

Article



DeepSeek-V3, GPT-4, Phi-4, and LLaMA-3.3 Generate Correct Code for LoRaWAN-Related Engineering Tasks

Daniel Fernandes ^{1,*}, João P. Matos-Carvalho ^{2,3}, Carlos M. Fernandes ^{1,3}, and Nuno Fachada ^{1,3}

- ¹ Copelabs, Lusófona University, 1749-024 Lisbon, Portugal; p7582@ulusofona.pt (C.M.F.); nuno.fachada@ulusofona.pt (N.F.)
- ² LASIGE, Departamento de Informática, Faculdade de Ciências, Universidade de Lisboa, 1749-016 Lisbon, Portugal; jpecarvalho@ciencias.ulisboa.pt
- ³ Center of Technology and Systems (UNINOVA-CTS) and Associated Lab of Intelligent Systems (LASI), 2829-516 Caparica, Portugal
- * Correspondence: daniel.fernandes@ulusofona.pt

Abstract: This paper investigates the performance of 16 Large Language Models (LLMs) in automating LoRaWAN-related engineering tasks involving optimal placement of drones and received power calculation under progressively complex zero-shot, natural language prompts. The primary research question is whether lightweight, locally executed LLMs can generate correct Python code for these tasks. To assess this, we compared locally run models against state-of-the-art alternatives, such as GPT-4 and DeepSeek-V3, which served as reference points. By extracting and executing the Python functions generated by each model, we evaluated their outputs on a zero-to-five scale. Results show that while DeepSeek-V3 and GPT-4 consistently provided accurate solutions, certain smaller models—particularly Phi-4 and LLaMA-3.3—also demonstrated strong performance, underscoring the viability of lightweight alternatives. Other models exhibited errors stemming from incomplete understanding or syntactic issues. These findings illustrate the potential of LLM-based approaches for specialized engineering applications while highlighting the need for careful model selection, rigorous prompt design, and targeted domain fine-tuning to achieve reliable outcomes.

Keywords: LoRaWAN; large language models; UAV placement; code generation; IoT

1. Introduction

The rapid expansion of Internet of Things (IoT) applications has led to increased attention to Low-Power Wide-Area Network (LPWAN) technologies, such as LoRa Wide Area Network (LoRaWAN), which provide long-range communication with low power consumption [1]. LoRaWAN networks are particularly appealing for rural areas, where infrastructure constraints can pose significant challenges to traditional wireless communication systems [2]. In this context, the integration of Unmanned Aerial Vehicles (UAVs) as mobile relays has emerged as a promising solution, enabling flexible deployments and extended coverage [3]. Determining the UAV position that minimizes signal propagation loss and assessing the corresponding received power are critical for ensuring reliable connectivity and resource-efficient operations in these rural scenarios [4].

Parallel to these developments in wireless communications, Large Language Models (LLMs) have shown rapid progress. Modern LLMs—including GPT-4 [5], recent open-source offerings locally installable with Ollama [6], and novel models such as DeepSeek [7]—have shown substantial capabilities in understanding complex tasks and



Academic Editors: Arkaitz Zubiaga, Panayiotis Bozanis and Leonidas Akritidis

Received: 19 February 2025 Revised: 26 March 2025 Accepted: 28 March 2025 Published: 1 April 2025

Citation: Fernandes, D.; Matos-Carvalho, J.P.; Fernandes, C.M.; Fachada, N. DeepSeek-V3, GPT-4, Phi-4, and LLaMA-3.3 Generate Correct Code for LoRaWAN-Related Engineering Tasks. *Electronics* 2025, 14, 1428. https://doi.org/10.3390/ electronics14071428

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/ licenses/by/4.0/). generating functional code for engineering problems [5]. Furthermore, these models demonstrate a broad applicability beyond code generation, including text clustering [8], text summarization [9], machine translation [10], and text classification/question answering [11]. However, despite these advancements, the effectiveness of lightweight, locally executed models in generating correct and efficient solutions for domain-specific engineering tasks remains an open question [12].

This study investigates whether lightweight and locally executed LLMs can generate correct Python code for UAV planning tasks in LoRaWAN environments. Specifically, we assess 16 different LLMs by evaluating their ability to generate Python functions that determine the optimal UAV position from a discrete set of candidate locations, minimizing propagation loss, and computing the corresponding received power (in dBm). Our primary goal is to compare the performance of locally run models, such as LLaMA-3.3 [13] and Phi-4 [14], against state-of-the-art large models such as GPT-4 [5] and DeepSeek-V3 [7], accessed via their online application programming interfaces (APIs). The inclusion of these larger models serves as a reference point to establish that such tasks can indeed be solved using advanced LLMs, allowing for a meaningful comparison with the performance of smaller, locally executed alternatives. The evaluation uses a zero-shot natural language prompt configuration, and correctness is measured through a scoring system based on function extraction and execution results.

Despite significant progress in AI-assisted UAV deployment, previous research has largely overlooked the unique communication and operational constraints inherent to Lo-RaWAN environments. LoRaWAN deployments pose distinct challenges such as stringent power limitations, specialized propagation characteristics at lower frequencies, and longrange communication requirements that differ fundamentally from scenarios commonly studied in existing UAV-AI literature. Existing approaches primarily focus on UAV trajectory planning, mission coordination, or visual scene understanding tasks, without explicitly addressing scenarios involving the low-power, wide-area network constraints and signal propagation peculiarities of LoRaWAN systems. This gap motivates our study, which specifically examines whether LLMs—particularly lightweight, locally executable variants—can effectively generate Python code to solve UAV placement and received power calculation tasks uniquely relevant to LoRaWAN environments.

The findings of this study are significant for two main reasons. First, they illustrate the extent to which lightweight, locally run LLMs can perform domain-specific engineering tasks, providing insight into their potential as cost-effective alternatives to proprietary, large-scale models [15]. Second, these findings may offer practical guidance not only for practitioners integrating LLM-generated code into IoT and UAV communication workflows but also for those in a wide range of other fields, as they highlight critical considerations such as reliability, correctness, and maintainability. The subsequent sections of this paper are organized as follows. Section 2 provides background information on the use of LLMs for human–UAV interaction and code generation, also discussing relevant aspects of prompt design. Section 3 describes the materials and methods employed, including the engineering problem context, prompt structure, model selection, and evaluation metrics. Results are presented in Section 4, followed by a detailed discussion in Section 5. Section 7 concludes the study's limitations and opportunities for future research. Finally, Section 7 concludes the paper with final remarks and recommendations.

2. Background

In this section, we start by addressing the general goal of integrating LLMs with UAVs to improve the behavior, organization, and communication of autonomous systems, as well as the specific implementation of UAVs as mobile relays and antennas in LoRoWAN envi-

ronments. In Section 2.2, we focus on the specific task of generating code for autonomous devices and on how LLMs are being used to incorporate code generation at different levels of workflow. Finally, in Section 2.3, we briefly discuss prompt engineering and its principles, the benefits and drawbacks of conversational and structured prompting, and how prompt design impacts code generation or task planning.

2.1. LLMs for Human–UAV Interaction

The nature of UAVs, namely their collective organization and communication requirements, strongly encourages integration with Artificial Intelligence (AI) algorithms. The recent emergence of LLM technologies in particular is inspiring new frameworks and prototypes for communication and design of several autonomous systems, and UAVs are no exception. As LLMs learning and adaptation capabilities in uncertain and dynamic environments grow and approach human-level proficiency, the scientific literature on the subject steadily increases [16,17]. Currently, there is a significant amount of knowledge on LLMs for human–UAV interaction. For a review on the state-of-the art literature on LLMs and UAVs, please refer to [16]. For a discussion of key areas where LLMs can impact UAVs, we urge the reader to refer to the paper by Phadke et al. [17]. In the following paragraphs, we discuss some recent developments on the usage of natural language models for controlling UAVs.

In [18], Aikins et al. present LEVIOSA, a framework for the generation of UAV trajectory based on text and speech. The authors use several LLMs to convert natural language prompts into sets of coordinates to guide the UAVs and low-level controllers to control each device in its path, aiming for accuracy, synchronization, and collision avoidance. LEVIOSA was tested on various scenarios with promising results.

Cui et al. [19] propose a Task Planning for Multi-UAV System (TPML) that uses LLMs as interfaces to translate UAV's operator instructions into executable codes. After validating the system in simulation environments and real-wold scenarios, the authors argue that TPML is able to control multiple UAVs in both synchronous and asynchronous missions with a single natural language input.

While most of the studies on natural language processing for UAVs focus on processing the user messages to program or optimize UAV behavior, others try to provide UVAs with scene descriptions skills in natural language, taking advantage of their capacity to acquire visual cues of the environment. In [20], the authors use LLMs and Visual Language Models (VLMs) to provide UAVs with the ability of scene description using natural language. The generated tests were subject to a readability test, some achieving a high school senior reading level (level 12 in the Gunning fog index).

In [21], the authors discuss a framework that integrates a novel factorization method—QTRAN—in a multi-agent reinforcement learning algorithm (MARL) [22] with an LLM to optimize UAV trajectories, overcoming limitations of value decomposition algorithms for trajectory planning, as they have difficulties in associating local observations with the global state of UAV swarms. Although QTRAN overcomes some of the limitations of standard MARLs, its performance can still be improved, namely by enhancing the representation network. For that purpose, the authors incorporate LLMs in the framework, boosting its overall performance in trajectory optimization and outperforming other reinforcement learning methods.

LPWAN-based systems are one of the emerging technologies in which UAVs are being tested and deployed. LPWANS, and LoROWANs in particular, rely on a set of fixed sensor stations, which measure and transmit a number of environmental data to a central unit. Traditionally, these stations are static, cover only very small areas and can be impaired by natural disasters. Due to their mobility, UAVs can act as moving communication nodes, which solves some of the limitations of static LoROWANs.

Several methods have been proposed to integrate UAVs in LoROWANs. In [23], UAVs are used to transfer information from ground-based LORAWAN nodes to the base station. The architecture of the systems thus consists of two layers, the first being the ground nodes that transmit data using LoRaWAN and the second the swarm of drones communicating over a WiFi ad hoc network. To enhance the performance of the systems, a distributed topology algorithm periodically adapts the UAV topology to the position of the ground nodes. In [24], the authors describe an air quality monitor system based on a LORAWAN and UAVs. In [25], a UAV emergency monitoring system using a LORAWAN is proposed to overcome the limitations of ground stations in disaster scenarios. Finally, Arroyo et al. [26] propose a UAV and LOROWAN system that enables data transfer from sensors to a central system and then use machine learning to classify the data. To the extent of our knowledge, there are no studies on the integration of LLMs and UAVs in a LoRaWAN environment.

2.2. Code Generation with LLMs

The landscape of AI-assisted programming has evolved significantly, with extensive research focusing on natural language generation and understanding of large codebases [27]. Shortly after their inception, some LLMs demonstrated capabilities in code assistance and code generation, even from natural language specifications. In the first models, those skills were somewhat limited and the output often required post-processing steps to improve the quality of the suggested code [28]. But LLMs quickly evolved, and their ability to provide executable code in due time improved significantly [29]. Furthermore, derivations of popular LLMs, like Open AI Codex [30], a descendant of ChatGPT-3, and Code Llama [13], Meta's programming tool, emerged as specialized models for coding. Nowadays, AI-assisted programming is a common practice in industry.

In the context of code generation for autonomous devices, Vemprala et al. [31] explore ChatGPT's ability on several robot-oriented tasks, including code synthesis. The authors present a framework for robot control that requires designing and implementing a library of APIs receptive to prompt engineering for ChatGPT. The proposed framework allows the generated code to be tested, verified, and validated by a user through simulation and manual inspection.

In [32], the authors adapt LLMs trained on code completion for writing robot policy code according to natural language prompts. The generated robot policies exhibit spatial-geometric reasoning and are able to prescribe precise values to ambiguous descriptions. By relying on a hierarchical prompting strategy, their approach is able to write more complex code and solve 39.8% of the problems on the HumanEval [30] benchmark.

Luo et al. [33] use LLMs to generate robot control programs, testing and optimizing the output in a simulation environment. After a number of optimization rounds, the robot control codes are deployed on a real robot for construction assembly tasks. The experiments show that their approach can improve the quality of the generated code, thus simplifying the robot control process and facilitating the automation of construction tasks.

2.3. Prompt Design

The piece of text or set of instructions that the user provides to an LLM to generate a specific response is called a prompt. Designing effective prompts is essential to take advantage of the potential of LLMs, and in a few years the craft established as a field of research and development of its own [34].

Prompting strategies can be broadly classified into structured and unstructured approaches. Structured prompting employs precise instructions with explicitly defined inputs, outputs, and constraints, often leading to more reliable and accurate code generation. However, structured prompts typically require a deeper understanding of both the problem domain and the underlying model, potentially limiting flexibility and accessibility. Conversely, unstructured prompting uses intuitive, conversational language, making it accessible to a broader audience, reflecting realistic scenarios where users may not possess specialized knowledge of prompt crafting. However, this can result in less consistent outputs due to inherent ambiguity.

Prompts may also be categorized based on the number of illustrative examples provided: zero-shot prompts provide no examples, one-shot prompts include a single example, and few-shot prompts incorporate multiple examples. Empirical research supports the trade-offs associated with different prompt styles; for instance, Liang et al. [32] demonstrate that structured, code-based prompts generally yield superior results for robot-related reasoning tasks compared to natural language prompts. However, advances in LLM technology continue to improve the viability of unstructured, natural language prompting in complex domains such as robotics [31]. Further improvements in output coherence have also been observed through structured reasoning techniques such as chain-of-thought (CoT) prompting [33,35].

In this study, we follow a natural language zero-shot prompt strategy, in which the request is performed in a relatively unstructured fashion without any examples. Nonetheless, established best practices for engineering-focused code generation were followed by explicitly specifying function inputs, expected return types, and required libraries, thus improving the clarity and reproducibility of the generated code [36].

3. Materials and Methods

This section starts with an overview of the theoretical context that informs our prompt design in Section 3.1. Next, Section 3.2 presents the proposed prompts and their respective scenarios. Section 3.3 describes and justifies the models analyzed in this study. Section 3.4 then outlines the prompting and response processing pipeline. The section concludes with a description of the experimental setup in Section 3.5, including all tested inputs for both the LLMs and the generated Python functions, the expected function results, and the evaluation metrics used.

3.1. Theoretical Context

The IoT paradigm refers to the interconnection of physical devices that collect, exchange, and process data over the Internet or other communication networks. According to Sanguesa et al. [37], it is estimated that by 2030, there will be approximately 125 billion IoT devices, ranging from simple temperature and humidity sensors to more complex sensors used in sectors such as agriculture and industry. The main goal of these sensors is to simplify and optimize daily activities. One of the challenges associated with this paradigm is the large volume of data generated and how it is processed. A potential solution for data collection is the use of UAVs, which can fly over (or carry) multiple sensors along a predefined path planning. These UAVs may or may not be capable of transmitting data in real time to a base station (BS). However, to use UAVs efficiently, it is often necessary to calculate their location and send control commands to adjust their position or even modify their flight path. Therefore, reliable communication between the UAV and a base station is crucial. One possible communication protocol for this purpose is LoRaWAN, which is based on LoRa (long-range) communication and enables effective long-distance data transmission [38,39]. Essentially, LoRa communication establishes a link between two points: the transmitter—in this case, the BS—and the receiver, i.e., the UAV. This communication is based on classical propagation models, such as those found in reference [40].

Regarding the modulation of a communication channel, the received power at the antenna (p_r) depends on factors such as the transmit power (p_t) , the gain of the antennas $(g_r \text{ and } g_t)$, the distance between the antennas (r), and the losses during transmission (free-space attenuation). Equation (1) represents the propagation loss l_F between the two points:

$$l_F = \frac{p_t \cdot g_r \cdot g_t}{p_r} = \left(\frac{4\pi r}{\lambda}\right)^2 = \left(\frac{4\pi r f}{c}\right)^2 \tag{1}$$

where λ represents the wavelength. In particular, $\lambda = \frac{c}{f}$, with *c* representing the speed of light and *f* the frequency, which in Europe is 868 MHz.

A lower propagation loss results in a stronger received signal. Propagation losses are typically expressed in dB units, and for a distance in meters and a frequency in Hz, Equation (1) can be rewritten as Equation (2), which represents the Free Space Path Loss formula. This formula is valid under free-space conditions, assuming a direct, unobstructed line of sight. In terms of notation, lowercase variables denote linear values, whereas uppercase variables denote logarithmic values.

$$L_F(dB) = 20\log(r_m) + 20\log(f_{Hz}) - 147.55$$
(2)

To estimate the received power, it is necessary to consider the transmitted power, the gain of the transmitting and receiving antennas, and the path losses that occur during transmission. Thus, Equation (3), derived from Equation (1), can be written as

$$P_{r(dBm)} = P_t + G_t + G_r - L_F \tag{3}$$

3.2. Scenarios and Prompts

To evaluate the LLM models, three zero-shot prompts with increasing levels of difficulty were designed—see Table 1. In this context, 'zero-shot' refers to prompts that do not provide any examples to the model being tested. Furthermore, these prompts use natural language, meaning that they are relatively unstructured and have undergone minimal refinement, apart from ensuring technical precision and clarity. This approach was chosen as it more closely follows real-world scenarios where domain experts may rely on direct, straightforward queries to achieve their goals.

The specific request posed by these prompts is for the LLM to identify, from a set of points, the point where the value of L_F is the lowest or to determine the received power at that point (i.e., the point with the lowest L_F). In all scenarios, a frequency of 868 MHz is considered, as well as a rural area where LoRa communication is possible up to 10 km. Both antennas are assumed to have a gain of 2.5 dBi each.

To simplify post-processing of responses, all prompts specify the available libraries, the expected indentation type, and that the return function should be self-contained—i.e., all required code including constants and auxiliary functions should be defined within the requested function.

Table 1. Prompts designed for this study, requiring the tested LLMs to generate Python functions that solve increasingly complex tasks related to LoRaWAN and UAVs.

Prompt 1

Consider that the LoRaWAN communication protocol is being used in a rural scenario where a base station communicates with a UAV at a communication frequency of 868 MHz. Assume a system with two axes (the x-axis and the y-axis) and that the base station is in position (0,0). Also, assume that all positions are in kilometers (km).

Create a Python function called `index_position()` which accepts a list of tuples, with each (x, y) tuple representing a possible position in which the UAV can be placed with respect to the base station. This function should return the list index of the tuple (i.e., UAV position) which minimizes the propagation loss. Assume that the math and numpy libraries are imported as follows, and no more libraries can be used:

import math
import numpy as np

Beyond importing these libraries, the `index_position()` function must be self-contained. In other words, all variables, constants, or helper functions must be defined within the `index_position()` function. Provide Python code with 4-space indentation following PEP 8.

Prompt 2

Consider that the LoRaWAN communication protocol is being used in a rural scenario where a base station communicates with a UAV at a communication frequency of 868 MHz. Assume a system with two axes (the latitude axis and the longitude axis) where each value is given in decimal degrees.

Create a Python function called `index_position()` which accepts a list of (latitude, longitude) tuples. The first tuple in this list represents the position of the base station, while the remaining tuples represent possible positions in which the UAV can be placed. This function should return the list index of the tuple which minimizes the propagation loss. Assume that the math and numpy libraries are imported as follows, and no more libraries can be used:

import math
import numpy as np

Beyond importing these libraries, the `index_position()` function must be self-contained. In other words, all variables, constants, or helper functions must be defined within the `index_position()` function. Provide Python code with 4-space indentation following PEP 8.

Prompt 3

Consider that the LoRaWAN communication protocol is being used in a rural scenario where a base station communicates with a UAV at a communication frequency of 868 MHz, with a transmission power of 27 dBm. Both the transmitter and UAV antennas have a gain of 2.5 dBi. Assume a system with two axes (the latitude axis and the longitude axis) where each value is given in decimal degrees.

Create a Python function called `power_received()` which accepts a list of (latitude, longitude) tuples. The first tuple in this list represents the position of the base station, while the remaining tuples represent possible positions in which the UAV can be placed. This function should return the power received (in dBm) by the UAV at the position that minimizes the propagation loss. Assume that the math and numpy libraries are imported as follows, and no more libraries can be used:

import math
import numpy as np

Beyond importing these libraries, the `power_received()` function must be self-contained. In other words, all variables, constants, or helper functions must be defined within the `power_received()` function. Provide Python code with 4-space indentation following PEP 8.

The first prompt is presented in the first row of Table 1. In this simpler scenario, the BS and the UAV's possible positions, measured in kilometers (km), are defined within a coordinate system with two axes: the *x*-axis and the *y*-axis. The BS is fixed at position (0,0), while the UAV's possible positions are provided as an input array to the function generated by the LLMs. To solve this problem, LLMs must generate a Python function that calculates the distance (e.g., Euclidean) between the BS and each possible UAV position, applies Equation (2) to compute power losses, and returns the index of the position with the lowest loss. The LLM must ensure that power losses maintain a one-to-one correspondence with the UAV positions to return the correct index.

Prompt 2, shown in the second row of Table 1, increases the complexity by considering geographical coordinates—latitude and longitude—instead of a simple (x, y) axis. LLMs must use a different method to calculate the distances between the UAV's position and the BS, such as Haversine's formula. This prompt further increases the difficulty by requiring that the UAV's position be given as the first element of the input array. Consequently, the generated functions must extract this information and return an index greater than zero, as index zero contains the UAV's position.

Prompt 3, presented in the last row of Table 1, closely resembles Prompt 2. However, instead of returning the index with the lowest loss, the generated function must return the value of that loss by applying Equation (3).

3.3. LLMs Considered

The LLMs models used in this paper were chosen based on their impact in AI research, innovative approaches, and performance across different domains such as programming, advanced reasoning, and computational efficiency. Table 2 lists and characterizes the LLMs selected for this study. For the remainder of this paper, the number of parameters associated with each model is expressed in billions or trillions with an uppercase B and T, respectively.

Table 2. Characteristics and main purpose of the LLMs tested in this study. 'Size' indicates the number of parameters in billions (B) or trillions (T). 'Tag' corresponds to the specific model version invoked in the respective API calls.

Family	Version	Size	Тад	Main Purpose
DeepSeek [7,41]	R1 R1	7B 70B	deepseek-r1:7b	Computationally efficient distilled reasoning model.
	KI	700	deepseek 11.700	computational efficiency.
	V3	671B	deepseek-v3	Mixture-of-Experts general-purpose model.
Gemma [42,43]	1.1	2B	gemma:2b	Lightweight model for dialogue, instruction-following and coding.
	2.0	2B	gemma2:2b	Compact general-purpose model trained with knowledge distillation.
GPT [5]	4	1.76T *	gpt-4-0613	Multimodal model optimized by OpenAI for text, audio, and image processing.
LLaMA [13]	3.2	3B	llama3.2:3b	Lightweight text-only model for multilingual dialogue and text summarization.
	3.3	70B	llama3.3:70b	Text-only model for deeper comprehension multilingual conversation.
	code	7B	codellama:7b	Code generation model.
Mistral [44]	0.3	7B	mistral:7b	Efficient model for text and code generation, supports function calling.
Phi [14]	4.0	14B	phi4:14b	Reasoning model trained using high-quality synthetic data.
Qwen [45,46]	2.5-coder	0.5B	qwen2.5-coder:0.5b	Code generation model.
	2.5-coder	1.5B	qwen2.5-coder:1.5b	Code generation model.
	2.5-coder	3B	qwen2.5-coder:3b	Code generation model.
	2.5	0.5B	qwen2.5:0.5b	General-purpose language model.
	qwq	32B	qwq:32b	Advanced reasoning model for complex problem-solving tasks.

* Unofficial estimate.

The DeepSeek family of models includes a range of architectures designed to balance performance and computational efficiency. DeepSeek-R1 (7B) and DeepSeek-R1 (70B) are distilled versions derived from the larger DeepSeek-R1 model (671B)—based on the Qwen and LLaMA architectures—to retain significant reasoning capabilities while reducing hardware demands [41]. In contrast, DeepSeek-V3 (671B) is a Mixture-of-Experts model designed to perform well in diverse tasks [7]. Considering these models is crucial due to their varied architectures and training methodologies, which offer insights into the trade-offs between model size, training techniques, and task-specific performance. The V3 671B model was selected over its more developed R1 counterpart, as initial trials demonstrated it was sufficiently accurate for the prompts presented in Section 3.2, providing a balance between performance and cost.

The Gemma model family [42,43], developed by Google DeepMind, comprises open models derived from the research and technology behind the Gemini models. While influenced by Gemini, Gemma is fully open-source and designed for efficient language understanding and reasoning. The lightweight Gemma v1.1 (2B) and Gemma2 (2B) implementations are optimized for resource-limited environments. Gemma2 (2B) incorporates knowledge distillation, improving efficiency and performance relative to its size. These models were included to assess the trade-offs in model scaling, particularly for the real-time and cost-sensitive applications associated with the tested prompts.

OpenAI's Generative Pre-trained Transformer (GPT) models are proprietary LLMs designed to understand and generate human-like text, facilitating tasks such as drafting documents, coding, and responding to queries [5]. Their popularity and advanced capabilities make them essential subjects in LLM comparison studies. In this context, GPT-4-0613 was selected over newer models such as GPT-40 and 01, as preliminary tests indicated its performance was sufficient for the presented prompts, therefore reducing costs.

The LLaMA series by Meta AI includes models optimized for various applications [13]. LLaMA-3.2 (3B) is a lightweight, multilingual model suited for mobile and edge devices, appropriate for text summarization and classification. LLaMA-3.3 (70B) is a larger, instruction-tuned model with superior performance in natural conversation and multilingual tasks. Code Llama (7B) specializes in code generation and understanding. Testing these three models is important for evaluating how model size, specialization, and efficiency in the LLaMA family impacts performance across the three implemented prompts.

The Mistral family of language models [44], developed by the French company Mistral AI, stands out for its efficient architecture and strong performance. Mistral models achieve high accuracy with fewer parameters, making them more accessible and computationally efficient compared to many large-scale models. The Mistral v0.3 (7B) model exemplifies this approach, demonstrating capabilities in text and code generation, conversation, and function calling, while effectively handling longer sequences. Its open-source nature offers a valuable option for research and application development, providing a European alternative to models predominantly from U.S.- and China-based companies.

The Phi model family [14], developed by Microsoft Research, is focused on the role of high-quality synthetic data for improving reasoning in compact language models. Phi-4, a 14-billion parameter model, prioritizes synthetic data to improve problem-solving in mathematics and coding, outperforming its teacher model, GPT-4, on several benchmarks. Unlike models that primarily scale with size, Phi-4 follows a distinct training approach, making it important to compare against other LLMs. Its relatively small size also makes it relevant for low-resource environments, where optimizing data efficiency can be a crucial factor in model deployment.

The Qwen model family, developed by Alibaba Cloud, includes general-purpose [45] and code-specialized [46] LLMs over a wide range of sizes. Their scalability, architectural optimizations, and strong reasoning capabilities make them valuable for benchmarking efficiency and specialization. Here, the most recent 2.5 versions are tested—namely the specialized coder implementations (0.5B, 1.5B, and 3B) and the general-purpose 0.5B model—as well as QwQ (Qwen with Questions) 32B model with advanced reasoning capabilities.

3.4. Implementation

The pipeline for submitting a prompt to an LLM, obtaining a response, extracting a Python function, and executing it is illustrated in Figure 1. The process begins by iterating through a predefined set of LLMs, seeds, temperatures, and prompts. Each prompt is submitted to the corresponding LLM, and its response is stored in a text file. Next, the function from each stored response is extracted by searching for the function definition (e.g., 'def requested_function():') and capturing all internal code up to the last properly indented 'return' statement. This ensures that functions defined within the external function do not prematurely terminate the extraction. The extracted function is then recorded in a Python file for execution. If the function is not successfully extracted—such as when the defined function name does not match the expected one—this information is logged in the results file, and a score of zero is assigned for that LLM, seed, temperature, and prompt combination.



Figure 1. Validation pipeline for the results of LLMs under study.

If the Python function is correctly generated and extracted, it is tested under Python 3.9.6 using the data provided for each scenario (presented in Section 3.5). One of three possible outcomes may occur:

- The code contains a syntax error and does not compile, in which case a score of 1 is recorded in the results file;
- The code executes but encounters a runtime error, resulting in an exception, in which case a score of 2 is stored in the results file;

• The code executes successfully and returns a result, in which case the score ranges from 3 to 5, as detailed below.

If the code executes successfully, the function's output is evaluated as follows: if the returned value is of a different type than expected (e.g., a float instead of an int), a score of 3 is recorded in the results file. This type check is performed broadly; for example, if an integer is expected, types such as int, np.int32, or np.int64 are considered valid (where np refers to the NumPy library). If the type is correct, the next step is to verify whether the returned value matches the expected value. For floating-point comparisons, a tolerance of 1% is allowed. If the result is incorrect, a score of 4 is assigned. Finally, if the returned value is correct, a score of 5 is recorded, indicating 100% functionally correct code. At the end of this process, a file containing all recorded scores is available for analysis.

In summary, scores between 0 and 5 are characterized as follows:

- 0. No Python file was generated—This indicates that the LLM did not generate a Python function or that the generated function does not have the name specified in the prompt.
- 1. Syntax error—The code does not compile.
- 2. Runtime error—The code is valid Python but has logic incongruencies and/or does not conform to the prompt requirements.
- 3. Code runs but returns an incorrect data type—For Prompts 1 and 2, it should return an integer (the index value), while in Prompt 3, it should return a float.
- 4. Code runs but returns an incorrect result.
- 5. Code runs and returns the correct result.

3.5. Experimental Setup

To thoroughly test the capabilities of the models listed in Section 3.3, the prompts presented in Section 3.2 were individually submitted to LLMs using six different pseudorandom number generator seeds across six temperature values, in a total of 36 submissions per prompt for each LLM. Temperatures were increased in 0.2 increments from 0.0 to 1.0 for locally executed LLMs via Ollama. Although Ollama accepts temperatures in the range of 0.0–1.0, both DeepSeek-V3 and GPT-4, executed through their online APIs, accept temperatures in the 0.0–2.0 range. Therefore, temperatures were doubled for these models. For example, and for the purpose of this study, a temperature of 0.6 in local models is doubled to 1.2 when submitting a prompt to online LLMs.

The LLM-generated Python functions were tested with the following input data, and return values for each prompt were expected:

- **Prompt 1** The input data are an array of four positions, namely [(2,5), (7,7), (1,8), (1,0.5)]. The expected return value is 3, corresponding to coordinate (1,0.5), which is the closest one to the BS, which is fixed at (0,0).
- **Prompt 2** The input data are an array containing the following coordinates:

BS / UAV	\rightarrow	(38.759297963817374, -9.154483012234662)
Position 1	\rightarrow	(38.749330295687805, -9.15304293547367)
Position 2	\rightarrow	(38.75727072916799, -9.157797377555926)
Position 3	\rightarrow	(38.737648166512336, -9.138660615310467)
Position 4	\rightarrow	(38.76841010033327, -9.160013961052972)

The expected return value is 2, corresponding to the index of Position 2, which minimizes the power loss.

Prompt 3 The input data are the same as in Prompt 2, but the expected value is -50.33 dBm, which is the minimal loss, obtained at Position 2.

As described in Section 3.4, the capabilities of the different LLMs in correctly answering Prompts 1–3 are assessed using a score between 0 and 5. For six submissions (one per seed) for each prompt–model–temperature combination, four summary statistics are calculated and presented: the mean score, a non-parametric 95% confidence interval around the mean, the percentage of perfect scores (score equal to 5), and a histogram of score distribution. These metrics allow for a detailed performance investigation of the capabilities of the 16 tested models to generate Python code to solve the three progressively complex LoRaWAN-related prompts.

In addition to these summary statistics, a formal statistical comparison between models is conducted using stratified permutation tests [47]. To account for varying prompt difficulty, model performance is stratified by prompt, allowing all three prompts to be included in a unified testing procedure. For each pairwise comparison between two models at a given temperature, scores are pooled by prompt (six scores per model per prompt, 12 in total), and a one-sided permutation test is applied. The test statistic is the sum of mean rank differences across prompts. All $\binom{12}{6} = 924$ possible permutations of model labels are precomputed per prompt, and 1000 stratified permutations are generated by randomly selecting one permutation per prompt and combining them. The resulting null distribution is used to estimate the probability of obtaining a test statistic as large or larger than the observed one under the null hypothesis of no difference. The tests are one-sided, since the goal is to determine whether one model significantly outperforms another—not whether it is worse. Finally, multiple testing correction is applied using the Benjamini–Hochberg procedure to control the false discovery rate (FDR) across all comparisons [48].

4. Results

Results for the simpler Prompt 1 are shown in Figure 2 and Table 3. While all models generated accurate code for certain seed/temperature combinations, DeepSeek-V3 and Phi-4 stood out, consistently providing correct answers across all seeds and temperatures. The three LLaMA models, the three Qwen coder models, and GPT-4 also demonstrated strong performance, reliably generating correct code for at least a subset of temperature values—typically at lower settings. Interestingly, GPT-4 exhibited a significant drop in answer quality at temperatures of 1.6 and higher (i.e., 2×0.8), with responses becoming essentially random at the highest temperature. In contrast, the DeepSeek-R1 models (7B and 70B), the Gemma models (2B), the Mistral model (7B), and the non-coder Qwen models (2.5–0.5B and QwQ-32B) failed to consistently produce correct answers.

Results for the slightly more complex Prompt 2, for which the UAV position is given as a function argument (i.e., it is not predefined within the function) and actual geographical coordinates are used, are shown in Figure 3 and Table 4. Only four models consistently generated accurate code: the larger online DeepSeek-V3 and GPT-4 models, as well as the smaller, locally tested LLaMA-3.3 and Phi-4. However, the drop in performance for GPT-4 at higher temperatures is even more pronounced for this prompt. Conversely, Gemma (2B), Mistral (7B), and both 0.5B Qwen models failed to produce a single correct answer.

Prompt 3, while similar to Prompt 2 in many respects, requires the requested function to return a concrete power loss value rather than merely the index of the position with the lowest loss. This distinction arguably makes it the most complex task for the models evaluated in this study. The results for this prompt are presented in Figure 4 and Table 5. The same four models continued to generate accurate code consistently, though within a more limited range of temperature settings. DeepSeek-V3 demonstrated the highest overall consistency, reliably producing correct code at temperatures of 0.8 (2×0.4) and $1.2 (2 \times 0.6)$, while maintaining a high percentage of accurate responses across the remaining temperatures. GPT-4 and Phi-4 achieved 100% accuracy when the temperature was set

to zero. However, while Phi-4 remained highly consistent at higher temperatures, GPT-4 exhibited a significant decline in performance. LLaMA-3.3 also demonstrated strong consistency, achieving 100% accuracy in all runs at temperatures of 0.2 and 0.4. None of the remaining models were able to successfully complete this task. The only exception was Qwen's QwQ (32B), which generated a single correct response at a temperature of 0.6. However, beyond this isolated instance, it predominantly produced code containing invalid syntax or runtime errors.

Figure 5 presents a pairwise significance heatmap based on *p*-values from a stratified permutation test, after FDR multiple testing correction, indicating which models (in rows) statistically outperformed others (in columns) across temperatures. Table 6 summarizes these results, showing the number of models each system significantly outperformed at each temperature, as well as the overall total across all temperatures. These results reinforce what was observed in the descriptive statistics—namely that DeepSeek-V3, GPT-4, Phi-4, and LLaMA-3.3 are the most consistent and competitive models in these engineering tasks. At nearly all temperature levels, these models significantly outperformed the majority of alternatives, with corrected *p*-values below the 0.05 threshold in a substantial number of pairwise comparisons. In particular, DeepSeek-V3, Phi-4, and LLaMA-3.3 achieved the highest number of significant wins at every temperature, while GPT-4 showed similarly strong performance at lower temperatures but exhibited a sharp decline in statistical superiority as temperature increased.

In contrast, the two DeepSeek-R1 models, as well as QwQ, registered very few significant wins at any temperature. Crucially, their only advantages were against GPT-4 at higher temperatures, where its output becomes increasingly random and unsuitable for these types of coding tasks. This further confirms their limited effectiveness, as already observed in previous results. An additional insight—less apparent in the descriptive statistics but clearly highlighted in Table 6—is the lack of correlation between model size and performance within the Qwen coder family. Specifically, the 1.5B Qwen coder model achieved the fourth highest total number of pairwise wins (47), surpassing even GPT-4 (45), while the larger 3B variant achieved roughly half as many.

Finally, Figure 6 presents the mean scores for the tested models across all three prompts, aggregating results from all seeds and temperature settings. While the initial assumption was that Prompts 1 to 3 increase in complexity, and the results thus far appear to support this hypothesis, Figure 6 provides a more comprehensive perspective. For most models, the mean score declines progressively with increasing prompt complexity, reinforcing this assumption. However, exceptions include both Gemma models and the non-coder Qwen-2.5 model, where the score reduction is not strictly monotonic. Another observation from this figure is that the highest performing models—DeepSeek-V3, GPT-4, LLaMA-3.3, and Phi-4—maintain consistent performance across prompts, with only a slight decline in mean score as complexity increases.



Figure 2. Prompt 1 mean answer score for the tested models over several temperatures. Each combination of model and temperature was tested with 6 different seeds. Error bars denote a 95% confidence interval. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Table 3. Prompt 1 answer statistics, namely the percentage of correct answers (score equal to 5) and histogram of scores (0–5) for the tested models over several temperatures. Each combination of model and temperature was tested with 6 different seeds. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Madal	Temperature									
widdel	0.0	0.2	0.4	0.6	0.8	1.0	Overall			
codellama:7b	100.0%	100.0%	50.0%	66.7%	100.0%	83.3%	83.3%			
deepseek-r1:70b	66.7% пП	33.3%	33.3% Пп	66.7% ⊓ .∏	16.7%	50.0%	44.4% пп			
deepseek-r1:7b	50.0% пп	50.0%	50.0%	66.7%	66.7% ⊓ .∏	66.7% пП	58.3%			
deepseek-v3	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%			
gemma2:2b	0.0%	16.7%	66.7%	33.3%	16.7%	16.7%	25.0%			
gemma:2b	0.0%	0.0%	16.7%	33.3%ПП	16.7%	33.3%	16.7%			
gpt-4-0613	100.0%	100.0%	100.0%	66.7%	0.0%	0.0%	61.1% ¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬¬			
llama3.2:3b	100.0%	100.0%	66.7%	66.7%	66.7%	66.7%	77.8%			
llama3.3:70b	100.0%	100.0%	100.0%	100.0%	83.3%	100.0%	97.2%			
mistral:7b	50.0%	16.7%	50.0%	16.7%	33.3%	33.3%	33.3%			
phi4:14b	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%			
qwen2.5-coder:0.5b	100.0%	83.3%	66.7%	50.0%	50.0%	66.7%	69.4%			
qwen2.5-coder:1.5b	100.0%	100.0%	100.0%	100.0%	50.0%	50.0%	83.3%			
qwen2.5-coder:3b	100.0%	83.3%	50.0%	50.0%	83.3%	16.7%	63.9%			
qwen2.5:0.5b	0.0%	50.0%	16.7%	0.0%	16.7%	16.7%	16.7%			
qwq:32b	0.0%	16.7%	0.0%	16.7%	0.0% 🗔	16.7%	8.3%			



iouci

Figure 3. Prompt 2 mean answer score for the tested models over several temperatures. Each combination of model and temperature was tested with 6 different seeds. Error bars denote a 95% confidence interval. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Table 4. Prompt 2 answer statistics, namely the percentage of correct answers (score equal to 5) and histogram of scores (0–5) for the tested models over several temperatures. Each combination of model and temperature was tested with 6 different seeds. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Madal	Temperature									
widdel	0.0	0.2	0.4	0.6	0.8	1.0	Overall			
codellama:7b	0.0%	16.7%	16.7%	16.7%	0.0%	0.0%	8.3%			
deepseek-r1:70b	0.0% ¬.Г	33.3% ПП	33.3%	33.3%	0.0%	0.0%	16.7%			
deepseek-r1:7b	50.0%	33.3%	83.3%	16.7% 🗔 🛶 🗖	16.7%	33.3%	38.9% ¬¬			
deepseek-v3	100.0%	100.0%	100.0%	100.0%	100.0%	83.3%	97.2%			
gemma2:2b	0.0%	0.0%	16.7%	0.0%7	16.7%	0.0%	5.6%			
gemma:2b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%			
gpt-4-0613	100.0%	100.0%	100.0%	66.7%	0.0%	0.0%	61.1% ¬П			
llama3.2:3b	0.0%	50.0%	0.0%	0.0%	0.0%	0.0%	8.3%			
llama3.3:70b	100.0%	100.0%	83.3%	100.0%	83.3%	83.3%	91.7%			
mistral:7b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%			
phi4:14b	100.0%	83.3%	100.0%	100.0%	100.0%	100.0%	97.2%			
qwen2.5-coder:0.5b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%			
qwen2.5-coder:1.5b	0.0%	0.0%	0.0%	0.0%	0.0%	33.3%	5.6%			
qwen2.5-coder:3b	50.0%	16.7%	0.0%	0.0%	0.0%	16.7%	13.9%			
qwen2.5:0.5b	0.0%	0.0%	0.0%٦	0.0%	0.0%	0.0%٦	0.0%			
qwq:32b	0.0%	0.0%	33.3%	0.0%	0.0%	0.0%	5.6%			



viouei

Figure 4. Prompt 3 mean answer score for the tested models over several temperatures. Each combination of model and temperature was tested with 6 different seeds. Error bars denote a 95% confidence interval. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Table 5. Prompt 3 answer statistics, namely the percentage of correct answers (score equal to 5) and histogram of scores (0–5) for the tested models over several temperatures. Each combination of model and temperature was tested with 6 different seeds. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Model	Temperature	!					
	0.0	0.2	0.4	0.6	0.8	1.0	Overall
codellama:7b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
deepseek-r1:70b	0.0%	0.0% [0.0%	0.0%	0.0%	0.0%	0.0%
deepseek-r1:7b	0.0%	0.0% [0.0% [0.0%	0.0%	0.0%	0.0% [
deepseek-v3	66.7%	83.3%	100.0%	100.0%	66.7%	83.3%	83.3%
gemma2:2b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
gemma:2b	0.0%	0.0%	0.0%	0.0%٦	0.0%7	0.0%	0.0%
gpt-4-0613	100.0%	50.0%	50.0%	0.0%	0.0%	0.0%	33.3%
llama3.2:3b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
llama3.3:70b	50.0%	100.0%	100.0%	83.3%	66.7%	16.7%	69.4%
mistral:7b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
phi4:14b	100.0%	66.7%	66.7%	66.7%п.П	66.7%	66.7%	72.2%
qwen2.5-coder:0.5b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
qwen2.5-coder:1.5b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
qwen2.5-coder:3b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
qwen2.5:0.5b	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
qwq:32b	0.0%	0.0% Г	0.0% Г	16.7%	0.0%	0.0%	2.8% Г



Figure 5. Pairwise significance heatmap of model performance comparisons for the three prompts across temperatures. Each colored block represents the *p*-value of a one-sided, rank-based stratified permutation test between two models (model in row vs. model in column) for a given temperature. Cells are colored based on statistical significance after Benjamini–Hochberg FDR multiple testing correction: dark green indicates a significant advantage of the model in the row against the model in the column (*p* < 0.01), light green indicates moderate significant advantage (*p* < 0.05), and light gray denotes no significant difference. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

Table 6. Number of statistically significant pairwise wins (corrected p < 0.05) per model across temperature settings. Bold values indicate the highest number of wins for each temperature column (including ties). Each cell represents how many times a given model significantly outperformed others at the corresponding temperature. Temperatures for online models, deepseek-v3 and gpt-4-0613, are twice the displayed values.

M- 1-1	Temperature						
Model	0.0	0.2	0.4	0.6	0.8	1.0	Overall
codellama:7b	5	5	2	5	9	4	30
deepseek-r1:70b	0	0	0	0	0	0	0
deepseek-r1:7b	0	0	0	0	0	1	1
deepseek-v3	12	12	12	13	13	13	75
gemma2:2b	2	3	2	3	3	3	16
gemma:2b	2	2	2	2	4	3	15
gpt-4-0613	12	12	12	9	0	0	45
llama3.2:3b	5	6	3	4	4	4	26
llama3.3:70b	12	12	12	13	13	13	75
mistral:7b	4	3	2	1	4	4	18
phi4:14b	12	12	12	13	13	13	75
qwen2.5-coder:0.5b	7	6	6	6	6	5	36
qwen2.5-coder:1.5b	7	7	8	8	9	8	47
qwen2.5-coder:3b	5	6	2	4	4	3	24
qwen2.5:0.5b	2	6	2	1	4	3	18
qwq:32b	0	0	0	0	1	1	2



Figure 6. Mean answer score for the tested models and the three prompts. Each combination of model and prompt was tested 36 times (6 seeds \times 6 temperatures). Error bars denote a 95% confidence interval.

5. Discussion

Within the DeepSeek model family, there was a surprising discrepancy between the well-performing DeepSeek-V3 and the underperforming DeepSeek-R1 models. The DeepSeek-R1 versions (7B and 70B), despite their larger parameter counts, rarely generated correct code. Interestingly, the DeepSeek-R1 models, as well as Qwen's QwQ (32B), tended to generate answers over five times longer than those from other models, yet without improved correctness. While these verbose outputs are particularly noticeable, we did not investigate the reasons behind them because this lies beyond the scope of this study. Nonetheless, the generated data—as well as further analyses on this matter—are available on Zenodo (https://doi.org/10.5281/zenodo.14888673) and may be addressed in future studies.

An important observation is that GPT-4 exhibits essentially random outputs when operating at higher temperatures. This behavior aligns with OpenAI's own documentation, which indicates that temperatures above 1.2 or 1.4 may lead to increasingly stochastic completions. In contrast, the other top-performing models in this study—DeepSeek-V3, LLaMA-3.3, and Phi-4—remain relatively robust under higher temperature settings. These considerations indicate that temperature influences each model differently. Differences in temperature scaling ranges (0–1 vs. 0–2) further complicate direct comparisons.

Although one might expect a clear correlation between model size and code generation quality, results support a more involved situation among locally run models. Larger models such as DeepSeek-R1 (70B) and QwQ (32B) do not necessarily outperform smaller alternatives: their answers were typically long yet largely incorrect. Conversely, some mid-to large-scale models, such as Phi-4 (14B) and LLaMA-3.3 (70B), consistently provided accurate solutions to all prompts. Another example, LLaMA-3.2 (3B), showed reasonable performance for simpler tasks but struggled with more complex prompts, highlighting a lower boundary for parameter count beyond which performance degrades. In contrast, Qwen's smaller coder models (0.5B, 1.5B, 3B) did not show any clear advantage with increasing size, confirming that raw parameter counts alone are insufficient to predict success across different tasks.

Within the Gemini-based lineage, Gemma-2 offered marginal improvements over its older v1.1 sibling, though neither model consistently produced correct outputs. On the other hand, LLaMA-3.3 (70B) clearly outperformed the related LLaMA-3.2 (3B), a result likely driven by its substantially larger parameter count. Phi-4 merits special mention for delivering accurate code across all tasks, seeds, and temperatures, while requiring considerably fewer parameters (14B) than the largest competitors. This affords Phi-4 a strong performance/size ratio among the locally executed models.

To support these observations, a stratified permutation test with FDR correction was applied across all model pairs and temperatures. The resulting significance heatmap and win counts showed strong agreement with the descriptive statistics. DeepSeek-V3, Phi-4, and LLaMA-3.3 consistently achieved the highest number of statistically significant wins, while GPT-4 also dominated at lower temperatures. These results reinforce that the observed differences in model performance are statistically meaningful and not artifacts of randomness or scoring variability.

From a broader perspective, these findings support the notion that carefully tuned, locally run models can achieve near-state-of-the-art performance in specialized Python code generation tasks without necessarily relying on proprietary solutions. Specifically, both Phi-4 and LLaMA-3.3 proved capable of reliably generating correct solutions for the type of UAV/LoRaWAN planning prompts tested in this work. Their consistency in providing accurate answers under varying seeds and temperature conditions places them among the top-performing models overall, comparable to GPT-4 and DeepSeek-V3. These results

address the central research question: lightweight and locally executed LLMs can, in fact, generate correct Python code for relatively simple LoRaWAN and UAV planning tasks, provided that their parameter counts and training procedures meet a certain threshold of quality and scale. The performance of Phi-4 was particularly impressive, especially considering it is a relatively lightweight model.

6. Limitations

Despite the insights gained from this study, several limitations should be acknowledged. First, the selection of models, while diverse, was not exhaustive. Only a subset of locally run lightweight models was evaluated, and online testing was limited to GPT-4 and DeepSeek-V3. Several potentially relevant models, such as Claude, Mistral (larger online versions), and specialized coding models (e.g., Gemma Coder or DeepSeek Coder), were not included. This restricted scope leaves open the possibility that other models may perform competitively or even outperform those tested in this study.

Second, model outputs were assessed solely based on functional correctness, without a detailed qualitative analysis of the responses. This introduces the risk that some answers classified as correct may not have been genuinely derived but instead relied on unintended memorization, dataset leakage, or other forms of 'cheating'. While this concern is most relevant for Prompts 1 and 2, where only an index is returned, Prompt 3 mitigates this issue by requiring a real-valued output. Nevertheless, a more rigorous analysis of response quality—including potential hallucinations, redundant reasoning, and incorrect assumptions—would strengthen future work.

Third, the study relied on a single test case per function, which limits the robustness of correctness assessments. A more comprehensive evaluation would include multiple test cases per function, ensuring that responses generalize beyond a specific input scenario. This is particularly relevant given the stochastic nature of LLM-generated code, where seemingly minor variations in the prompt or execution conditions can lead to significant changes in output validity.

Fourth, all evaluations were conducted using zero-shot natural language prompts, without fine-tuning or explicit prompt engineering. While this choice aligns with practical use cases where domain experts may rely on straightforward instructions, further experimentation with prompt optimization strategies—such as chain-of-thought prompting or few-shot learning—could provide deeper insights into model capabilities.

Additionally, the study focused on relatively simple UAV/LoRaWAN planning tasks. While these scenarios are relevant to real-world applications, they do not necessarily capture the full complexity of autonomous UAV coordination, network interference, or real-time decision-making in dynamic environments. The strong performance of top models suggests they may be capable of handling more complex scenarios, but this remains an open question for future research.

A final limitation concerns the use of statistical significance testing. While stratified permutation tests confirmed the robustness of performance differences, they do not account for the magnitude or practical implications of those differences. Moreover, the use of discrete, ordinal scores simplifies model outputs and may obscure subtle qualitative distinctions. Although multiple testing correction was applied to reduce false positives, this also reduces sensitivity to borderline effects. Additionally, comparisons at non-zero temperatures should be interpreted with caution, as temperature scaling is handled differently across models, potentially resulting in varying degrees of output randomness for the same nominal value. These tests therefore complement, but do not replace, the broader descriptive analysis presented earlier.

These limitations do not diminish the validity of the study's conclusions but highlight areas for refinement in subsequent investigations. A broader model selection, more rigorous evaluation metrics, and extended task complexity would further improve the understanding of LLMs' capabilities in UAV and LoRaWAN-related computational tasks.

7. Conclusions

This paper analyzed the capabilities of 16 LLMs to generate Python functions for practical LoRaWAN-related engineering tasks involving UAV placement and signal propagation. By progressively increasing the complexity of prompts, we evaluated each model's ability to return valid and correct solutions under a standardized scoring system. The findings indicate that several recent models—particularly DeepSeek-V3, GPT-4, LLaMA-3.3, and Phi-4—consistently generated accurate and executable functions. Particularly, Phi-4 displayed exceptional performance despite its relatively lightweight architecture, demonstrating that well-optimized, smaller-scale models can be highly effective for specialized engineering applications. Models that did not achieve high scores often struggled with prompt interpretation, code syntax, or domain-specific computations, underlining the need for careful prompt engineering and model fine-tuning in similar applications.

The demonstrated viability of lightweight and locally executed LLMs for specialized engineering tasks such as UAV planning in LoRaWAN environments suggests that these models could significantly lower computational barriers and costs, allowing for broader and more flexible integration of AI-driven code generation into practical engineering workflows.

While this study highlighted the strong potential of LLMs in engineering workflows, certain limitations must be acknowledged, including the constrained model selection, the single test case per function, and the absence of qualitative analysis of responses. However, these limitations present opportunities for future research. Expanding test sets, incorporating more complex domain requirements, and evaluating additional models—particularly other lightweight alternatives—could further enrich our understanding of LLM-driven code generation in wireless communications and related fields. Future research could also explore the incorporation of reinforcement learning with human feedback to further improve the code generation capabilities of lightweight LLMs [49].

Author Contributions: Conceptualization, D.F., J.P.M.-C. and N.F.; methodology, D.F., J.P.M.-C. and N.F.; software, D.F., J.P.M.-C. and N.F.; validation, D.F., J.P.M.-C. and C.M.F.; formal analysis, N.F.; investigation, D.F. and C.M.F.; resources, D.F. and J.P.M.-C.; data curation, N.F.; writing—original draft preparation, D.F., J.P.M.-C., C.M.F. and N.F.; writing—review and editing, D.F., J.P.M.-C., C.M.F. and N.F.; visualization, N.F. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by: Fundação para a Ciência e a Tecnologia (FCT, https://ror.org/00snfqn58, accessed on 26 March 2025) under Grants Copelabs ref. UIDB/04111/2020, Centro de Tecnologias e Sistemas (CTS) ref. UIDB/00066/2020, LASIGE Research Unit ref. UIDB/00408/2025, and COFAC ref. CEECINST/00002/2021/CP2788/CT0001; Instituto Lusófono de Investigação e Desenvolvimento (ILIND, Portugal) under Project COFAC/ILIND/COPELABS/1/2024; and, Ministerio de Ciencia, Innovación y Universidades (MICIU/AEI/10.13039/501100011033, https://ror.org/05r0vyz12, accessed on 26 March 2025) under Project PID2023-147409NB-C21.

Data Availability Statement: The data generated by this study and respective analysis are available at https://doi.org/10.5281/zenodo.14888673 under the CC-BY license.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations and Symbols

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
API	Application Programming Interface
BS	Base Station
СоТ	Chain of Thought
FDR	False Discovery Rate
GPT	Generative Pre-trained Transformer
IoT	Internet of Things
LLM	Large Language Model
LoRa	Long-range
LoRaWAN	LoRa Wide Area Network
LPWAN	Low-Power Wide-Area Network
MARL	Multi-Agent Reinforcement Learning algorithm
TPML	Task Planning for Multi-UAV System
UAV	Unmaned Aerial Vehicle
VML	Visual Language Models

The following symbols are used in this manuscript:

- λ Wavelength
- c Speed of light
- f Frequency
- G_r Gain of the receiving antenna
- G_t Gain of the transmitting antenna
- *L_F* Propagation loss
- *P_r* Received power
- *P_t* Transmit power
- *r* Distance between the antennas

References

- Petajajarvi, J.; Mikhaylov, K.; Roivainen, A.; Hanninen, T.; Pettissalo, M. On the coverage of LPWANs: Range evaluation and channel attenuation model for LoRa technology. In Proceedings of the 2015 14th International Conference on ITS Telecommunications (ITST), Copenhagen, Denmark, 2–4 December 2015; IEEE: New York, NY, USA, 2015; pp. 55–59. [CrossRef]
- Augustin, A.; Yi, J.; Clausen, T.; Townsley, W.M. A study of LoRa: Long range & low power networks for the internet of things. Sensors 2016, 16, 1466. [CrossRef] [PubMed]
- 3. Mozaffari, M.; Saad, W.; Bennis, M.; Debbah, M. Wireless communication using unmanned aerial vehicles (UAVs): Optimal transport theory for hover time optimization. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 8052–8066. [CrossRef]
- 4. Sanchez-Iborra, R.; Sanchez-Gomez, J.; Ballesta-Viñas, J.; Cano, M.D.; Skarmeta, A.F. Performance evaluation of LoRa considering scenario conditions. *Sensors* 2018, *18*, 772. [CrossRef] [PubMed]
- 5. Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F.L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. GPT-4 Technical Report. *arXiv* 2024, arXiv:2303.08774. [CrossRef]
- 6. Morgan, J.; Chiang, M. Ollama: Get Up and Running with Large Language Models. GitHub. 2023. Available online: https://ollama.com/ (accessed on 10 February 2025).
- Liu, A.; Feng, B.; Xue, B.; Wang, B.; Wu, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; et al. DeepSeek-V3 Technical Report. arXiv 2024, arXiv:2412.19437. [CrossRef]
- 8. Petukhova, A.; Matos-Carvalho, J.P.; Fachada, N. Text clustering with large language model embeddings. *Int. J. Cogn. Comput. Eng.* **2025**, *6*, 100–108. [CrossRef]
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; Zettlemoyer, L. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; Jurafsky, D., Chai, J., Schluter, N., Tetreault, J., Eds.; Association for Computational Linguistics: Cambridge, MA, USA, 2020; pp. 7871–7880. [CrossRef]

- 10. Alves, D.M.; Pombal, J.; Guerreiro, N.M.; Martins, P.H.; Alves, J.; Farajian, A.; Peters, B.; Rei, R.; Fernandes, P.; Agrawal, S.; et al. Tower: An Open Multilingual Large Language Model for Translation-Related Tasks. *arXiv* 2024, arXiv:2402.17733. [CrossRef]
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R.; Le, Q.V. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In Proceedings of the 33rd International Conference on Neural Information Processing Systems, 2019, NeurIPS, Vancouver, BC, Canada, 8–14 December 2019; pp. 5753–5763.
- 12. Gu, X.; Chen, M.; Lin, Y.; Hu, Y.; Zhang, H.; Wan, C.; Wei, Z.; Xu, Y.; Wang, J. On the effectiveness of large language models in domain-specific code generation. *ACM Trans. Softw. Eng. Methodol.* **2024**, *34*, 1–22. [CrossRef]
- 13. Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. The Llama 3 Herd of Models. *arXiv* 2024, arXiv:2407.21783. [CrossRef]
- 14. Abdin, M.; Aneja, J.; Behl, H.; Bubeck, S.; Eldan, R.; Gunasekar, S.; Harrison, M.; Hewett, R.J.; Javaheripi, M.; Kauffmann, P.; et al. Phi-4 technical report. *arXiv* 2024, arXiv:2412.08905. [CrossRef]
- 15. Ling, C.; Zhao, X.; Lu, J.; Deng, C.; Zheng, C.; Wang, J.; Chowdhury, T.; Li, Y.; Cui, H.; Zhang, X.; et al. Domain specialization as the key to make large language models disruptive: A comprehensive survey. *arXiv* **2024**, arXiv:2305.18703. [CrossRef]
- 16. Javaid, S.; Fahim, H.; He, B.; Saeed, N. Large Language Models for UAVs: Current State and Pathways to the Future. *IEEE Open J. Veh. Technol.* **2024**, *5*, 1166–1192. [CrossRef]
- 17. Phadke, A.; Hadimlioglu, A.; Chu, T.; Sekharan, C.N. Integrating large language models for UAV control in simulated environments: A modular interaction approach. *arXiv* **2024**, arXiv:2410.17602. [CrossRef]
- 18. Aikins, G.; Dao, M.P.; Moukpe, K.J.; Eskridge, T.C.; Nguyen, K.D. LEVIOSA: Natural Language-Based Uncrewed Aerial Vehicle Trajectory Generation. *Electronics* **2024**, *13*, 4508. [CrossRef]
- Cui, J.; Liu, G.; Wang, H.; Yu, Y.; Yang, J. TPML: Task Planning for Multi-UAV System with Large Language Models. In Proceedings of the 2024 IEEE 18th International Conference on Control & Automation (ICCA), Reykjavik, Iceland, 18–21 June 2024; IEEE: New York, NY, USA, 2024; pp. 886–891. [CrossRef]
- 20. de Curtò, J.; de Zarzà, I.; Calafate, C.T. Semantic Scene Understanding with Large Language Models on Unmanned Aerial Vehicles. *Drones* **2023**, *7*, 114. [CrossRef]
- 21. Zhu, F.; Huang, F.; Yu, Y.; Liu, G.; Huang, T. Task Offloading with LLM-Enhanced Multi-Agent Reinforcement Learning in UAV-Assisted Edge Computing. *Sensors* 2024, 25, 175. [CrossRef]
- 22. Son, K.; Kim, D.; Kang, W.J.; Hostallero, D.E.; Yi, Y. QTRAN: Learning to Factorize with Transformation for Cooperative Multi-Agent Reinforcement Learning. In Proceedings of the 36th International Conference on Machine Learning, PMLR, Long Beach, CA, USA, 9–15 June 2019; Volume 97, pp. 5887–5896.
- Saraereh, O.A.; Alsaraira, A.; Khan, I.; Uthansakul, P. Performance Evaluation of UAV-Enabled LoRa Networks for Disaster Management Applications. *Sensors* 2020, 20, 2396. [CrossRef] [PubMed]
- Chen, L.Y.; Huang, H.S.; Wu, C.J.; Tsai, Y.T.; Chang, Y.S. A LoRa-Based Air Quality Monitor on Unmanned Aerial Vehicle for Smart City. In Proceedings of the 2018 International Conference on System Science and Engineering (ICSSE), New Taipei City, Taiwan, 28–30 June 2018; pp. 1–5. [CrossRef]
- Pan, M.; Chen, C.; Yin, X.; Huang, Z. UAV-Aided Emergency Environmental Monitoring in Infrastructure-Less Areas: LoRa Mesh Networking Approach. *IEEE Internet Things J.* 2022, *9*, 2918–2932. [CrossRef]
- Arroyo, P.; Herrero, J.L.; Lozano, J.; Montero, P. Integrating LoRa-Based Communications into Unmanned Aerial Vehicles for Data Acquisition from Terrestrial Beacons. *Electronics* 2022, 11, 1865. [CrossRef]
- 27. Wong, M.F.; Guo, S.; Hang, C.N.; Ho, S.W.; Tan, C.W. Natural language generation and understanding of big code for AI-assisted programming: A review. *Entropy* **2023**, *25*, 888. [CrossRef]
- Jain, N.; Vaidyanath, S.; Iyer, A.; Natarajan, N.; Parthasarathy, S.; Rajamani, S.; Sharma, R. Jigsaw: Large language models meet program synthesis. In Proceedings of the 44th International Conference on Software Engineering, Pittsburgh, PA, USA, 21–29 May 2022; ICSE '22, pp. 1219–1231. [CrossRef]
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. In Proceedings of the 34th Conference on Neural Information Processing Systems, NeurIPS 2020, Vancouver, BC, Canada, 6–12 December 2020; Volume 33, pp. 1877–1901.
- 30. 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* **2021**, arXiv:2107.03374. [CrossRef]
- 31. Vemprala, S.; Bonatti, R.; Bucker, A.; Kapoor, A. ChatGPT for robotics: Design principles and model abilities. *arXiv* 2023, arXiv:2306.17582. [CrossRef]
- Liang, J.; Huang, W.; Xia, F.; Xu, P.; Hausman, K.; Ichter, B.; Florence, P.; Zeng, A. Code as Policies: Language Model Programs for Embodied Control. In Proceedings of the 2023 IEEE International Conference on Robotics and Automation (ICRA), London, UK, 29 May–2 June 2023; IEEE: New York, NY, USA, 2023; pp. 9493–9500. [CrossRef]
- 33. Luo, H.; Wu, J.; Liu, J.; Antwi-Afari, M.F. Large language model-based code generation for the control of construction assembly robots: A hierarchical generation approach. *Dev. Built Environ.* **2024**, *19*, 100488. [CrossRef]

- 34. Amatriain, X. Prompt design and engineering: Introduction and advanced methods. arXiv 2024, arXiv:2401.14423. [CrossRef]
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.H.; Le, Q.V.; Zhou, D. Chain-of-thought prompting elicits reasoning in large language models. In Proceedings of the 36th International Conference on Neural Information Processing Systems, New Orleans, LA, USA, 28 November–9 December 2022; Curran Associates Inc.: Newry, UK, 2022; Volume 35, pp. 24824–24837.
- Li, Y.; Shi, J.; Zhang, Z. An approach for rapid source code development based on ChatGPT and prompt engineering. *IEEE Access* 2024, 12, 53074–53087. [CrossRef]
- Sanguesa, J.A.; Torres-Sanz, V.; Serna, F.; Martinez, F.J.; Garrido, P.; Calafate, C.T. Improving LoRaWAN Connectivity in Smart Agriculture Contexts Using Aerial IoT. In Proceedings of the 2023 IEEE Globecom Workshops (GC Wkshps), Kuala Lumpur, Malaysia, 4–8 December 2023; IEEE: New York, NY, USA, 2023; pp. 1027–1032. [CrossRef]
- Raimundo, A.; Fernandes, D.; Gomes, D.; Postolache, O.; Sebastião, P.; Cercas, F. UAV GNSS Position Corrections based on IoT LoRaWAN Communication Protocol. In Proceedings of the 2018 International Symposium in Sensing and Instrumentation in IoT Era (ISSI), Shanghai, China, 6–7 September 2018; IEEE: New York, NY, USA, 2018; pp. 1–5. [CrossRef]
- Ghazali, M.H.M.; Teoh, K.; Rahiman, W. A Systematic Review of Real-Time Deployments of UAV-Based LoRa Communication Network. *IEEE Access* 2021, 9, 124817–124830. [CrossRef]
- 40. Saunders, S.R.; Aragón-Zavala, A. Antennas and Propagation for Wireless Communication Systems, 2nd ed.; J. Wiley & Sons: Hoboken, NJ, USA, 2024.
- 41. Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. DeepSeek-R1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv* 2025, arXiv:2501.12948. [CrossRef]
- 42. Mesnard, T.; Hardin, C.; Dadashi, R.; Bhupatiraju, S.; Pathak, S.; Sifre, L.; Rivière, M.; Kale, M.S.; Love, J.; Tafti, P.; et al. Gemma: Open models based on Gemini research and technology. *arXiv* 2024, arXiv:2403.08295. [CrossRef]
- 43. Riviere, M.; Pathak, S.; Sessa, P.G.; Hardin, C.; Bhupatiraju, S.; Hussenot, L.; Mesnard, T.; Shahriari, B.; Ramé, A.; Ferret, J.; et al. Gemma 2: Improving open language models at a practical size. *arXiv* **2024**, arXiv:2408.00118. [CrossRef]
- Jiang, A.Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D.S.; Casas, D.d.l.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. Mistral 7B. arXiv 2023, arXiv:2310.06825. [CrossRef]
- 45. Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; et al. Qwen2.5 technical report. *arXiv* 2025, arXiv:2412.15115. [CrossRef]
- Hui, B.; Yang, J.; Cui, Z.; Yang, J.; Liu, D.; Zhang, L.; Liu, T.; Zhang, J.; Yu, B.; Lu, K.; et al. Qwen2.5-Coder Technical Report. *arXiv* 2025, arXiv:2409.12186. [CrossRef]
- 47. Good, P.I. *Permutation, Parametric and Bootstrap Tests of Hypotheses: A Practical Guide to Resampling Methods for Testing Hypotheses,* 3rd ed.; Springer Series in Statistics; Springer: Berlin/Heidelberg, Germany, 2004. [CrossRef]
- 48. Benjamini, Y.; Hochberg, Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B (Methodol.)* **1995**, *57*, 289–300. [CrossRef]
- 49. Wong, M.F.; Tan, C.W. Aligning Crowd-Sourced Human Feedback for Reinforcement Learning on Code Generation by Large Language Models. *IEEE Trans. Big Data* **2024**. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.