;
    int32_t var_34_1 = 0;
    int64_t var_30 = arg2;
    int32_t var_38_1 = 0;
    do {
        int32_t rax_3 = get_next_format(&var_20, var_34_1);
        var_34_1 = rax_3;
        if (rax_3 == 0xe) return arg1;
        if (var_34_1 == 0xf) return arg2;
        var_38_1 = get_next_format(&var_30, var_38_1);
    } while (var_34_1 == var_38_1);
    return arg2;
}

long fatcheck(long param_1, long param_2)
{
    int local_38;
    int local_34;
    long local_30;
    long local_20;
    long local_28;
    long local_18;
    long local_10;
    local_10 = param_2;
    if (param_1 != 0) {
        local_34 = 0;
        local_30 = 0;
        local_38 = param_2;
        local_20 = param_1;
        local_28 = param_2;
        local_18 = param_1;
        do {
            local_34 = get_next_format(&local_20, local_34);
            if (local_34 == 0xe) {
                return local_18;
            }
            if (local_34 == 0xf) {
                return local_20;
            }
            local_38 = get_next_format(&local_30, local_38);
        } while (local_34 == local_30);
        local_10 = local_20;
    }
    return local_10;
}

```

Figure A.21: Code — task-file_fmtcheck-00 (Reason for inclusion: Binary Ninja vs Ghidra comparison)

A.3 Quiz results

Table A.1: Quiz responses (R = Right, L = Left, - = blank).

| Participant | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 |
|-------------|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|
| P1 | R | R | R | R | R | R | R | R | L | R | R | R | R | R | L | L | L | R | L | L | L |
| P2 | L | R | R | R | R | R | R | L | L | R | R | R | R | R | L | L | R | R | L | R | L |
| P3 | L | - | R | R | R | R | R | R | L | R | R | R | L | R | L | L | R | R | L | - | L |
| P4 | R | L | R | R | R | R | L | R | L | L | R | R | L | R | R | R | R | R | R | L | L |
| P5 | L | L | R | R | R | R | R | L | L | L | R | R | - | R | L | L | R | - | L | R | R |
| P6 | R | L | R | R | R | R | L | R | R | R | R | R | R | R | R | R | R | R | R | R | R |
| P7 | L | L | R | R | R | R | R | R | L | R | R | R | L | R | L | R | L | R | R | L | R |
| P8 | R | L | R | R | R | R | R | R | L | R | L | R | R | R | R | R | R | R | R | L | L |
| P9 | L | R | R | R | R | R | R | R | L | R | R | R | L | R | R | L | L | R | L | L | L |
| P10 | L | - | R | R | R | R | - | R | L | R | R | R | L | R | - | L | R | R | L | R | R |
| P11 | L | R | R | L | R | R | R | R | L | R | R | R | R | R | R | R | R | R | L | R | R |

Appendix B

Future Ghidra work

B.1 Alternative decompilation flows

A possible future direction, based on the “groups” mechanism described in the Section 3.3.2, is to define a standard way to enable alternative action pipelines in the decompiler. In practice, existing actions could include guarded branches (e.g., if-based checks, like line 7 of this example) to activate different transformation flows and generate multiple candidate C outputs for the same function.

Listing B.1: Example of different branch of code if a specific group is set

```
1 void BlockWhileDo::finalTransform(Funcdata &data){
2     // Simplification style
3     const string &currentAction = data.getArch()->allacts.getCurrentName();
4     BlockGraph::finalTransform(data);
5     if (!data.getArch()->analyze_for_loops) return;
6     if (hasOverflowSyntax())
7         if (currentAction == "different"){// LukeSerne example of PR-specific
8             transformation
9             // Apply a different set of transformations to generate an alternative
10            decompilation variant
11            [...]
12        }else
13            return; // Still too complex
14        FlowBlock *copyBl = getFrontLeaf();
15        [...] // Original transformation flow for the standard decompilation variant
16    }
```

On top of this, a dedicated Ghidra extension could orchestrate the process by:

- collecting all generated decompilation variants,
- sending them to an LLM through an API call for ranking/selection,
- and showing to the analyst only the variant selected as the most readable or human-like.

We have prototyped this architecture in a separate project, and the following Figures B.1 and B.2 show a possible code structure for the extension and a mockup of the user interface.

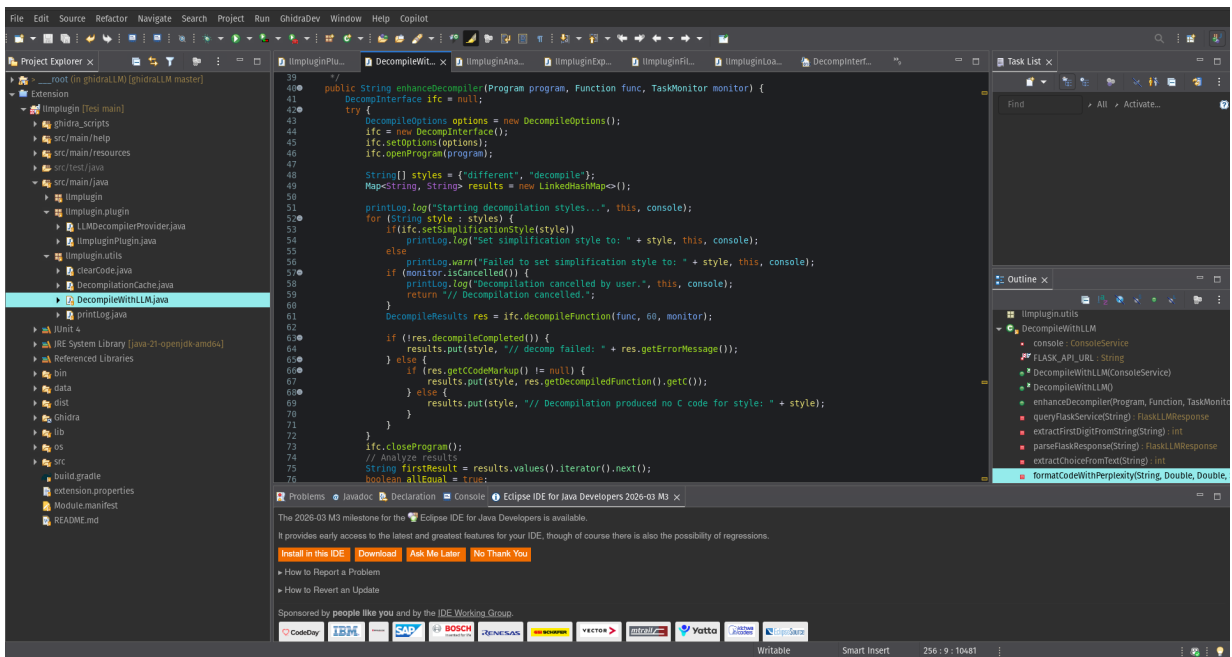


Figure B.1: A possible code architecture for a Ghidra extension that manages multiple decompilation variants and uses an LLM to select the best one.

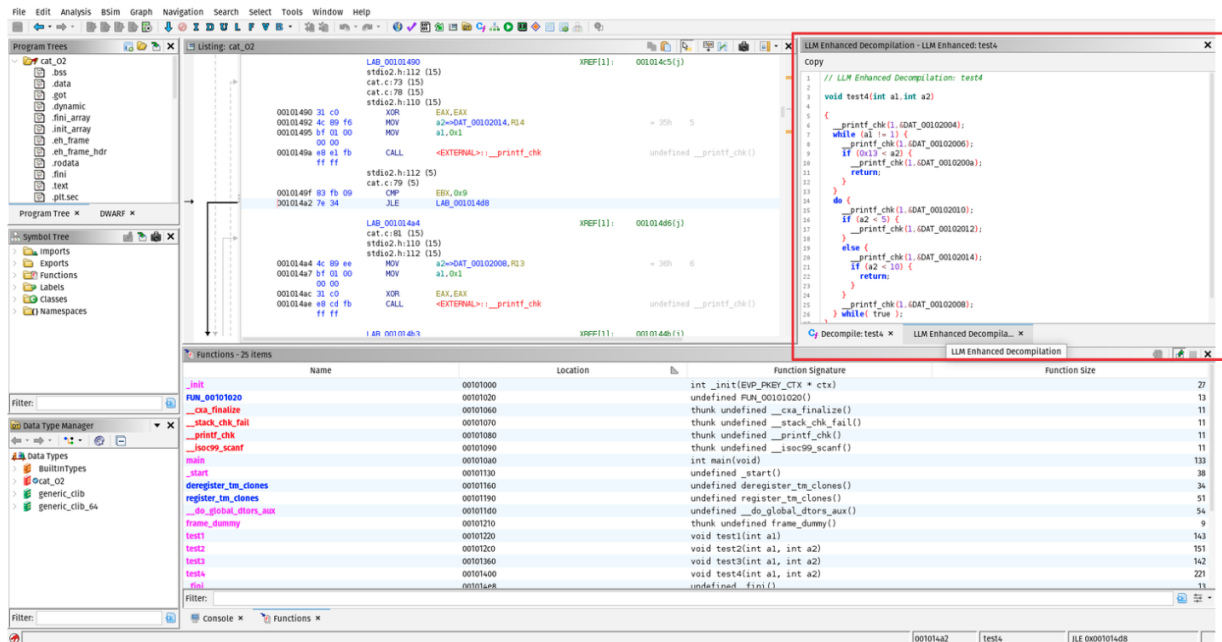


Figure B.2: A possible user interface for the Ghidra extension, showing the selected decompilation variant.

This would provide a modular framework for experimenting with decompilation strategies while keeping the final user interface simple.

B.2 Implementing in GitHub Actions

Another future direction is to implement the LLM-based evaluation pipeline in a GitHub Action that can be triggered on pull requests. This would allow the decompilation quality of new PRs to be automatically evaluated and compared against the base version, providing immediate feedback to developers and maintainers about the readability and human-likeness of their changes. This will need to be carefully designed to manage the computational cost and ensure that the evaluation process is efficient and does not introduce significant delays in the PR workflow.

It must be not a blocking check, but rather a complementary tool that provides insights and suggestions to developers without hindering the development process.

Appendix C

Prompts

These are the prompts that were used to evaluate the decompiled code snippets with the LLM-as-a-Judge approach.

C.1 LLM-as-a-Judge prompt templates

Listing C.1: Prompt templates used in the LLM-as-a-Judge evaluation.

```
def get_quality_prompt_s(diff_text, source_code):
    return (
        "You are a Senior Compiler Engineer evaluating decompilation
        fidelity.\n"
        "Your goal: Determine if the changes in the Diff move the code closer
        to the Source Code structure** or further away.\n\n"

        "### HOW TO READ THE COMPARISON\n"
        "- Lines starting with '%' belong ONLY to Candidate A.\n"
        "- Lines starting with '&' belong ONLY to Candidate B.\n"
        "- Lines starting with a space belong to BOTH (shared context).\n\n"

        "### EVALUATION CRITERIA\n"
        "1. Structural Isomorphism: Which version (A or B) matches the
        control flow structures of the Source Code?\n"
        "  - Source 'switch' -> Winner must have 'switch'.\n"
        "  - Source 'for' -> Winner must have 'for'.\n\n"

        "### CHAIN OF VERIFICATION\n"
```

```

"Before making your final decision, perform a step-by-step verification
  to check for potential errors:\n"
"1. **Isolate Modifications**: Explicitly identify what AST structural
  nodes were actually added or removed in the diff.\n"
"2. **Fact-Check**: Verify that your assessment is based only on
  these specific modifications and not on unverified assumptions.\n"
"3. **Error Check**: Double-check that your choice strictly follows the
  'Human vs Machine' hierarchy without logical contradictions.\n\n"

"### INPUT DATA\n"
"--- GROUND TRUTH (SOURCE CODE) ---\n"
f"{'c\n{source_code}\n'\n\n"
"--- STRUCTURAL DIFF (A vs B) ---\n"
f"{'diff\n{diff_text}\n'\n\n"

"### OUTPUT FORMAT\n"
"Output ONLY valid JSON:\n"
"{\n"
'  "motivation": "Briefly describe which version aligns with Source
  structure.",\n'
'  "winner": "A" | "B"\n'
"}"
)

```

```

def get_ast_prompt_s(diff_text, source_ast):
    return (
        "You are a Decompilation Architect comparing AST changes against a
        Ground Truth.\n"
        "Your goal: Determine which version in the Diff (A or B) mirrors the
        **Source AST topology**.\n\n"

        "### HOW TO READ THE COMPARISON\n"
        "- Lines starting with '%' belong ONLY to **Candidate A**.\n"
        "- Lines starting with '&' belong ONLY to **Candidate B**.\n"
        "- Lines starting with a space belong to BOTH (shared context).\n\n"

        "### EVALUATION CRITERIA\n"
        "1. **Topology Match**: Which version (A or B) preserves the Source AST
        node types and structure?\n"
        "2. **Complexity**: Does the one version of code match better the
        nesting depth of the Source?\n\n"
    )

```

```

#### CHAIN OF VERIFICATION\n"
"Before making your final decision, perform a step-by-step verification
  to check for potential errors:\n"
"1. Isolate Modifications: Explicitly identify what AST structural
  nodes were actually added or removed in the diff.\n"
"2. Fact-Check: Verify that your assessment is based only on
  these specific modifications and not on unverified assumptions.\n"
"3. Error Check: Double-check that your choice strictly follows the
  'Human vs Machine' hierarchy without logical contradictions.\n\n"

#### INPUT DATA\n"
f"--- GROUND TRUTH (SOURCE AST) ---\n{source_ast}\n\n"
f"--- AST DIFF (A vs B) ---\n"
f"``diff\n{diff_text}\n``\n"

#### OUTPUT FORMAT\n"
"Output ONLY valid JSON:\n"
"{\n"
'  "motivation": "Briefly describe which version aligns with Source
  structure.",\n'
'  "winner": "A" | "B"\n'
"}"
)

```

```

def get_quality_prompt(diff_text):
    return (
        "You are a Lead C Code Auditor performing a blind review of a patch
        (Diff).\n"
        "Your goal is to decide which version in the diff (the version A '-'
        lines or the version B '+' lines) represents better Human
        Engineering Practices.\n"
        **CRITICAL RULE**: Focus ONLY on the logic flow and readability
        changes shown in the diff.\n\n"

        ### HOW TO READ THE COMPARISON\n"
        "- Lines starting with '%' belong ONLY to Candidate A.\n"
        "- Lines starting with '&' belong ONLY to Candidate B.\n"
        "- Lines starting with a space belong to BOTH (shared context).\n\n"

        ### THE ONLY HIERARCHY: HUMAN ENGINEERING & READABILITY\n"
        "We prioritize code that looks like it was authored by an experienced
        human programmer over raw, algorithmic output.\n\n"
    )

```

```

"1. Human-like Traits (Winner):\n"
"  - Natural, idiomatic choices (e.g., using a clean 'for' loop instead
    of a clunky 'while' with manual increments).\n"
"  - Simplified, elegant conditionals and optimized operations.\n"
"  - Clean visual structure that prioritizes readability and clear
    developer intent.\n\n"
"2. Machine-like Artifacts (Loser):\n"
"  - Over-complicated, mechanical translations of simple logic.\n"
"  - Redundant operations, unnecessary casts, or awkward block
    structures.\n"
"  - Literal, step-by-step logic that a human would naturally abstract
    away.\n\n"

```

```

"### DIFFERENTIAL ANALYSIS\n"

```

```

Judge: Does the difference make the code more like a Human
(Abtract) or more like a Machine (Concrete)?\n\n"

```

```

"### FORCED DECISION RULES\n"

```

```

- NO NEUTRALITY: Pick a winner.\n"

```

```

- Tie-Breaker: If logic is identical, choose the representation
with less artificial nesting/depth.\n\n"

```

```

"### CHAIN OF VERIFICATION\n"

```

```

Before making your final decision, perform a step-by-step verification
to check for potential errors:\n"

```

```

1. Isolate Modifications: Explicitly identify what AST structural
nodes were actually added or removed in the diff.\n"

```

```

2. Fact-Check: Verify that your assessment is based only on
these specific modifications and not on unverified assumptions.\n"

```

```

3. Error Check: Double-check that your choice strictly follows the
'Human vs Machine' hierarchy without logical contradictions.\n\n"

```

```

"### INPUT DATA (UNIFIED DIFF)\n"

```

```

f"diff\n{diff_text}\n""\n\n"

```

```

"### OUTPUT FORMAT\n"

```

```

Output ONLY valid JSON:\n"

```

```

{\n"

```

```

' "motivation": "Why the winner is more human-like.",\n'

```

```

' "winner": "A" | "B"\n'

```

```

}"

```

```

)

```

```

def get_ast_prompt(diff_text):
    return (
        "You are a Static Analysis Expert evaluating changes in Control Flow
        Skeletons (AST).\n"
        "Your goal: Decide if the structural changes in the Diff improve the
        **Idiomatic Design**.\n\n"

        "### HOW TO READ THE COMPARISON\n"
        "- Lines starting with '%' belong ONLY to **Candidate A**.\n"
        "- Lines starting with '&' belong ONLY to **Candidate B**.\n"
        "- Lines starting with a space belong to BOTH (shared context).\n\n"

        "### THE ONLY HIERARCHY: HUMAN ENGINEERING & READABILITY\n"
        "We prioritize code that looks like it was authored by an experienced
        human programmer over raw, algorithmic output.\n\n"
        "1. **Human-like Traits (Winner)**:\n"
        "  - Natural, idiomatic choices (e.g., using a clean 'for' loop instead
        of a clunky 'while' with manual increments).\n"
        "  - Simplified, elegant conditionals and optimized operations.\n"
        "  - Clean visual structure that prioritizes readability and clear
        developer intent.\n\n"
        "2. **Machine-like Artifacts (Loser)**:\n"
        "  - Over-complicated, mechanical translations of simple logic.\n"
        "  - Redundant operations, unnecessary casts, or awkward block
        structures.\n"
        "  - Literal, step-by-step logic that a human would naturally abstract
        away.\n\n"

        "### DIFFERENTIAL ANALYSIS\n"
        "**Judge**: Does the difference make the code more like a Human
        (Abstract) or more like a Machine (Concrete)?\n\n"

        "### FORCED DECISION RULES\n"
        "- **NO NEUTRALITY**: Pick a winner.\n"
        "- **Tie-Breaker**: If logic is identical, choose the representation
        with less artificial nesting/depth.\n\n"

        "### CHAIN OF VERIFICATION\n"
        "Before making your final decision, perform a step-by-step verification
        to check for potential errors:\n"
        "1. **Isolate Modifications**: Explicitly identify what AST structural
        nodes were actually added or removed in the diff.\n"
    )

```

```
"2. **Fact-Check**: Verify that your assessment is based *only* on
  these specific modifications and not on unverified assumptions.\n"
"3. **Error Check**: Double-check that your choice strictly follows the
  'Human vs Machine' hierarchy without logical contradictions.\n\n"
```

```
#### INPUT DATA (AST DIFF)\n"
f"''diff\n{diff_text}\n'''\n\n"
```

```
#### OUTPUT FORMAT\n"
"Output ONLY valid JSON:\n"
"{\n"
'  "motivation": "Why the winner is more human-like.",\n'
'  "winner": "A" | "B"\n'
"}"
```

```
)
```

```
#####Dogbolt version#####
```

```
def Cget_quality_prompt_s(code_a, code_b, source_code):
    return (
        "You are a Senior Compiler Engineer evaluating decompilation
        fidelity.\n"
        "Your goal: Determine which candidate version (A or B) is **closer to
        the Source Code structure**.\n\n"

        "### EVALUATION CRITERIA\n"
        "1. **Structural Isomorphism**: Which version matches the control flow
        structures of the Source Code?\n"
        "  - Source 'switch' -> Winner should have 'switch'.\n"
        "  - Source 'for' -> Winner should have 'for'.\n"
        "2. **Logical Fidelity**: Which version correctly represents the logic
        of the source without introducing assembly-level artifacts (like
        unnecessary gotos)?\n\n"

        "### CHAIN OF VERIFICATION\n"
        "Before making your final decision, perform a step-by-step verification
        to check for potential errors:\n"
        "1. **Isolate Modifications**: Explicitly identify what AST structural
        nodes were actually added or removed in the diff.\n"
        "2. **Fact-Check**: Verify that your assessment is based *only* on
        these specific modifications and not on unverified assumptions.\n"
```

```
"3. **Error Check**: Double-check that your choice strictly follows the
  'Human vs Machine' hierarchy without logical contradictions.\n\n"
```

```
### INPUT DATA\n"
```

```
--- GROUND TRUTH (SOURCE CODE) ---\n"
```

```
f"'\n{source_code}\n'\n\n"
```

```
--- CANDIDATE A ---\n"
```

```
f"'\n{code_a}\n'\n\n"
```

```
--- CANDIDATE B ---\n"
```

```
f"'\n{code_b}\n'\n\n"
```

```
### OUTPUT FORMAT\n"
```

```
Output ONLY valid JSON:\n"
```

```
{\n"
```

```
' "motivation": "Briefly describe which candidate aligns better with
  the Source structure.",\n'
```

```
' "winner": "A" | "B"\n'
```

```
}"
```

```
)
```

```
def Cget_ast_prompt_s(ast_a, ast_b, source_ast):
```

```
    return (
```

```
        "You are a Decompilation Architect comparing AST topologies against a
          Ground Truth.\n"
```

```
        "Your goal: Determine which candidate (A or B) mirrors the **Source AST
          topology**.\n\n"
```

```
        ### EVALUATION CRITERIA\n"
```

```
        "1. **Topology Match**: Does one Candidate restore a node type (e.g.
```

```
          'SwitchStatement') present in the Source but missing in the other?\n"
```

```
        "2. **Complexity**: Does the candidate match the nesting depth and
          statement hierarchy of the Source?\n\n"
```

```
        ### CHAIN OF VERIFICATION\n"
```

```
        "Before making your final decision, perform a step-by-step verification
          to check for potential errors:\n"
```

```
        "1. **Isolate Modifications**: Explicitly identify what AST structural
          nodes were actually added or removed in the diff.\n"
```

```
        "2. **Fact-Check**: Verify that your assessment is based only on
          these specific modifications and not on unverified assumptions.\n"
```

```
        "3. **Error Check**: Double-check that your choice strictly follows the
          'Human vs Machine' hierarchy without logical contradictions.\n\n"
```

```

#### INPUT DATA\n"
---- GROUND TRUTH (SOURCE AST) ----\n"
f"{source_ast}\n\n"
---- CANDIDATE A AST ----\n"
f"{ast_a}\n\n"
---- CANDIDATE B AST ----\n"
f"{ast_b}\n\n"

#### OUTPUT FORMAT\n"
Output ONLY valid JSON:\n"
"{\n"
  ' "motivation": "Identify specific node types or nesting levels that
    match the Ground Truth.",\n'
  ' "winner": "A" | "B"\n'
}"
)

```

```

def Cget_quality_prompt(code_a, code_b):
    return (
        "You are a Lead C Code Auditor performing a blind review of two
        decompilation candidates.\n"
        "Your goal: Decide which version represents better **Human Engineering
        Practices**.\n\n"

        "### THE HIERARCHY OF IDIOMATIC CONTROL FLOW\n"
        "Prioritize structures that map to high-level human thinking over raw
        machine output:\n"
        "1. **Semantic Structure (Preferred)**: 'for' loops, 'do-while',
        'switch' statements, and clean scoping.\n"
        "2. **Graph Artifacts (Avoid)**: 'goto' spaghetti, deep 'if-else'
        cascades where 'switch' applies, or artificial '{ { ... } }' wrapper
        blocks.\n\n"

        "### FORCED DECISION RULES\n"
        "- **NO NEUTRALITY**: Pick a winner.\n"
        "- **Tie-Breaker**: If logic is identical, choose the version with less
        artificial nesting depth.\n\n"

        "### CHAIN OF VERIFICATION\n"
        "Before making your final decision, perform a step-by-step verification
        to check for potential errors:\n"
    )

```

```

1. **Isolate Modifications**: Explicitly identify what AST structural
   nodes were actually added or removed in the diff.\n"
2. **Fact-Check**: Verify that your assessment is based only on
   these specific modifications and not on unverified assumptions.\n"
3. **Error Check**: Double-check that your choice strictly follows the
   'Human vs Machine' hierarchy without logical contradictions.\n\n"

#### INPUT DATA\n"
--- CANDIDATE A ---\n"
f"'"'c\n{code_a}\n'"'\n\n"
--- CANDIDATE B ---\n"
f"'"'c\n{code_b}\n'"'\n\n"

#### OUTPUT FORMAT\n"
Output ONLY valid JSON:\n"
{\n"
' "motivation": "Describe why the winner is more
  human-readable/idiomatic.",\n'
' "winner": "A" | "B"\n'
"}"
)

```

```

def Cget_ast_prompt(ast_a, ast_b):
    return (
        "You are a Static Analysis Expert evaluating Control Flow Skeletons
          (AST).\n"
        "Your goal: Decide which structural representation (A or B) is more
          **Idiomatic**.\n\n"

        "#### CRITERIA\n"
        "We prioritize high-level human abstractions over raw assembly-derived
          graphs:\n"
        "- **Winner**: Natural loops ('for'/'while'), 'switch' cases, and
          logical nesting.\n"
        "- **Loser**: Conditional jumps to labels ('goto'), excessive 'if-else'
          cascades, and redundant wrapper blocks.\n\n"

        "#### CHAIN OF VERIFICATION\n"
        "Before making your final decision, perform a step-by-step verification
          to check for potential errors:\n"
        "1. **Isolate Modifications**: Explicitly identify what AST structural
          nodes were actually added or removed in the diff.\n"
    )

```

2. ****Fact-Check****: Verify that your assessment is based *only* on these specific modifications and not on unverified assumptions.\n"
3. ****Error Check****: Double-check that your choice strictly follows the 'Human vs Machine' hierarchy without logical contradictions.\n\n"

```
### INPUT DATA\n
```

```
--- CANDIDATE A AST ---\n
```

```
f"{ast_a}\n\n"
```

```
--- CANDIDATE B AST ---\n
```

```
f"{ast_b}\n\n"
```

```
### OUTPUT FORMAT\n
```

```
Output ONLY valid JSON:\n
```

```
{\n
```

```
  "motivation": "Explain the structural advantages of the winner.",\n
```

```
  "winner": "A" | "B"\n
```

```
}"
```

```
)
```
