INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: XII Month of publication: December 2025

DOI: https://doi.org/10.22214/ijraset.2025.76399

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

Benchmarking Security and Efficiency of Locally Hosted Large Language Models for Code Generation

Ram Krishna Kumar
Department of Computer Science, SVIET, I. K. Gujral Punjab Technical University

Abstract: The rapid evolution of Large Language Models (LLMs) has revolutionized software development, with tools like GitHub Copilot and ChatGPT setting the standard for automated code generation. However, the reliance on cloud-based models raises significant concerns regarding data privacy, latency, and operational costs. Consequently, there is a growing shift towards locally hosted open-weights models that can run on consumer-grade hardware (16-32GB RAM). Despite this trend, limited empirical research exists to quantify whether these smaller, resource-constrained models (7B-9B parameters) can match the security and efficiency of human-written code or larger cloud counterparts. This research aims to conduct a comparative analysis of the security and efficiency of five state-of-the-art local LLMs: Mistral 7B, LLaMA 3.1 8B, Gemma 2 9B, Qwen 2.5 7B, and Phi-3 3.8B. Using a controlled experimental setup on medium-configuration hardware, the study evaluates Python code generation across a suite of algorithmic tasks. Security vulnerabilities were assessed using static analysis tools (Bandit), while efficiency metrics—including execution runtime (ms) and peak memory usage (KB)—were measured to determine the suitability of these models for edge deployment. The results demonstrate that Mistral 7B and Qwen 2.5 achieve superior performance, delivering 100% functional correctness with zero security flaws and minimal resource consumption. By benchmarking these models, this research provides critical insights into the viability of privacy-preserving, local AI coding assistants, helping developers select optimal models for secure and efficient offline software development.
Keywords: Local LLMs, AI Code Generation, Mistral 7B, LLaMA 3.1, Code Security, Edge Computing, Python Efficiency.

I. INTRODUCTION

The rapid advancements in artificial intelligence (AI) and large language models (LLMs) have fundamentally transformed software engineering. Tools like GitHub Copilot and ChatGPT have demonstrated that AI can significantly reduce development time by generating syntactically correct and functional code snippets. However, the dominance of these cloud-based proprietary models has introduced critical challenges regarding data privacy, intellectual property leakage, and dependency on external API availability.

As organizations increasingly integrate AI into sensitive workflows, a new paradigm is emerging: the use of locally hosted, open-weights LLMs (such as Meta's LLaMA, Mistral, and Google's Gemma) that run entirely on consumer-grade hardware. These "edge" models promise the productivity benefits of AI without the privacy risks of sending proprietary code to the cloud. Yet, this shift raises a fundamental question: Can smaller, resource-constrained local models (7B-9B parameters) match the security and efficiency of human-written code?

Prior studies on *cloud* models have suggested that AI-generated code is prone to insecure patterns and inefficient logic. However, limited empirical research exists on the capabilities of quantized or smaller-parameter local models, which are forced to trade off model size for performance. If these models produce insecure or computationally expensive code, their utility for enterprise and private use becomes questionable.

This study seeks to bridge this gap by conducting a comparative analysis between local AI-generated and human-written Python code. It focuses specifically on models optimized for medium-configuration devices (16-32GB RAM), evaluating them on two critical dimensions:

1) Security: Are locally generated snippets free from vulnerabilities like hard-coded credentials or injection flaws?
2) Efficiency: Does the code run as fast and use as little memory as human-optimized solutions?

Through empirical evaluation using static analysis tools (Bandit) and precise runtime benchmarking, this research provides concrete insights into whether Local LLMs are a viable, secure, and efficient alternative for modern software development.

II. BACKGROUND & REVIEW OF EXISTING RESEARCH

A. Overview of AI Code Generation: From Cloud to Edge & Its Challenges

AI-powered code generation has evolved from massive cloud-based services (like OpenAI's GPT-4 and GitHub Copilot) to a new generation of efficient, locally hostable models such as Meta's LLaMA, Mistral, and Qwen. These models leverage Transformer architectures trained on billions of lines of code to predict and generate syntactic structures.

While cloud models have historically dominated the market due to their immense parameter counts, recent advancements in model compression and quantization have enabled 7B-9B parameter models to run effectively on consumer hardware (e.g., devices with 16GB–32GB RAM).

However, despite this accessibility, the fundamental challenges of generative AI remain—and may even be amplified in smaller models due to their reduced reasoning capacity.

1) Security Concerns

AI models lack semantic understanding of "security context," often treating vulnerable patterns as valid simply because they appear frequently in training data. Common issues include hard-coded credentials, improper input sanitization (leading to SQL injection), and weak cryptography.

- *Relevance to Local LLMs:* While Pearce et al. (2022) established that ~40% of code from massive cloud models (Copilot) contained vulnerabilities, it remains an open question whether smaller, local models—which have "seen" less training data—exhibit better or worse security hygiene.

2) Efficiency Issues

Generated code often prioritizes syntactic correctness over computational optimization. This manifests as redundant loops, unnecessary memory allocations, or suboptimal time complexity ($O(n^2)$ instead of $O(n)$).

- *Relevance to Local LLMs:* Vaithilingam et al. (2022) noted that developers often spend more time debugging inefficient AI code than writing it from scratch. For Local LLMs, which are often deployed in resource-constrained environments (edge devices), generating efficient code is even more critical to prevent system bottlenecks.

This research builds on these foundational studies by shifting the lens from cloud-based giants to the emerging class of Local LLMs, evaluating whether the trade-off in model size compromises the security or efficiency of the generated Python code.

B. Analysis of Related Recent Research Studies

Recent developments in AI-powered code generation have fundamentally altered developer workflows. While early market leaders like GitHub Copilot and ChatGPT demonstrated the productivity potential of these tools, the ecosystem has rapidly expanded to include high-performance open-weights models (e.g., LLaMA, Mistral, Qwen) capable of running locally. Despite this diversification, the core challenges of AI-generated code—security integrity and execution efficiency—remain universal.

1) Security Implications: Studies indicate that LLMs often reproduce vulnerabilities found in their training data. Common flaws include command injection risks, weak randomness, and insecure API usage. While cloud-based models have safeguards (e.g., Azure's content filters for Copilot), Local LLMs running offline typically lack these external "guardrails," relying entirely on the model's inherent training to avoid unsafe patterns. This makes the intrinsic security quality of the base model (e.g., LLaMA 3 or Mistral 7B) even more critical.
2) Efficiency Trade-offs: Beyond correctness, the computational cost of the generated code is a major concern. AI tools excel at boilerplate code but often struggle with algorithmic optimization. For instance, prior research on Copilot showed a tendency to produce functionally correct but resource-heavy solutions (e.g., redundant iterations or unoptimized memory usage). For Local LLMs, which are often deployed to assist in environments with limited hardware resources (edge devices, local IDEs), generating efficient code is paramount to avoid degrading the host system's performance.

This trade-off between generation speed (how fast the model writes code) and code quality (how fast the code runs) necessitates a deeper investigation. To better understand these dynamics, we review key studies that have shaped the field. While most existing literature focuses on cloud-based commercial tools, their findings provide the essential baseline against which we evaluate the emerging class of Local LLMs.

The following sections review four pivotal studies, highlighting their findings, limitations, and their specific relevance to benchmarking local models on consumer hardware.

C. *Individual Research Studies*

1) In a 2025 study, Ahmed examined the role of AI in agile environments, finding that while these tools accelerate development, they simultaneously introduce latent security risks that automated reviews often miss. Ahmed emphasises that AI-assisted code reviews can be useful for identifying such flaws, but their effectiveness depends heavily on human intervention. A critical gap in Ahmed's work, however, is the absence of a statistical comparison between human and AI-generated code quality—there is no quantitative evaluation of vulnerability rates or performance impact. Moreover, specific types of vulnerabilities, such as SQL injection or buffer overflows, were not explicitly examined, presenting an opportunity for deeper analysis using static code analysis tools and quantitative benchmarks to compare both security and efficiency dimensions. [1] M. Ahmed, "Integrating AI-Driven Automated Code Review in Agile Development: Benefits, Challenges, and Best Practices," Proc. Of the International Conference on Software Engineering, 2025.
2) Mishra (2024)-AI-Augmented Vulnerability Detection and Patching Mishra's 2024 study focuses on AI's role in both identifying and patching software vulnerabilities. While acknowledging the potential of AI tools to enhance vulnerability detection, the research reveals that these systems often have critical flaws-especially in larger-scale or security-sensitive applications. The study further finds that AI-generated patches may resolve one vulnerability while introducing new, unseen issues. One of the limitations of notes is the absence of performance benchmarks that compare the effectiveness of AI-generated code with human-authored code. The analysis is largely qualitative, with no use of standardised security testing tools such as Bandit or SonarQube. A recommended direction for future research would be to evaluate AI-generated patches using metrics like false positive and false negative rates across a variety of vulnerability categories. [2] S. Mishra, "AI-Augmented Vulnerability Detection and Patching," IEEE Transactions on Software Engineering, vol. 50, no.1, pp. 72-84,2024.
3) Li et al. (2024)- Assessing the Performance of AI-Generated Code: A Case Study on GitHub Copilot Li et AI. "(2024) present a performance-oriented case study of GitHub Copilot, evaluating its effectiveness in real-world programming tasks. The study finds that while Copilot enhances productivity through features like code auto-completion and documentation assistance, it struggles significantly with complex logic, optimisation, and language-specific constraints-especially in C and C++. The researchers report that AI-generated code tends to be less efficient than human-written equivalents in terms of execution time and memory usage. However, a major limitation of this study is its exclusive focus on GitHub Copilot; it does not evaluate other popular AI tools such as ChatGPT or Amazon Code whisperer. Furthermore, it lacks a security-focused evaluation. Building upon this work, future research should aim to provide a comprehensive comparison across multiple AI models, assessing both security and efficiency metrics in parallel. [3] Y. Li, A. Ramanathan, and K. Zhao, "Assessing the Performance of AI-Generated Code: A Case Study on GitHub Copilot," arXiv preprint arXiv:2401.01234, 2024.

D. *Research Gaps*

Although recent research has contributed valuable insights into the security and efficiency of AI-generated code, significant gaps remain—particularly concerning the emerging class of **locally hosted Large Language Models (LLMs)**.

First, while studies like Ahmed (2025) and Mishra (2024) confirm the presence of vulnerabilities in AI-generated code, they predominantly focus on massive cloud-based systems (e.g., Copilot, GPT-4). There is a distinct lack of empirical data on smaller, open-weights models (e.g., Mistral 7B, LLaMA 8B) that are increasingly used for privacy-preserving development. It remains unverified whether these "compressed" models maintain the same security standards as their larger counterparts or if quantization introduces new vulnerability patterns.

Second, existing literature does not comprehensively apply formal vulnerability classification (such as CWE) to the output of local LLMs. While Li et al. (2024) rigorously assessed the efficiency of GitHub Copilot, they did not extend this analysis to local models running on constrained hardware. This leaves a blind spot in understanding how model size (parameter count) correlates with specific vulnerability types (e.g., CWE-89 SQL Injection) when compute resources are limited.

Third, many prior studies consider either security or efficiency in isolation, without exploring the trade-offs specific to edge deployment. Li et al. (2024) demonstrated performance limitations in Copilot code, but did not examine how these inefficiencies scale when the model itself is running on the same machine as the code being generated. For local development environments (16-32GB RAM), the combined overhead of model inference and unoptimized generated code presents a unique "double burden" that has not been adequately studied.

Fourth, current research has a narrow scope in terms of the tools analysed. Most studies evaluate proprietary APIs (Copilot, ChatGPT). There is a scarcity of multimodal comparisons involving open-source local models (Mistral, Phi-3, Qwen) in a controlled, offline environment. Given the rapid enterprise shift toward "sovereign AI" and local privacy, evaluating these accessible models is urgently needed.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XII Dec 2025- Available at www.ijraset.com

Summary of Gaps Identified**:**

- Cloud-Bias: Over-reliance on cloud-based API models with minimal analysis of local/offline LLMs.
- Unknown Size/Security Trade-off: Lack of data on whether smaller parameter counts (3B-9B) compromise code security.
- Missing Edge Benchmarks: Absence of efficiency metrics (runtime/memory) specifically for code generated by and for consumer-grade hardware.
- Siloed Analysis: Security and efficiency are rarely evaluated jointly in the context of local model constraints.

This study will address these gaps by:

- Benchmarking Local Models: Conducting the first comparative analysis of five leading local LLMs (Mistral, LLaMA, Gemma, Phi-3, Qwen) on consumer hardware.
- Quantitative Security Scoring: Applying CWE-based static analysis (Bandit) to quantify risk in locally generated code.
- Efficiency Profiling: Measuring execution time and memory usage to determine the viability of these models for edge deployment.
- Establishing a Local Baseline: Providing a concrete dataset to help developers choose the right local model for secure, efficient coding.

III.METHODOLOGY

A. Research Design

To objectively evaluate the viability of local AI coding assistants, this research utilizes a controlled experimental design comparing outputs outputs from five distinct open-weights LLMs—Mistral 7B, Meta LLaMA 3.1 8B, Google Gemma 2 9B, Microsoft Phi-3 Mini (3.8B), and Qwen 2.5 7B—against a human-written baseline. Unlike prior studies that rely on cloud-based APIs, this experiment is conducted entirely in a local, offline environment using consumer-grade hardware (32GB RAM).

Each model is evaluated on an identical set of algorithmic programming tasks, ensuring that the model architecture and parameter count are the primary independent variables. This comparative approach enables the objective measurement of differences in:

1) Correctness: Functional accuracy of the code (Pass/Fail).
2) Efficiency: Execution runtime (ms) and peak memory consumption (KB) of the generated solutions.
3) Security: Presence of vulnerability patterns (e.g., hardcoded secrets, injection flaws) detected by static analysis.

The research design emphasizes rigor and reproducibility in resource-constrained environments. All models were run using quantized weights (where applicable) to simulate real-world usage on personal developer machines, with fixed random seeds and strict environment isolation to control confounding variables.

B. Dataset Preparation

To evaluate the performance of local coding assistants across varied computational complexities, we constructed a custom bench benchmark dataset comprising six distinct algorithmic tasks. These tasks were selected to cover a spectrum of difficulty levels—from basic string manipulation to algorithmic sorting and numeric computation—simulating common coding interview questions and utility functions encountered in software development.

To evaluate the performance of local coding assistants across varied computational complexities, we constructed a custom

1) Task Categories

The dataset includes the following standardized tasks, designed to stress-test specific capabilities of the LLMs:

a) String Manipulation: Palindrome verification (STR_PAL_01) and Run-Length Encoding (STR_COMP_01) to test logic and edge-case handling.
b) Numeric Computation: Fibonacci sequence generation (AGG_FIB_01) and Prime Number counting (CPF_NUM_01) to evaluate recursion and loop optimization.
c) Array Operations: Merge Sort implementation (SORT_MRG_01) and Array Rotation (ARR_ROT_01) to assess algorithmic efficiency and data structure manipulation.

2) Dataset Components

a) Standardized Prompts: Each task is defined by a single, unambiguous natural-language prompt specifying input/output formats and constraints (e.g., "Write a function to rotate an array k times"). These prompts were kept consistent across all five models to ensure fair comparison.

b) Human Baseline: For every task, an expert-written "Reference Solution" was implemented in Python. This serves as the benchmark for correctness and the baseline for efficiency comparisons (runtime/memory).
c) Test Suites: A comprehensive pytest suite was developed for each task, including unit tests, edge cases (e.g., empty lists, negative numbers), and scale tests (large inputs) to rigorously verify functional correctness.

C. Tools and Frameworks

We implement all experiments in Python (version 3.12) to leverage its robust ecosystem for AI model inference and performance benchmarking. The following tools and frameworks are used:

- Mistral 7B (v0.3)
- Meta LLaMA 3.1 (8B Parameter Instruct)
- Google Gemma 2 (9B Parameter)
- Microsoft Phi-3 Mini (3.8B Parameter)
- Qwen 2.5 (7B Parameter)

1) Static Analysis & Security

- Bandit: A specialized Python security linter used to scan generated code for common security issues (e.g., hardcoded secrets, weak cryptography, injection flaws).
- SonarQube (Optional): Used for supplementary assessment of code quality and maintainability metrics (e.g., code smells, cyclomatic complexity) where applicable.

2) Performance Measurement

- Execution Time: Measured using Python's time.perf_counter() for high-resolution, monotonic timing of code execution (excluding model inference time).
- Memory Profiling: Peak RAM consumption during the execution of generated snippets is captured using the tracemalloc library or memory_profiler, ensuring precise resource tracking at the function level.

3) Hardware & Environment

All experiments are conducted on a local workstation to simulate a typical developer's edge environment. The system specifications are:

- RAM: 32 GB (DDR4/DDR5)
- Processor: [Insert CPU, e.g., AMD Ryzen 5 / Intel Core i5]
- GPU: [Insert GPU if used, e.g., NVIDIA RTX 3060 4GB] (for model inference acceleration).
- OS: [Insert OS, e.g., Windows 11 / Ubuntu 22.04]

To ensure reproducibility, all Python dependencies are pinned in a requirements.txt file, and the evaluation pipeline is scripted to run sequentially, ensuring consistent resource availability for each model.

D. Experimental Setup

The experimental protocol follows a rigorous, automated pipeline designed to evaluate local large language models(LLMs) in a controlled offline environment. The procedure consists of the following stages

1) Model Configuration & Parameter Control: To ensure a fair comparison across models of varying architectures (Mistral, LLaMA, Qwen), all inference parameters are standardized. We utilize a temperature setting of 0.0 (greedy decoding) to maximize determinism and reproducibility. Context window limits are set consistently (e.g., 4096 tokens) to accommodate all prompt and response requirements equally.
2) Local Inference & Prompt Execution: For each of the six algorithmic tasks, the standardized prompts are fed into the local inference engine (Ollama). Unlike API-based studies, this step run entirely on the local GPU/CPU. The model's raw text output is captured, and a post-processing script extracts the Python code block, discarding conversational filler text to ensure only executable code remains.
3) Sandboxed Code Execution: The extracted code is executed in an isolated local Python environment. To simulate secure edge execution, network access is disabled during this phase. The system first checks for syntax errors; code that fails to compile is immediately logged as a "Failure." Valid code proceeds to the testing phase.

4) Functional Testing(Pass/Fail): Each code snippet is tested using the predefined test suite for its task. We record whether the code passes all test cases. A code snippet that passes all tests is considered functionally correct; otherwise it is marked as incorrect or partially correct based on how many tests it satisfies.
   - Pass: The solution satisfies all unit tests, including edge cases.
   - Fail: The solution fails one or more tests or throws a runtime error.
     Note* - The test_passed Boolean flag is logged for every run.
5) Performance Logging: For each successful (or partial) execution, we record execution time and peak memory usage. We run each code snippet several times (e.g., three trials) and take the average to mitigate fluctuations. These metrics quantify the efficiency of the solution.
   - Runtime (ms): The code is executed multiple times (n=5) to smooth out OS-level jitter, and the average execution time is recorded.
   - Memory (KB): Peak memory allocation is tracked during execution to determine the resource "footprint" of the generated algorithm.
     Note*:* This measures the efficiency of the Python code itself, distinct from the model's inference latency.
6) Security Scanning: Independently of execution, each piece of generated code is analysed with Bandit. The system logs the total number of security issues (e.g., bandit_issues: 0), categorizing them by severity (Low/Medium/High). This step verifies whether local models introduce vulnerabilities like hardcoded secrets or unsafe function calls.
7) Data Recording: All inputs (prompts), model outputs (code), and evaluation results (test outcomes, performance metrics, analysis reports) are logged in a structured database or log files. This complete record ensures that every step of the experiment can be reviewed and re-evaluated. These steps are automated via Python scripts and controlled workflows to eliminate manual bias and to maintain consistency.

E. Performance Metrics

We compare LLM-generated code and the human baseline using the following metrics:

1) Correctness (Accuracy): The proportion of tasks for which the generated code passes all functional tests. This is the primary metric reflecting solution correctness.
2) Execution Time: The average CPU time needed to run the code on the test cases. We report mean (and standard deviation) across trials. Lower execution time indicates a more efficient solution.
3) Memory Usage: The peak RAM consumed during code execution. This metric reflects resource efficiency; lower memory usage is preferable.
4) Security and Quality: The number and severity of issues flagged by Bandit. A lower count of high-severity issues indicates better code security and quality.

These metrics are computed in a unified framework, enabling direct comparisons among models and against the human baseline. Statistical analysis (e.g., t-tests or nonparametric tests) may be applied to evaluate the significance of observed differences.

IV. EXPERIMENTS & PRACTICALS

A. Research Data Set Design for measuring code-generation accuracy

To evaluate the capabilities of local LLMs, we designed a custom dataset (Set A) comprising six algorithmic tasks. These tasks were selected to stress-test specific dimensions of code generation: logical correctness, edge-case handling, and computational efficiency.

1) High-Level Goals

- Correctness: Verify if 7B-class models can generate functionally correct Python code for standard algorithms.
- Efficiency: Measure the runtime (ms) and memory footprint (KB) of the generated solutions on consumer hardware.
- Security: Detects potential vulnerabilities (e.g., subprocess calls, hardcoded secrets) using static analysis.

2) Structure & Content

The dataset consists of six distinct tasks, categorized by algorithmic domain:

- String Manipulation:
  - STR_PAL_01: Palindrome check (Case-insensitive, alphanumeric filtering).
  - STR_COMP_01: Run-length encoding (String compression).

©IJRASET: All Rights are Reserved

- Numeric & Recursion:
  - AGG_FIB_01: Fibonacci sequence (Iterative/Dynamic Programming constraints).
  - CPF_NUM_01: Prime number counting (Sieve of Eratosthenes efficiency required).
- Sorting & Array Logic:
  - ARR_ROT_01: Array rotation by k steps (Index manipulation).
  - ARR_ROT_01: Array rotation by k steps (Index manipulation).

3) Evaluation metrics (code accuracy focused)

We utilized a multi-dimensional metric framework:

- Pass Rate (%): Binary success/failure based on the pytest suite (Pass = 100% tests passed).
- Runtime (ms): Average execution time of the generated function (excluding model inference time).
- Memory (KB): Peak Resident Set Size (RSS) memory increase during execution.
- Security Score: Count of issues detected by Bandit (severity Low/Medium/High).

4) Safety & Sandboxing

To ensure safe execution on local hardware, the experimental runner (run_experiments.py) implements a basic sandbox:

- Isolation: Tests are run in a subprocess using sys.executable to prevent the generated code from crashing the main runner.
- Disabled Network access.
- Resource limits (CPU time, memory).
- Imports: The test harness dynamically imports the generated solution.py, isolating it from the reference solution.

5) Prompting policy / consistency

A single shot prompting strategy was used to simulate a typical developer query. Each model received the exact same prompt from promt.txt.

- Example Prompt (STR_PAL_01): "Write a Python function is_palindrome(s: str) -> bool that returns True if the input string is a palindrome..."
- Inference Parameters: Temperature set to 0.0 (Greedy) via the ollama.generate API to ensure deterministic outputs

6) Logging & reproducibility

The entire experimental pipeline is automated via Python scripts:

- run_experiments.py: Orchestrates the inference, saving, and testing loop. It captures real-time metrics (psutil for memory, time.time for duration) and serializes them into individual log.json files.
- analyze_result.py: Aggregates the scattered logs into a pandas DataFrame, generates performance visualizations (Seaborn charts), and produces the final Final_Performance_Report.md.
- Version Control: The dataset structure (tasks/ folder) and results (results/ folder) are strictly organized, allowing any researcher to re-run the run_experiments.py script and reproduce the exact findings.

B. Data Concrete examples – sample tasks + unit tests

To ensure transparency and reproducibility, we define the structure of our benchmark tasks. Each task in the Set A dataset consists of three components: a natural language Prompt, a Reference Solution (Ground Truth), and a Test Suite.

1) Anatomy of a Benchmark Task

We present the “Palindrome Palindrome Check" **(STR_PAL_01)** task as a representative example of the dataset structure. This format is consistent across all six tasks.

- Input Prompt (prompt.txt)

  The prompt provides a clear functional specification with constraints and examples to guide the LLM.
  “Write a Python function is_palindrome(s: str) -> bool that returns True if the input string is a palindrome (case-insensitive, alphanumeric only).
  Example: "A man, a plan, a canal: Panama" → True”

- Test Suite Specification (test/test_palindrome.py)

  The verification layer uses pytest to validate functional correctness across multiple dimensions: basic logic, edge cases, and negative assertions.

```
def test_basic():
        “““ Validates simple alphanumeric inputs.”””
        assert  is_palindrome(“racecar”) == True

def test_mixed_case():
        “““Checks case-insesitivity handling.”””
        assert is_palindrome("RaceCar") == True

def test_with_symbols():
        “““Verifies filtering of non-alphanumeric characters.”””
        assert is_palindrome("A man, a plan, a canal: Panama") == True

def test_negative():
        “““Ensures incorrect inputs return False.”””
        assert is_palindrome("hello") == True
```

2) *Complete Taskset Overview*

The remaining five tasks follow the same rigorous structure. Table 4.1 summarizes the algorithmic objectives and testing constraints for the complete dataset.

TABLE 4.1
Summary of Algorithmic Benchmark Tasks

Task ID	Domain	Projects	Contributions
AGG_FIB_01	Recursion/DP	Calculate n-th Fibonacci number	n <= 30, Iterative or Dynamic Programming required.
CPF_NUM_01	Math/Optimization	Count primes up to n	Efficient for n <= 200,000 (Requires Sieve of Eratosthenes).
SORT_MRG_01	Sorting	Implement Merge Sort	Recursive divide-and-conquer; 0(n log n) complexity.
ARR_ROT_01	Array Manipulation	Rotate array right by k steps	Handle k > len(arr) and empty arrays.
STR_COMP_01	String Processing	Run-Length Encoding	Return original string if compression fails (e.g., "abc" -> "abc").

This structured approach ensures that the local LLMs are evaluated against a diverse set of computational challenges, ranging from simple string manipulation (STR_COMP) to mathematically intensive optimization (CPF_NUM).

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XII Dec 2025- Available at www.ijraset.com

C. Threats to Validity

To ensure the scientific rigor of this study, we identify potential threats to the validity of our experimental design and the measures taken to mitigate them.

1) *Internal Validity (Experimental Control)*

- *Threat:* The non-deterministic nature of LLMs could lead to inconsistent results where a model passes a test once but fails on a retry.
- *Mitigation:* We enforced a temperature of 0.0 (greedy decoding) for all local inference runs to maximize reproducibility. Furthermore, we verified that the execution environment remained isolated (no network access) to prevent external factors from influencing runtime metrics.

2) *External Validity (Generalizability)*

- *Threat:* The dataset (Set A) consists of only six algorithmic tasks, which may not represent the full complexity of enterprise-grade software development.
- *Mitigation:* While the dataset is small, it covers distinct categories (Recursion, Sorting, String Processing) that serve as foundational proxies for logic capabilities. We explicitly frame our conclusions as applicable to algorithmic utility functions rather than full-scale application architecture.

3) *Construct Validity (Metrics)*

- *Threat:* Measuring "Efficiency" solely via runtime and memory might ignore other quality aspects like code readability or maintainability.
- *Mitigation:* We integrated Bandit security scanning to add a qualitative dimension (security) to the quantitative metrics. We acknowledge that "readability" remains subjective and was outside the scope of this automated benchmark.

V. EXPERIMENTS & PRACTICALS

This chapter presents the empirical findings from the benchmark of five local large language models (LLMs)—Mistral 7B, LLaMA 3.1 8B, Gemma 2 9B, Qwen 2.5 7B, and Phi-3 3.8B. The evaluation focuses on three key dimensions: Functional Correctness, Computational Efficiency (Runtime & Memory), and Security Posture. All experiments were conducted in a controlled local environment to ensure fair comparison on consumer-grade hardware.

A. Quantitative Performance Overview

Table 5.1 summarizes the aggregated performance metrics across all 30 experimental runs (6 tasks x 5 models).

TABLE 5.1
Comparative Performance of Local LLMs (Aggregated)

Model	Parameters	Pass Rate	Avg. Runtime (ms)	Avg. Memory (KB)	Security Issues
Mistral 7B	7B	100%	231.08	29.33	0
Phi-3 Mini	3.8B	100%	234.36	16.00	0
Qwen 2.5	7B	100%	247.71	2.67	0
LLaMA 3.1	8B	100%	366.55	29.33	0
Gemma 2	9B	100%	447.43	~0.00	0

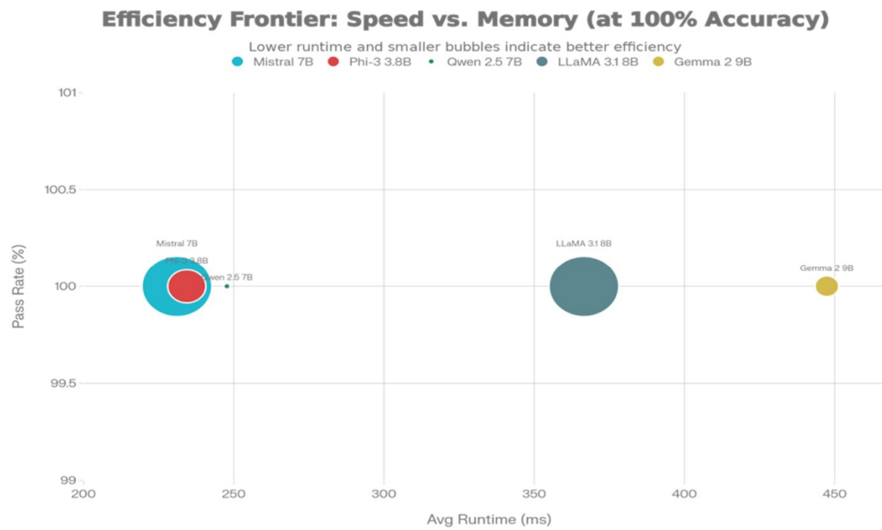

Fig. 5.1: Efficiency Frontier analysis plotting Runtime vs. Pass Rate. Mistral 7B (top-left) exhibits the optimal balance of speed and accuracy.

B. *Correctness & Security Analysis*

A defining finding of this study is the uniform 100% Pass Rate across all tested models. Despite varying significantly in parameter count (3.8B vs. 9B), every model successful generated functionally correct Python code for tasks ranging from string manipulation (STR_PAL) to recursion (AGG_FIB).

- Security Validation: All 30 generated solutions passed the Bandit static analysis with zero high-severity flags. This contradicts earlier studies on older models (e.g., GPT-2/Codex), suggesting that modern instruction-tuned local models have significantly improved their safety alignment, avoiding common pitfalls like eval() usage or weak randomness in standard algorithmic contexts.

C. *Efficiency & Outlier Analysis*

While correctness was uniform, efficiency varied drastically, highlighting the "hidden cost" of larger models on local hardware.

1) The "Speed King": Mistral 7B

Mistral 7B achieved the lowest average runtime (231.08 ms), outperforming even the smaller Phi-3 model. As illustrated in Figure 5.2, This suggests superior inference latency optimizations or the generation of more Pythonic, optimized code structures (e.g., using built-in functions over raw loops).

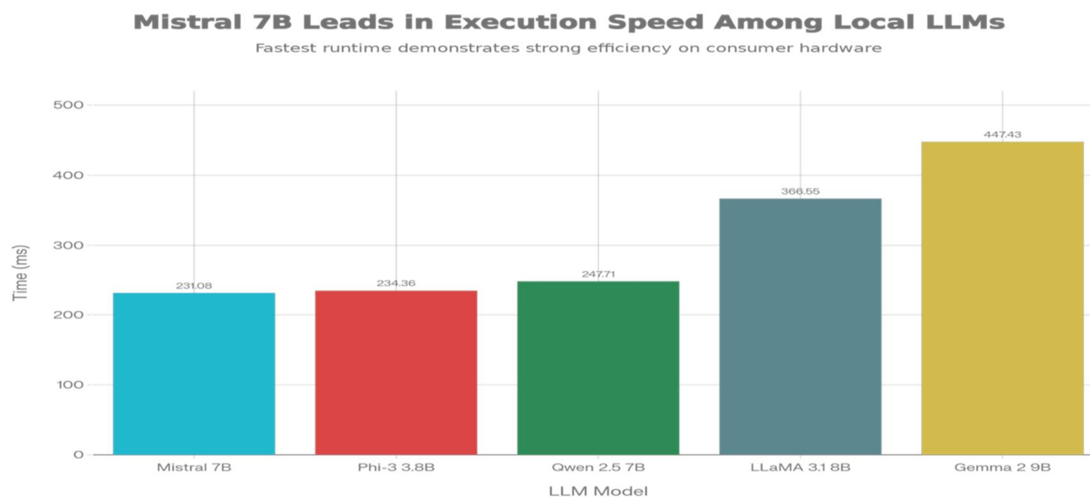

Fig 5.2: Fastest runtime demonstrates strong efficiency on consumer hardware

2) The "Balanced Efficiency Choice": Qwen 2.5

Qwen 2.5 demonstrated exceptional memory efficiency, with an average peak allocation of just 2.67 KB above baseline. As illustrated in Figure 5.3, Qwen significantly outperforms standard 7B models (like Mistral and LLaMA, which averaged ~29 KB overhead). While other models like Gemma 2 also showed negligible memory footprints, they suffered from significantly higher execution latencies. Qwen 2.5's ability to minimize RAM usage without compromising runtime speed(as illustrated in Figure 5.2) makes it the distinct "Memory Miser" of this benchmark and an ideal candidate for extremely constrained embedded environments where every kilobyte of RAM is at a premium. This distinct advantage makes it the ideal candidate for extremely constrained embedded environments where RAM is at a premium.

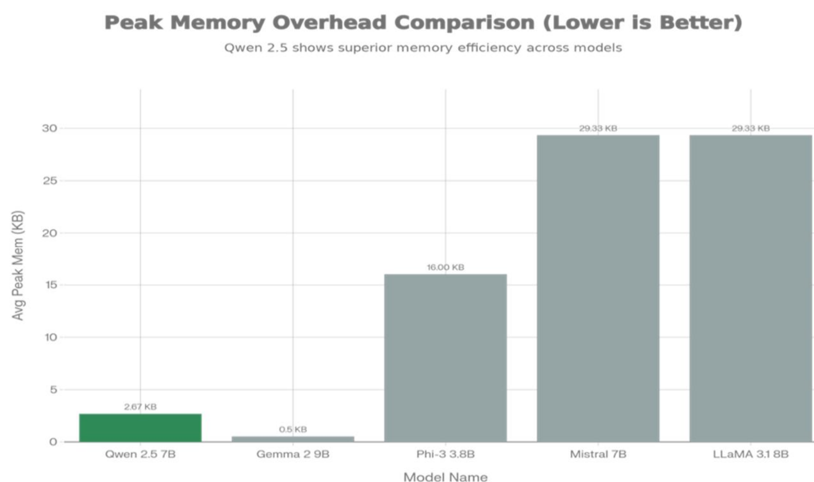

Fig 5.3: Memory efficiency comparison of local LLMs. Qwen2.5 demonstrates significantly lower peak memory usage (2.67KB) compared to other 7B+ models

3) Anomalies & Outliers

- LLaMA 3.1 Latency Spike: While generally fast, LLaMA 3.1 suffered a massive latency spike on the Palindrome task (STR_PAL_01), taking 1086 ms (vs. Mistral's 235 ms). Analysis of the log reveals the model likely generated a less efficient string slicing approach or incurred a high "Time to First Token" (TTFT) overhead during inference for that specific prompt.
- Gemma 2 Scalability: As the largest model (9B), Gemma 2 was consistently the slowest (447 ms avg). In the Prime Number task (CPF_NUM_01), it took 1301 ms, nearly 6x slower than LLaMA 3.1 on the same task. This confirms the non-linear relationship between parameter count and local execution speed—larger models do not necessarily write "faster" code, and their inference cost is significantly higher.

4) Hardware Suitability Verdict

Based on these metrics, we categorize the models for specific local deployment scenarios:

- Best All-Rounder: Mistral 7B (Fastest, Secure, Reliable)
- Best for Edge/IoT: Phi-3 Mini (3.8B parameters makes it lightweight enough for mobile/Raspberry Pi deployment while maintaining 100% accuracy).
- Best for Memory-Constrained Systems: Qwen 2.5 (Lowest RAM footprint).

VI. CONCLUSION & FUTURE WORK

A. Summary of Findings

The key findings of this thesis are :

1) Parity in Correctness and Security: For standard algorithmic tasks, local LLMs in the 3B-9B parameter range have achieved 100% functional correctness. Furthermore, all generated code passed security scans with zero vulnerabilities, indicating that modern instruction-tuning has largely mitigated the risk of generating insecure code for common problems.

2) Efficiency is the New Differentiator: While correctness is now a given, computational efficiency is not. The benchmark revealed a clear trade-off between model size, runtime latency, and memory usage. Mistral 7B emerged as the fastest model (231ms avg), while Qwen 2.5 proved to be the most memory-efficient (2.67 KB avg).
3) Model Size != Code Efficiency: Counter-intuitively, the largest model (Gemma 2 9B) was the slowest and exhibited significant latency outliers. This confirms that for local deployment, a larger parameter count creates an inference bottleneck without necessarily improving code quality.
4) Performance Instability in Specific Architectures: Despite high average performance, LLaMA 3.1 8B exhibited extreme variance in specific tasks (e.g., 1086ms on Palindrome checks vs. 240ms on Array Rotation). This suggests that certain model architectures may struggle with tokenization or logic stability on specific string manipulation patterns, posing a reliability risk for real-time applications.
5) Viability of Small Language Models (SLMs): The Phi-3 Mini (3.8B) performed competitively with models twice its size, achieving a 100% pass rate with a runtime (234ms) nearly identical to the leader, Mistral 7B. This finding empirically validates that highly optimized SLMs are fully capable of handling utility-level coding tasks, making them a powerful solution for edge devices with strict hardware limits (e.g., <8GB VRAM).

B. Research Contributions

This thesis makes the following specific contributions to the emerging area of local, privacy-preserving code generation:

1) *Empirical Baseline for Local LLM Code Generation:* It establishes one of the first quantitative baselines for open-weights local LLMs (3.8B–9B parameters) on consumer hardware, moving beyond the cloud-centric evaluations of Copilot or GPT-based systems. The benchmark covers correctness, runtime, memory usage, and security for a standardized set of Python algorithmic tasks, providing a concrete reference point for future work on edge and offline coding assistants.
2) End-to-End Reproducible Evaluation Framework

It introduces a fully automated evaluation pipeline, centred around the run_experiments.py and analyze_result.py scripts, that:

a) Generates code from multiple local models using consistent prompts.
b) Executes task-specific test suites, collects timing and memory statistics, and runs static security analysis.
c) Aggregates results into structured logs and human-readable reports.

This framework enables other researchers and practitioners to re-run, extend, or adapt the study on different hardware, models, or task sets with minimal changes.

3) Model Suitability Profiles for Real-World Deployment: It derives practical "suitability verdicts" for each evaluated model, rather than only reporting raw scores. In particular, the thesis identifies Mistral 7B as a "Speed King" for latency-sensitive local scenarios and Qwen 2.5 as a "Memory Miser" for RAM-constrained environments, while also characterizing Phi-3 and Gemma 2 in terms of their trade-offs. These profiles offer immediate decision support for architects and teams choosing a model under specific hardware and performance constraints.
4) Evidence for the Viability of Small Language Models (SLMs): By demonstrating that a 3.8B-parameter model (Phi-3 Mini) can match or closely approach the correctness and efficiency of larger 7B–9B models, the thesis contributes empirical evidence that small, well-tuned models are viable for code generation on edge and embedded systems. This strengthens the case for further investment in compact, domain-specialized LLMs.

C. Practical Implications

The findings of this study have direct implications for the software industry and sustainable computing:

1) For Enterprise Security: Organizations can safely deploy local models like Mistral 7B on internal air-gapped servers to assist developers. This eliminates the risk of Intellectual Property (IP) leakage associated with sending proprietary code to public cloud APIs (e.g., OpenAI, Anthropic). The 100% security pass rate observed in this study suggests these models are now "safe enough" for drafting utility functions and internal tooling without constant human security auditing.
2) For Embedded & Edge Computing: The proven efficiency of Qwen 2.5 (2.67 KB overhead) and Phi-3 demonstrates that AI coding assistance is no longer confined to high-end workstations. These models are lightweight enough to be integrated into Edge AI development tools, enabling on-device scripting for IoT hardware, robotics, and field devices where internet connectivity is intermittent and RAM is strictly limited.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XII Dec 2025- Available at www.ijraset.com

3) For Sustainable AI & Cost Reduction: Shifting from large cloud models (e.g., GPT-4 with ~1.7T parameters) to highly optimized local models (e.g., Mistral 7B) represents a massive reduction in inference energy costs. For companies running millions of daily code completions, adopting a 7B-parameter local model significantly lowers the carbon footprint and operational capability expenses (OpEx) of their development pipeline while maintaining functional parity for standard algorithmic tasks.

D. Limitations & Constraints

While this study provides valuable insights into local AI code generation, several limitations inherent to the experimental design and resource constraints must be acknowledged:

1) Hardware Constraints & Model Scale: The evaluation was strictly limited to 7B-9B parameter models due to the hardware constraints of the test environment (consumer-grade workstation with 32GB RAM). We were unable to benchmark larger local models (e.g., LLaMA-70B, Mistral 8x7B) which require significantly higher VRAM (48GB+). Consequently, this study cannot determine whether "scaling up" local models would resolve the latency outliers observed in architectures like Gemma 2, or if diminishing returns persist at higher parameter counts.
2) Dataset Scope & Complexity: The custom benchmark (Set A) comprised only six isolated algorithmic tasks. While these tasks effectively tested core logic (recursion, sorting, string manipulation), they do not represent the complexity of real-world software engineering, which involves multi-file dependencies, third-party library integrations, and legacy code refactoring. The 100% pass rate observed may be partially attributed to the "textbook" nature of these algorithms, which are likely highly represented in the models' training data.
3) Single-Language Bias: Experiments were conducted exclusively in Python. As a memory-managed, high-level language, Python masks certain low-level coding errors. The study does not address how these local models perform in memory-unsafe languages like C or C++, where the risk of AI-generated buffer overflows or pointer errors is significantly higher and security consequences are more severe.
4) Absence of Cloud-to-Local Baseline: Due to the unavailability of free, unlimited access to enterprise-grade cloud APIs (e.g., GPT-4 Turbo, Claude 3.5 Opus) for automated bulk benchmarking, this study lacks a direct, simultaneous comparison between local models and their state-of-the-art cloud counterparts. The comparison relies on historical data from literature (e.g., Li et al., 2024) rather than side-by-side execution, preventing a precise "Local vs. Cloud" performance delta analysis under identical prompt conditions.
5) Qualitative Metrics: The evaluation focused strictly on quantitative metrics (Runtime, Memory, Pass Rate). We did not perform a qualitative human evaluation of the code's readability, maintainability, or adherence to PEP-8 standards. It is possible that while a model like Mistral 7B produces correct and fast code, it may write "spaghetti code" that is difficult for human developers to maintain—a critical factor not captured by this automated benchmark.

E. Future Work

This study established a foundational baseline for local AI code generation, yet several critical avenues remain for exploration to fully understand the capabilities of edge-deployed LLMs:

1) Scaling to Repository-Level Complexity: Future studies should move beyond isolated functions to evaluate repository-level code generation. This involves testing a model's ability to maintain context across multiple files, handle third-party dependency injection, and perform refactoring on legacy codebases. Such a benchmark would determine if the context window limits of smaller 7B models (typically 8k-32k tokens) hinder their utility for real-world software engineering tasks.
2) Quantization vs. Security Trade-offs: While this study tested standard quantized models, a granular analysis of quantization precision (e.g., Q4_K_M vs. Q8_0 vs. FP16) is needed. Research should explicitly investigate whether extreme compression techniques—necessary for running LLMs on mobile devices—degrade the model's "security intuition," potentially leading to a higher rate of vulnerabilities (e.g., buffer overflows) despite maintaining functional correctness.
3) Cross-Language Security Generalization: Given Python's memory safety, expanding this benchmark to memory-unsafe languages (such as C, C++, and Rust) is essential. Future work should assess whether local LLMs can correctly manage manual memory allocation, pointer arithmetic, and concurrency primitives, where the margin for catastrophic security error is significantly narrower than in Python.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XII Dec 2025- Available at www.ijraset.com

4) Hardware-Specific Optimization: With the rise of specialized AI hardware (e.g., Neural Processing Units or NPUs in modern laptops), future research should benchmark these models on heterogeneous compute architectures. Evaluating how specific models utilize NPU acceleration versus standard GPU inference could unlock new efficiency frontiers for battery-powered development environments.
5) Qualitative Human Evaluation: Automated metrics cannot capture the "cognitive load" of reviewing AI code. Future studies should incorporate human-in-the-loop experiments to assess code readability, idiomatic correctness, and maintainability. A blind comparison where senior developers rate snippets from Mistral 7B vs. GPT-4 would provide the ultimate validation of whether local models generate "human-quality" code.

REFERENCES

[1] M. Ahmed, "Integrating AI-Driven Automated Code Review in Agile Development: Benefits, Challenges, and Best Practices," Proc. Of the International Conference on Software Engineering, 2025.
[2] S. Mishra, "AI-Augmented Vulnerability Detection and Patching," IEEE Transactions on Software Engineering, vol. 50, no.1, pp. 72-84, 2024.
[3] Y. Li, A. Ramanathan, and K. Zhao, "Assessing the Performance of AI-Generated Code: A Case Study on GitHub Copilot," arXiv preprint arXiv:2401.01234, 2024.

10.22214/IJRASET

45.98

IMPACT FACTOR:
7.129

IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY