Privacy-Layered Model Agents: Architecting Locally Executed, Speech Cognitive Automation for Multi-Domain Intelligent Systems

Dr. Ian Cox, Dr. Vincent Wang, Dr. Ivy Zhang
Personal Intelligence Technology
Hong Kong, Shenzhen, Beijing, Taipei
cox@ai-agent.com, wang@ai-agent.com, zhang@ai-agent.com

ABSTRACT

Autonomous agents will inevitably operate at speed of human thinking and be embedded into our lives with increasing far-reaching consequences, enhancing efficiency and exacerbating systemic risks if deployed without control.

LAPSCA (Local Autonomous Speech Cognitive Automation) is a reference architecture for multi-domain, intelligent, task-specific AI agents that operate locally on consumer-grade hardware, processing sensitive data without cloud dependency. It integrates local speech processing, cognitive reasoning, and tool-based execution within a privacy-layered, multi-agent architecture designed for reliability, auditability, and human-centered control.

LAPSCA is not a generic automation tool. It is a composable framework for building task-specific agents across domains. Whether you are a trader optimizing a quant strategy, a professional streaming workflows, or someone simplifying personal tasks, LAPSCA provides the architecture to build, deploy, and control agents that work for you locally, privately, and reliably. This paper presents LAPSCA's design principles, layered architecture, and a reference implementation for multi-domain deployment.

The LAPSCA Framework highlights five critical findings. Modern hardware enables sophisticated AI inference locally, eliminating cloud dependency for latency-sensitive and privacy-critical applications. Speech-to-text has matured to the point where natural language becomes a viable interface for agent control. Model Context Protocol (MCP) provides a standardized, extensible interface between agents and tools. Layered safety architecture with kill-switch governance is essential for preventing runaway agent behaviors. Task-specific agents outperform general-purpose assistants in reliability, latency, and user trust within their domains.

While some proposed platforms attempt to aggregate tools, they often fail to deliver actionable insights or facilitate secure execution. These tools provide the interface but cannot synthesize data into meaningful strategies or ensure reliable execution. LAPSCA addresses these gaps by integrating knowledge retrieval, flexible reasoning, and secure execution into a single, extensible deployment. Furthermore, it incorporates a layered safety architecture and human-in-the-loop controls, ensuring that automation remains transparent, auditable, and under the user's control.

Imagine a world where every user, every person, and every digital interaction can be an operative semi-autonomously executing, analyzing and learning 24/7 to achieve specific objectives with speed, precision, and adaptability. This paper explores the architecture, design principles, and implementation of such a system, LAPSCA (Local Autonomous Speech Cognitive Automation), a framework for building reliable, auditable, and human-centered AI agents that operate locally on consumer hardware. We demonstrate the framework's applicability across three distinct domains: quantitative trading, professional streaming workflows, and personal task automation. These systems demonstrate that, contrary to prevailing cloud-centric AI paradigms, complex cognitive automation can be achieved with high reliability and strong privacy guarantees when architected for local execution and human-in-the-loop control.

2. Firm-Funded AI for Blockchain Decision-Making

2.1 Blockchain Specific Intelligence & Voice Controlled Operations

LAPSCA combines the speed of automated decision-making in crypto markets with human oversight and strategic refinement that traditional AI systems lack.

At the core of LAPSCA is a layered architecture that processes data through specialized agents operating at different speeds and with different scopes. The system combines local speech processing, cognitive reasoning, tool-based execution, and a multi-layered safety framework to deliver a system that is responsive, auditable, and under the user's control. The architecture is designed to be modular and extensible, allowing new tools, new reasoning patterns, and new domain-specific agents to be added without disrupting the core system.

The data layer has a persistent through-line of short-term memory.

Supervised Data Processing The model can ingest a wealth of structured or unstructured data in a repeatable, scalable manner. Complex, analysing critical transactions, and compiling on-chain data. The model generates insights and recommendations, which are then presented to the user for review and approval before any action is taken.

Human-Supervised Strategy Refinement The system facilitates a collaborative loop between the AI and the user. The AI proposes strategies, analyzes performance, and suggests adjustments to parameters or logic. The user provides high-level guidance, approves strategy changes, and intervenes when the AI encounters ambiguity or novel situations. This ensures that the system leverages the AI's analytical power while maintaining human oversight and control over the strategy's evolution.

Strategic Offline Job Queue. LAPSCA allows users to define complex, multi-step tasks, monitoring wallet transactions, and compiling on-chain data. The system can then execute these jobs autonomously when the user is offline, delivering the final analysis or a summary upon their return. This feature transforms the AI from a tool requiring constant supervision into a true digital assistant that can work in the background.

2.2 Full Decision Control, Limitless Automation at Your Fingertips

LAPSCA aims to be an enabling ecosystem rather than a closed, walled garden. The core framework is open-source, allowing advanced users to inspect, modify, and contribute to the system. The architecture is designed to be extensible, with clear interfaces for adding new tools, new reasoning patterns, and new domain-specific agents. The system can be deployed on various hardware configurations, from a standard laptop to a dedicated server, and can be configured to prioritize local processing for maximum privacy or leverage cloud APIs for enhanced performance when appropriate.

Imagine, While you sleep, LAPSCA scrapes the entire crypto market, analyzing trends, identifying alpha opportunities, and preparing a ready-to-review strategy by morning. This transforms the user from a manual operator into a strategic director, leveraging the system's tireless analysis and execution capabilities to maintain a competitive edge in the fast-paced crypto market.

Complete Control, Customizable Strategies The user retains ultimate control over the system's operation. The system allows for the creation and management of multiple, independent strategies, each with its own set of rules, data sources, and execution parameters. The user can define the AI's role in each strategy, from simple data aggregation to complex, multi-step analysis and execution.

Seamless Integration with Trading Platforms The system is designed to interface with various cryptocurrency exchanges and trading platforms, enabling automated execution of trades based on the AI's analysis and the user's predefined rules. This integration is crucial for leveraging the speed and precision of algorithmic trading while maintaining human oversight.

Secure, Local Execution Prioritizing Privacy The system can be configured to run entirely on local hardware, ensuring that sensitive trading data and strategies never leave the user's control. This is a significant advantage over cloud-based AI trading platforms, providing enhanced security and privacy.

Transparent and Auditable Decision Making The system should provide clear logs and explanations of the AI's reasoning for each trade or action, allowing the user to audit the system's performance and understand the rationale behind its decisions.

2.3 Multi-Source Data Integration

LAPSCA integrates numerous data sources to provide a comprehensive view of the market. The system can ingest data from various exchanges, blockchain explorers, social media platforms, and news aggregators. The AI then synthesizes these disparate data streams to identify patterns, correlations, and potential trading opportunities that may be overlooked by human analysts or traditional algorithms.

2.4 Risk Management Integration

LAPSCA integrates sophisticated risk management capabilities to safeguard capital. The system can be configured with various risk parameters, such as maximum position sizes, stop-loss orders, and exposure limits across different assets or strategies. The AI can also be trained to recognize patterns indicative of high-risk market conditions and to dynamically adjust its trading behavior accordingly, for example, reducing position sizes or temporarily halting trading during periods of extreme volatility.

2.5 Multi-Chain Knowledge Graphs for Real-Time Market Synthesis

LAPSCA leverages multi-chain knowledge graphs to provide the AI with a deep understanding of the cryptocurrency ecosystem. These graphs can include information on token relationships, smart contract interactions, wallet activity, governance proposals, and social sentiment. The AI can then use this knowledge to identify complex patterns and relationships that may be indicative of market movements or emerging trends.

2.6 Proxy-Pool Composition

LAPSCA mitigates rate limits and proxy bottlenecks by intelligently distributing requests across a pool of proxy servers. The system can dynamically allocate requests to different proxies based on their performance, geolocation, and the specific requirements of the target website. This ensures uninterrupted data collection and maximum efficiency, even when dealing with rate-limited APIs or geo-restricted content.

The platform's leverage is rather large, compiling paradigms and fundamental trading frameworks in market high-volume arbitrage opportunities. The system can identify and exploit arbitrage opportunities across different exchanges or between different trading pairs, capitalizing on price discrepancies or inefficiencies efficiently and leaving no information or potential arbitrageurs. These capabilities extend beyond simple price arbitrage to more complex, multi-leg arbitrage strategies that may involve interacting with various DeFi protocols. To begin with, global data production is exploding (the DGDPA and CSTAI, CAGR of 25% and 23% respectively) and Generative AI is rapidly transforming how data is created and consumed. The sheer volume and complexity of this data, coupled with the need for real-time analysis to maintain a competitive edge, necessitate the development of sophisticated, AI-driven systems capable of processing and synthesizing information at scale. Traditional, manual approaches are simply not sufficient to keep pace with the speed and complexity of the modern data landscape. The LAPSCA framework is designed to address this challenge by leveraging the power of AI to automate the entire data processing pipeline, from data ingestion and cleaning to analysis and visualization, enabling users to extract actionable insights from vast and complex datasets with unprecedented speed and efficiency.

Additionally, LAPSCA can also implemented strategy analysis such as Interflipnetwork The System LAPSCA and Knowledge-based Multi-Agent Systems (KBMAS) for Strategic Decision Support in Algorithmic Trading. The framework can be adapted to analysis and synthesize vast amounts of research and contextual information, such as news articles, social media sentiment, and on-chain data, to inform trading decisions. This capability can provide a significant edge in understanding market dynamics and anticipating potential price movements.

2.6 Use Cases

LAPSCA is a universal technology capable of automating and optimizing tasks across Retail and Retail Activities. Its flexible agentic nature lets firms prioritize gathering, summarizing, and analyzing trading-relevant information in a trading workflow automation.

Some of the use cases are as follows (not exhaustive):

Automated Trading and Arbitrage Rapidly scan order book depth, liquidity, exchange opportunities, and optimize price strategies with near-instantaneous execution.

Market Sentiment Analysis Seamlessly monitor trades accounts across platforms like Pinterest, Wordpress, and Dubtoo, consolidating engagement data into actionable insights.

Data Feeds Trading Agents Access real-time, high-frequency market data, including order book information, trade data, and sentiment analysis. LAPSCA can aggregate and process this data to provide the trading agent with a comprehensive view of the market.

Portfolio Management Automate the process of rebalancing portfolios based on predefined rules or AI-driven recommendations, analyzing market trends and risk tolerance to optimize asset allocation and maximize returns.

Risk Management Implement real-time monitoring of market conditions and automatically adjust trading parameters to mitigate potential losses. This includes setting stop-loss orders, limiting exposure to specific assets, and dynamically adjusting trade sizes based on volatility and other risk factors.

Sentiment Analysis Integrate with social media and news feeds to gauge market sentiment. LAPSCA can analyze keywords, hashtags, and mentions to identify potential market movements driven by news events or social trends.

Backtesting and Optimization Leverage historical data to backtest trading strategies and optimize their performance. LAPSCA can automate the process of running simulations and analyzing the results to identify the most profitable strategies.

Customer Support Automate responses to frequently asked questions regarding trading strategies, platform features, and technical issues. This can free up human support agents to focus on more complex and critical tasks.

Data Analysis Perform sentiment analysis on customer feedback, identify trends in customer behavior, and generate reports on key performance indicators (KPIs) to inform business decisions.

Content Generation Generate marketing materials, such as blog posts, social media updates, and email newsletters, to promote trading services and attract new customers.

Figure 1 High Level Architecture Diagram

a. Functional Architecture

The functional blueprint serves as the primary passing point for users, equipped with both voice-enabled and text-based inputs, utilizing Automated Speech Recognition (ASR) to process commands. The system's core is managed by a central orchestrator, which routes tasks to specialized agents for data ingestion, analysis, and execution. The architecture also incorporates a feedback loop for continuous learning, where the system learns from past runs and adapts to market changes with minimal delay. The critical control action integrates all trading functionalities, enabling the system to execute trades, manage portfolios, and interact with various exchanges. This end-to-end functionality is crucial for automating the entire trading workflow, from data gathering to order execution.

b. Interactive Agent & Back-to-Manual

The core interface with LAPSCA is designed to be user-friendly for both technically savvy users and non-traders. Its accessibility is paramount, ensuring a wider adoption. The system supports voice commands for hands-free operation, allowing users to initiate tasks, monitor performance, and receive alerts while multitasking. The system provides real-time feedback on its actions and performance, keeping the user informed and in control. The architecture is designed to be modular, allowing for the easy integration of new tools and functionalities as required by different trading strategies.

c. Data & Trading Layer

The data and trading layer serves LAPSCA by aggregating data from Centralized Exchanges (CEXs), Decentralized Exchanges (DEXs), and on-chain sources. This layer is crucial for enhancing decision-making capabilities by providing a comprehensive view of the market. The system leverages the Model Context Protocol (MCP) to standardize tool integrations, ensuring seamless and secure communication between the AI agents and various data sources and execution venues. The architecture is designed to be modular and extensible, allowing for the easy integration of new data sources and trading venues as required by different trading strategies.

The layer is built on a modular, microservice-based architecture, where each blockchain interaction (reading data or executing transactions) is encapsulated within a dedicated microservice. This design ensures scalability, flexibility, and fault tolerance, allowing for the addition of new blockchains or protocols without disrupting the core system. The system also incorporates a robust error handling and retry mechanism to ensure reliable execution even in the face of network issues or exchange outages.

d. Multi-Chain & Cross-Market Support

LAPSCA's architecture supports both chain-level and asset-level operations, allowing users to manage assets across Ethereum, Solana, Bitcoin, and Base, with potential for future expansion. It utilizes multi-chain frameworks like Ethereum Attestation Service (EAS) and cross-chain messaging protocols like Wormhole to enable seamless interoperability between different blockchain networks.

Users can execute cross-chain trades or bridge assets by querying data from through integrated APIs. For example, a user might instruct the system to find the best price for USDC across Ethereum, Solana, and Base, and then execute the trade on the chain with the most favorable conditions, accounting for gas fees and bridging costs.

The platform's layer architecture supports advanced cross-chain strategies, such as monitoring liquidity across different chains to identify potential arbitrage opportunities or assessing the risk associated with bridging assets between different networks. The system can also be configured to prioritize certain chains or exchanges based on factors like liquidity, security, or transaction fees.

e. Human-in-the-Loop Governance

While LAPSCA aims for high levels of automation, it also incorporates a human-in-the-loop governance model to ensure that the system's actions align with the user's intent and risk tolerance. The system can be configured to require user approval for certain actions, such as large trades or interactions with new protocols. The user can also set predefined rules and boundaries for the AI agents, such as maximum trade sizes, allowed trading pairs, and stop-loss limits.

The platform's governance framework allows users to define and manage different trading strategies, each with its own set of rules, data sources, and execution parameters. The user can monitor the performance of each strategy and intervene if necessary to adjust the parameters or pause the strategy altogether.

3. Verifiable Intelligence & Voice Controlled Operations

LAPSCA prioritizes the creation of auditable and verifiable AI agents, ensuring that their actions are transparent and accountable. This is crucial for building trust and mitigating the risks associated with autonomous decision-making. The framework incorporates several mechanisms to achieve this, including detailed logging of all agent actions, the generation of cryptographic proofs of execution, and the integration with blockchain technology for immutable record-keeping.

Voice-controlled AI agents offer a more intuitive and accessible interface for interacting with complex systems. However, this also introduces new security and privacy considerations. The LAPSCA framework addresses these concerns by implementing robust voice authentication and command verification mechanisms, ensuring that only authorized users can access and control the system. The framework also supports the use of local, on-device speech processing to minimize the reliance on cloud services and enhance privacy.

3.1 Data Flow

LAPSCA is a universal technology capable of automating and optimizing tasks across Retail and Retail Activities. Its flexible agentic nature lets firms prioritize gathering, summarizing, and analyzing trading-relevant information in a trading workflow automation.

Some of the use cases are as follows (not exhaustive):

Automated Trading and Arbitrage Rapidly scan order book depth, liquidity, exchange opportunities, and optimize price strategies with near-instantaneous execution.

Market Sentiment Analysis Seamlessly monitor trades accounts across platforms like Pinterest, Wordpress, and Dubtoo, consolidating engagement data into actionable insights.

Data Feeds Trading Agents Access real-time, high-frequency market data, including order book information, trade data, and sentiment analysis. LAPSCA can aggregate and process this data to provide the trading agent with a comprehensive view of the market.

Portfolio Management Automate the process of rebalancing portfolios based on predefined rules or AI-driven recommendations, analyzing market trends and risk tolerance to optimize asset allocation and maximize returns.

Risk Management Implement real-time monitoring of market conditions and automatically adjust trading parameters to mitigate potential losses. This includes setting stop-loss orders, limiting exposure to specific assets, and dynamically adjusting trade sizes based on volatility and other risk factors.

Sentiment Analysis Integrate with social media and news feeds to gauge market sentiment. LAPSCA can analyze keywords, hashtags, and mentions to identify potential market movements driven by news events or social trends.

Backtesting and Optimization Leverage historical data to backtest trading strategies and optimize their performance. LAPSCA can automate the process of running simulations and analyzing the results to identify the most profitable strategies.

Customer Support Automate responses to frequently asked questions regarding trading strategies, platform features, and technical issues. This can free up human support agents to focus on more complex and critical tasks.

Data Analysis Perform sentiment analysis on customer feedback, identify trends in customer behavior, and generate reports on key performance indicators (KPIs) to inform business decisions.

Content Generation Generate marketing materials, such as blog posts, social media updates, and email newsletters, to promote trading services and attract new customers.

Figure 2 Data Flow Diagram

a. Functional Architecture

The functional blueprint serves as the primary passing point for users, equipped with both voice-enabled and text-based inputs, utilizing Automated Speech Recognition (ASR) to process commands. The system's core is managed by a central orchestrator, which routes tasks to specialized agents for data ingestion, analysis, and execution. The architecture also incorporates a feedback loop for continuous learning, where the system learns from past runs and adapts to market changes with minimal delay. The critical control action integrates all trading functionalities, enabling the system to execute trades, manage portfolios, and interact with various exchanges. This end-to-end functionality is crucial for automating the entire trading workflow, from data gathering to order execution.

b. Interactive Agent & Back-to-Manual

The core interface with LAPSCA is designed to be user-friendly for both technically savvy users and non-traders. Its accessibility is paramount, ensuring a wider adoption. The system supports voice commands for hands-free operation, allowing users to initiate tasks, monitor performance, and receive alerts while multitasking. The system provides real-time feedback on its actions and performance, keeping the user informed and in control. The architecture is designed to be modular, allowing for the easy integration of new tools and functionalities as required by different trading strategies.

c. Data & Trading Layer

The data and trading layer serves LAPSCA by aggregating data from Centralized Exchanges (CEXs), Decentralized Exchanges (DEXs), and on-chain sources. This layer is crucial for enhancing decision-making capabilities by providing a comprehensive view of the market. The system leverages the Model Context Protocol (MCP) to standardize tool integrations, ensuring seamless and secure communication between the AI agents and various data sources and execution venues. The architecture is designed to be modular and extensible, allowing for the easy integration of new data sources and trading venues as required by different trading strategies.

The layer is built on a modular, microservice-based architecture, where each blockchain interaction (reading data or executing transactions) is encapsulated within a dedicated microservice. This design ensures scalability, flexibility, and fault tolerance, allowing for the addition of new blockchains or protocols without disrupting the core system. The system also incorporates a robust error handling and retry mechanism to ensure reliable execution even in the face of network issues or exchange outages.

d. Multi-Chain & Cross-Market Support

LAPSCA's architecture supports both chain-level and asset-level operations, allowing users to manage assets across Ethereum, Solana, Bitcoin, and Base, with potential for future expansion. It utilizes multi-chain frameworks like Ethereum Attestation Service (EAS) and cross-chain messaging protocols like Wormhole to enable seamless interoperability between different blockchain networks.

Users can execute cross-chain trades or bridge assets by querying data from through integrated APIs. For example, a user might instruct the system to find the best price for USDC across Ethereum, Solana, and Base, and then execute the trade on the chain with the most favorable conditions, accounting for gas fees and bridging costs.

The platform's layer architecture supports advanced cross-chain strategies, such as monitoring liquidity across different chains to identify potential arbitrage opportunities or assessing the risk associated with bridging assets between different networks. The system can also be configured to prioritize certain chains or exchanges based on factors like liquidity, security, or transaction fees.

e. Human-in-the-Loop Governance

While LAPSCA aims for high levels of automation, it also incorporates a human-in-the-loop governance model to ensure that the system's actions align with the user's intent and risk tolerance. The system can be configured to require user approval for certain actions, such as large trades or interactions with new protocols. The user can also set predefined rules and boundaries for the AI agents, such as maximum trade sizes, allowed trading pairs, and stop-loss limits.

The platform's governance framework allows users to define and manage different trading strategies, each with its own set of rules, data sources, and execution parameters. The user can monitor the performance of each strategy and intervene if necessary to adjust the parameters or pause the strategy altogether.

4. Security and Privacy Architecture

LAPSCA prioritizes a comprehensive security framework that provides users with fine-grained control over their data and agent actions.

Local Data Processing The system leverages local processing for sensitive data, minimizing reliance on cloud services. Automated Speech Recognition (ASR) and Text-to-Speech (TTS) models can be deployed on local hardware, ensuring that voice commands and trading strategies remain within the user's control.

Local Encrypted Vault Sensitive data, including API keys and trading strategies, is stored in an encrypted vault on the user's device. This ensures that even if the system is compromised, the attacker will not gain access to the user's sensitive information.

Granular Permissions The system allows users to define granular permissions for each agent, specifying the data it can access and the actions it can perform. This prevents unauthorized access and limits the potential damage from a compromised agent.

Human-in-the-Loop Approval For critical actions, such as large trades or interactions with new protocols, the system can be configured to require user approval. This provides an additional layer of security and allows users to maintain control over the system's behavior.

Phase 2: Advanced Automation and Scalability (M 30-M 60)

The goal of the phase is to stream LAPSCA into a fully autonomous, privacy-centric platform with advanced capabilities. The plan emphasizes enhancing the system's intelligence, optimizing performance, and expanding its ecosystem through community contributions and partnerships.

Automated Order Execution The system will gain the ability to execute trades on various exchanges, with trade routing strategies since local latency is not a critical factor. The system will prioritize cost-effective execution and leverage aggregators where beneficial.

On-Chain Composability The system will be designed to interact with various DeFi protocols, enabling automated yield farming, liquidity provision, and other on-chain activities.

Complex Task Aggregation Advanced tools will be developed to automate high-performing trading strategies and solve more complex, multi-step problems. The system will be able to chain together various tools and data sources to achieve sophisticated objectives.

Offline Research Assistant The system will be enhanced to function as a fully autonomous research assistant, capable of conducting deep dives into market trends, analyzing competitive landscapes, and generating comprehensive reports.

Multi-Modal Analysis The system will be enhanced to process and synthesize information from various sources, including text, audio, and video, providing a more holistic understanding of market dynamics.

Community Contribution The system will be designed to facilitate community contributions, with clear guidelines and tools for developing and integrating new tools, strategies, and data sources.

Phase 3: Ecosystem Maturity (M 60-M 90)

The final phase focuses on expanding the capabilities established in Phase 1 by integrating more sophisticated tools and enhancing system interoperability. The primary objective is to transform the system into a robust, adaptable, and community-driven platform.

Data Scraping and Ethical Engineering This initiative, including an open call after the initial phase, aims to safeguard and make ethically sourced data available for training and fine-tuning models, ensuring high-quality, ethical data.

LoRA Fine-Tuning and RAG The LoRA (Low-Rank Adaptation) and Retrieval-Augmented Generation (RAG) framework will enable personalized, cost-effective, and high-fidelity agent adaptation.

Tool Ecosystem The system will leverage the Model Context Protocol (MCP) to dynamically connect AI agents with a vast ecosystem of tools, including APIs, databases, and other software, enhancing their capabilities and interoperability.

Phase 4: Advanced Automation and Scalability (M 90-M 180)

The goal of the phase is to stream LAPSCA into a fully autonomous, privacy-centric platform with advanced capabilities. The plan emphasizes enhancing the system's intelligence, optimizing performance, and expanding its ecosystem through community contributions and partnerships.

Automated Order Execution The system will gain the ability to execute trades on various exchanges, with trade routing strategies since local latency is not a critical factor. The system will prioritize cost-effective execution and leverage aggregators where beneficial.

On-Chain Composability The system will be designed to interact with various DeFi protocols, enabling automated yield farming, liquidity provision, and other on-chain activities.

Complex Task Aggregation Advanced tools will be developed to automate high-performing trading strategies and solve more complex, multi-step problems. The system will be able to chain together various tools and data sources to achieve sophisticated objectives.

Offline Research Assistant The system will be enhanced to function as a fully autonomous research assistant, capable of conducting deep dives into market trends, analyzing competitive landscapes, and generating comprehensive reports.

Multi-Modal Analysis The system will be enhanced to process and synthesize information from various sources, including text, audio, and video, providing a more holistic understanding of market dynamics.

Community Contribution The system will be designed to facilitate community contributions, with clear guidelines and tools for developing and integrating new tools, strategies, and data sources.

5. Unified Framework & Product Roadmap

5.1 Operational Framework

An end-workflow is structured around four key stations:

Control Protocol LAPSCA integrates gaming-inspired keyboard macros or high-performance voice dictation that makes inter-process communication (IPC) with the agent feel instantaneous and natural.

Sensory Tool Integration While the GUI Grounder and tool aggregation are systems to interact across the main stack, adding capabilities, such as Optical Character Recognition (OCR) and Automated Speech Recognition (ASR) are critical for a full sensory suite.

Result Synthesis Raw machine outputs are not actionable. Agents must synthesize data into a human-readable format, such as a summary, a chart, or a direct command to an API.

Execution Automation Upon user confirmation, LAPSCA can then operationalize the overall strategy or strategy through its integrated crypto wallet and DeFi. This automation optimizes transactions for speed, cost, and efficiency.

By integrating these components into a unified, standard framework, LAPSCA empowers users to interact with decentralized systems with unprecedented speed and control, transforming complex, multi-step processes into a seamless, conversational experience.

5.2 Product Roadmap

The development of LAPSCA is guided by a structured roadmap that builds on its specialized framework, with each milestone designed to validate the core hypothesis and expand the system's capabilities.

Phase 1: Foundation Building (Week 1-8)

The goal during this phase is to establish the foundational components that will enable LAPSCA to deliver its core functionality. This includes setting up the core infrastructure, integrating essential tools for data ingestion and execution, and developing the initial version of the GUI for monitoring and control. The focus is on ensuring a stable and reliable platform for future development.

Tasks Deadline: The GUI will be developed to provide real-time visualization of agent activity, logs, and historical trends, ensuring users can maintain control and oversight of the system's operations.

Control Middle Manager A crucial component responsible for the orchestration, managing the lifecycle of agents, allocating resources, and ensuring the overall stability and security of the system.

Initial Job Dispatcher Vision Call-ahead. Backgrounds Tool to retrieve blockchain data and execute accounts. Unknown Initially, the system will support a limited set of predefined jobs, focusing on core functionalities like data retrieval and simple transactions.

Phase 2: Expansion and Intelligence (Week 9-16)

The focus shifts to expanding the capabilities established in Phase 1 by integrating more sophisticated tools and enhancing system interoperability. The primary objective is to transform the system into a robust, adaptable, and community-driven platform.

Data Scraping and Ethical Engineering This initiative, including an open call after the initial phase, aims to safeguard and make ethically sourced data available for training and fine-tuning models, ensuring high-quality, ethical data.

LoRA Fine-Tuning and RAG The LoRA (Low-Rank Adaptation) and Retrieval-Augmented Generation (RAG) framework will enable personalized, cost-effective, and high-fidelity agent adaptation.

Tool Ecosystem The system will leverage the Model Context Protocol (MCP) to dynamically connect AI agents with a vast ecosystem of tools, including APIs, databases, and other software, enhancing their capabilities and interoperability.

Phase 3: Advanced Automation and Scalability (M 30-M 60)

The goal of the phase is to stream LAPSCA into a fully autonomous, privacy-centric platform with advanced capabilities. The plan emphasizes enhancing the system's intelligence, optimizing performance, and expanding its ecosystem through community contributions and partnerships.

Automated Order Execution The system will gain the ability to execute trades on various exchanges, with trade routing strategies since local latency is not a critical factor. The system will prioritize cost-effective execution and leverage aggregators where beneficial.

On-Chain Composability The system will be designed to interact with various DeFi protocols, enabling automated yield farming, liquidity provision, and other on-chain activities.

Complex Task Aggregation Advanced tools will be developed to automate high-performing trading strategies and solve more complex, multi-step problems. The system will be able to chain together various tools and data sources to achieve sophisticated objectives.

Offline Research Assistant The system will be enhanced to function as a fully autonomous research assistant, capable of conducting deep dives into market trends, analyzing competitive landscapes, and generating comprehensive reports.

Multi-Modal Analysis The system will be enhanced to process and synthesize information from various sources, including text, audio, and video, providing a more holistic understanding of market dynamics.

Community Contribution The system will be designed to facilitate community contributions, with clear guidelines and tools for developing and integrating new tools, strategies, and data sources.

6. Verifiable Intelligence & Voice Controlled Operations

LAPSCA prioritizes the creation of auditable and verifiable AI agents, ensuring that their actions are transparent and accountable. This is crucial for building trust and mitigating the risks associated with autonomous decision-making. The framework incorporates several mechanisms to achieve this, including detailed logging of all agent actions, the generation of cryptographic proofs of execution, and the integration with blockchain technology for immutable record-keeping.

Voice-controlled AI agents offer a more intuitive and accessible interface for interacting with complex systems. However, this also introduces new security and privacy considerations. The LAPSCA framework addresses these concerns by implementing robust voice authentication and command verification mechanisms, ensuring that only authorized users can access and control the system. The framework also supports the use of local, on-device speech processing to minimize the reliance on cloud services and enhance privacy.

6.1 Data Flow

LAPSCA is a universal technology capable of automating and optimizing tasks across Retail and Retail Activities. Its flexible agentic nature lets firms prioritize gathering, summarizing, and analyzing trading-relevant information in a trading workflow automation.

Some of the use cases are as follows (not exhaustive):

Automated Trading and Arbitrage Rapidly scan order book depth, liquidity, exchange opportunities, and optimize price strategies with near-instantaneous execution.

Market Sentiment Analysis Seamlessly monitor trades accounts across platforms like Pinterest, Wordpress, and Dubtoo, consolidating engagement data into actionable insights.

Data Feeds Trading Agents Access real-time, high-frequency market data, including order book information, trade data, and sentiment analysis. LAPSCA can aggregate and process this data to provide the trading agent with a comprehensive view of the market.

Portfolio Management Automate the process of rebalancing portfolios based on predefined rules or AI-driven recommendations, analyzing market trends and risk tolerance to optimize asset allocation and maximize returns.

Risk Management Implement real-time monitoring of market conditions and automatically adjust trading parameters to mitigate potential losses. This includes setting stop-loss orders, limiting exposure to specific assets, and dynamically adjusting trade sizes based on volatility and other risk factors.

Sentiment Analysis Integrate with social media and news feeds to gauge market sentiment. LAPSCA can analyze keywords, hashtags, and mentions to identify potential market movements driven by news events or social trends.

Backtesting and Optimization Leverage historical data to backtest trading strategies and optimize their performance. LAPSCA can automate the process of running simulations and analyzing the results to identify the most profitable strategies.

Customer Support Automate responses to frequently asked questions regarding trading strategies, platform features, and technical issues. This can free up human support agents to focus on more complex and critical tasks.

Data Analysis Perform sentiment analysis on customer feedback, identify trends in customer behavior, and generate reports on key performance indicators (KPIs) to inform business decisions.

Content Generation Generate marketing materials, such as blog posts, social media updates, and email newsletters, to promote trading services and attract new customers.

a. Functional Architecture

The functional blueprint serves as the primary passing point for users, equipped with both voice-enabled and text-based inputs, utilizing Automated Speech Recognition (ASR) to process commands. The system's core is managed by a central orchestrator, which routes tasks to specialized agents for data ingestion, analysis, and execution. The architecture also incorporates a feedback loop for continuous learning, where the system learns from past runs and adapts to market changes with minimal delay. The critical control action integrates all trading functionalities, enabling the system to execute trades, manage portfolios, and interact with various exchanges. This end-to-end functionality is crucial for automating the entire trading workflow, from data gathering to order execution.

b. Interactive Agent & Back-to-Manual

The core interface with LAPSCA is designed to be user-friendly for both technically savvy users and non-traders. Its accessibility is paramount, ensuring a wider adoption. The system supports voice commands for hands-free operation, allowing users to initiate tasks, monitor performance, and receive alerts while multitasking. The system provides real-time feedback on its actions and performance, keeping the user informed and in control. The architecture is designed to be modular, allowing for the easy integration of new tools and functionalities as required by different trading strategies.

c. Data & Trading Layer

The data and trading layer serves LAPSCA by aggregating data from Centralized Exchanges (CEXs), Decentralized Exchanges (DEXs), and on-chain sources. This layer is crucial for enhancing decision-making capabilities by providing a comprehensive view of the market. The system leverages the Model Context Protocol (MCP) to standardize tool integrations, ensuring seamless and secure communication between the AI agents and various data sources and execution venues. The architecture is designed to be modular and extensible, allowing for the easy integration of new data sources and trading venues as required by different trading strategies.

The layer is built on a modular, microservice-based architecture, where each blockchain interaction (reading data or executing transactions) is encapsulated within a dedicated microservice. This design ensures scalability, flexibility, and fault tolerance, allowing for the addition of new blockchains or protocols without disrupting the core system. The system also incorporates a robust error handling and retry mechanism to ensure reliable execution even in the face of network issues or exchange outages.

d. Multi-Chain & Cross-Market Support

LAPSCA's architecture supports both chain-level and asset-level operations, allowing users to manage assets across Ethereum, Solana, Bitcoin, and Base, with potential for future expansion. It utilizes multi-chain frameworks like Ethereum Attestation Service (EAS) and cross-chain messaging protocols like Wormhole to enable seamless interoperability between different blockchain networks.

Users can execute cross-chain trades or bridge assets by querying data from through integrated APIs. For example, a user might instruct the system to find the best price for USDC across Ethereum, Solana, and Base, and then execute the trade on the chain with the most favorable conditions, accounting for gas fees and bridging costs.

The platform's layer architecture supports advanced cross-chain strategies, such as monitoring liquidity across different chains to identify potential arbitrage opportunities or assessing the risk associated with bridging assets between different networks. The system can also be configured to prioritize certain chains or exchanges based on factors like liquidity, security, or transaction fees.

e. Human-in-the-Loop Governance

While LAPSCA aims for high levels of automation, it also incorporates a human-in-the-loop governance model to ensure that the system's actions align with the user's intent and risk tolerance. The system can be configured to require user approval for certain actions, such as large trades or interactions with new protocols. The user can also set predefined rules and boundaries for the AI agents, such as maximum trade sizes, allowed trading pairs, and stop-loss limits.

The platform's governance framework allows users to define and manage different trading strategies, each with its own set of rules, data sources, and execution parameters. The user can monitor the performance of each strategy and intervene if necessary to adjust the parameters or pause the strategy altogether.

7. Key Takeaways and Final Insights

Both Kimi-Dev-72B agents from the Kimi-X technology combining to handle tasks that are relevant with digital incentives. Meanwhile Kimi-Dev-72B is a specialized 72B parameter model for software engineering tasks, achieving state-of-the-art performance on SWE-bench Verified (60.4%). Its capabilities include automated bug fixing and feature implementation, optimized for complex, multi-file codebases. LAPSCA provides the infrastructure needed to integrate Kimi-Dev-72B, empowering users with intuitive, privacy-first tools to execute coding tasks locally.

Kimi-Dev-72B's performance in software engineering benchmarks demonstrates its potential to automate complex coding tasks. The model's ability to understand and modify multi-file codebases, coupled with its high accuracy in bug fixing and feature implementation, makes it a valuable asset for developers and organizations seeking to enhance their software development processes. By integrating Kimi-Dev-72B into the LAPSCA framework, users can leverage its capabilities while maintaining control over their data and workflows.

The combination of Kimi-Dev-72B and LAPSCA offers a compelling proposition for organizations prioritizing data privacy and security. The ability to run these models locally on consumer-grade hardware ensures that sensitive data never leaves the user's control, mitigating the risks associated with cloud-based AI services. This approach is particularly relevant for organizations in regulated industries or those with strict data governance policies.

The LAPSCA framework's modular design allows for the seamless integration of other specialized models and tools, enabling users to create a tailored AI ecosystem that addresses their specific needs and workflows. This flexibility is crucial for maximizing the value of AI investments and ensuring that the technology remains aligned with evolving business requirements.

The development roadmap for LAPSCA indicates a commitment to continuous improvement and the addition of new features and functionalities. The planned integration of advanced tools for data scraping, LoRA fine-tuning, and RAG (Retrieval-Augmented Generation) will further enhance the system's capabilities and its value proposition.

The success of this approach will be measured by the adoption rate of the Kimi-Dev-72B agent within the LAPSCA framework, the performance of the agent in real-world coding tasks, and the overall cost-effectiveness of the system compared to traditional cloud-based AI solutions. The ability to demonstrate significant cost savings while maintaining high performance will be a key factor in the framework's success.

References

El Mellouki, Oussama, Mohammad Amin Khenkeh, and Walid Gaaloul. "Blockchain Oracle Design Patterns." Proceedings of the Blockchain and Trustworthy Systems: First International Conference, BTS 2019, Guangzhou, China, December 7–8, 2019. Springer Singapore, 2020.

Ferdinand, Jacob, Ashish Kundu, Kiran Kate, Stavros Shiokas, Dinesh Garg, and Giorgio Cecchinato. "Fast Verification of Llm Responses to Blockchain Smart Contracts." ICLR 2025 Workshop: Foundation Models in the Wild, 2025.

Gao, Yunfan, Yun Xiong, Xinyu Gao, Kangxiang J