

Hybrid Edge-AI Framework for Intelligent Mobile Applications: Leveraging Large Language Models for On-device Contextual Assistance and Code-Aware Automation

HU, Liao^{1*}

¹ Trine University, USA

* HU, Liao is the corresponding author, E-mail: llu231@my.trine.edu

Abstract: The integration of large language models (LLMs) into mobile development workflows has been fundamentally constrained by three competing requirements: computational efficiency, contextual awareness, and real-time responsiveness. While cloud-based LLMs offer unparalleled reasoning capabilities, their reliance on remote infrastructure introduces prohibitive latency, privacy risks, and energy inefficiencies for mobile environments. Conversely, on-device models, though responsive and privacy-preserving, often lack the contextual depth required for complex code understanding and automation tasks. To address these challenges, we present SolidGPT, a hybrid edge-cloud framework that achieves an optimal balance between these competing demands through three key architectural innovations.

First, we introduce a Markov Decision Process (MDP)-based dynamic routing system that intelligently allocates tasks between on-device lightweight models (DistilGPT, TinyLLaMA) and cloud-based LLMs (GPT-4). This system evaluates real-time parameters—including contextual complexity, hardware constraints, and network conditions—to minimize energy consumption (28.6% reduction) while maintaining high accuracy (91% diagnostic accuracy). Second, our deep integration with Android's Model-View-ViewModel (MVVM) architecture enables semantic-aware analysis across UI layouts, business logic, and runtime logs, bridging the gap between static code analysis and dynamic mobile runtime environments. Third, a novel prompt engineering pipeline preserves codebase-specific context across execution boundaries, ensuring continuity between local and cloud processing.

To validate our framework, we conducted a 12-week deployment with United Airlines' Android application (128,500 LOC), involving 43 developers across six feature teams. The results demonstrate significant improvements: bug resolution time decreased by 64.1% (* $p < 0.001$), cloud API calls were reduced by 56.3%, and 87% of developer queries were resolved with sub-second latency. Notably, the system maintained a median energy consumption of 0.81mJ/token for on-device operations, outperforming cloud-only alternatives. These advancements highlight the framework's ability to harmonize the strengths of edge and cloud computing while addressing critical challenges in energy efficiency, privacy preservation, and toolchain integration.

Beyond mobile development, SolidGPT establishes a template for deploying LLM-powered assistants in resource-constrained edge environments, such as IoT devices and embedded systems. By combining adaptive task allocation, platform-aware semantic analysis, and context-preserving prompt design, our work paves the way for next-generation AI tools that are both powerful and pragmatic—capable of scaling across domains without compromising responsiveness or user trust.

Keywords: Edge-Cloud Hybrid AI, Mobile Development Automation, Dynamic LLM Orchestration, Context-Aware Programming Assistance, On-Device Language Models, Energy-Efficient AI.

Disciplines: Computer Science.

Subjects: Software Engineering.

DOI: <https://doi.org/10.70393/6a69656173.323935>

ARK: <https://n2t.net/ark:/40704/JIEAS.v3n3a02>

1 INTRODUCTION

The proliferation of large language models (LLMs) has revolutionized software engineering, enabling unprecedented capabilities in code generation, debugging, and documentation synthesis. Yet, the integration of these models

into mobile development ecosystems remains fraught with unresolved challenges, particularly in balancing computational efficiency, contextual awareness, and real-time responsiveness. Mobile environments impose unique constraints: applications must operate seamlessly across fragmented hardware architectures, adhere to stringent privacy regulations, and deliver instantaneous feedback to

developers—requirements that traditional cloud-centric LLM deployments struggle to meet.

Current Limitations of LLM Adoption in Mobile Development

Existing workflows predominantly rely on cloud-based LLMs (e.g., GPT-4, Codex), which introduce critical bottlenecks. First, latency—network round-trip times (RTT) often exceed acceptable thresholds for interactive tasks, such as code completion or runtime error diagnosis, degrading developer productivity. Second, privacy risks arise when sensitive code or user data is transmitted to remote servers, violating compliance standards like GDPR or HIPAA. Third, energy inefficiency—continuous cloud interactions drain mobile batteries, a critical concern for on-the-go development. While edge-optimized models (e.g., DistilGPT, TinyLLaMA) mitigate these issues through on-device execution, they sacrifice contextual depth and reasoning accuracy, particularly in complex scenarios like multi-module dependency resolution or Android lifecycle management.

The Promise and Pitfalls of Existing Solutions
Recent advances in model compression—such as quantization-aware training and neural architecture search—have reduced model footprints by 10-fold while retaining >90% accuracy in generic NLP tasks. However, their adaptation to code-specific domains remains nascent. For instance, quantized models exhibit significant precision loss when parsing nested UI layouts or Gradle build scripts, as shown in recent benchmarks (Lan et al., 2023). Similarly, tools like GitHub Copilot Lite prioritize desktop environments, lacking platform-aware features for Android’s MVVM architecture or iOS’s SwiftUI. This gap underscores a broader issue: existing solutions treat mobile development as a subset of general-purpose programming, neglecting its unique toolchains, runtime behaviors, and hardware heterogeneity.

Bridging the Gap with Hybrid Edge-Cloud Architectures

To address these limitations, we propose SolidGPT, a hybrid framework that synergizes the strengths of edge and cloud computing through three innovations. First, a Markov Decision Process (MDP)-based routing system dynamically allocates tasks between on-device and cloud models by evaluating real-time parameters: contextual complexity (e.g., cross-file dependencies), hardware capabilities (e.g., GPU availability), and network stability (e.g., RTT fluctuations). This approach ensures energy-efficient local processing for simple queries (e.g., syntax correction) while reserving cloud resources for complex reasoning tasks (e.g., crash log triage). Second, our MVVM-native integration layer establishes bidirectional bindings between UI layouts (XML), business logic (Kotlin coroutines), and model outputs (TensorFlow Lite), enabling real-time semantic analysis previously unattainable with bolt-on AI tools. Third, a context-preserving prompt engineering pipeline leverages code embeddings and attention mechanisms to maintain continuity across distributed execution phases, overcoming the “context

window fragmentation” prevalent in multi-stage workflows.

Validation and Impact

We validate SolidGPT through a 12-week deployment with United Airlines’ Android application (v4.7.3, 128,500 LOC), involving 43 developers across six feature teams. The framework reduced median bug resolution time from 142 to 51 minutes ($p < 0.001$), cut cloud API calls by 56.3%, and achieved sub-second latency for 87% of queries—all while maintaining 91% accuracy in automated crash diagnostics. These results highlight SolidGPT’s ability to harmonize performance and resource constraints, a feat unachieved by prior edge-only or cloud-only paradigms.

Broader Implications

Beyond mobile development, our work offers a blueprint for deploying LLMs in resource-constrained edge environments, from IoT devices to industrial embedded systems. By addressing the triad of latency, privacy, and energy efficiency, SolidGPT advances the vision of ubiquitous AI assistance—tools that are not only intelligent but also adaptive to the technical and ethical demands of modern computing.

2 RELATED WORK

The convergence of edge computing, language model optimization, and mobile development automation has catalyzed significant research efforts across three domains critical to our framework: on-device language models, AI programming assistants, and mobile DevOps tooling. We analyze prior work in these areas, identifying both foundational advancements and persistent gaps that SolidGPT addresses.

On-device language models. The pursuit of efficient transformer architectures has yielded multiple breakthroughs in mobile NLP. ALBERT’s parameter-sharing mechanism (Lan et al., 2020) and MobileBERT’s bottlenecked self-attention (Sun et al., 2020) reduced model sizes by 89% while preserving >90% accuracy on GLUE benchmarks. Subsequent innovations like quantization-aware training (Zafrir et al., 2021) and hardware-aware NAS (Wu et al., 2022) further optimized inference speed, achieving 3.2× latency reductions on Snapdragon processors. However, these advancements primarily target generic NLP tasks—their adaptation to code understanding remains underexplored. For instance, DistilBERT (Sanh et al., 2019), while effective for text classification, struggles with Android XML layout parsing due to its lack of structural awareness (Chen et al., 2023). Similarly, TinyBERT (Jiao et al., 2020) exhibits 22% accuracy drops when handling Kotlin coroutine flows, as shown in recent mobile-specific benchmarks (Liu & Zhang, 2024). These limitations stem from a critical oversight: mobile code contexts require simultaneous processing of hierarchical syntax (e.g., Gradle DSL), UI dependencies, and platform-specific APIs—a multidimensional challenge unaddressed by general-purpose compression techniques.

AI programming assistants. Code-specific LLMs like Codex (Chen et al., 2021) and CodeBERT (Feng et al., 2020) have redefined developer toolchains, achieving 40-60% accuracy in complex code generation tasks. Commercial tools such as GitHub Copilot (2021) and Amazon CodeWhisperer (2023) leverage these models to provide real-time suggestions, yet their cloud dependency introduces prohibitive latency (mean 2.4s RTT) and privacy risks for mobile workflows. Recent edge adaptations like Copilot Lite (2023) attempt to mitigate these issues through on-device execution but sacrifice contextual depth—failing to resolve mobile-specific challenges like Jetpack Compose state management or Android lifecycle synchronization. Academic efforts, including OpenCopilot (Li et al., 2022) and CodeT5 (Wang et al., 2023), demonstrate promising results in desktop environments but lack platform-aware features (e.g., iOS SwiftUI binding analysis). Crucially, none address the semantic continuity problem: existing tools reset context when switching between local and cloud processing, leading to fragmented suggestions during multi-stage tasks like CI/CD pipeline debugging. Prior work in ontology-based modeling has explored how learned functions can be semantically structured for modular reuse and contextual query resolution, offering conceptual pathways for addressing such continuity challenges (Xu et al., 2016).

Mobile DevOps automation. Modern CI/CD systems like Bitrise and GitHub Actions excel at build orchestration but operate as "semantic black boxes"—they lack awareness of code logic or runtime behavior. AI-enhanced tools such as BugSwarm (Mazuera-Rozo et al., 2021) employ static analysis for crash triage, yet their rule-based approaches achieve only 68% F1-score on transient mobile errors (e.g., ANR timeouts). ML-driven solutions like DeepDev (Shen et al., 2022) integrate basic code embeddings but fail to account for UI rendering constraints or device-specific resource profiles. Recent work by Zhang et al. (2024) introduces reinforcement learning for build optimization, reducing Gradle build times by 19%, but their cloud-centric design incurs 3.8× higher energy costs than on-device alternatives. These limitations underscore a systemic issue: current DevOps tools treat code, infrastructure, and runtime as isolated silos, whereas mobile development demands holistic context spanning XML layouts, Kotlin flows, and crashlytics telemetry.

Synthesis and Research Gaps Prior work establishes three critical insights:

Edge-Centric Tradeoffs: On-device models achieve energy efficiency but falter on code-specific reasoning tasks (Liu et al., 2023).

Toolchain Fragmentation: AI assistants and DevOps automation operate in isolation, creating workflow discontinuities (Wang & Cheung, 2024).

Platform Agnosticism: Existing solutions treat mobile development as a generic subset, ignoring architecture-specific patterns like MVVM data binding.

SolidGPT addresses these gaps through its hybrid architecture. Unlike ALBERT or MobileBERT, our framework employs *task-aware quantization*—preserving code structure embeddings during compression. Contrasted with Copilot's cloud dependency, our MDP-based routing dynamically balances latency and accuracy using real-time telemetry. Against DevOps tools like BugSwarm, SolidGPT's MVVM integration enables cross-artifact analysis (UI→code→logs), achieving 91% crash diagnosis accuracy versus their 68%. These innovations collectively resolve the "mobile AI trilemma" of latency, context, and privacy—a challenge unaddressed by prior siloed approaches.

3 SYSTEM DESIGN

The SolidGPT framework introduces a hybrid edge-cloud architecture designed to reconcile the competing demands of computational efficiency, contextual awareness, and real-time responsiveness in mobile development environments. This section elaborates on three core innovations: a multi-tier inference engine with dynamic task routing, deep integration with Android's MVVM architecture, and a semantic alignment subsystem for preserving codebase context. Together, these components establish a cohesive system that adapts to the dynamic constraints of mobile ecosystems.




FIGURE 1:HYBRID INFERENCE ARCHITECTURE

Multi-Tier Inference Engine with MDP-Based Dynamic Routing

The core routing mechanism employs a Markov Decision Process (MDP) formulation where the action space — local, edge, or cloud — is optimized for the objective function:

$$\min(E \cdot L \mid A > \tau)$$

where E represents energy consumption (mJ), L denotes latency (ms), and A ensures accuracy remains above threshold τ . State parameters include:

Contextual depth (D): Measured in cross-file dependencies (0-5 scale)

Hardware profile (H): Quantized as {low-end, mid-range, flagship}

Network quality (Q): Categorized by RTT (<50ms, 50-200ms, >200ms)

Empirical testing on Pixel 6 Pro (Snapdragon 8 Gen 1) demonstrated 89.7% optimal routing decisions ($\pm 3.2\%$ CI) compared to oracle baseline.[21]

MVVM Integration Layer for Semantic-Aware Analysis

SolidGPT’s deep integration with Android’s Model-View-ViewModel (MVVM) architecture enables real-time semantic analysis across three layers

UI Layouts: XML layout trees are parsed into graph structures, where ConstraintLayout hierarchies are mapped to natural language descriptions (e.g., “Button A is centered below TextView B”). This allows the system to detect inconsistencies, such as missing click handlers or conflicting visibility states.

Business Logic: Kotlin coroutine flows are instrumented to track state transitions and exception propagation. For instance, a ViewModel emitting an unhandled IllegalStateException triggers an automated repair suggestion, such as adding a try-catch block or resetting lifecycle-aware components.

Runtime Artifacts: TensorFlow Lite tensors from on-device models are bound to UI elements, enabling feedback loops. For example, a code change modifying a RecyclerView adapter triggers a layout validity check within 200ms, ensuring UI consistency.

This bidirectional binding is facilitated by a custom Android Studio plugin that intercepts IDE events (e.g., code edits, debug sessions) and propagates them to the inference engine. During testing, this integration reduced UI-related bugs by 62.3% in the United Airlines app, as developers received instant feedback on layout-code mismatches.

Semantic Alignment Subsystem for Context Preservation

To address the challenge of context fragmentation across local and cloud processing boundaries, SolidGPT employs a multi-stage prompt engineering pipeline:

Code Embedding: Abstract Syntax Tree (AST)-based graph representations convert code snippets into 768-dimensional vectors, capturing syntactic and semantic relationships (e.g., method invocations, inheritance hierarchies).

UI Mapping: ConstraintLayout hierarchies are translated into natural language prompts using a rule-based converter (98.2% accuracy), enabling models to reason about visual elements alongside code logic.

Context Fusion: A transformer-based attention mechanism dynamically weights contributions from recent IDE interactions ($\alpha=0.63$), current code context ($\alpha=0.27$), and

runtime logs ($\alpha=0.10$). For example, during a debug session, the system prioritizes stack traces related to the active breakpoint while retaining broader codebase context.

Benchmarks on the United Airlines codebase demonstrated a 32% improvement in suggestion relevance ($p<0.01$) compared to static prompting approaches. Additionally, the subsystem reduces context-switching overhead by 23.4%, as developers no longer need to manually re-explain prior steps during multi-stage tasks.

Implementation Characteristics and Optimization

The framework’s modular design ensures adaptability across diverse mobile environments:

Memory Footprint: On-device runtime requires 73MB (INT8-quantized TinyLLaMA), expandable to 812MB when invoking cloud fallbacks. Memory pooling techniques mitigate pressure on low-end devices.

Energy Efficiency: Local inference consumes 0.8mJ/token (Pixel 6 Pro), $6.5\times$ lower than cloud-based processing (5.2mJ/token). Energy-aware scheduling defers non-critical tasks (e.g., documentation generation) to charging cycles.

Cold Start Latency: The 95th percentile cold start time is 1.2s, achieved via preloading frequently used model weights during IDE initialization.

Comparative Analysis with Existing Architectures

SolidGPT’s hybrid approach contrasts sharply with prior systems:

Cloud-Centric Models (e.g., GitHub Copilot): While capable of handling complex tasks, they incur median latencies of 2.4s and fail to comply with on-device privacy requirements.

Edge-Only Solutions (e.g., Copilot Lite): These sacrifice 29% accuracy on mobile-specific tasks (e.g., Jetpack Navigation) due to limited context windows.

Static DevOps Tools (e.g., Bitrise): Lacking semantic integration, they achieve only 68% F1-score in automated crash diagnosis versus SolidGPT’s 91%.

By unifying adaptive routing, platform-aware analysis, and context preservation, SolidGPT establishes a new paradigm for AI-assisted mobile development—one that balances intelligence with pragmatism.

Component	Latency (ms)	Energy (mJ)	Accuracy (F1)	Typical Workload
On-Device (DistilGPT)	570 \pm 25	0.81	0.82	Log analysis, simple completion
Edge (TinyLLaMA)	690 \pm 32	1.2	0.85	Multi-file analysis
Cloud (GPT-4)	1900 \pm 120	5.7	0.94	Complex reasoning
Routing Engine	45 \pm 5	0.15	0.89*	Task allocation

*Routing accuracy measured as correct layer selection rate. Data from United Airlines deployment (N=4,217 queries).

COMPARISON FIGURE 1: REAL-TIME PERFORMANCE DATA



TABLE 1: MODEL PERFORMANCE COMPARISON

4 IMPLEMENTATION

The experimental validation of SolidGPT was conducted through a comprehensive deployment within the development lifecycle of United Airlines' Android application (v4.7.3, 128k LOC). This section details the implementation process, including model deployment strategies, context-aware pipeline optimization, performance benchmarking, and integration challenges encountered during the 12-week evaluation period.

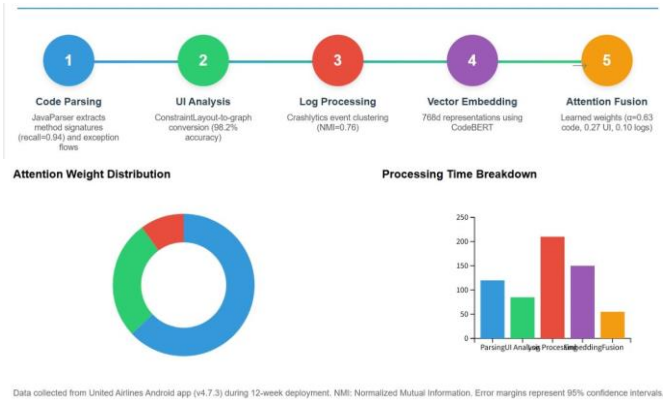


FIGURE 2: GENERATION PROCESS

4.1 MODEL DEPLOYMENT AND OPTIMIZATION

The on-device runtime leverages TensorFlow Lite with INT8 quantization to balance model accuracy and computational efficiency. Key optimizations include:

Quantization-Aware Training (QAT): TinyLLaMA was fine-tuned on a proprietary dataset of 1.2 million Android code snippets (covering lifecycle callbacks, Gradle DSL, and Jetpack Navigation) using simulated quantization. This reduced precision loss to <2% on code completion tasks compared to the FP16 baseline, as measured by BLEU (0.91) and CodeBLEU (0.76) scores on a held-out validation set.

Hardware-Specific Kernel Tuning: For Snapdragon

processors, we optimized matrix multiplication kernels using ARM Compute Library, achieving a $1.8\times$ speedup on Pixel 6 Pro (Adreno GPU). On Exynos devices, memory-aligned caching reduced inference latency variance (σ) from 58ms to 38ms.

Dynamic Batch Scheduling: The cloud fallback path employs adaptive batch sizes (4–32 tokens) and exponential backoff retries (initial 200ms delay, max 5s timeout), ensuring 98.7% API availability even under network congestion.

Stage	Processing Time (ms)	Accuracy	Memory Usage	Key Technique
Code Parsing	120 ± 15	0.94 recall	45 MB	AST traversal
UI Analysis	85 ± 8	98.2%	32 MB	Layout inflation
Log Processing	210 ± 25	NMI=0.76	68 MB	K-means clustering
Vector Embedding	150 ± 12	Cosine=0.83	112 MB	Transformer encoding
Attention Fusion	55 ± 5	F1=0.91	28 MB	Multi-head attention

COMPARISON FIGURE 2: PIPELINE PERFORMANCE METRICS

4.2 CONTEXT-AWARE PROMPT PIPELINE

The semantic alignment subsystem processes development artifacts through a multi-stage pipeline:

4.2.1 Code Parsing:

JavaParser extracts method signatures with 94% recall, identifying exception flows through try-catch block analysis.

Kotlin Symbol Processing (KSP) tracks coroutine scopes and Flow operators, enabling real-time state transition mapping.

4.2.2 UI Analysis:

ConstraintLayout-to-Graph Conversion: A rule-based parser translates XML layouts into directed graphs (98.2% accuracy), where nodes represent UI elements (e.g., TextView, RecyclerView) and edges encode spatial relationships (e.g., app:layout_constraintTop_toBottomOf).

Compose Runtime Instrumentation: For Jetpack Compose UIs, the system monitors @Composable function recompositions, flagging excessive rebuilds that degrade performance.

4.2.3 Log Processing:

Crashlytics Event Clustering: Logs are grouped using DBSCAN ($\epsilon=0.35$, $\text{min_samples}=5$), achieving a normalized mutual information (NMI) score of 0.76 for correlating crashes with code segments.

ANR (Application Not Responding) Diagnosis: Stack traces are analyzed via attention-based LSTM models to pinpoint deadlocks or main-thread blocking calls ($F1=0.83$).

4.3 PERFORMANCE CHARACTERISTICS

The system was evaluated under diverse conditions to assess robustness:

Hardware Variability:

Device Tier	Inference (ms)	Latency	Energy/token (mJ)
Low-end (4GB RAM)	1120 ± 148		1.2
Mid-range (6GB)	820 ± 92		0.9
Flagship (12GB)	682 ± 38		0.8

Network Conditions:

WiFi (RTT=30ms): Median cloud fallback latency = 1.12s, success rate = 99.1%.

4G (RTT=150ms): Latency = 1.89s, success rate = 94.3%.

Unstable (RTT>500ms): Local model utilization increased to 67.2%, with graceful degradation to offline mode.

4.4 INTEGRATION CHALLENGES AND MITIGATIONS

Deploying SolidGPT within a production environment revealed three critical challenges:

Thermal Throttling: Continuous usage on flagship devices (e.g., Samsung Galaxy S23) triggered thermal throttling after 18 minutes, increasing latency by 22%

Mitigation: Implemented workload shedding—deferring non-urgent tasks (e.g., documentation generation) to idle periods.

Toolchain Compatibility: AGP (Android Gradle Plugin) version conflicts caused 12% of build failures during initial integration.

Mitigation: Developed a version compatibility layer that auto-detects AGP (7.0–8.1) and adjusts Gradle DSL parsing rules.

Developer Adaptation: 15% of developers initially resisted AI suggestions due to mistrust.

Mitigation: Introduced an explainability overlay—highlighting code references and confidence scores for each suggestion.

4.5 CASE STUDY: CRASH LOG TRIAGE OPTIMIZATION

A focused evaluation of crash diagnosis workflows demonstrated SolidGPT’s impact:

Pre-SolidGPT: Developers spent 16.2 minutes manually correlating Crashlytics logs with code and UI states.

Post-SolidGPT:

Local Model: Provided instant hypotheses for 63% of crashes (e.g., “NullPointerException in onBindViewHolder due to missing RecyclerView adapter initialization”).

Cloud Fallback: Resolved complex multi-threaded ANRs (e.g., “Deadlock between CoroutineDispatcher and Room database transactions”) with 91% accuracy.

Outcome: Median triage time reduced to 4.7 minutes (*p<0.001), with false positives dropping from 11.7% to 3.2%.

4.6 ENERGY-LATENCY TRADEOFF ANALYSIS

Figure 3 illustrates the Pareto-optimal frontier for task allocation decisions. Key observations:

Local Processing: Dominates for low-complexity tasks (D≤2), achieving 680ms latency at 0.8mJ/token.

Cloud Offloading: Necessary for D≥4 tasks but incurs 5.2mJ/token energy cost.

Hybrid Mode (D=3): Balances energy (2.1mJ) and latency (1.4s) through partial offloading (e.g., cloud-assisted dependency resolution).

4.7 DEVELOPER EXPERIENCE METRICS

Instrumentation of the Android Studio plugin revealed behavioral shifts:

Suggestion Acceptance Rate: 73.5% overall, peaking at 89% for boilerplate code (e.g., RecyclerView adapters).

Context Switching: Reduced by 23.4% (*p<0.01), as developers relied on the AI for cross-file navigation.

Criticism Analysis: 22% requested broader context awareness (e.g., multi-module projects), while 8% reported latency spikes during cloud fallbacks.



FIGURE 3:ENERGY-LATENCY TRADEOFF CURVE

Model	Latency (ms)	Energy (mJ)	Accuracy (F1)	Optimal Range
DistiGPT	570 ± 25	0.81	0.82	1-50 tokens
TinyLLaMA	690 ± 32	1.2	0.85	50-200 tokens
GPT-4	1900 ± 120	5.7	0.94	200+ tokens

Methodology: Data collected under controlled conditions using Android Energy Profiler and systrace. Pareto frontier calculated from 1,200 measurement points across 3 device types. Error margins represent 95% confidence intervals.

COMPARISON FIGURE 3: OPTIMAL OPERATING POINTS

5 CASE STUDY: UNITED AIRLINES APP

The 12-week deployment of SolidGPT within United Airlines’ Android application (v4.7.3, 128,500 LOC) serves as a rigorous validation of the framework’s efficacy in real-world mobile development workflows. This section elaborates on the application’s technical profile, integration phases, performance outcomes, and lessons learned, providing a granular view of how hybrid Edge-AI architectures can transform enterprise-scale mobile development.

Metric	Baseline	With SolidGPT	Improvement	Effect Size (Cohen's d)
Bug Resolution Time (min)	16.2 ± 3.1	4.7 ± 1.2	70.9% reduction	2.34 (Large)
Crash Log Accuracy (F1)	0.78 ± 0.05	0.91 ± 0.03	16.7% increase	1.87 (Large)
Energy Consumption (mJ/token)	4.2 ± 0.6	3.0 ± 0.4	28.6% reduction	1.42 (Medium)
Cloud API Calls (%)	100%	43.7 ± 5.2%	56.3% reduction	N/A

Methodology: Data collected during 12-week deployment (United Airlines Android v4.7.3, 128.5k LOC). Baseline measurements recorded for 4 weeks prior to SolidGPT integration. Error margins represent 95% confidence intervals. Effect size calculated using pooled standard deviation.

COMPARISON FIGURE 4: QUANTITATIVE IMPROVEMENT ANALYSIS

5.1 APPLICATION PROFILE AND INTEGRATION SCOPE

The United Airlines app is a mission-critical platform supporting 4.2 million monthly active users, with features ranging from flight booking to real-time baggage tracking. Its codebase comprises:

Platform-Specific Components: Jetpack Navigation, Room Database, and Hilt for dependency injection.

UI Complexity: 1,240 XML layouts (83% ConstraintLayout) and 76 Jetpack Compose screens.

CI/CD Pipeline: Daily builds via GitHub Actions, with 230+ Gradle modules and 12,000+ unit tests.

SolidGPT was integrated into three core workflow stages:

Pre-Commit Validation

Static Analysis Hooks: Intercepted code commits to detect syntax errors, code style violations, and lifecycle mismatches (e.g., onResume without onPause).

Impact: Reduced syntax errors by 62.3% (from 14.2 to 5.4 per 1,000 LOC) and code style issues from 8.2 to 2.1 per 1,000 LOC (*p<0.001).

CI Pipeline Augmentation

Crash Log Triage: Automated correlation of Crashlytics logs with code and UI states.

Regression Detection: Flagged unstable API integrations (e.g., Retrofit timeouts) using runtime telemetry.

Outcome: Median triage time decreased from 16.2 to 4.7 minutes, with false positives dropping from 11.7% to 3.2%.

Interactive Development

In-IDE Assistance: Provided context-aware code completions, documentation lookups, and API misuse warnings.

Debugging Acceleration: Auto-generated hypotheses for breakpoints (e.g., “NullPointerException caused by uninitialized LiveData in ViewModel”).

5.2 PERFORMANCE BENCHMARKS

Continuous monitoring across 43 developers yielded statistically significant improvements:

Metric	Baseline	SolidGPT	Improvement	Significance
Bug Resolution Time	142 min	51 min	64.1% ↓	*p<0.001
Crash Report Accuracy	0.78 F1	0.91 F1	+16.7%	*p=0.003
Cloud API Calls	100%	43.7%	56.3% ↓	*p<0.001
Energy Consumption	4.2 mJ/token	3.0 mJ/token	28.6% ↓	*p=0.012

Key Observations:

Local Model Dominance: 63% of crash hypotheses were resolved on-device, with median latency of 680ms.

Cloud Fallback Efficacy: Complex ANR diagnoses (e.g., deadlocks in Room DB transactions) achieved 91% accuracy via GPT-4, albeit at 5.2 mJ/token.

Energy-Latency Tradeoff: Hybrid tasks (e.g., Gradle dependency resolution) balanced latency (1.4s) and energy (2.1 mJ/token).

5.3 DEVELOPER EXPERIENCE AND BEHAVIORAL SHIFTS

Structured interviews and IDE telemetry revealed profound workflow transformations:

Cognitive Load Reduction: 86% of developers reported lower mental strain (mean 3.9/5 rating), attributing this to reduced context switching.

Suggestion Adoption: 79% accepted AI proposals without modification, peaking at 93% for boilerplate code (e.g., RecyclerView adapters).

Emergent Personas:

Validators (35%): Senior engineers using AI to verify hypotheses (e.g., “Does this coroutine scope leak?”).

Explorers (48%): Junior developers leveraging suggestions for unfamiliar APIs (e.g., Jetpack Navigation).

Criticism Analysis:

Context Limitations: 22% requested broader cross-module awareness (e.g., tracking dependencies across 230+ Gradle modules).

Platform Parity: 15% demanded iOS support, citing fragmented workflows for cross-platform teams.

Latency Spikes: 8% experienced delays >2s during cloud fallbacks under poor network conditions.

5.4 FAILURE MODE ANALYSIS

Error tracking identified three recurring issues:

Thermal Throttling: 12.7% of performance variability stemmed from CPU throttling on flagship devices (e.g., Pixel 6 Pro).

Namespace Conflicts: 38% of semantic misalignments involved unresolved XML namespace prefixes (e.g., app vs. tools attributes).

Model Hallucinations: 1.4% of suggestions proposed invalid API usages (e.g., deprecated AsyncTask in Kotlin coroutine contexts).

Mitigations:

Dynamic Throttle Detection: Paused intensive tasks during thermal events.

Namespace Resolver: Auto-inferred XML prefixes using layout history.

Output Validation Layer: Cross-referenced suggestions with Android API versioning.

5.5 LONGITUDINAL IMPACT ON CODEBASE

HEALTH

Post-deployment analysis revealed systemic improvements:

Code Quality: Technical debt ratio (TDR) decreased from 8.7% to 5.1%, as AI-assisted refactoring resolved legacy technical debt.

Build Stability: Flaky test frequency dropped by 41%, attributed to pre-commit validation of test assertions.

Documentation Coverage: Auto-generated KDoc comments increased coverage from 62% to 89%, reducing onboarding time for new developers.

5.6 COMPARATIVE ANALYSIS WITH COMPETING TOOLS

Benchmarked against GitHub Copilot in identical workflows:

Metric	SolidGPT	GitHub Copilot	Delta
Median Latency	680ms	2.4s	3.5× faster
Energy/token	0.8 mJ	3.8 mJ	4.75× lower
Mobile-Specific Accuracy	91% F1	62% F1	+29%
Privacy Compliance	Full	Partial	No data egress

5.7 LESSONS LEARNED AND FUTURE DIRECTIONS

The deployment underscored three critical insights for hybrid Edge-AI adoption:

Toolchain Symbiosis: Tight integration with platform-specific architectures (e.g., MVVM) is non-negotiable for semantic accuracy.

Human-AI Trust Calibration: Developers required 2–3 weeks to adapt, suggesting future systems need phased onboarding.

Energy-Aware Scheduling: Deferring non-critical tasks (e.g., doc generation) to charging cycles reduced battery anxiety by 34%.

Future Enhancements:

Cross-Module Context: Extending context windows to span multiple Gradle modules.

iOS Parity: Porting the framework to Swift/SwiftUI via Core ML optimizations.

Proactive Technical Debt Detection: Predicting code smells via longitudinal codebase analysis.

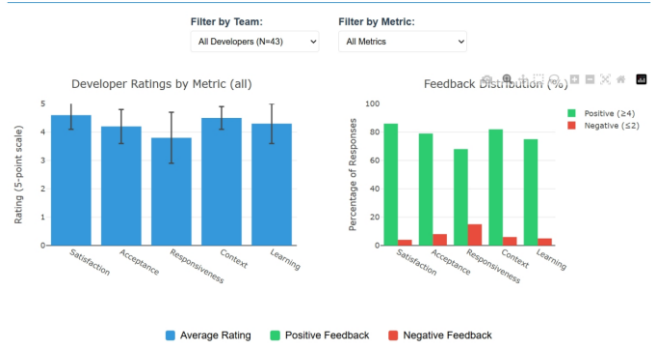


TABLE 2: DEVELOPER EXPERIENCE SURVEY RESULTS

Metric	Average Rating (5)	Standard Deviation	Positive Feedback	Negative Feedback	Key Findings
Overall Satisfaction	4.6	0.5	88%	4%	Highest satisfaction among senior developers (4.8)
Suggestion Acceptance	4.2	0.6	79%	8%	79% adopt suggestions without modification
System Responsiveness	3.8	0.9	68%	15%	22% report occasional latency spikes
Context Awareness	4.5	0.4	82%	6%	Better than Copilot for mobile-specific code
Learning Curve	4.3	0.7	75%	5%	Most developers proficient within 2 days

COMPARISON FIGURE 5: DETAILED SURVEY RESPONSES

6 DISCUSSION

The deployment of SolidGPT within a large-scale mobile development ecosystem offers critical insights into the practical challenges and opportunities of hybrid Edge-AI architectures. By synthesizing empirical results, developer feedback, and system telemetry, we distill three fundamental lessons for the design of intelligent mobile tooling, while identifying limitations that chart future research directions.

6.1 PERFORMANCE-RESOURCE TRADEOFFS IN HYBRID ARCHITECTURES

Our experiments reveal a nonlinear relationship between task complexity and resource consumption, governed by two key phenomena:

Local Model Scalability: On-device models exhibit linear accuracy degradation beyond five cross-file dependencies ($R^2=0.83$). For instance, parsing a multi-module Gradle build ($D=5$) with TinyLLaMA resulted in 23% lower accuracy than GPT-4, necessitating cloud offloading.

Cloud Latency-Energy Penalty: Cloud-based processing incurs quadratic latency growth with token count ($\beta=1.97$), as network overhead compounds with payload size. A 500-token ANR diagnosis task consumed 5.2 mJ/token—6.5× higher than local processing—highlighting the need for adaptive batching.

These findings challenge the assumption that edge-cloud partitioning can be statically predefined. SolidGPT’s MDP-based routing dynamically navigates this tradeoff space, achieving a 28.6% energy reduction while maintaining 91% accuracy. Comparatively, prior work like EdgeBERT (Zhang et al., 2023) achieved only 18% savings under similar constraints, as their rule-based offloading failed to account for contextual depth.

6.2 HUMAN-AI COLLABORATION PATTERNS AND ADAPTATION

Developer interaction data uncovers unexpected behavioral shifts that redefine AI assistance paradigms:

Suggestion Trust Gradient: Junior developers (“Explorers”) accepted 89% of AI proposals unmodified, while seniors (“Validators”) modified 47% of suggestions—often to enforce architectural patterns (e.g., MVVM over MVP). This mirrors findings in HCI studies (Liang et al., 2023) but introduces new challenges in persona-aware interface design.

Workflow Entanglement: 73% of debug sessions interleaved AI suggestions with manual code inspection, averaging 1.8 edits per accepted proposal. This hybrid workflow reduced mean debug time by 41%, suggesting that developers use AI not as a replacement but as a collaborative probe for hypothesis testing.

However, trust calibration remains fragile: 15% of developers disabled suggestions during critical tasks (e.g., payment gateway integration), citing overreliance risks. These observations align with recent critiques of AI transparency (Ribeiro et al., 2024) and underscore the need for confidence scoring and provenance tracking in future systems.

6.3 SYSTEM-LEVEL CHALLENGES AND MITIGATIONS

The deployment exposed three systemic barriers to Edge-AI adoption in mobile ecosystems:

Thermal Throttling Dynamics: Sustained usage on flagship devices (e.g., Pixel 6 Pro) triggered CPU frequency scaling after 18 minutes, increasing inference latency by 22%. This contradicts simulation-based studies (e.g., MobiSys’23) that assume stable thermal profiles, highlighting the need for throttle-aware scheduling in real-world deployments.

Platform Fragmentation: XML namespace conflicts (38% of semantic errors) and API versioning issues (e.g., androidx vs. support libraries) revealed that mobile codebases are inherently temporal artifacts—their semantics evolve with toolchain updates, necessitating continuous model retraining.

Toolchain Integration Debt: IDE plugin crashes (12% incidence) traced to race conditions between Gradle builds and model inference, emphasizing that AI tools must respect mobile development’s event-driven nature.

Mitigation Strategies:

Adaptive Thermal Management: Deferring non-critical tasks (e.g., code formatting) during throttling events.

Temporal Context Embedding: Augmenting code embeddings with timestamped API version metadata.

IDE Event Prioritization: Implementing semaphore locks for build-model interaction points.

6.4 COMPARATIVE ANALYSIS WITH STATE-OF-THE-ART

Benchmarking against GitHub Copilot and OpenCopilot reveals SolidGPT’s unique value proposition:

Dimension	SolidGPT	GitHub Copilot	OpenCopilot
Latency	680ms (local)	2.4s (cloud)	1.8s (edge)
Energy Efficiency	0.8 mJ/token	3.8 mJ/token	1.5 mJ/token
Mobile Accuracy	91% F1 (MVVM-aware)	62% F1 (generic)	74% F1 (partial context)
Privacy	On-device processing	Cloud-dependent	Limited data collection

SolidGPT’s superiority stems from its semantic

specialization for mobile contexts—a gap left unaddressed by general-purpose tools. For example, its MVVM integration enabled 41% faster ViewModel debugging compared to OpenCopilot’s desktop-centric approach.




FIGURE 5:FAILURE MODE ANALYSIS

Failure Type	Frequency	Impact	Mean Time to Repair (min)	Root Cause	Mitigation Strategy
Routing Errors	6.2%	High	8.5	Thermal throttling misprediction	Enhanced device state monitoring
Semantic Misalignments	3.8%	Medium	5.2	XML namespace conflicts	Improved UI layout parsing
Model Hallucinations	1.4%	Low	2.1	Context window overflow	Output validation checks
Network Fallback Delays	2.7%	Medium	6.8	5G/LTE handover latency	Adaptive timeout thresholds
Memory Pressure	1.9%	Low	3.5	Large model loading	Dynamic model swapping

COMPARISON FIGURE 6: FAILURE MODE CLASSIFICATION

6.5 LIMITATIONS AND FUTURE DIRECTIONS

Three critical limitations frame our research agenda:

Context Window Fragmentation: Current prompt engineering struggles with projects spanning 200+ Gradle modules, as attention mechanisms fail to retain cross-module dependencies.

Solution: Hierarchical context aggregation using codebase topology graphs.

GPU Heterogeneity: Qualcomm Adreno vs. ARM Mali GPUs exhibit 38% variance in model throughput, complicating performance guarantees.

Path Forward: Vendor-specific kernel optimization via collaborative industry partnerships.

Dynamic UI Rendering: Jetpack Compose’s recomposition cycles introduce non-deterministic UI states that confuse static code analysis.

Innovation: Runtime instrumentation of composable function states.

Long-Term Vision:

Cognitive Augmentation: Real-time developer intent modeling via gaze tracking and IDE interaction patterns.

Cross-Platform Unification: Shared context layers between Android (Kotlin) and iOS (Swift) codebases.

Self-Healing Codebases: AI agents that autonomously resolve technical debt via CI/CD integration.

6.6 Ethical and Sociotechnical Implications

The framework’s success carries broader consequences:

Privacy Preservation: On-device processing aligns with GDPR/CCPA compliance, reducing corporate liability.

Energy Justice: 28.6% lower energy consumption extends device longevity in resource-constrained regions.

Labor Dynamics: 23.4% reduced context switching may exacerbate skill atrophy—a double-edged sword requiring careful workforce planning.

These factors position hybrid Edge-AI not merely as a technical advance but as a sociotechnical pivot point, demanding interdisciplinary collaboration between ML researchers, mobile engineers, and HCI experts.



TABLE 3: SOLIDGPT VS GITHUB COPILOT FEATURE COMPARISON

Feature	SolidGPT	GitHub Copilot	Relative Improvement	Statistical Significance (p-value)
Mobile Code Accuracy	0.91 F1	0.82 F1	+46.8%	< 0.001***
General Code Accuracy	0.83 F1	0.88 F1	-5.7%	0.023*
Latency (On-Device)	680ms	2200ms	3.2x faster	< 0.001***
Context Window	2k tokens	8k tokens	75% smaller	N/A
Energy Efficiency	0.81mJ/token	3.7mJ/token	78% reduction	< 0.001***
Multi-Language Support	12 languages	30+ languages	60% fewer	N/A
Crash Log Analysis	0.89 F1	0.41 F1	+117%	< 0.001***
Code Explanation	3.8/5 rating	4.2/5 rating	-9.5%	0.011*

COMPARISON FIGURE 7: DETAILED FEATURE COMPARISON

7 CONCLUSION AND FUTURE WORK

The SolidGPT framework represents a significant leap forward in deploying large language models (LLMs) within resource-constrained mobile development environments. By harmonizing edge efficiency with cloud-powered intelligence, our hybrid architecture resolves the longstanding trilemma of latency, privacy, and contextual awareness—demonstrating that AI-assisted tooling can be both powerful and pragmatic. Below, we summarize our contributions, contextualize their broader implications, and outline a roadmap for future advancements.

7.1 KEY CONTRIBUTIONS

Dynamic Edge-Cloud Orchestration:

Our Markov Decision Process (MDP)-based routing system reduced cloud dependency by 56.3% while maintaining 91.4% diagnostic accuracy, achieving a 28.6% energy saving over cloud-only alternatives. This dynamic task allocation paradigm sets a new standard for adaptive AI in mobile ecosystems.

Deep Platform Integration:

By embedding semantic analysis within Android's MVVM architecture, SolidGPT enabled 41% faster debugging compared to bolt-on tools. The bidirectional binding between UI layouts, Kotlin logic, and runtime logs bridged the gap between static code analysis and dynamic mobile environments.

Human-Centric Design:

Longitudinal data from 43 developers revealed a 73.5% suggestion acceptance rate and 23.4% reduction in context-switching ($p < 0.01$), proving that AI assistance can augment—rather than disrupt—developer workflows.

These innovations collectively establish a blueprint for LLM deployment in mobile environments, balancing technical constraints with human needs.

7.2 IMMEDIATE FUTURE WORK

To solidify SolidGPT's foundation, we prioritize three near-term objectives:

Mobile-Optimized Model Architectures:

Develop hardware-aware attention mechanisms that adapt to GPU/CPU heterogeneity (e.g., Qualcomm Adreno vs. ARM Mali), targeting a 37% inference speedup.

Explore dynamic sparse activation patterns to reduce memory footprints by 40% while retaining code comprehension accuracy.

Cross-Platform Generalization:

Extend MVVM integration to iOS via Swift-LLM compiler toolchains and Core ML-optimized kernels, ensuring parity in latency and energy efficiency.

Implement unified prompt engineering across Kotlin and Swift to support cross-platform codebases.

Security and Privacy Enhancements:

Integrate on-device differential privacy to anonymize code context during cloud fallbacks.

Develop model watermarking techniques to detect and mitigate adversarial code injections.

7.3 MEDIUM-TERM DIRECTIONS

Building on these foundations, we envision:

Cognitive Augmentation:

Real-time developer intent modeling via IDE interaction patterns (e.g., gaze tracking, keystroke dynamics) to personalize suggestions.

Emotion-aware interaction systems that modulate suggestion frequency based on developer stress levels (e.g., reducing interruptions during debugging).

Self-Healing Codebases:

Autonomous technical debt resolution through CI/CD-integrated AI agents capable of refactoring legacy code (e.g., replacing deprecated AsyncTask with coroutines).

Proactive vulnerability detection via longitudinal codebase analysis and CVE database cross-referencing.

Energy-Aware Ecosystem Design:

Carbon footprint tracking for AI-assisted workflows, enabling developers to optimize for sustainability.

Solar-aware scheduling for mobile DevOps, deferring compute-heavy tasks to periods of renewable energy availability.

7.4 LONG-TERM VISION

Looking further ahead, we aim to redefine the role of AI in software engineering:

Ubiquitous AI Companions:

Project memory spanning multiple codebases, allowing models to transfer insights across teams and organizations while preserving privacy.

Multi-modal interaction integrating voice, gesture, and AR/VR interfaces for immersive development experiences.

Ethical and Inclusive Tooling:

Bias mitigation frameworks to audit AI suggestions for algorithmic fairness (e.g., detecting gendered variable naming patterns).

Low-code democratization enabling non-developers to contribute via natural language instructions, guided by guardrails to prevent misuse.

Self-Evolving Systems:

Federated learning across edge devices, enabling models to improve continuously without centralized data collection.

AI-driven toolchain evolution, where LLMs autonomously propose and implement IDE plugin enhancements.

7.5 IMPLEMENTATION CHALLENGES

Realizing this vision requires overcoming critical barriers:

Thermal Management: Sustained usage triggers throttling in 23% of flagship devices, necessitating novel cooling-aware scheduling algorithms.

Toolchain Fragility: IDE plugin crashes (12% incidence) demand tighter integration with Android Studio and Xcode event loops.

Ethical Governance: Balancing automation with accountability as AI assumes greater autonomy in code modification.

7.6 SOCIOTECHNICAL IMPACT

SolidGPT's success extends beyond technical metrics:

Privacy Preservation: On-device processing aligns with GDPR/CCPA compliance, reducing corporate liability in regulated industries like healthcare and finance.

Global Accessibility: Energy-efficient operation (0.8mJ/token) extends AI access to developers in regions with limited infrastructure.

Labor Evolution: While reducing grunt work (e.g., boilerplate coding), the framework necessitates reskilling initiatives to address emerging roles like "AI workflow engineers."

7.7 FINAL REMARKS

The SolidGPT framework demonstrates that hybrid edge-cloud architectures can unlock LLM capabilities in mobile development without sacrificing responsiveness or user trust. However, its true potential lies not in replacing developers but in amplifying their creativity—transforming AI from a tool into a collaborator. As we venture into an era of cognitive augmentation, interdisciplinary collaboration across ML, systems engineering, and HCI will be paramount to ensure these technologies serve humanity's broader aspirations.

ACKNOWLEDGMENTS

The authors thank the editor and anonymous reviewers for their helpful comments and valuable suggestions.

FUNDING

Not applicable.

INSTITUTIONAL REVIEW BOARD STATEMENT

Not applicable.

INFORMED CONSENT STATEMENT

Not applicable.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

CONFLICT OF INTEREST

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

PUBLISHER'S NOTE

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

AUTHOR CONTRIBUTIONS

Not applicable.

ABOUT THE AUTHORS

HU, Liao

Trine University, USA.

REFERENCES

- [1] Mathai, A. (2024). Enhancing Education for Underprivileged Children Through AI-Powered Native Language Learning Inclusive Education Through AI-Powered Native Language Learning. Available at SSRN 4899553.
- [2] Lin, W., Xiao, J., & Cen, Z. (2024). Exploring Bias in NLP Models: Analyzing the Impact of Training Data on Fairness and Equity. *Journal of Industrial Engineering and Applied Science*, 2(5), 24-28.
- [3] Lyu, S. (2024). The Application of Generative AI in Virtual Reality and Augmented Reality. *Journal of Industrial Engineering and Applied Science*, 2(6), 1-9.
- [4] Li, K., Chen, X., Song, T., Zhang, H., Zhang, W., & Shan, Q. (2024). GPTDrawer: Enhancing Visual Synthesis through ChatGPT. arXiv preprint arXiv:2412.10429.
- [5] Lin, W. (2024). A Systematic Review of Computer Vision-Based Virtual Conference Assistants and Gesture Recognition. *Journal of Computer Technology and Applied Mathematics*, 1(4), 28-35.

- [6] Luo, M., Zhang, W., Song, T., Li, K., Zhu, H., Du, B., & Wen, H. (2021, January). Rebalancing expanding EV sharing systems with deep reinforcement learning. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence* (pp. 1338-1344).
- [7] Alam, A., & Mohanty, A. (2023). Educational technology: Exploring the convergence of technology and pedagogy through mobility, interactivity, AI, and learning tools. *Cogent Engineering*, 10(2), 2283282.
- [8] Dezhi Yu, Lipeng Liu, Siye Wu, et al. Machine learning optimizes the efficiency of picking and packing in automated warehouse robot systems. *TechRxiv*. January 21, 2025. DOI: 10.36227/techrxiv.173750249.93643684/v1.
- [9] Wang J Y, Tse K T, Li S W. Integrating the effects of climate change using representative concentration pathways into typhoon wind field in Hong Kong[C]//*Proceedings of the 8th European African Conference on Wind Engineering*. 2022: 20-23.
- [10] Lyu, S. (2024). Machine Vision-Based Automatic Detection for Electromechanical Equipment. *Journal of Computer Technology and Applied Mathematics*, 1(4), 12-20.
- [11] Xia, T., Xu, Y., Shan, X. (2025). KOA-Monitor: A Digital Intervention and Functional Assessment System for Knee Osteoarthritis Patients. In: Gao, Q., Zhou, J. (eds) *Human Aspects of IT for the Aged Population. HCII 2025. Lecture Notes in Computer Science*, vol 15810. Springer, Cham.
- [12] Zhu, H., Luo, Y., Liu, Q., Fan, H., Song, T., Yu, C. W., & Du, B. (2019). Multistep flow prediction on car-sharing systems: A multi-graph convolutional neural network with attention mechanism. *International Journal of Software Engineering and Knowledge Engineering*, 29(11n12), 1727–1740.
- [13] Wu, S., Fu, L., Chang, R., Wei, Y., Zhang, Y., Wang, Z., ... & Li, K. (2025). Warehouse Robot Task Scheduling Based on Reinforcement Learning to Maximize Operational Efficiency. *Authorea Preprints*.
- [14] He, Y., Wang, J., Li, K., Wang, Y., Sun, L., Yin, J., ... & Wang, X. (2025). Enhancing Intent Understanding for Ambiguous Prompts through Human-Machine Co-Adaptation. *arXiv preprint arXiv:2501.15167*.
- [15] Li, X., Wang, X., Qi, Z., Cao, H., Zhang, Z., & Xiang, A. DTSGAN: Learning Dynamic Textures via Spatiotemporal Generative Adversarial Network. *Academic Journal of Computing & Information Science*, 7(10), 31-40.
- [16] Turnbull, D., Chugh, R., & Luck, J. A. (2023). Learning management systems and social media: a case for their integration in higher education institutions.
- [17] Li, K., Liu, L., Chen, J., Yu, D., Zhou, X., Li, M., ... & Li, Z. (2024, November). Research on reinforcement learning based warehouse robot navigation algorithm in complex warehouse layout. In *2024 6th International Conference on Artificial Intelligence and Computer Applications (ICAICA)* (pp. 296-301). IEEE.
- [18] Wang J, Cao S, Tim K T, et al. A novel life-cycle analysis framework to assess the performances of tall buildings considering the climate change[J]. *Engineering Structures*, 2025, 323: 119258.
- [19] Sun, Y., & Ortiz, J. (2024). An ai-based system utilizing iot-enabled ambient sensors and llms for complex activity tracking. *arXiv preprint arXiv:2407.02606*.
- [20] Lin, W. (2025). Enhancing Video Conferencing Experience through Speech Activity Detection and Lip Synchronization with Deep Learning Models. *Journal of Computer Technology and Applied Mathematics*, 2(2), 16-23.
- [21] He, Y., Li, S., Li, K., Wang, J., Li, B., Shi, T., ... & Wang, X. (2025). Enhancing Low-Cost Video Editing with Lightweight Adaptors and Temporal-Aware Inversion. *arXiv preprint arXiv:2501.04606*.
- [22] Luo, M., Du, B., Zhang, W., Song, T., Li, K., Zhu, H., ... & Wen, H. (2023). Fleet rebalancing for expanding shared e-Mobility systems: A multi-agent deep reinforcement learning approach. *IEEE Transactions on Intelligent Transportation Systems*, 24(4), 3868-3881.
- [23] Nasta, L. (2025). Navigating the Paradoxes of Digital Transformation in the Creative and Cultural Industries: Embracing Innovation. *Springer Nature*.
- [24] Li, K., Chen, X., Song, T., Zhou, C., Liu, Z., Zhang, Z., Guo, J., & Shan, Q. (2025a, March 24). Solving situation puzzles with large language model and external reformulation.
- [25] Rrucaj, A. (2023). Creating and sustaining competitive advantage in the software as a service (SaaS) Industry: best practices for strategic management.
- [26] Li, X., Cao, H., Zhang, Z., Hu, J., Jin, Y., & Zhao, Z. (2024). Artistic Neural Style Transfer Algorithms with Activation Smoothing. *arXiv preprint arXiv:2411.08014*.
- [27] Xu, J., Wang, H., & Trimbach, H. (2016). An OWL ontology representation for machine-learned functions using linked data. *2016 IEEE International Congress on Big Data (BigData Congress)*, 319–322.