Trim My View: An LLM-Based Code Query System for Module Retrieval in Robotic Firmware

Sima Arasteh*, Pegah Jandaghi[†], Nicolaas Weideman[†], Dennis Perepech*,

Mukund Raghothaman*, Christophe Hauser[‡] and Luis Garcia[§]

*University of Southern California, {arasteh,perepech,raghotha}@usc.edu † USC Information Sciences Institute, jandaghi@usc.edu, nhweideman@gmail.com [‡]Dartmouth College, christophe.hauser@dartmouth.edu [§]University of Utah, la.garcia@utah.edu

Abstract—The software compilation process has a tendency to obscure the original design of the system and makes it difficult both to identify individual components and discern their purpose simply by examining the resulting binary code. Although decompilation techniques attempt to recover higherlevel source code from the machine code in question, they are not fully able to restore the semantics of the original functions. Furthermore, binaries are often stripped of metadata, and this makes it challenging to reverse engineer complex binary software.

In this paper we show how a combination of binary decomposition techniques, decompilation passes, and LLM-powered function summarization can be used to build an economical engine to identify modules in stripped binaries and associate them with high-level natural language descriptions. We instantiated this technique with three underlying open-source LLMs—CodeQwen, DeepSeek-Coder and CodeStral—and measured its effectiveness in identifying modules in robotics firmware. This experimental evaluation involved 467 modules from four devices from the ArduPilot software suite, and showed that CodeStral, the bestperforming backend LLM, achieves an average F1-score of 0.68 with an online running time of just a handful of seconds.

I. INTRODUCTION

Knowledge of the high-level organization of software—its functions, files and modules—and how its functionality is divided among these units is crucial to productive programming, reverse engineering, and other software engineering practices. While at least some of this information is evident while examining the *source code*, it is much more challenging to recover when working with *binaries*.

Crucially, the compilation process tends to obscure the organizational units of software, including information about its modules, comments, data types, and identifier names (when working with stripped binaries). Although this is sometimes desirable (for example, to protect intellectual property), it

Workshop on Binary Analysis Research (BAR) 2025 28 February 2025, San Diego, CA, USA ISBN 979-8-9919276-4-2 https://dx.doi.org/10.14722/bar.2025.23021 www.ndss-symposium.org severely hinders the understandability of the final binary. Furthermore, despite impressive advances in decompilation techniques [1], they are unable to completely recover semantic details of the original source code [2].

Recent breakthroughs in the development of large language models (LLMs) have prompted researchers to investigate their application in various problems related to the semantic understanding of (binary) code: i.e., binary code summarization [3], [4], source-to-binary matching [5], and function identification [6]. Several studies have also evaluated their effectiveness for function summarization tasks [3], [4], [7], [8], [9].

Despite their impressive accuracy, we argue that function summarization is insufficient for a number of problems related to understanding the high-level structure of programs. For example, engineers may wish to identify portions of the codebase responsible for implementing communication protocols such as the serial peripheral interface (SPI), or for managing the universal asynchronous receiver/transmitter (UART). These functionalities may be deeply embedded in the code and/or divided among multiple functions. This limits the effectiveness of function summarization tools for many reverse engineering activities, which instead require identifying and navigating between larger units of the programs, such as its modules.

Another challenge that arises while using LLMs is economic in nature: Well-known commercial LLMs such as GPT-X [10], Copilot,¹ and Gemini [11] charge money when used on a large scale. As an example, the authors of BinSum [3] report incurring an expense of \$11,000 when evaluating GPT-4. One could use open-source LLMs instead [12], [13]. However, GPUs with large amounts of memory are prohibitively expensive, and this would restrict us to use mid-size language models. This raises an additional worry that these smaller locally-run models would be less accurate or effective at discharging their tasks.

In this paper, we describe a system—called ChatCPS—in which the user requests parts of the code related to some

¹https://copilot.github.com/

high-level functionality of interest. The code takes the form of stripped binaries. In response, the system automatically identifies modules/groups of related functions related to the desired functionality and presents them to the user for examination. As indicated by the name $(\cdots - CPS)$, our primary focus is on reverse engineering firmware for robot controllers.

Our system relies on a combination of binary decomposition techniques, function decompilation, and multiple levels of LLM application. In the first LLM pass, we obtain summaries of individual functions with the binary. In parallel, we apply BCD [14], an algorithm for decomposing binary code into components/modules. (Because the source code of the original BCD implementation was unavailable, and because we were working with ARM rather than x86-64 binaries, we had to reimplement this algorithm). We then group together summaries for related functions, and use a second pass with the LLM to obtain a categorization of the module as a whole. We summarize this workflow in Figure 1. We focus on four categories in which CPS experts frequently express interest: data transfer, navigation, controllers, and safety checks. We describe these categories in more detail in Section III-C.

In our experience, this two-pass approach—first summarize individual decompiled functions, and then summarize groups of related function summaries—improves classification accuracy when compared to a hypothetical single-pass procedure (which would presumably attempt to directly summarize entire modules that were previously identified by BCD). This improvement in accuracy allows us to use less expensive (i.e., smaller and open-source) LLMs, rather than expensive commercial offerings.

We instantiated our system with three mid-sized LLMs drawn from the BigCodeBench leaderboard [15], [16]: CodeQwen [17], DeepSeek-Coder [13] and CodeStral [18]. Our experimental evaluation with 467 modules from the ArduPilot software suite [19] indicates that CodeStral provides the most accurate high-level categorization, with an average F1-score of 0.68. As a secondary measurement, we also separately calculated the accuracy of the first LLM pass, i.e., we measured the accuracy of the generated function summaries. This measurement required us to build a source code parser that normalized function bodies by removing comments, variable and function names. We describe our experimental setup and our results in Sections IV and V respectively.

Our hope is that ChatCPS would streamline the reverse engineering process and enable users to efficiently identify and distinguish parts of the codebase related to different functionalities. We also hope that module categorization would enhance the user's understanding of the binary's underlying structure. Finally, we expect that this top-level categorization would form a basis for more fine-grained categorization of the functions within each module. This would help to distinguish different kinds of algorithms, including Kalman filters [20] from PID controllers [21], both of which may be found in modules related to navigation and control. At an even lower level, after groups of functions related to individual algorithms have been identified, users could run more specific queries. To our knowledge, this is the first such code query-based system that utilizes LLMs to retrieve specific functionalities within stripped binaries. We summarize our contributions:

- 1) We have developed a query system that defines and retrieves module categories using LLM-generated summaries.
- 2) We implement a new variant of the BCD algorithm adjusted for our needs and evaluate its effectiveness in decomposing robotic firmware into modular components within cyber-physical systems.
- We evaluate the effectiveness of three open-source LLMs (DeepSeek-Coder, CodeQwen and CodeStral), in summarizing decompiled functions within modules for stripped firmware.
- 4) We establish the ground truth categories for 467 modules across four devices in the ArduPilot [19] dataset, including QuadCopter, HeliCopter, Rover, and Submarine.

The ChatCPS implementation may be downloaded from https: //github.com/SimaArasteh/chatcps.

II. RELATED WORK

Code semantics analysis has been widely studied by researchers. Traditionally, researchers have relied on signaturebased methods to detect specific, well-known functions, such as mathematical operations. For example, advanced binary analysis tools like IDA Pro [22], Ghidra [23], and angr [24] are designed to recognize the usage of common libraries within binaries, such as standard math functions. More sophisticated techniques employ semantic pattern matching to identify known functions. Studies such as those by Kim et al. [25] and Xu et al. [26] have focused on extracting and recognizing semantic patterns associated with specific mathematical functions, including cryptographic operations.

However, these traditional approaches are generally constrained to a predefined set of patterns and are incapable of detecting unknown functions.

To address these limitations, recent research has begun leveraging LLMs. LLMs have demonstrated remarkable effectiveness across a variety of tasks in the area of software engineering, including function identification and code summarization. For example, BinaryAI [5] leverages LLMs to correlate binary code snippets with their original source code, aiding in tasks like detecting reused third-party libraries. Similarly, FoC [6] focuses on identifying known cryptographic functions within stripped binaries via LLMs.

Moreover, LLMs can extract code semantics and software functionality. Before the widespread adoption of LLMs, researchers used smaller transformer models such as CodeT5 [27], BinT5 [7] and CodeBERT [28] for code summarization tasks to enhance comprehension of both binary and source code [29], [8], [30], [31], [7]. Smaller transformer models like CodeT5, BinT5, and CodeBERT may offer less detailed code summaries due to their limited parameters, which can restrict their ability to handle diverse coding languages and scenarios effectively. Larger LLMs, with their extensive training and advanced architectures, provide more accurate and



B. Decompilation Process

Fig. 1. The Workflow of the ChatCPS system. A) ChatCPS decomposes the binary firmware into its modules. B) Shows the decompilation process. F_n indicates binary functions, while F_n^* indicates the corresponding decompiled functions. C) ChatCPS summarizes each function in a module using three open-source LLMs (CodeQwen, CodeStral, DeepSeek-Coder) and D) categorizes modules based on function descriptions.

comprehensive summaries. As a result, researchers have begun to explore the capabilities of larger LLMs in understanding the code semantics. For example, [3], [9], [32], [33], [34] and [30] conducts evaluations of different LLMs such as ChatGPT, Llama2 [35], and Code-Llama [36] for the task of code summarization.

While code semantic summarization is useful for understanding individual components of a software system, it does not offer a comprehensive view of the software's overall architecture and interactions. In reverse engineering, especially with large and complex software, analyzing every function can be overwhelming and unnecessary. Instead, focusing on specific functionalities that are crucial can be more effective. Developing a system that selectively highlights these key functionalities would greatly enhance the efficiency of the analysis.

III. THE DESIGN OF CHATCPS

Reverse engineering complex software without access to the source code is a challenging and labor-intensive task. To facilitate this process, we introduce ChatCPS, a code-based query system designed specifically for cyber-physical systems. This system aims to assist reverse engineers in focusing on critical software functionalities. ChatCPS consists of a threelayer pipeline. First, we decompose binary firmware into its software components called modules. Then, we summarize the decompiled version of functions in the modules using LLMs. Finally, we categorize modules based on the function



Fig. 2. Illustration of how the BCD algorithm uses information about the call graph, data references, and code locality to decompose the binary into groups of functions. F_n indicates functions in the binary. The image shows a final weighted graph generated by the BCD algorithm. Green edges show the function call relationship between functions. Red edges show the code locality, and the purple edges indicate the data references.

descriptions by leveraging LLMs. Figure 1 shows the overall workflow of ChatCPS.

A. Example of a Decompiled code

C. Examples of summaries of different LLMs



Fig. 3. The process of summarizing decompiled functions using LLMs. Section A- presents a decompiled example of the function *_mav_finalize_message_chan_send* from the ArduPilot project, extracted from the stripped binary of a copter device compiled for the fmuv3 board. Section B- depicts the prompt used to summarize decompiled functions. Section C- displays the function descriptions of *_mav_finalize_message_chan_send*, as returned by three distinct LLMs.

A. Binary Modularization

We start by decomposing binary firmware into its constituent modules using a modified version of the BCD method as described by Karande et al. [14]. While the original source code was unavailable and required us to re-implement the technique, the modifications were made specifically to address the incompatibility of the original BCD method with ARM binaries. In this section, we outline the original workflow of the BCD method and detail the changes we introduced to adapt it for ARM binaries.

To decompose a binary into its modules, the BCD method constructs a directed graph from the binary program, where nodes represent functions and edges capture three types of relationships between functions: code locality, data references, and function calls. It then applies the Newman algorithm [37] to cluster the graph into smaller components, referred to as modules.

BCD relies on three facts to decompose a binary program.

1) Code Locality: Functions within a module typically share similar functionalities and are physically located close to each other. This proximity is maintained by the compiler in the binary program, where functions are organized sequentially from lower to higher memory addresses. BCD utilizes this characteristic to construct a directed graph known as SG (for *sequence graph*), which illustrates these relationships. In this graph, an edge extends from a function at a lower memory

address to a function at a higher address, effectively mapping the sequential organization of functions within the binary.

2) Data References: Functions within a module often access and share the same data. BCD examines this access for global and static variables, as well as constant string literals located in the *.data*, *.bss*, and *.rodata* sections. In this process, BCD constructs a graph known as the Data Reference Graph (*DRG*). In this graph, an edge is drawn between two functions if they access at least one common variable. Consequently, this shared access results in a mutual edge being established between these functions within the graph.

3) Function Calls: The principle underlying BCD is the observation that functions within a module tend to call each other more frequently than they do functions outside the module. To leverage this insight, BCD constructs a graph called the Call Graph CG, wherein it assigns a weight to each edge based on the number of calls between functions. Additionally, BCD synthesizes a comprehensive weighted graph by integrating the CG with two other graphs: (SG) and (DRG).

In this composite graph, the weight of each edge is determined by a linear combination of its weights from SG, DRG, and CG. During the final stage of the process, BCD employs an algorithm known as the Newman algorithm to cluster the graph into smaller graphs refer as a module. Based on this algorithm, each node is considered an independent cluster. Subsequently, nodes are merged based on the optimal connectivity indicated



Fig. 4. Module Retrieval Process. We provide function summaries and module categories as an input prompt to an LLM. Then, we instruct the LLM to decide based on the function descriptions and ranked the categories.

by their edges, ultimately forming the final clusters. For a detailed explanation of this methodology, we refer the reader to the original paper [14]. Figure 2 shows the workflow of the BCD algorithm.

B. Function Summarization in Binary Modules

In this step, we leverage open-source LLMs to summarize decompiled functions in each module. Within our codebase, some functions are relatively small, primarily serving as interfaces that call numerous other functions. These smaller functions typically reveal less about the program's semantic structure. To focus on more substantively informative functions, we exclude functions with a line count below a specified threshold. This approach ensures that our analysis concentrates on functions that provide significant insights into the software's architecture. Figure 3 shows the prompt and the process that we used for summarizing functions. In this Figure, you can observe summaries generated by three opensource LLMs(Codeqwen, Codestral, and DeepSeek-Coder) for function *mav_finalize_message_chan_send* within a module of a helicopter device.

C. Module Retrieval

In this step, our goal is to identify modules with specific functionalities within cyber-physical systems. Given their wide range of applications and unique capabilities, these systems serve as our primary focus, though our methodology is applicable to various types of applications. We define five distinct module categories based on common functionalities across cyber-physical systems (e.g., robots): 1) Data transfer (crucial for real-time communication between vehicle sensors and systems), 2) navigation (essential for route optimization and precise positioning), 3) control (manages operational capabilities such as steering and braking), and 4) safety checks (monitors systems and environment to prevent accidents). We also consider the modules that do not align with these defined categories are not mutually exclusive, as some modules

belong to multiple categories, adding complexity to the module retrieval process.

To identify the module category using LLMs, we construct a prompt for each module. The prompt contains three sections to provide the LLM with sufficient information to accurately infer the category of the module:

Category Definition/Knowledge: This section includes the precise definition of module categories, which are curated by an expert. We include this section to ensure that LLM uses the intended semantics of categories rather than relying on their semantic prior.

Functions Summaries: In this section, we include the summaries of the functions contained in the module, which are generated in the previous stage of the pipeline.

Prompt Instruction: This section provides the LLM with specific instructions, including the context of the code and a clear description of the task it is expected to perform. Specifically, these instructions guide the LLM to rank the categories based on the likelihood of the module belonging to each. This approach also allows the LLM to retrieve multiple categories when a module aligns with more than one. Figure 4 shows the prompt and the process we use for this phase.

IV. EXPERIMENTAL SETUP

In this section, we describe the benchmarks and experimental setup, including detailed information on the configurations and methods applied in our experiments.

A. Benchmark

To evaluate our approach, we use the open-source and well-known autopilot dataset called ArduPilot. This dataset is a comprehensive collection of data and configurations that supports the development and testing of autonomous vehicle software across various platforms. Developers created ArduPilot to boost research and development in unmanned vehicle systems, offering a solid base for simulation, testing, and the practical application of autonomous control algorithms. The

TABLE I THE EFFECTIVENESS OF THE BCD ALGORITHM IN DECOMPOSING BINARY FIRMWARE INTO ITS MODULES IN ARDUPILOT.

Device	# of Modules	# of Functions	$\mathbf{P}_{\mathbf{w}}$	$\mathbf{R}_{\mathbf{w}}$	$F1_w$
QuadCopter	121	8556	0.77	0.69	0.72
HeliCopter	121	7583	0.75	0.73	0.74
Rover	112	7258	0.72	0.72	0.72
Submarine	113	6994	0.76	0.68	0.71

dataset includes a wide range of vehicles, such as Quad-Copters, HeliCopters, Planes, Rovers, and Submarines. As an open-source software, ArduPilot is widely used to innovate and improve unmanned vehicle technology.

We compile the ArduPilot dataset for four QuadCopters, devices—including HeliCopters, Rovers. and Submarines- specifically for the FMUV3 board. The FMUV3 is a critical flight management unit version that supports various processors and sensors. This versatility enhances its utility across different unmanned vehicle platforms within the ArduPilot framework. Though Ardupilot provides many different boards, we only focus on building the ArduPilot dataset for just one board due to the inherent challenges of establishing ground truth for all boards. Also, Different boards share common functionalities and modules. We plan to extend our analysis to additional boards in future work.

B. Binary Modularization Setup

We decompose binaries into modules using the BCD algorithm [14]. As the original source code was unavailable, we adapted the approach outlined in the paper with some modifications. Originally, BCD was based on IDA-Pro [22], which is not free. Therefore, we opted to implement BCD using angr [24]. Also, the original BCD algorithm was designed for x86 binaries. However, as ArduPilot binaries utilize the ARM architecture, we modify the BCD implementation to align with our dataset. This adaptation specifically addresses the differences in CPU architecture, which affect the identification and processing of shared data references between functions. Due to the distinct ways ARM and x86 architectures handle these references, our implementation requires customization to effectively manage these architectural variations.

C. Decompilation Process Setup

We decompile binary functions extracted from stripped binaries in ArduPilot using Ghidra [23], a free and opensource reverse engineering tool developed by the National Security Agency (NSA). Ghidra offers a wide range of features for analyzing compiled code across multiple platforms and architectures.

For each module, we select functions that contain at least 15 lines of code to ensure they are large enough to provide meaningful semantic insights. Some functions in a module are limited to calling other internal or external functions, providing limited semantic insight by themselves. As a result, we filter out such functions in our evaluation. Furthermore, we focus solely on the body of the function itself without delving into the bodies of callee functions. This approach is justified for several reasons. Firstly, larger functions typically contain enough content to disclose their semantic purpose. Secondly, a single function might invoke numerous other functions with entirely different semantics. For example, a controller function might call a data-transfer function to send a message. Including the bodies of these callee functions could skew the results of our categorization and potentially mislead our analysis.

D. LLM-based Unit Setup

To retrieve modules with specific functionalities, we evaluate three open-source LLMs: Deepseek-Coder[13], CodeQwen[17], and CodeStral [18]. DeepSeek-Coder, CodeQwen, and CodeStral are specialized versions of well-known language models designed for code analysis and development tasks. DeepSeek excels at interpreting both language and code, significantly aiding in the review of complex code and the creation of documentation. CodeQwen, built to improve coding help, provides features such as real-time debugging and code optimization. CodeStral builds on the capabilities of the Mistral model family and specializes in analyzing code in multiple languages. It supports a variety of programming environments and helps ensure that code works well across different languages.

Based on their good performance on the CodeBench leaderboard [15], we selected these LLMs for our study and plan to analyze additional LLMs in future work. To ensure a fair comparison, we chose models with similar parameter counts, specifically Qwen2.5-Coder-32B-Instruct, DeepSeek-Coder-33B-Instruct, and Codestral-22B-v0.1. The context window for these models is 128000, 16000, and 32000 tokens, respectively. Our experiments were conducted on a CentOSbased server equipped with two GPUs, using parallel processing across both GPUs and featuring 48 GB of RAM. We take advantage of the versions of these open-source LLMs available through HuggingFace². For inference, we use the widely adopted Im-eval-harness library [38]. This evaluation framework provides a streamlined and flexible interface for benchmarking LLM performance across diverse tasks, enabling systematic and reproducible evaluations.

Our implementation of ChatCPS including the modified version of BCD contains approximately 2000 lines of Python.

V. EXPERIMENTAL RESULTS

Our assessment of ChatCPS addresses the following research questions:

- **RQ1.** How effectively can BCD decompose binaries into modules within the autopilot dataset?
- **RQ2.** How effectively do CodeQwen, Codestral, and DeepSeek-Coder create function summaries that demonstrate their proficiency in understanding code?

²https://huggingface.co/

TABLE II LLM FUNCTION SUMMARIZATION EVALUATION. THE NUMBERS INDICATE THE AVERAGE COSINE SIMILARITY OF THE ENCODED SUMMARY OF THE NORMAL CODE AND THE DECOMPILED CODE OF THE FUNCTION GENERATED BY THE LLMS IN THE COLUMNS.

Device	Deepseek-Coder	CodeStral	CodeQwen	
QuadCopter	0.76 ± 0.16	0.62 ± 0.20	0.68 ± 0.20	
HeliCopter	0.75 ± 0.18	0.62 ± 0.20	0.68 ± 0.21	
Rover	0.75 ± 0.17	0.62 ± 0.21	0.67 ± 0.21	
Submarine	0.76 ± 0.16	0.62 ± 0.21	0.68 ± 0.21	

- **RQ3.** How proficiently does ChatCPS retrieve module categories?
- **RQ4.** How long does it take for each LLM to complete function summarization?

A. RQ1: Binary modularization Effectiveness

We evaluate how well the BCD algorithm can decompose Ardupilot binaries into its corresponding modules. Table I presents statistics for four devices in ArduPilot, including the number of modules created in the binaries, the number of functions, and the performance of modularization. We compare the predicted modularization with the ground truth. We then define the true positives (TP) as the fraction of functions that were correctly placed in the current module, the false positives (FP) as the fraction of functions that were incorrectly placed in the current module, and the false negatives (FN) as the fraction of functions that were not included in the current module but were expected to be. From here, similar to the original paper, we calculate a weighted version of precision, recall, and F_1 score shown in equations 1, 2 and 3 respectively.

$$P_w = \frac{\sum_{i=1}^{N_c} P_i \cdot n_i}{N_f} \tag{1}$$

$$R_w = \frac{\sum_{i=1}^{N_c} R_i \cdot n_i}{N_f} \tag{2}$$

$$F_{1w} = \frac{\sum_{i=1}^{N_c} F_{1i} \cdot n_i}{N_f}$$
(3)

In these formulas, P_i , R_i , and F_{1i} are the precision, recall, and F_1 score for module C_i , n_i is the number of functions in module C_i , N_f is the total number of functions in all modules, and N_c is the total number of modules. We use weighted metrics due to the significant variation in the number of functions each module contains. This method assigns the appropriate importance to each module based on its size, ensuring that modules with a larger number of functions have a greater impact on the overall metrics. This approach is vital as it prevents smaller modules from disproportionately influencing the results, which could distort the system's perceived effectiveness. By adjusting scores relative to the number of functions, the evaluation more precisely reflects the system's performance across modules of varying sizes and complexities, providing a balanced and detailed assessment of its binary decomposition capabilities.

As suggested by Table I, the modified version of the BCD algorithm achieved acceptable performance in decomposing binaries into modules. It is important to note that modularizing large binary firmware is a challenging task, yet our implementation of the BCD algorithm has attained a satisfactory level of effectiveness. Even though recovering modules from binary code may be imperfect, it does not significantly impact the effectiveness of categorization. For instance, two different modules in ArduPilot, such as AC_circle and AC_WPNav, can both fall under the same category, namely navigation. Furthermore, upon manual inspection, we find that this imperfection in modularization does not substantially cause functions from different modules that belong to distinct categories to be grouped into the same module.

B. RQ2: Function Summarization Effectiveness

Module categories are established based on function summaries, underscoring the importance of their reliability. Therefore, we evaluate the function summarization generated by LLMs separately. To evaluate the effectiveness of function summaries, we compare the similarity between decompiled function summaries and the summaries derived from the normalized version of the original source code. The normalized version of the source code contains the original code, but function names and variables are stripped. Removing function names and variable names before using LLMs to summarize code helps isolate the LLM's ability to understand and summarize code based purely on its structure and logic rather than relying on potentially descriptive names. Function names and variables can give significant hints about what the code does, which may lead LLMs to rely more on these cues rather than analyzing the underlying code patterns. To evaluate an LLM's raw capability to understand code, we minimize external influences like descriptive names. Ideally, LLM creates a similar textual summary for the decompiled code and the normalized code of the same function.

To obtain the normalized version of the source code, we develop a source code parser specifically designed to extract function bodies directly from the source code. Then, we parse the ArduPilot source code [39] to extract function bodies, using the Clang [40] tool to identify the function and variable instances. We replace these identifiers with anonymized names. Additionally, to align the source code with the decompiled functions, we extract mappings between function names and their addresses from binaries that contain debug symbols (DWARF). Using these mapped names and addresses, we identify the corresponding function bodies in the source code. Lastly, we remove all comments from the source code to focus exclusively on evaluating the code understanding capabilities.

To compute the textual summary similarity between the normalized code and decompiled code for each function, we encode the text using an embedding model. The embedding model generates a vector for each of the textual summaries. Then, we compare the similarity of generated embedding vec-

TABLE III

THE EFFECTIVENESS OF RETRIEVING MODULE CATEGORIES BY THREE OPEN-SOURCE LLMS IN THE CHATCPS DESIGN. NUMBERS BETWEEN PARENTHESES INDICATE THE UPPER BOUND OF MODULE CATEGORIZATION. THESE NUMBERS SHOW THE RESULT OF MODULE CATEGORIZATION OBTAINED FROM NORMALIZED SOURCE CODE.

Device	Category	DeepSeek-Coder		CodeStral			Code-Qwen			
		Precision	Recall	F1 Score	Precision	Recall	F1 Score	Precision	Recall	F1 Score
QuadCopter	Data Transfer	0.33 (0.52)	0.14 (0.42)	0.20 (0.46)	0.78 (0.76)	0.73 (0.92)	0.75 (0.83)	0.61 (0.72)	0.55 (0.55)	0.58 (0.62)
	Navigation	0.24 (0.25)	0.68 (0.86)	0.36 (0.39)	0.63 (0.73)	0.70 (0.80)	0.66 (0.76)	0.24 (0.30)	0.82 (0.93)	0.37 (0.46)
	Controller	0.93 (0.93)	0.65 (0.67)	0.77 (0.77)	0.92 (0.94)	0.92 (0.82)	0.92 (0.87)	0.87 (0.88)	0.44 (0.52)	0.59 (0.64)
	Safety Check	0 (0)	0 (0)	0 (0)	0.25 (0.5)	0.16 (0.5)	0.2 (0.5)	0 (0)	0 (0)	0 (0)
HeliCopter	Data Transfer	0.5 (0.66)	0.36 (0.54)	0.42 (0.6)	0.73 (0.77)	0.73 (0.73)	0.73 (0.75)	1.0 (1.0)	0.25 (0.33)	0.4 (0.5)
	Navigation	0.05 (0.13)	0.5 (1)	0.09 (0.23)	0.82 (0.9)	0.83 (0.9)	0.82 (0.9)	0.37 (0.65)	0.75 (0.9)	0.5 (0.75)
	Controller	0.8 (0.88)	0.4 (0.53)	0.53 (0.66)	0.92 (0.92)	0.9 (0.9)	0.91 (0.91)	0.81 (0.84)	0.9 (0.91)	0.85 (0.87)
	Safety Check	0 (0)	0 (0)	0 (0)	0.38 (0.5)	0.33 (0.56)	0.35 (0.53)	0 (0)	0 (0)	0 (0)
Rover	Data Transfer	0.23 (0.71)	0.42 (0.45)	0.3 (0.55)	0.72 (0.83)	0.81 (0.83)	0.76 (0.83)	0.5 (0.55)	0.5 (0.71)	0.5 (0.62)
	Navigation	0.21 (0.4)	0.8 (1.0)	0.33 (0.57)	0.68 (0.70)	0.68 (0.73)	0.68 (0.71)	0.55 (0.6)	0.83 (0.75)	0.66 (0.66)
	Controller	0.75 (1.0)	0.23 (0.53)	0.35 (0.69)	0.78 (0.0.89)	0.92 (0.92)	0.85 (0.91)	0.83 (1.0)	0.71 (0.75)	0.76 (0.85)
	Safety Check	0 (0)	0 (0)	0 (0)	0.38 (0.44)	0.50 (0.67)	0.43 (0.53)	0 (0)	0 (0)	0 (0)
Submarine	Data transfer	0.38 (0.44)	0.45 (0.8)	0.41 (0.57)	0.77 (0.78)	0.83 (0.85)	0.80 (0.81)	0.53 (0.75)	0.41 (0.69)	0.46 (0.71)
	Navigation	0.17 (0.23)	1.0 (1.0)	0.29 (0.37)	0.73 (0.8)	0.78 (0.81)	0.75 (0.80)	0.6 (0.73)	0.6 (0.58)	0.6 (0.64)
	Controller	1.0 (0.66)	0.28 (0.42)	0.44 (0.52)	0.78 (0.81)	0.96 (0.96)	0.86 (0.88)	0.77 (0.9)	0.87 (0.52)	0.82 (0.66)
	Safety Check	0 (0)	0 (0)	0 (0)	0.30 (0.44)	0.43 (0.80)	0.35 (0.57)	0 (0)	0 (0)	0 (0)

tors via cosine similarity. We use NV-Embed-v2 [41] due to its high performance in the leaderboard³ of embedding model quality[42]. The results of this experiment are represented in Table II.

C. RQ3: Module Retrieval Effectiveness

For the evaluation of module retrieval, we establish the ground truth for 476 modules from four devices in ArduPilot: QuadCopter, HeliCopter, Submarine, and Rover. We establish ground truth for all modules within the binaries, which requires approximately 100 human hours. We compare the module categories retrieved by the LLMs using this ground truth. In personal communication, the ArduPilot developers expressed their belief that module categorization was fundamentally hard, even for developers with expert insight.

In all cyber-physical systems, certain functionalities are essential and typically implemented within components known as modules. For clarity and structured analysis, we define four common module categories within this domain. Autopilot devices, for instance, include modules dedicated to communication protocols such as MAVLink, which facilitates the exchange of data between the device and ground stations or other Unmanned Aerial Vehicles (UAVs). These communication modules ensure seamless integration and coordination across the network.

Navigation modules play a crucial role in determining and maintaining the position and orientation of the vehicle, integrating sensors and GPS data to guide its path accurately. Additionally, control modules are fundamental in defining the vehicle's behavior and managing the dynamics and responses based on operational inputs and environmental conditions.

Safety modules are designed to monitor the integrity and operational status of the vehicle, implementing fail-safes and checks that prevent malfunctions. Lastly, some modules do not neatly fit into these categories, encompassing functionalities that are unique or less common. Thus, our categories are defined as data transfer, navigation, safety check, controller, and others. It is crucial to explicitly define each category for the LLM, recognizing that each model may have different understandings and interpretations of these categories, which influences their functionality.

It is important to note that some modules may overlap across categories, such as those fitting both the navigation and controller categories, which manage route planning and vehicle control simultaneously. To evaluate the module retrieval performance of the LLM based on the provided ranking, we select the top k categories from the LLM's output, where k corresponds to the number of categories to which the module belongs. To assess the effectiveness of the module retrieval unit, we employed precision, recall, and F1 score as our evaluation metrics. Table III displays the performance results for each LLM. For each module category, we show the number of modules and the results of the evaluation metrics for each LLM. We calculate the precision, recall, and F_1 score for module categorization based on summaries from both decompiled functions and normalized source code. The results from the normalized code serve as the upper bounds for each evaluation metric, as indicated in parentheses in Table III. Our experimental results show that Codestral could outperform DeepSeek-Coder and CodeQwen. Recall that, according to

³https://huggingface.co/spaces/mteb/leaderboard

Table III, the evaluation metrics for safety check modules are not high due to the limited number of these modules.

D. RQ4: LLM Time Analysis

Evaluating the time performance of the LLMs DeepSeek-Coder, CodeQwen, and CodeStral is crucial due to the significant time required for tasks such as summarizing decompiled functions and retrieving modules. As these LLMs are free to use, the main cost involved is the computational and operational time, which can be substantial. This is particularly important in environments where efficient resource utilization affects project timelines and costs.

We calculate the total seconds each LLM requires to generate function summaries for all decompiled functions per device. We present this analysis in hours in Figure 5. This evaluation indicates that CodeStral takes an average of 20 hours, while DeepSeek-Coder is the fastest, requiring approximately 5 hours to complete the summarization task. We believe that the speed efficiency in DeepSeek-Coder is because of our use of the structured (user vs. assistant) chat template style while prompting the LLM. In our experience, this usually yields shorter outputs from the language model. Note that this template was used to resolve some complications in DeepSeek-Coder in which inconsistencies in the special end token resulted in extremely long and repeated outputs.

The module categorization process is notably efficient, typically completed in a relatively short period and not more than two hours in the most extended cases. Given the minimal variability and the generally swift completion time, we have opted not to include a detailed analysis of time differences in this paper.



Fig. 5. Comparison of time required for function summarization by three open-source LLMs (CodeStral, CodeQwen, and DeepSeek-Coder) across four ArduPilot devices.

VI. CHALLENGES AND LIMITATIONS

This section highlights the limitations of our approach.

- ArduPilot encompasses a comprehensive dataset with numerous functions and modules that often overlap across multiple categories. This overlap between module categories complicates our analysis because it can lead to inconsistencies in how modules are categorized.
- Given the complexity and the specific nature of our dataset, we have deliberately limited the number of categories. Despite this limitation, the categories we have defined are highly useful and provide essential insights that researchers can leverage for effective analysis.
- The categories we choose for our evaluation in this paper are tailored specifically to the domain of cyber-physical systems. Two objections might be raised in this context: First, that these categories are overly coarse-grained and that developers might want to see a more fine-grained categorization of functions within the binary. We chose these categories because they were of broad interest to CPS researchers, and also because of ease of establishing ground truth. The second objection is that they have a narrow focus on robotics firmware. This choice was purely a historical accident. However, we believe that the concept can be readily adapted to other domains. For example, malware analysts might be interested in categories such as network communication, key-logging, or process monitoring.
- We also did not perform prompt engineering to optimize our results. It is possible that with more effort put into prompt design, one or more LLMs might achieve better accuracy.

VII. FUTURE WORK

In this section, we outline our future plans for expanding both our evaluation and the conceptual framework of our research.

- We intend to broaden our evaluation by incorporating well-known and proprietary LLMs such as GPT-4, Copilot, and Claude into our assessment, comparing their effectiveness in function summarization and module retrieval.
- As part of our future plans, we aim to enhance our evaluation by identifying the most effective prompts through a prompt engineering process.
- We aim to enrich the semantic context of modules by using some tools such as SensorLoader [43]. These frameworks could potentially label modules with sensorspecific data flows or peripheral access patterns, enhancing the module categorization process.
- Our current system handles a limited number of general queries. Moving forward, we plan to refine our approach to accommodate more specialized queries, particularly those targeting specific functions within modules. For example, we aim to enable queries that identify well-known mathematical functions commonly utilized in cyber-physical systems.

• We aim to further validate our approach by conducting evaluations on the PX4 dataset, widely recognized in the field of autopilot systems.

VIII. CONCLUSION

In this research, we design a code-based query system called ChatCPS to retrieve different module categories from stripped binary firmware. ChatCPS first decomposes binary firmware into its constituent modules using the BCD algorithm. It then categorizes these modules based on function summaries from decompiled firmware functions, generated by three opensource LLMs: CodeQwen, CodeStral, and DeepSeek-Coder. This system can isolate specific functionalities within binary firmware, facilitating the analysis and understanding of its architecture. Our experimental results demonstrate that Code-Stral surpasses other LLMs in module retrieval, achieving the best performance in our comparisons.

REFERENCES

- Z. Liu and S. Wang, "How far we have come: Testing decompilation correctness of c decompilers," in *Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2020, pp. 475–487.
- [2] Y. Cao, R. Zhang, R. Liang, and K. Chen, "Evaluating the effectiveness of decompilers," in *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2024, pp. 491–502.
- [3] X. Jin, J. Larson, W. Yang, and Z. Lin, "Binary code summarization: Benchmarking chatgpt/gpt-4 and other large language models," *arXiv* preprint arXiv:2312.09601, 2023.
- [4] Z. Song, J. Chen, and K. Zhang, "Bin2summary: Beyond function name prediction in stripped binaries with functionality-specific code embeddings," *Proceedings of the ACM on Software Engineering*, vol. 1, no. FSE, pp. 47–69, 2024.
- [5] L. Jiang, J. An, H. Huang, Q. Tang, S. Nie, S. Wu, and Y. Zhang, "Binaryai: Binary software composition analysis via intelligent binary source code matching," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.
- [6] G. Chen, X. Shang, S. Cheng, Y. Zhang, W. Zhang, and N. Yu, "Foc: Figure out the cryptographic functions in stripped binaries with llms," arXiv preprint arXiv:2403.18403, 2024.
- [7] A. Al-Kaswan, T. Ahmed, M. Izadi, A. A. Sawant, P. Devanbu, and A. van Deursen, "Extending source code pre-trained language models to summarise decompiled binaries," in 2023 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER). IEEE, 2023, pp. 260–271.
- [8] A. Mastropaolo, M. Ciniselli, M. Di Penta, and G. Bavota, "Evaluating code summarization techniques: A new metric and an empirical characterization," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.
- [9] T. Ahmed and P. Devanbu, "Few-shot training llms for project-specific code-summarization," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 2022, pp. 1–5.
- [10] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [11] R. Anil, S. Borgeaud, Y. Wu, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican *et al.*, "Gemini: A family of highly capable multimodal models," *arXiv preprint arXiv:2312.11805*, vol. 1, 2023.
- [12] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, "Qwen technical report," *arXiv preprint arXiv:2309.16609*, 2023.
- [13] Q. Zhu, D. Guo, Z. Shao, D. Yang, P. Wang, R. Xu, Y. Wu, Y. Li, H. Gao, S. Ma *et al.*, "Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence," *arXiv preprint arXiv:2406.11931*, 2024.
- [14] V. Karande, S. Chandra, Z. Lin, J. Caballero, L. Khan, and K. Hamlen, "Bcd: Decomposing binary code into components using graph-based clustering," in *Proceedings of the 2018 on Asia Conference on Computer* and Communications Security, 2018, pp. 393–398.

- [15] "Bigcodebench leaderboard," Hugging Face Spaces, 2023, accessed: 2023-01-22. [Online]. Available: https://huggingface.co/spaces/bigcode/ bigcodebench-leaderboard
- [16] T. Y. Zhuo, M. C. Vu, J. Chim, H. Hu, W. Yu, R. Widyasari, I. N. B. Yusuf, H. Zhan, J. He, I. Paul *et al.*, "Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions," *arXiv preprint arXiv:2406.15877*, 2024.
- [17] B. Hui, J. Yang, Z. Cui, J. Yang, D. Liu, L. Zhang, T. Liu, J. Zhang, B. Yu, K. Lu *et al.*, "Qwen2. 5-coder technical report," *arXiv preprint arXiv:2409.12186*, 2024.
- [18] D. S. Chaplot, "Albert q. jiang, alexandre sablayrolles, arthur mensch, chris bamford, devendra singh chaplot, diego de las casas, florian bressand, gianna lengyel, guillaume lample, lucile saulnier, lélio renard lavaud, marie-anne lachaux, pierre stock, teven le scao, thibaut lavril, thomas wang, timothée lacroix, william el sayed," arXiv preprint arXiv:2310.06825, 2023.
- [19] A. D. Team, "ArduPilot Dataset," Online, 2025. [Online]. Available: http://ardupilot.org
- [20] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [21] A. Visioli, Practical PID control. Springer Science & Business Media, 2006.
- [22] C. Eagle, "The ida pro book: the unofficial guide to the world's most popular disassembler. 2008."
- [23] Ghidra, "Ghidra reference," https://ghidra-sre.org/, 2018.
- [24] Y. Shoshitaishvili, R. Wang, C. Salls, N. Stephens, M. Polino, A. Dutcher, J. Grosen, S. Feng, C. Hauser, C. Kruegel *et al.*, "Sok:(state of) the art of war: Offensive techniques in binary analysis," in 2016 IEEE symposium on security and privacy (SP). IEEE, 2016, pp. 138–157.
- [25] D. Kim, E. Kim, S. K. Cha, S. Son, and Y. Kim, "Revisiting binary code similarity analysis using interpretable feature engineering and lessons learned," arXiv preprint arXiv:2011.10749, 2020.
- [26] D. Xu, J. Ming, and D. Wu, "Cryptographic function detection in obfuscated binaries via bit-precise symbolic loop mapping," in 2017 IEEE Symposium on Security and Privacy (SP). IEEE, 2017, pp. 921– 937.
- [27] Y. Wang, W. Wang, S. Joty, and S. C. Hoi, "Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation," arXiv preprint arXiv:2109.00859, 2021.
- [28] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang *et al.*, "Codebert: A pre-trained model for programming and natural languages," *arXiv preprint arXiv:2002.08155*, 2020.
- [29] L. Huanzhen, "Improve code summarization via prompt-tuning codet5," *Wuhan University Journal of Natural Sciences*, vol. 28, no. 6, pp. 474– 482, 2023.
- [30] S. Arakelyan, R. J. Das, Y. Mao, and X. Ren, "Exploring distributional shifts in large language models for code analysis.(3 2023)," *CoRR abs/2303.09128*, 2023.
- [31] D. Mondal, A. Lodha, A. Sahoo, and B. Kumari, "Understanding code semantics: An evaluation of transformer models in summarization," arXiv preprint arXiv:2310.16314, 2023.
- [32] W. Sun, Y. Miao, Y. Li, H. Zhang, C. Fang, Y. Liu, G. Deng, Y. Liu, and Z. Chen, "Source code summarization in the era of large language models," arXiv preprint arXiv:2407.07959, 2024.
- [33] R. Haldar and J. Hockenmaier, "Analyzing the performance of large language models on code summarization," arXiv preprint arXiv:2404.08018, 2024.
- [34] V. Lomshakov, A. Podivilov, S. Savin, O. Baryshnikov, A. Lisevych, and S. Nikolenko, "Proconsul: Project context for code summarization with llms," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, 2024, pp. 866–880.
- [35] G. Rejithkumar, P. R. Anish, and S. Ghaisas, "Towards understanding contracts grammar: A large language model-based extractive questionanswering approach," in 2024 IEEE 32nd International Requirements Engineering Conference (RE). IEEE, 2024, pp. 310–320.
- [36] B. Roziere, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, R. Sauvestre, T. Remez *et al.*, "Code llama: Open foundation models for code," *arXiv preprint arXiv:2308.12950*, 2023.
- [37] M. E. Newman, "Fast algorithm for detecting community structure in networks," *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, vol. 69, no. 6, p. 066133, 2004.
- [38] L. Gao, J. Tow, B. Abbasi, S. Biderman, S. Black, A. DiPofi, C. Foster, L. Golding, J. Hsu, A. Le Noac'h, H. Li, K. McDonell, N. Muennighoff, C. Ociepa, J. Phang, L. Reynolds, H. Schoelkopf,

A. Skowron, L. Sutawika, E. Tang, A. Thite, B. Wang, K. Wang, and A. Zou, "A framework for few-shot language model evaluation," 07 2024. [Online]. Available: https://zenodo.org/records/12608602

- [39] ArduPilot, "Ardupilot open source autopilot," https://github.com/ ArduPilot/ardupilot, 2023.
- [40] Clang, "A C language family frontend for LLVM," http://clang.llvm.org, 2013, [Accessed 27-August-2013].
- [41] C. Lee, R. Roy, M. Xu, J. Raiman, M. Shoeybi, B. Catanzaro, and W. Ping, "Nv-embed: Improved techniques for training llms as generalist embedding models," 2024. [Online]. Available: https: //arxiv.org/abs/2405.17428
- [42] N. Muennighoff, N. Tazi, L. Magne, and N. Reimers, "Mteb: Massive text embedding benchmark," 2023. [Online]. Available: https://arxiv.org/abs/2210.07316
- [43] A. Dasbach-Prisk, C. Dewitt, and L. Garcia, "Sensorloader: Bridging the gap in cyber-physical reverse engineering across embedded peripheral devices," in *Proceedings of the First International Workshop on Security* and Privacy of Sensing Systems, 2023, pp. 30–36.