

# Towards Specification-Driven LLM-Based Generation of Embedded Automotive Software

Minal Suresh Patil<sup>1,2\*</sup>, Gustav Ung<sup>2</sup> and Mattias Nyberg<sup>2</sup>

<sup>1</sup> Umeå Universitet, 90187 Umeå, Sweden  
minalsp@cs.umu.se

<sup>2</sup> Scania, Granparksvägen 10, 15148 Södertälje, Sweden  
{gustav.ung, mattias.nyberg}@scania.com

**Abstract.** The paper studies how code generation by LLMs can be combined with formal verification to produce critical embedded software. The first contribution is a general framework, `spec2code`, in which LLMs are combined with different types of critics that produce feedback for iterative backprompting and fine-tuning. The second contribution presents a first feasibility study, where a minimalistic instantiation of `spec2code`, without iterative backprompting and fine-tuning, is empirically evaluated using three industrial case studies from the heavy vehicle manufacturer Scania. The goal is to automatically generate industrial-quality code from specifications only. Different combinations of formal ACSL specifications and natural language specifications are explored. The results indicate that formally correct code can be generated even without the application of iterative backprompting and fine-tuning

**Keywords:** Code Generation · Formal verification · Large Language Models · Automated Software Engineering

## 1 Introduction

Recent advancements in Large Language Models (LLMs) have shown promising, and sometimes astonishing, results in code generation [35,29]. However, from several studies [33,41], it is also clear that it is hard to guarantee code correctness and quality. In the area of automotive embedded systems, correctness and quality of the software are crucial. To be more specific, by *correctness* we here mean functional correctness with respect to functional specifications and also absence of errors that may cause safety and cybersecurity issues. By *quality*, we mean all other properties typically expected in embedded code, as defined in coding standards and guidelines such as MISRA-C [23] and “the power of 10” rules [14].

In the present paper, we consider the problem of using LLMs to generate source code for critical embedded software. We make the following two contributions:

---

\* Work was done while the author was at Scania

- We introduce the `spec2code` framework. This framework is an adaptation of the previously presented LLM-Modulo Framework from the position paper [17] in the area of planning. In `spec2code`, LLMs are combined iteratively with *critics*, such as software verifiers, to produce high-quality, correct software.
- Based upon industrial case studies, we investigate the feasibility of the `spec2code` framework by considering a minimalistic instantiation of it, in which we use non-iterative prompting, exclude fine-tuning, and focus mainly on formal functional correctness of the code.

In the second contribution, within the context of the `spec2code` framework, we explore if it is at all feasible to have LLMs generate automotive safety-critical embedded C code from specifications. We consider specifications being given in natural language (NL) and in the formal ANSI/ISO C Specification Language (ACSL) [4]. To evaluate the correctness of the code generated, we apply deductive verification and the tool Frama-C [7] to verify whether the code implements given ACSL specifications. Since this is our first experiments with `spec2code`, and the focus is exploring its fundamental feasibility, we have, as stated above, chosen to just implement a version without non-iterative prompting.

For the evaluation, we use three industrial case studies from the heavy vehicle manufacturer Scania. Each case study consists of a single software *module*, i.e. a pair of a `.c`-file and a `.h`-file, from real production software. From the specifications of each such module, we try to generate functionally equivalent code by using two different LLMs, namely `gpt-3.5` and `gpt-4-turbo`. We aim to address the following research questions:

- **RQ1:** How can the integration of formal verification tools, such as Frama-C, with off-the-shelf LLMs be combined to automatically generate formally verified C code from ACSL specifications, along with high- and low-level natural language software specifications?
- **RQ2:** How can specifications be effectively translated into prompts or inputs for LLMs to ensure the generation of correct and verified C code?
- **RQ3:** How do we assess the functional correctness and quality of the generated code?

The `spec2code` framework is, to the best of our knowledge, the first presented general framework focusing on using LLM’s to generate critical embedded software from specifications. A related work is the so-called TriCo methodology [2] envisioning how LLMs and formal methods can be exploited to *co-pilot* the simultaneous development of the three artifacts implementation, specification, and tests. There are however noticeable differences between TriCo and `spec2code`. Firstly, TriCo aims at *supporting* the programmer in writing code by giving suggestions, while `spec2code` aims at *replacing* the programmer with LLM- and formal-based technology, and thus, only relies on having developers creating specifications. Secondly, while TriCo emphasizes the triple including test cases, `spec2code` relies on only formal formal verification to guarantee correct code, thus eliminating the need for testing.

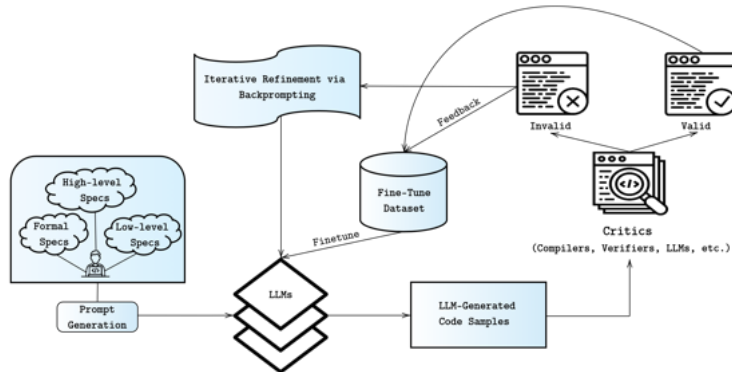


Fig. 1: Overview of `spec2code` framework

The paper is organized as follows. In Section 2, background information on deductive verification of C programs and LLMs are presented. Then, in Section 3, we present the first contribution, i.e. the `spec2code` framework. Section 4, 5, and 6 then describe the second contribution, i.e. the case study on using the minimalist instantiation of the `spec2code` framework applied to real industrial code. Finally, Section 7 and 8 contains a discussion and the conclusions of the paper.

## 2 Background

### 2.1 Deductive Verification of C Programs

FRAMA-C [7] is an open-source, extensively developed framework that verifies C programs annotated with ACSL [4] specifications. Its WP plug-in enables users to verify that C code meets ACSL specifications through deductive verification [1,13]. Specifically, WP utilizes an advanced form of weakest precondition calculus [3,20], which inspired its name. ACSL specifications are placed within special comments `/*@ ...*/` and primarily include function contracts and code annotations. Function contracts feature pre-conditions (**requires** clauses) and post-conditions (**ensures** clauses), which are pure logical formulas that need to be verified before and after any function call. The **assigns** clause specifies the frame conditions of the function, i.e. the set of memory locations, including global variables and pointers, that are allowed to be modified by the function. Code annotations, such as **assert** clauses, are pure logical formulas linked to specific points in the program that must be validated on every execution path that passes through these points. *Ghost variables* are variables that are evaluated in a static runtime environment disjoint from the normal heap and stack. These can be related to the normal program variables by *ghost statements*, but may not affect the program.

## 2.2 Large Language Models

In recent years, Large Language Models (LLMs) have shown impressive progress in a wide range of tasks. Notably, there has been a significant increase in the number of LLMs designed specifically for coding tasks, particularly in the area of code generation. LLMs utilize Transformers [36], which are the most advanced neural architecture in natural language processing. Additionally, Transformers have proven highly effective in addressing classic problems in verification [9,11], reasoning [21], and in the auto-formalization of mathematics and formal specifications [12,8]. OpenAI’s `gpt-3.5` and `gpt-4-turbo` build on the pre-trained GPT-3 model, with additional fine-tuning using Reinforcement Learning with Human Feedback (RLHF) [27]. Although they are not specifically optimized for code generation, both models have shown impressive performance on various related tasks [24,30].

## 2.3 Prompting LLMs

Prompts are user-provided inputs, such as queries and instructions that guide LLMs and instruct their behavior for specific tasks. Previous work [38,37] have demonstrated that the quality of input prompts plays a crucial role in the performance of LLMs, significantly influencing the quality of the output. A prompt template is often used to generate prompts in a consistent and reproducible way. Chain-of-Thought (CoT) prompting [31] has significantly enhanced zero- and few-shot performance in various complex reasoning tasks, such as arithmetic, commonsense, symbolic, and logical reasoning. A crucial feature of CoT prompting is the use of rationales, which demonstrate the step-by-step reasoning process. In this work, we design our prompts based on Zero-shot-CoT [19], which involves merely introducing the rationale-triggering sentence “*Let’s think step by step*” to the LLMs, which leads to significant improvements in zero-shot performance.

## 3 Spec2code Framework

Kambhampati et al. [17] examines the role of LLMs in planning tasks known as *LLM-Modulo Framework*. It concludes that while LLMs cannot independently plan or verify, they can assist within a hybrid framework that integrates LLMs with external *critics*.

Building on this core principle, we introduce `spec2code` which aims to leverage the capabilities of LLMs to generate code from both informal and formal specifications via prompting. `spec2code` aims to make use of the generative power of LLMs while ensuring the functional correctness and quality of the code generated through software verifiers.

Figure 1 shows the `spec2code` framework that involves a systematic method for generating code from specifications using LLMs. The process begins with the verification engineer writing both high-level and low-level specifications in

natural language, as well as formal specifications in ACSL. These specifications describe the desired functionality and behavior of the code to be generated.

High-level specifications are typically abstract and outline the module’s objectives and functionalities, whereas formal specifications describe the precise semantics in a formal language of software behavior. Low-level specifications describe the implementation and specific requirements the module must achieve to realize the high-level specifications. Examples of these specifications are provided in Section 5.1.

By using a carefully crafted prompt that employs one or a combination of prompting strategies such as in-context learning, zero-shot/few-shot prompting, or Chain-of-Thought, these specifications are used to generate initial programs using LLMs. These programs are first checked for successful compilation by a compiler, and then a software verifier assesses the functional correctness of these programs. Here, we term the compiler and the software verifier as critics. If a program fails to compile successfully or fails the verification process, the feedback provided by the compiler or the software verifier is used as part of the fine-tuning dataset. This dataset is continuously updated along with the successfully verified programs to further fine-tune the LLMs. Furthermore, we employ iterative backprompting to refine our initial prompt based on feedback from various sources that may come from the compiler, counterexamples produced by software verifiers such as Frama-C, code quality assessments provided by other LLMs (which can be further optimized using prompt optimization techniques such as TextGrad and DSPy [40,18]), or human evaluators. By continuously incorporating feedback from these critics, we enhance the accuracy and quality of the prompts. This refinement process ensures that each iteration of both the LLM-generated code and the prompts aligns closely with the desired code quality and standards.

Building upon this iterative refinement process, we can further fine-tune the LLMs using a supervised approach known as Supervised Fine-Tuning (SFT) and then perform preference optimization to produce preferred outputs [27]. This process requires data from different sources and feedback mechanisms. During the SFT stage, we collect data that includes paired datasets of specifications (both natural language and ACSL), and verified C code implementations; functional safety requirements (e.g., ISO 26262 compliance [32]); industry-standard coding guidelines (e.g., MISRA C [23]); and successfully verified C code with common verification properties such as memory safety and absence of runtime errors. This step ensures the LLMs generate appropriate programs for given inputs. Feedback data will consist of failed verification attempts with counterexamples along with explanations of why the verification failed, human expert reviews or comments via iterative backprompting, compiler warnings and errors, and any feedback from other LLMs. We then construct preference pairs, i.e., preferred (verified, correct) vs. non-preferred (failed verification, incorrect) programs, which are optimized to produce preferred outputs using optimization techniques such as Direct Preference Optimization (DPO)[28]. This approach aligns the LLM to produce verified code by leveraging multiple feedback sources

and domain knowledge. This process necessitates a robust pipeline for code generation, verification, feedback collection, and model updating, with careful balancing of feedback sources to prevent overfitting.

The main objective of the `spec2code` framework is to produce high-quality, specification-compliant code through a combination of initial LLM-based code generation from specifications and through iterative refinement via backprompting. Our framework integrates high-level, low-level and formal specifications with LLM-generated code, iterative refinement via backprompting, and a fine-tuning dataset. The iterative nature and continuous improvement of the fine-tuning dataset ensure that the generated code progressively aligns more closely with the intended specifications, resulting in reliable, safe, and trustworthy code.

## 4 Case Studies

Our case study focuses on three automotive control modules obtained from Scania: the Steering Fluid Level Detection (SFLD) module, the Brake Light Activation (BRAK) module, and the Power Steering Backup (STEE) module. We assess the correctness and quality of the generated code. This work not only evaluate the feasibility of using LLMs for such tasks but also aims to identify potential limitations and areas for improvement in the context of automotive software engineering.

### 4.1 Code Style of Application Modules in Embedded/Automotive Software

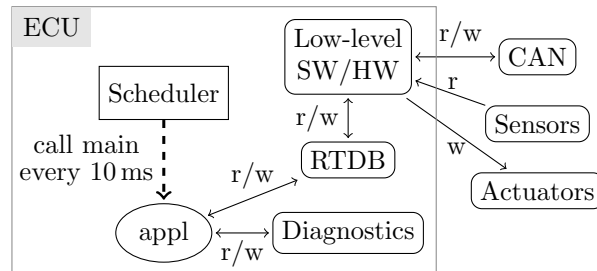


Fig. 2: ECU System architecture

All case study modules in this study are written in C [15], and are developed according to internal coding guidelines largely identical to those of MISRA-C [23]. The modules are executed in two of several Electronic Control Units (ECUs) in the embedded system of Scania trucks. Each ECU has a static scheduler which repeatedly calls the main entry point of each application module once every 10 ms. The control flow is simple, containing no recursion or loops. The

module (in isolation) is strictly sequential, although as part of the larger system it is not. Other than the scheduler, the application modules communicates directly with other application modules and infrastructure software. The two main mechanism for internal communication within the ECU are a diagnostics module, and a real-time database (RTDB) module which facilitates communication between the different application level modules as shown in Figure 2. The external communication between ECUs is done by CAN [16], and the application can indirectly interact with the physical environment by sensors and actuators. The implementation is generally structured as follows: when the main function is called, all inputs from the system are read by calling functions in RTDB, then several computations are made using these inputs and static variables within the module, before finally writing all outputs to RTDB. Reads and writes to the diagnostics module occur throughout computing results, and may be part of, or the result of, those computations.

## 4.2 Working with the case studies

All three case studies are taken from Scania’s development repository. Each case study has a set of natural language specifications of varying quality associated with it. From these natural language specifications, the ACSL specifications were derived by hand.

We restrict all case studies to one single translation unit each, meaning that the interaction with library code has been abstracted away. For example, the interaction with RTDB has been removed, and instead the relevant signals are assumed to be globally defined static variables. Importantly, the actual code functionality has not been changed, and therefore the modified code follows the original code.

## 4.3 Case study 1: Oil Level Warning (SFLD)

The oil level warning module is a software responsible for emitting diagnostics warnings whenever the oil level has been low for a specific amount of time. The original module has 200 LoC, not including declarations and type definitions from header files. There are 1 high-level natural language specification, 11 low-level natural language specifications, and 10 derived ACSL specifications for this module.

## 4.4 Case study 2: Brake Light Activation (BRAK)

The brake light module is a software responsible for activation of brake lights. The original module has 400 LoC not including declarations and type definitions from header files. In total, there are 1 high-level natural language specification, 17 low-level natural language specifications, and 17 derived ACSL specifications for this module

### 4.5 Case study 3: Power Steering Backup (STEE)

The power steering backup module is software responsible for activation of a backup system in case the primary power steering system fails. This module and its specifications has been thoroughly studied in a previous paper [34]. In the present paper, a modified version has been used. There are 1 high-level natural language specification, 5 low-level natural language specifications, and 5 derived ACSL specifications for this module. This version of the module has 150 LoC.

## 5 Experimental Setup

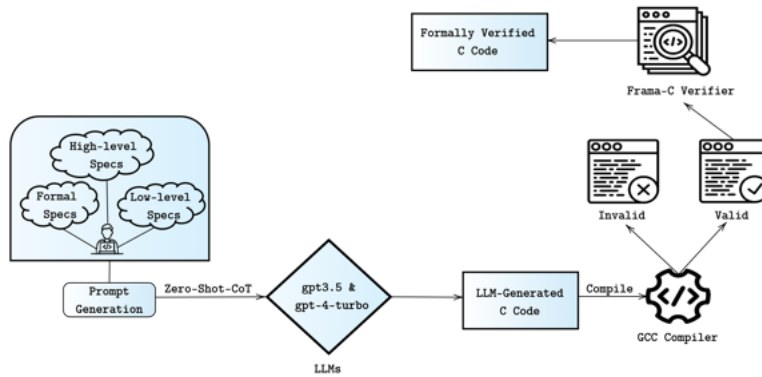


Fig. 3: Overview of the minimalist version of spec2code framework

As shown in Figure 3, we evaluate a minimalist instantiation of spec2code without iterative prompting and fine-tuning. With a minimalistic approach, we aim to explore how the latest State-of-the-art (SotA) Large Language Models (LLMs) perform in a real-world industrial setting. This particular instantiation includes GCC as the C compiler and Frama-C as the software verifier, which serve as the critics. For our evaluation, we also use the *diffkemp* [22] program equivalence tool, and human code quality inspection.

### 5.1 Specification Types

1. **High-Level Natural Language (HLNL)**: Set of high-level specifications that captures the overall functionality and constraints of the system focusing on what the system needs to achieve. A high-level specification is shown below.

*High-Level Natural Language specification 1*

If there is a brake light request, then the truck lights shall be activated.



2. **Low-Level Natural Language (LLNL)**: Set of low-level specifications which combined aims to realize the high-level specifications. A low-level specification related to the high-level specification above is shown below.

<i>Low-level Natural Language specification 1</i>
If the operational state is normal or emergency with limited functionality, and the supply voltage is not low, then the brake lights should be enabled.

3. **ACSL**: Set of formal specifications in written in ACSL. An ACSL specification derived from the low-level specification above is shown below.

<i>ACSL specification 1</i>
<pre>((gh_operationalState == NORMAL_OPERATION    gh_operationalState == EMERGENCY_STOP_LIMITED) &amp;&amp; gh_supplyVoltageLevel != VOLTAGE_LOW) ==&gt; gh_brakeLightEnabled == TRUE;</pre>

In our experiments, we use all possible combinations of specifications to thoroughly evaluate the LLM’s capability to generate functionally correct code with respect to the specifications. By incorporating different specification types, we aim to evaluate the effectiveness in producing functionally correct code across various specification categories.

## 5.2 LLM Selection

In our experiments, we evaluate our framework on current SotA LLMs namely `gpt-4-turbo` [26] and `gpt-3.5` [25]. These models were chosen because they offer API access, which allows us to leverage their capabilities without the need for high memory requirements. For `gpt-4-turbo` and `gpt-3.5` we configure the parameters of the LLMs as follows: `max_token` : 4096, `temperature` : 0.8, `presence_penalty` : 0.5, `frequency_penalty` : 1, and `top-p` : 1. A temperature of 0.8 is preferred for our experiments because it introduces controlled randomness, allowing for more diverse and creative programs while still maintaining coherence.

## 5.3 Prompt Design

We designed a prompt based on the Zero-shot-CoT prompting technique to guide the LLM in generating code from ACSL specifications, as well as high-level and low-level specifications provided in natural language. We first specify the *system prompt* to define the intended role, capabilities, and behavior of the LLM. Next, we specify the *user prompt* that instructs the LLM to complete the given task, i.e., completing the function definition from specifications. We provided the structure of the C program, including the header, global variables, and function header, as part of the prompt (See [Appendix A](#) for prompt template).

## 5.4 Evaluation of Generated Code

In all case studies, each combination of model and specification produces a generated program which is then compiled. If the compilation is successful, then the program is evaluated in three different ways: functional correctness, program equivalence, and code quality.

**Functional correctness.** The functional correctness of the program is verified with respect to its ACSL specification by using the weakest-preconditions calculus plugin (WP) in Frama-C. We choose Frama-C as a verification framework because we deem it to be mature and of state-of-the-art. In our evaluation we use the Frama-C version 27.1 with the backend solvers Z3 [10] and Alt-Ergo [6]. The results can be seen in section 6.

**Program equivalence.** The functional correctness verification shows whether the implementation satisfies the specification, but it is possible that the specification itself might be incomplete. By incompleteness, we mean that the set of program behaviors specified by the specification might not cover the intended program behaviors of the code. We evaluate the *relative correctness* of the original code from the Scania repository and the code generated by the LLM using program equivalence based on program semantics.

All case studies include one handwritten program each, which is verified w.r.t. its formal specification, and is considered to be the reference program. The generated programs are compared to the relevant reference program, and any semantic deviation is captured. Concretely, the evaluation is done by using the tool *diffkemp* [22]. This tool can show program equivalence on instruction-level, for programs written in LLVM byte-code. A set of semantics preserving transformations is used to accomplish this task. Hence, the tool can fail to show equivalence either because the programs are not equivalent, or because the programs are too different structurally.

**Code Quality:** The code quality is evaluated quantitatively by counting the lines of code (LoC), and qualitatively by evaluating its conformance with “the power of 10” rules [14]. Only the lines of code in the implemented function are counted, and all empty lines and comments are removed. Generally a longer program implementing the same functionality as a short program can be considered to be more complex, and of lower quality. However, analysis of LoC is surely not sufficient for evaluating the code quality. Hence, we also manually inspect the generated code and evaluate its conformance with “the power of 10” rules. A summary of the rules is the following:

1. Avoid complex flow constructs, such as goto and recursion.
2. All loops must have fixed bounds.
3. Do not use dynamic memory allocation after initialization.
4. Restrict size of function to around 60 LOC.
5. (Use a minimum of two runtime assertions per function.) **Not used.**
6. Restrict the scope of data to the smallest possible.

7. Check the return value of all non-void functions, or cast to void to indicate the return value is useless.
8. Use the preprocessor sparingly.
9. Limit pointer use to a single dereference, and do not use function pointers.
10. Compile with all possible warnings active; all warnings should then be addressed before release of the software.

As indicated in the list above, we do not use Rule 5. We argue that the use of this rule is questionable, at least in the context of automotive embedded software. In C/C++, it is in fact generally recommended to not use assertions in production code, e.g. see [39]. Rule 8 is deemed to be followed when each instance of preprocessor usage can be justified, and is judged to be simpler than an alternative approach. If a file is not compilable, then no code quality analysis is performed.

## 6 Results

In this section, we discuss the results of our research by placing them in the context of the three research questions:

1. Capability of current SotA LLMs to generate code based on formal, high-level, and low-level specifications (**RQ1**).
2. Effectiveness of input prompt design on the LLM’s performance (**RQ2**).
3. Evaluation of both functional correctness and overall quality of the LLM-generated code (**RQ3**).

It is worth noting that all compilable programs produced by the LLM adhered to the “power of 10” rules, and no warnings were issued by the GCC compiler. Consequently, our analysis tables focus solely on the Lines of Code (LoC) metric.

Model	Specification Type	Compiled	Eq. Check	Verified (Proved Goals)	LoC
GPT-4-turbo	ACSL	Yes	Not Eq	33 / 33	35
	HLNL	Yes	Not Eq	5 / 23	25
	LLNL	Yes	Not Eq	23 / 29	54
	ACSL + HLNL + LLNL	Yes	Not Eq	23 / 29	35
	HLNL + LLNL	Yes	Not Eq	23 / 29	40
	ACSL + LLNL	Yes	Not Eq	18 / 25	36
	ACSL + HLNL	Yes	Not Eq	24 / 29	37
GPT-3.5	ACSL	Yes	Not Eq	21 / 27	37
	HLNL	Yes	Not Eq	11 / 22	25
	LLNL	No	n/a	n/a	n/a
	ACSL + HLNL + LLNL	No	n/a	n/a	n/a
	HLNL + LLNL	No	n/a	n/a	n/a
	ACSL + LLNL	No	n/a	n/a	n/a
	ACSL + HLNL	Yes	Not Eq	18 / 27	35

Table 1: Steering Fluid Level Detection (SFLD) module

Model	Specification Type	Compile	Eq. Check	Verified (Proved Goals)	LoC
GPT-4-turbo	ACSL	Yes	Not Eq	41 / 50	48
	HLNL	Yes	Not Eq	13 / 26	21
	LLNL	Yes	Not Eq	39 / 40	48
	ACSL + HLNL + LLNL	Yes	Not Eq	40 / 52	64
	HLNL + LLNL	Yes	Not Eq	30 / 30	32
	ACSL + LLNL	Yes	Not Eq	39 / 40	65
	ACSL + HLNL	Yes	Not Eq	41 / 52	30
GPT-3.5	ACSL	Yes	Not Eq	41 / 43	36
	HLNL	Yes	Not Eq	11 / 24	20
	LLNL	No	n/a	n/a	n/a
	ACSL + HLNL + LLNL	Yes	Not Eq	39 / 40	44
	HLNL + LLNL	No	n/a	n/a	n/a
	ACSL + LLNL	Yes	Not Eq	13 / 30	35
	ACSL + HLNL	Yes	Not Eq	38 / 42	33

Table 2: Brake Light Activation (BRAK) module

Model	Specification Type	Compile	Eq. Check	Verified (Proved Goals)	LoC
GPT-4-turbo	ACSL	Yes	Not Eq	6 / 18	27
	HLNL	Yes	Not Eq	3 / 8	6
	LLNL	Yes	Not Eq	8 / 8	22
	ACSL + HLNL + LLNL	Yes	Not Eq	8 / 8	26
	HLNL + LLNL	Yes	Not Eq	8 / 8	24
	ACSL + LLNL	Yes	Not Eq	8 / 8	38
	ACSL + HLNL	Yes	Not Eq	8 / 8	36
GPT-3.5	ACSL	Yes	Not Eq	8 / 8	17
	HLNL	Yes	Not Eq	3 / 8	9
	LLNL	Yes	Not Eq	8 / 8	16
	ACSL + HLNL + LLNL	Yes	Not Eq	8 / 8	16
	HLNL + LLNL	No	n/a	n/a	n/a
	ACSL + LLNL	Yes	Not Eq	8 / 8	16
	ACSL + HLNL	Yes	Not Eq	8 / 8	16

Table 3: Power Steering Backup (STEE) module

## 7 Discussion

**LLM comparison.** To ensure a fair comparison between `gpt-3.5` and `gpt-4-turbo`, we evaluate and report only the initial code generation response, known as *pass@1* [5]. The *pass@1* metric offers a direct measure of an LLM’s capability to produce correct code on the first attempt, i.e., producing syntactically correct code that is also successfully verified by a software verifier. By focusing on the first generated code output, we highlight the LLM’s efficiency and reliability in producing correct solutions without requiring further iterations thus providing an unbiased assessment of each LLM’s performance under the same parameter settings, making it easier to compare their effectiveness. However, it worth noting that the *pass@1* is a strict metric because it does not account for the LLM’s performance on subsequent attempts or its ability to learn from feedback.

In Tables 1, 2 and 3 we compare the performance of `gpt-4-turbo` and `gpt-3.5` models across various specification types (ACSL, HLNL, LLNL, and their combinations) for each of the case study.

In Table 1, `gpt-4-turbo` successfully produces code that compiles for all specification types, while `gpt-3.5` is unable to do so with several. In terms of verification, `gpt-4-turbo` produces code that is successfully verified with ACSL specifications where Frama-C proves all 33 goals, whereas `gpt-3.5` performs less consistently, with fewer proved goals and failing to compile for multiple specification types.

In Table 2, `gpt-4-turbo` successfully compiles all specification types, while `gpt-3.5` fails to compile LLNL and HLNL + LLNL specification types. `gpt-4-turbo` generally achieves higher numbers of verified (proved) goals compared to `gpt-3.5`, with the HLNL + LLNL combination for `gpt-4-turbo` achieving 30/30 proved goals.

In Table 3, `gpt-4-turbo` successfully compiles all specification types, while `gpt-3.5` fails only for the HLNL + LLNL combination. In terms of verified proved goals, `gpt-4-turbo` successfully verifies most combined specifications except for ACSL and HLNL specification types, whereas `gpt-3.5` fails to compile the HLNL + LLNL specification type but is able to verify all specification types except the HLNL specification type.

From the results, we see that relying solely on high-level natural language specifications often leads to suboptimal results. This is likely because high-level specifications are inherently ambiguous and lack technical specificity, causing LLMs to misinterpret the specifications. Consequently, the generated code may not align with the user’s actual intent.

**Ghost variables.** Ghost variables in ACSL play an important role in formal verification by serving as auxiliary variables to help prove program properties. However, we observe that their use can significantly impact the verification process in code generation by LLMs since they are highly context-dependent and if the LLM misinterprets the role of ghost variables or if they are not properly specified in the prompt, the LLM tends to generate code that misuses or incorrectly implements these variables, leading to failed verification instances. Furthermore, the additional overhead or added complexity of specifications involving ghost variables can inhibit the capability of the LLMs to produce code, especially if they have not been adequately trained on similar examples.

Thus, ensuring consistency and correct usage of ghost variables in both the specification and the generated code is crucial for successful verification of the code. Due to the mentioned challenges in using ghost variables for code generation, in all our case studies, we converted the ghost variables in the specification to concrete variables in order to ensure the LLMs produce at least that code that is able to compile successfully.

**Equivalence checking.** The program equivalence checking tool was not successful in showing equivalence in any case. By doing manual inspection, we found that this is due to several reasons. First, we must recognize that the LLM was tasked with generation of code according to the specification, and nothing more. However, while the specification is supposed to be complete for the generated program (by construction), this specification does not necessarily need to specify the complete behavior of the reference program. Therefore, the specified behavior is equivalent for both programs, but the unspecified behavior is not.

Another reason for the program inequality is due to ghost variables in the specification. As previously stated, all generated programs treat ghost variables as concrete program variables. In the reference programs, the corresponding

behavior is instead implemented in terms of local variables which are not visible in the function contract due to scoping rules. Therefore, ghost annotations are used in the implementation, relating the local variables with the ghost variables. Consequently, the semantic equivalence can not be shown, because the generated program additionally affects the concrete (formerly ghost) variables, i.e. the set of concrete global variables involved in side effects is larger for the generated program.

While the results show that equivalence could not be shown between the reference program and the generated programs, equivalence could be shown between generated programs in some cases. In particular, five of the `gpt-4-turbo` generated programs for the STEE case study, all of which could be functionally verified, could be divided into three different equivalence classes. These variations clearly illustrate that multiple interpretations were allowed by the specification.

**Code quality** As stated in the results, all code modules adhere to “the power of 10” rules. Despite having temporal specifications, i.e. specifications that must be interpreted over time and multiple executions of the code modules, the generated code did not contain any use of loops or recursion, in adherence to rule 1. This was probably prevented by giving the LLM context about the function execution environment such as details about the scheduler.

Importantly, the generated code did not feature any dynamic memory allocations or pointers, in adherence to rules 3 and 9. It is possible that pointers were avoided due to the usage of a single function for each module, this removes the need for transferring struct data between functions, which is a normal use-case for pointers.

The LLM was provided with the module interface, and the intended execution context, namely automotive software, and this is likely to have influenced the style of the generated code.

While the general style of code produced is deemed suitable for usage in automotive embedded devices, the code style was inconsistent in some cases. For example, in some cases input validation was done by using normal if-statements mixed with ternary if-statements. It is possible that the usage of ternary-statements were influenced by preprocessor macros that are defined in terms of ternary-statements.

## 7.1 Threats to Validity

**Internal Validity.** A potential threat to the internal validity of our method is the natural non-deterministic behavior of LLMs, which leads to unpredictable and non-repeatable code generation. We experimented with various temperature settings and ultimately set the LLM’s temperature to 0.8 to control the randomness in the code produced. Another, a potential threat to the internal validity, when evaluating performance of LLMs for code generation, is that the code examples used in the evaluation have been part of the training data. This threat is completely avoided in the experiments of the present paper since all case studies use proprietary, not published, code from the manufacturer Scania.

**External Validity.** While it is technically feasible to substitute `gpt-4-turbo` and `gpt-3.5` with other LLMs used in our experiments, this may yield different results, indicating that our findings might not apply universally. It is essential to consider whether our findings are valid for other codebases at Scania and in different companies. The variability in contexts, coding standards, and organizational requirements could affect the applicability of our findings, thus necessitating further validation to ensure broader relevance.

## 8 Conclusion

The paper has studied how code generation by LLMs can be combined with formal verification with the aim of producing critical embedded software. The first contribution is a general framework `spec2code`, where LLMs are combined with different types of critics. In the second contribution, as a first feasibility study, a minimalistic instantiation of `spec2code`, without iterative backprompting and fine-tuning, was applied to three industrial case studies from the heavy vehicle manufacturer Scania. The goal was to, from specifications only, automatically generate industrial quality code that could be hypothetically integrated in the Scania production code.

Despite not using iterative backprompting and any fine-tuning, we were able to produce successfully compiled code in all three case studies. The code was also formally verified, for some combination of specifications, in two of the three case studies. We view this result as quite promising, although the case studies were all relatively small and of low complexity. However, we expect that with the addition of iterative backprompting and fine-tuning, which we plan as future work, the approach will also be much more powerful and become capable of handling more challenging coding problems. Additionally, we intend to include the  $pass@k$  in future work as this will offer a more comprehensive understanding of the LLMs ability to generate code by accounting for the presence of a correct program among the top  $k$  candidate programs. Furthermore, while we acknowledge that closed-source models for scientific evaluation may be sub-optimal, we aim to address this concern in the future work by evaluating our framework using open-source LLMs.

An important conclusion of the work is that we note that, while a human programmer utilizes specifications *and* context knowledge, the LLMs in `spec2code` only have access to specifications. Thus, the correctness of the generated code is with respect to the specifications only. This implies that to have *complete* specifications becomes crucial. This we judge to be rare in real industrial software development. This raises the question if the extra efforts needed to make specifications complete, is more or less compared to the effort saved by delegating the programming to LLMs. This is also something we plan to investigate in future work.



## References

1. Ahrendt, W., Beckert, B., Bubel, R., Hähnle, R., Schmitt, P.H., Ulbrich, M.: Deductive software verification—the key book. Lecture notes in computer science **10001** (2016)
2. Ahrendt, W., Gurov, D., Johansson, M., Rümmer, P.: Trico—triple co-piloting of implementation, specification and tests. In: International Symposium on Leveraging Applications of Formal Methods. pp. 174–187. Springer (2022)
3. Baudin, P., Bobot, F., Correnson, L., Dargaye, Z., Blanchard, A.: WP Plug-in Manual. CEA LIST, Inria, <https://www.frama-c.com/download/frama-c-wp-manual.pdf>
4. Baudin, P., Filiâtre, J.C., Marché, C., Monate, B., Moy, Y., Prevosto, V.: Acsl: Ansi/iso c specification. URL <https://frama-c.com/html/acsl.html> (2021)
5. Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H.P.d.O., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., et al.: Evaluating large language models trained on code. arXiv preprint arXiv:2107.03374 (2021)
6. Conchon, S., Coquereau, A., Iguernlala, M., Mebsout, A.: Alt-ergo 2.2. In: SMT Workshop: International Workshop on Satisfiability Modulo Theories (2018)
7. Correnson, L., Cuoq, P., Kirchner, F., Maroneze, A., Prevosto, V., Pucetti, A., Signoles, J., Jakobowski, B.: Frama-C User Manual. CEA LIST, Inria, <http://frama-c.com/download/frama-c-user-manual.pdf>
8. Cosler, M., Hahn, C., Mendoza, D., Schmitt, F., Trippel, C.: nl2spec: Interactively translating unstructured natural language to temporal logics with large language models. In: International Conference on Computer Aided Verification. pp. 383–396. Springer (2023)
9. Cosler, M., Schmitt, F., Hahn, C., Finkbeiner, B.: Iterative circuit repair against formal specifications. arXiv preprint arXiv:2303.01158 (2023)
10. De Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: International conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340. Springer (2008)
11. Hahn, C., Schmitt, F., Kreber, J.U., Rabe, M.N., Finkbeiner, B.: Teaching temporal logics to neural networks. arXiv preprint arXiv:2003.04218 (2020)
12. Hahn, C., Schmitt, F., Tillman, J.J., Metzger, N., Siber, J., Finkbeiner, B.: Formal specifications from natural language. arXiv preprint arXiv:2206.01962 (2022)
13. Hähnle, R., Huisman, M.: Deductive software verification: from pen-and-paper proofs to industrial tools. Computing and Software Science: State of the Art and Perspectives pp. 345–373 (2019)
14. Holzmann, G.J.: The power of 10: Rules for developing safety-critical code. Computer **39**(6), 95–99 (2006)
15. International Organization for Standardization: Programming languages—C (1999), ISO/IEC 9899:1999
16. International Organization for Standardization: Road vehicles controller area network (CAN) (2015), ISO 11898-1:2015
17. Kambhampati, S., Valmeekam, K., Guan, L., Stechly, K., Verma, M., Bhambri, S., Saldyt, L., Murthy, A.: Llms can’t plan, but can help planning in llm-modulo frameworks. arXiv preprint arXiv:2402.01817 (2024)
18. Khattab, O., Singhvi, A., Maheshwari, P., Zhang, Z., Santhanam, K., Vardhamanan, S., Haq, S., Sharma, A., Joshi, T.T., Moazam, H., et al.: Dspy: Compiling declarative language model calls into self-improving pipelines. arXiv preprint arXiv:2310.03714 (2023)

19. Kojima, T., Gu, S.S., Reid, M., Matsuo, Y., Iwasawa, Y.: Large language models are zero-shot reasoners. *Advances in neural information processing systems* **35**, 22199–22213 (2022)
20. Leino, K.R.M.: Efficient weakest preconditions. *Information Processing Letters* **93**(6), 281–288 (2005)
21. Lewkowycz, A., Andreassen, A., Dohan, D., Dyer, E., Michalewski, H., Ramasesh, V., Slone, A., Anil, C., Schlag, I., Gutman-Solo, T., et al.: Solving quantitative reasoning problems with language models. *Advances in Neural Information Processing Systems* **35**, 3843–3857 (2022)
22. Malík, V., Vojnar, T.: Automatically checking semantic equivalence between versions of large-scale c projects. In: 2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST). pp. 329–339. IEEE (2021), <https://github.com/diffkemp/diffkemp>
23. MIRA Ltd: MISRA-C:2004 Guidelines for the use of the C language in critical systems (Oct 2004), <https://misra.org.uk/misra-c>
24. Olausson, T.X., Inala, J.P., Wang, C., Gao, J., Solar-Lezama, A.: Demystifying gpt self-repair for code generation. arXiv preprint arXiv:2306.09896 (2023)
25. OpenAI: gpt-3.5-turbo-0125. <https://www.openai.com> (2022), accessed: 25/04/2024
26. OpenAI: gpt-4-turbo. <https://www.openai.com> (2023), accessed: 25/04/2024
27. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *Advances in neural information processing systems* **35**, 27730–27744 (2022)
28. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems* **36** (2024)
29. Ross, S.I., Martinez, F., Houde, S., Muller, M., Weisz, J.D.: The programmer’s assistant: Conversational interaction with a large language model for software development. In: Proceedings of the 28th International Conference on Intelligent User Interfaces. pp. 491–514 (2023)
30. Schäfer, M., Nadi, S., Eghbali, A., Tip, F.: An empirical evaluation of using large language models for automated unit test generation. *IEEE Transactions on Software Engineering* (2023)
31. Shi, F., Suzgun, M., Freitag, M., Wang, X., Srivats, S., Vosoughi, S., Chung, H.W., Tay, Y., Ruder, S., Zhou, D., et al.: Language models are multilingual chain-of-thought reasoners. arXiv preprint arXiv:2210.03057 (2022)
32. for Standardization (ISO), I.O.: Iso 26262-1: 2018—road vehicles—functional safety. Geneva, Switzerland (2018)
33. Tambon, F., Dakhel, A.M., Nikanjam, A., Khomh, F., Desmarais, M.C., Antoniol, G.: Bugs in large language models generated code. arXiv preprint arXiv:2403.08937 (2024)
34. Ung, G., Amilon, J., Gurov, D., Lidström, C., Nyberg, M., Palmkog, K.: Post-hoc formal verification of automotive software with informal requirements: An experience report. In: Accepted at 2024 IEEE 32nd International Requirements Engineering Conference (RE). IEEE (2024)
35. Vaithilingam, P., Zhang, T., Glassman, E.L.: Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In: Chi conference on human factors in computing systems extended abstracts. pp. 1–7 (2022)

36. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
37. Wang, B., Min, S., Deng, X., Shen, J., Wu, Y., Zettlemoyer, L., Sun, H.: Towards understanding chain-of-thought prompting: An empirical study of what matters. *arXiv preprint arXiv:2212.10001* (2022)
38. Wei, X., Cui, X., Cheng, N., Wang, X., Zhang, X., Huang, S., Xie, P., Xu, J., Chen, Y., Zhang, M., et al.: Zero-shot information extraction via chatting with chatgpt. *arXiv preprint arXiv:2302.10205* (2023)
39. Wikipedia contributors: Assertion (software development) — Wikipedia, the free encyclopedia. [https://en.wikipedia.org/w/index.php?title=Assertion\\_\(software\\_development\)&oldid=1179241560](https://en.wikipedia.org/w/index.php?title=Assertion_(software_development)&oldid=1179241560) (2023), [Online; accessed 27-May-2024]
40. Yuksekgonul, M., Bianchi, F., Boen, J., Liu, S., Huang, Z., Guestrin, C., Zou, J.: Textgrad: Automatic” differentiation” via text. *arXiv preprint arXiv:2406.07496* (2024)
41. Zhong, L., Wang, Z.: Can llm replace stack overflow? a study on robustness and reliability of large language model code generation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 38, pp. 21841–21849 (2024)

## Appendix A

### Prompt Template for HLNL for Specification Type

**System Prompt:** You are an experienced verification engineer with expertise in safe embedded C programming and writing ANSI/ISO C Specification Language (ACSL) specifications for safety-critical systems. Your role is to analyze C programs with accompanying ACSL and natural language specifications. Produce complete C functions that satisfy the given specifications by following these guidelines:

- Do not alter the provided specifications.
- Do not explain or comment on the code you produce.
- Do not modify any provided header files.

When given a task, focus only on implementing the required C function that meet the specifications. Prioritize safety and correctness in your implementations, ensuring that your code not only meets the given specifications but also adheres to best practices for safety-critical systems by following the 10 Rules for Developing Safety-Critical Code.

- Avoid complex flow constructs, such as goto and recursion.
- All loops must have fixed bounds.
- Do not use dynamic memory allocation after initialization.
- Restrict size of function to around 60 LOC.
- (Use a minimum of two runtime assertions per function.)
- Restrict the scope of data to the smallest possible.
- Check the return value of all non-void functions, or cast to void to indicate the return value is useless.
- Use the preprocessor sparingly.
- Limit pointer use to a single dereference, and do not use function pointers.
- Compile with all possible warnings active; all warnings should then be addressed before release of the software.

**User Prompt:** Generate the C code for a function that implements the following high-level specification.

- If there is a brake light request, then the truck and trailer lights shall be activated.

```
//Header
typedef unsigned char    tB;
...
```

```

typedef struct {
    tB val;
    tU08 ss_U08;
} tBS;
...

#define TRUE                1
#define FALSE              0
...

/*@
    assigns *outSig;
    ensures \valid(outSig);
    ...
*/
extern void validateInputBool(const tBS * const inSig,
                             tBS * const outSig,
                             tB defaultNotExist,
                             tB defaultNotGood);

//Input variables
static tU08S rtdb_state;
static tBS rtdb_req;
...

//Output variables
static tBS rtdb_truck;
...

//Concrete variables
tU08 gh_operationalState;
tU08 gh_supplyVoltageLevel;
tU08 gh_brakeLightEnabled;
...

//Function
void Brak_10ms(void);

```