

AI-Powered Unit Testing using Large Language Models (LLMs)

“If you’re not doing unit testing, you’re not doing software engineering; you’re doing software experimentation.”

— By Kent Beck.

Unit testing is the cornerstone of maintainable software development.

In today’s competitive digital landscape, software quality is a direct reflection of a company's brand and reputation. However, maintaining high-quality code is increasingly challenging as systems grow larger and more complex. Traditional unit testing processes often struggle to keep up, leading to incomplete coverage, human error, and delivery delays.




Fig 1: Integrating AI into Test environments

In our previous article [UnitTesting AI ML](#) we explored the growing need for **AI in Unit Testing**, especially in the context of modern software development. We discussed the challenges in traditional unit test automation, such as incomplete test coverage, time-consuming manual efforts, and lack of adaptability to code changes. We then introduced how machine learning (ML) can play a transformative role in overcoming these limitations by learning patterns, understanding code structures, and intelligently generating tests.

The article also presented a foundational workflow for integrating AI into unit testing and introduced key capabilities, including AI-Driven Test Case Generation (leveraging static and dynamic analysis), Code Coverage Analysis, Intelligent Test Prioritization, Anomaly Detection, and Automated Test Maintenance and Refactoring.

Additionally, it highlighted how the AI-powered solution can be seamlessly embedded within CI/CD pipelines to enable continuous, automated testing with every code change facilitating early detection and resolution of issues in the development lifecycle.

In this follow-up article, we take a deeper dive into two core features outlined in our previous roadmap:




Fig2: Core Features of Unit Testing

These two features form the backbone of an intelligent testing framework one that not only automates the generation of test cases but also evaluates their effectiveness with precision. In the sections that follow, we'll explore the inner workings of feature, demonstrate how they seamlessly integrate into existing development workflows.

AI-Driven Unit Test Case Generation: A Smarter Approach to Software Quality

Writing comprehensive unit tests has always been a crucial but time-consuming task for developers. It requires a deep understanding of complex code logic, manual effort to cover all scenarios including edge cases and adherence to consistent testing patterns. Under tight deadlines, this process can often be overlooked or rushed, leading to gaps in quality and coverage.

With the advent of Large Language Models (LLMs), we now have a transformative solution that changes the way we approach unit testing.

By leveraging LLMs trained on vast amounts of programming data and capable of code understanding, we can intelligently automate the generation of unit tests. Here's how:

- **Code Comprehension:** LLMs can parse and understand the structure, intent, and flow of source code, including functions, classes, and dependencies.
- **Scenario Detection:** They can identify key execution paths, edge cases, and potential failure points, and use this insight to propose meaningful test scenarios.

- **Test Generation:** Using natural language and code prompts, LLMs can generate high-quality test cases in real time, following best practices and framework-specific conventions.
- **Developer Control:** While the heavy lifting is automated, developers retain full control. They can review, edit, and refine generated tests to ensure alignment with project requirements and coding standards.

This approach drastically reduces the manual effort traditionally involved in test creation while improving speed, consistency, and coverage. More importantly, it frees developers to focus on higher-value tasks like feature development and architectural improvements.

LLM-driven unit testing is not just a productivity boost it's a strategic shift towards smarter, AI-assisted software engineering. As we continue integrating LLMs into the development lifecycle, the future of testing looks faster, more reliable, and significantly more intelligent.

Generated Unit Tests

```
void test_factorial() {
    assert(factorial(0) == 1);
    assert(factorial(1) == 1);
    assert(factorial(5) == 120);
    assert(factorial(10) == 3628800);
}

void test_isPrime() {
    assert(isPrime(-1) == 0);
    assert(isPrime(0) == 0);
    assert(isPrime(1) == 0);
}

void test_reverseString() {
    char str1[] = "hello";
    reverseString(str1);
    assert(strcmp(str1, "olleh") == 0);

    char str2[] = "";
    reverseString(str2);
    assert(strcmp(str2, "") == 0); // Edge case: empty string

    char str3[] = "a";
    reverseString(str3);
    assert(strcmp(str3, "a") == 0); // Edge case: single character

    char str4[] = "madam";
    reverseString(str4);
    assert(strcmp(str4, "madam") == 0); // Palindrome

    char str5[] = "123 456";
    reverseString(str5);
    assert(strcmp(str5, "654 321") == 0); // String with space

    char str6[] = "A man, a plan, a canal, Panama";
    reverseString(str6);
    assert(strcmp(str6, "amanaP ,lanac a ,nalp a ,nam A") == 0); // String with punctuation
}
```

Fig3: AI Generated Sample Unit Test Cases

Ensuring Quality with Code Coverage Analysis

While generating unit tests using Large Language Models (LLMs) significantly accelerates development, it's only half the equation. The real value of testing lies in understanding **how much of your code is being exercised**—because you can't ensure reliability or release with confidence without visibility into what's tested and what's not.

That's where **Code Coverage Analysis** becomes essential. To complement AI-generated tests, integrating comprehensive coverage evaluation provides actionable insights across key dimensions:

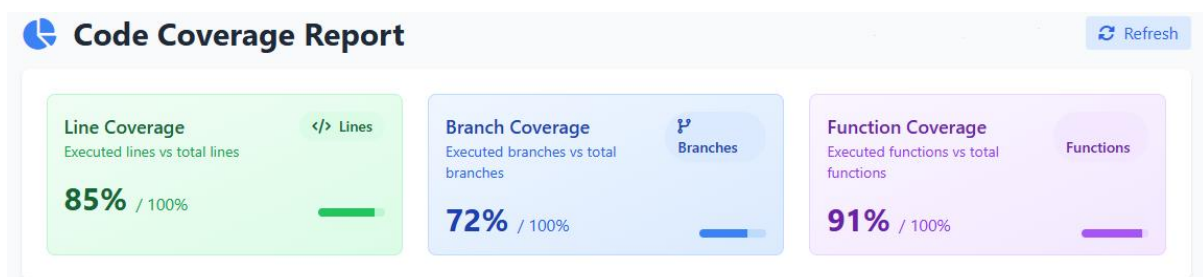


Fig4: Sample Unit Test Coverage Report

Line Coverage: Which lines of code are executed during tests.

Branch Coverage: Whether all conditional branches (if-else, switch, etc.) are explored.

Function Coverage: Whether all functions and methods are invoked during test execution.

By training LLMs with awareness of these coverage metrics, we can go beyond generating generic test cases—they begin to **prioritize untested logic, validate critical business paths, and optimize test strategies for high-risk areas**.

This synergy between AI-generated tests and coverage-driven insights enables development teams to:

- Identify blind spots and unreachable code.
- Improve test effectiveness and reduce technical debt.
- Release with confidence, knowing that the most critical logic is truly validated.

The future of testing isn't just automated—it's intelligent, strategic, and coverage-aware.

Process of Unit Testing using LLM

Unit test generation using Large Language Models (LLMs) follows a structured pipeline that leverages machine learning and natural language understanding to automate test creation. The process involves several key stages, each designed to optimize the model's ability to understand code and produce relevant, high-quality unit tests.

Secure Model Acquisition and Confidential Fine-Tuning

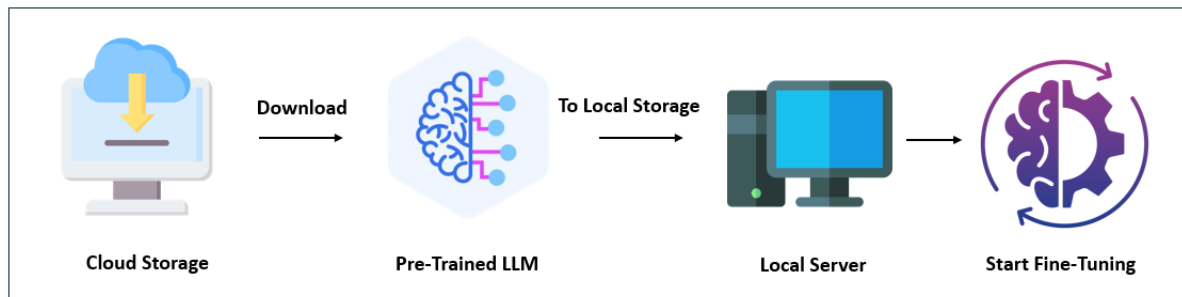


Fig5: Model Transfer and Localization

To ensure both the integrity of the source model and the confidentiality of sensitive data, we adopted a secure workflow for acquiring and fine-tuning our Large Language Model (LLM). The base model was downloaded from a trusted cloud repository and stored in our internal, access-controlled server environment. This step was performed over encrypted HTTPS connections, ensuring protection against tampering and man-in-the-middle attacks. Once securely downloaded, the model was disconnected from public access and moved into an **isolated internal environment** for further processing.

The **fine-tuning process was executed entirely on-premises**, within our secure infrastructure, without any outbound data transfer. This approach ensures that proprietary data, training artifacts, and domain-specific modifications remain completely confidential and protected under our internal data governance policies. By retaining full control over the model lifecycle, from acquisition to fine-tuning and deployment, we uphold rigorous standards of **data security, auditability, and regulatory compliance**.

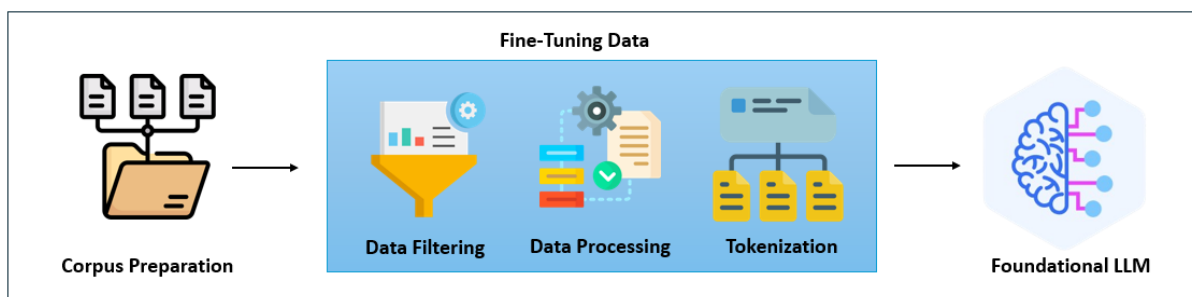


Fig6: Fine-Tuning Process of Dataset

Corpus Preparation

The first step involves building a robust dataset composed of source code and their corresponding unit tests. These examples are sourced from open-source repositories, developer benchmarks, and curated test suites. Each function is paired with a meaningful, well-structured test case to help the model learn the logical relationship between implementation and validation.

Fine-Tuning Data

Collected data undergoes preprocessing to ensure consistency and quality. This includes aligning each function with its respective test, standardizing formatting and naming conventions, and filtering out noisy or incomplete samples. The data is also adapted to match the LLM's tokenization process, allowing it to learn effectively from clean and structured input.

Code Representation for LLMs

Unlike traditional ML, LLMs do not require manual feature engineering. Instead, code is represented as token sequences enriched with contextual elements such as docstrings, type annotations, and logical indicators. This structured representation enhances the model's understanding of function intent and helps guide accurate test generation.

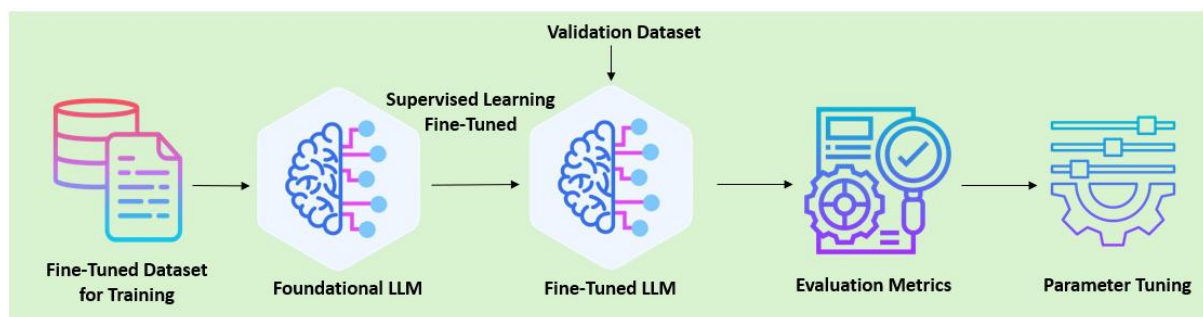


Fig7: Training Process of LLM

Model Training

The LLM is fine-tuned on the pre-processed dataset, enabling it to learn how to generate unit tests based on function definitions. Training involves exposure to diverse coding patterns, test strategies (e.g., assertions, edge cases), and best practices. Through this, the model learns not just to mirror syntax, but to apply logical reasoning to testing scenarios.

Model Validation

After training, the model is validated on unseen examples to assess its effectiveness. It is tested for syntactic correctness, behavioural accuracy, coverage impact, and runtime success.



Fig8: Model Inference

Deployment via Web Interface

To make the system accessible, the fine-tuned LLM is deployed within a chat-based web interface. Developers can input source code directly, and the system responds with generated test cases in real time. The interactive UI allows users to review and refine output, while providing insights like expected coverage making AI-powered unit testing intuitive and developer-friendly.

Enhancing Unit Testing with LLMs and AI

Integrating Large Language Models (LLMs) into unit testing marks a major shift from traditional manual and scripted automation. While manual test writing and maintenance can be time-consuming and error-prone, and even traditional automation requires significant setup and coding, LLMs dramatically streamline the process.

By enabling natural language-to-test generation, intelligent test suggestions, and seamless CI/CD integration, LLMs can **reduce testing time by 60–80%** compared to manual methods and up to **40–60%** compared to conventional automation. This not only accelerates test coverage and code quality but also improves developer productivity, enhances legacy system support, and ensures consistency across teams.

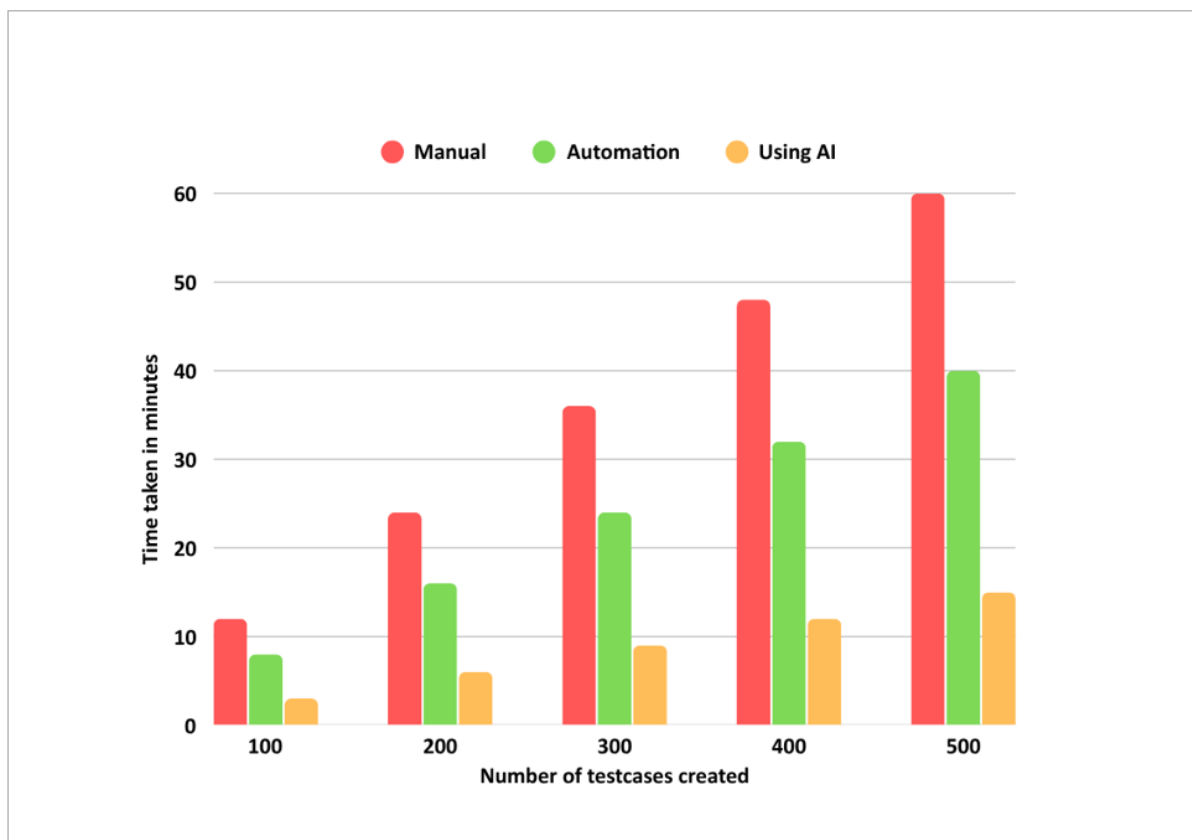


Fig9: Time-Effectiveness




Fig10: Key Advantages

Limitations of AI in Unit Testing

Data Dependency - AI models rely on historical data to generate tests, but if the data is incomplete or biased, the resulting tests may be suboptimal.

Difficulty Handling Complex Logic - AI can find patterns but might struggle to generate effective test cases for highly complex logic or algorithms.

Insufficient Mocking or Setup – AI struggles to handle asynchronous or multi-threaded code properly.

Our LLM is not limited to unit test generation, it is designed to support comprehensive test case creation, including integration, functional, regression, and edge-case testing. Although initially tailored for unit test automation, the model's architecture is inherently flexible and extensible, enabling it to adapt to a wide range of testing requirements across various software modules and environments. Its adaptability ensures it can meet diverse testing needs, regardless of application complexity or domain. This versatility empowers QA and development teams to scale their testing processes, enhance test coverage, and ensure high code quality while accelerating release cycles.

DISCLAIMER

The implementation of AI-powered unit testing using Large Language Models (LLMs) is an evolving area of research and practice. While the techniques described aim to enhance software testing efficiency and coverage, they should be used with human oversight and integrated into development workflows responsibly. Results may vary based on model configurations, datasets, and programming environments. Connect with Jasmin today for a personalized consultation and discover how AI-driven unit testing can strengthen your development process and accelerate software quality with confidence.

Contact Details

Jasmin Infotech Private Limited (HQ)
Plot 119 Velachery Tambaram Road
Pallikaranai, Chennai 600100
India

MS. JEYASUDHA / MR. NARENDRAN OVI

M: +91 89251 09996
narendran.ovi@jasmin-infotech.com