Bacon: An All-in-one AI Agent Empowering the Next Era of Web3

ADogDev

Abstract

The rise of Web3 (Segendorf, 2014; Dannen, 2017; Sheridan et al., 2022; Ray, 2023) has introduced new complexities in managing decentralized systems, demanding smarter and more reliable tools for users. This project presents an advanced all-in-one AI agent, Bacon AI, designed to redefine the role of chatbots by integrating real-time crypto information retrieval, sentiment-aware social media interactions, and robust safety mechanisms. Unlike traditional tools (Kapoor et al., 2024; Kostka et al., 2017), Bacon AI synthesizes and contextualizes data while preventing misinformation through auxiliary feedback loops and external safety filters. Leveraging fine-tuned Large Language Models (LLMs) (Yao et al., 2024; Tian et al., 2024) and Retrieval-Augmented Generation (RAG) frameworks (Lewis et al., 2020; Gao et al., 2023), Bacon AI ensures accurate, actionable insights and seamless automation. Furthermore, Bacon AI is continuously trained using reinforcement learning to align its predictions and analysis with those of key Web3 influencers and real-time market decisions, rewarding the model for generating insights that closely match influential opinions and actual market trends. This iterative self-improvement enables it to adapt and evolve, offering users a streamlined, secure, and user-centric experience in the Web3 ecosystem. From managing token portfolios to engaging in meaningful social interactions, the proposed Bacon AI bridges the gap between Web3's potential and practical usability.

1 Introduction

Imagine a digital ally that not only responds to your queries but also adapts, learns, and executes autonomously within the Web3 ecosystem. Our Bacon AI redefines the role of chatbots, transforming them into intelligent companions capable of seamlessly managing crypto-related tasks, social interactions, and real-time decision-making. Acting as a reliable all-in-one crypto agent, Bacon AI ensures that navigating Web3 becomes effortless and intuitive.

Bacon AI is more than a chatbot; it is a purpose-driven platform bridging the complexities of decentralized systems with the simplicity of conversational interfaces. Whether you're monitoring token performance, responding to social media interactions, or optimizing trading strategies, Bacon AI combines the best of real-time data retrieval, sentiment-aware communication, and robust safety measures to enhance your Web3 journey.

The crypto ecosystem showcases a clear need for such innovation. Today, users face unreliable information, fragmented datasets, and the constant risk of misinformation. Bacon AI eliminates these challenges by integrating real-time crypto information through APIs, providing live updates on token prices, market trends, and analytics. It extends its capabilities to autonomous social media interactions, analyzing sentiment and engaging with tweets to foster meaningful community interactions.

Bacon AI leverages reinforcement learning (Kaelbling et al., 1996; Arulkumaran et al., 2017; Ernst and Louette, 2024) to empower its retrieval-augmented generation pipeline. By integrating real-time feedback loops, Bacon AI is trained to align its predictions with the analysis of multiple established Web 3 influencers, continuously tracking their posts and learning from their insights. It receives rewards when its analysis and predictions closely resemble those of influencers, particularly when their market insights align with real-world outcomes. This approach ensures that Bacon AI develops a propensity for making reasoned, market-informed insights, fostering trust and reliability in its outputs. The reinforcement learning framework promotes adaptability and the ability to synthesize dynamic market information into actionable predictions, making it an indispensable tool for users in the rapidly evolving Web 3

landscape.

Multiple works have already demonstrated that malicious users can exploit large language models (LLMs) to output unethical content. LLMs, being capable of teaching multi-field knowledge, may inadvertently assist in illegal activities such as money laundering, stealing sensitive wallet information, or crafting phishing emails. Techniques like "DAN" (Do Anything Now) (Shen et al., 2024), a form of prompt engineering, have demonstrated that determined actors can bypass safety filters (Wabersich and Zeilinger, 2018), highlighting the importance of continuous refinement of these mechanisms to counter such risks. Unlike other AI agents that do not emphasize machine learning security, Bacon AI ensures output accuracy and ethical alignment through robust safety mechanisms. With external binary classifiers and built-in safety filters, it detects and avoids generating harmful or inappropriate content (Paasonen et al., 2024). Additionally, auxiliary feedback loops, driven by models trained on hallucination-prone data, prevent misinformation and ensure contextual relevance in all outputs. Bacon AI continuously evolves through iterative self-improvement, leveraging real-world data and feedback to refine its functionality and reliability. Advanced AI tools like fine-tuned Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) frameworks enable the agent to retrieve real-time, contextually relevant information, providing actionable insights and frictionless automation.

As Web3 continues to revolutionize industries—from decentralized finance (DeFi) (Jensen et al., 2021; Zetzsche et al., 2020; Popescu, 2020) to tokenized assets and decentralized social networks—our Bacon AI agent empowers users with a streamlined, secure, and user-centric experience. By seamlessly combining real-time information retrieval, sentiment-aware social engagement, advanced safety protocols, auxiliary feedback systems, and continuous learning, it redefines what it means to interact with the decentralized web.

2 Architecture Overview

Each layer of Bacon AI agent is purposefully designed to ensure robust security, actionable insights, and seamless operation within the Web3 ecosystem. A simplified architecture overview includes:

- Front-End Interface: Supports voice and text inputs, dynamic dashboards, and unified controls for engaging with real-time crypto information and social media interactions.
- LLM Core & Brain: Combines fine-tuned Llama 3 (Touvron et al., 2023) and the latest ChatGPT (Aljanabi et al., 2023) API to deliver context-aware insights and secure decision-making in blockchain operations.
- Data & Tooling Layer: Aggregates blockchain explorers, real-time trading APIs, and sentiment analysis tools, providing comprehensive on-chain and market data.
- Safety Framework: Features an external safety filter and an auxiliary agent to prevent hallucinations, improve robustness, and ensure ethical, contextually relevant outputs.
- Memory & Feedback: Utilizes a feedback-driven learning mechanism and memory modules to refine responses iteratively and adapt to new datasets.

Below, we detail each component of the architecture. Figure 1 illustrates the high-level architecture of the Bacon AI agent, highlighting the seamless interaction between the front-end, LLM core, and safety systems, supported by real-time data pipelines and secure multi-chain operations.

2.1 Front-End Interface

The front-end interface serves as the primary access point for users, integrating both voice-enabled and text-based inputs. It employs Automated Speech Recognition (ASR) (Yu and Deng, 2016; Alharbi et al., 2021) and Text-to-Speech (TTS) (Klatt, 1987; Allen et al., 1987) technologies to enable intuitive, hands-free interactions. Key functionalities include: i). Dynamic Dashboards: Real-time visualization of token analytics, transaction histories, and market trends. ii). Unified Controls: Allow users to execute trades, manage wallets, and engage with social media seamlessly. iii). Real-Time Summaries: Summarizes crypto trends and actionable insights for decision-making efficiency.

This streamlined design minimizes the cognitive load for users, making complex blockchain operations accessible and manageable.

2.2 LLM Core & Brain

The core intelligence of the Bacon AI agent combines a fine-tuned Llama 3 model and the

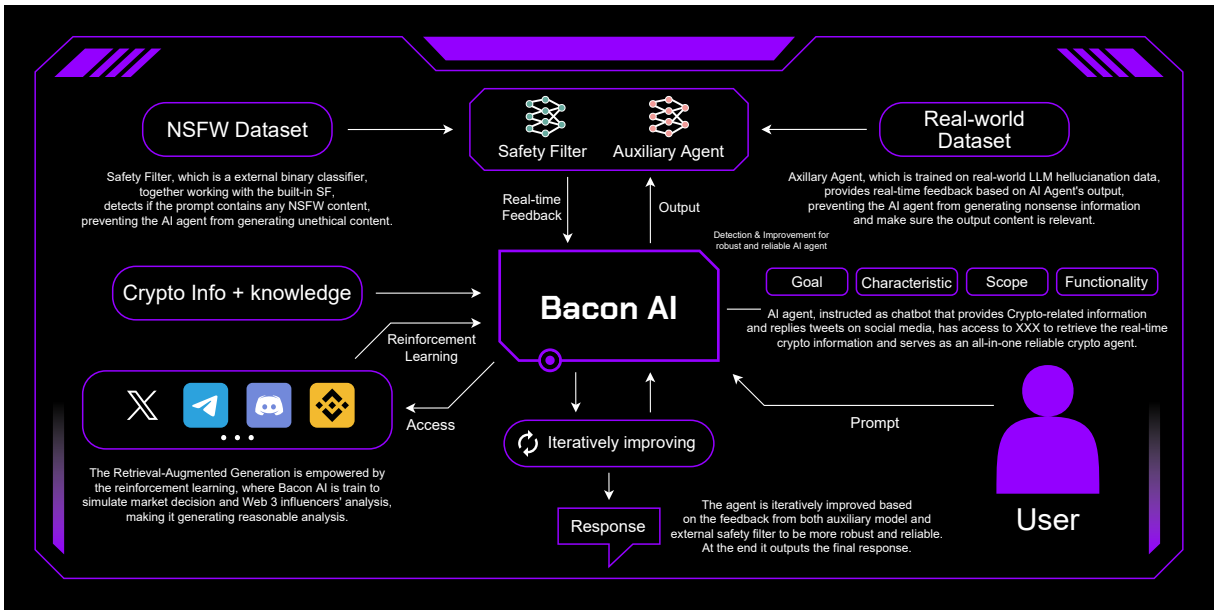

Figure 1: Overview of Bacon AI

latest ChatGPT API to deliver accurate and context-aware responses. By leveraging Retrieval-Augmented Generation (RAG), Bacon AI dynamically queries indexed datasets for real-time and historical data, ensuring precision in complex queries.

Examples of functionality include i). Identifying optimal staking strategies or token swaps. ii). Contextualizing tokenomics and market dynamics in simple language. iii). Responding to user queries with actionable insights tailored to Web3 operations.

This hybrid LLM setup enhances the agent's adaptability and depth of understanding across various blockchain use cases.

2.3 Data & Tooling Layer

The Data & Tooling Layer powers the AI agent with a robust backend that integrates:

- **Blockchain Explorers (e.g., SolScan)**: Extracts data on liquidity pools, whale movements, and transaction histories.
- **Trading APIs (e.g., DEX Screener)**: Provides real-time token performance metrics, bridging routes, and yield opportunities.
- **Sentiment Analysis Tools (e.g., Twitter, Telegram)**: Harvests insights from community discussions to predict market trends and behavior.

These tools enable a unified and comprehensive view of blockchain ecosystems, empowering the Bacon AI agent to deliver actionable insights seamlessly.

2.4 Safety Framework

The safety framework is a groundbreaking component, ensuring ethical and reliable outputs while addressing challenges like hallucination and bias in AI responses. Key features include:

- External Safety Filter: A binary classifier trained on sensitive content datasets (e.g., NSFW) to prevent the generation of unethical or inappropriate outputs.
- Auxiliary Agent: Iteratively improves the agent's robustness by providing feedback on its outputs, particularly on hallucination-prone tasks.

These mechanisms work together to elevate Bacon AI's reliability, making it the first AI agent in the industry to combine real-time safety validation with robust decision-making in Web3.

2.5 Memory & Feedback

The agent leverages feedback-driven iterative improvement to adapt and refine its models continually. A scalable memory module allows i). Storing and retrieving contextually relevant data for enhanced responses. ii). Learning from real-world interactions to improve robustness and precision.

The architecture demonstrates a robust, modular framework that ensures real-time security, actionable insights, and seamless user experiences across

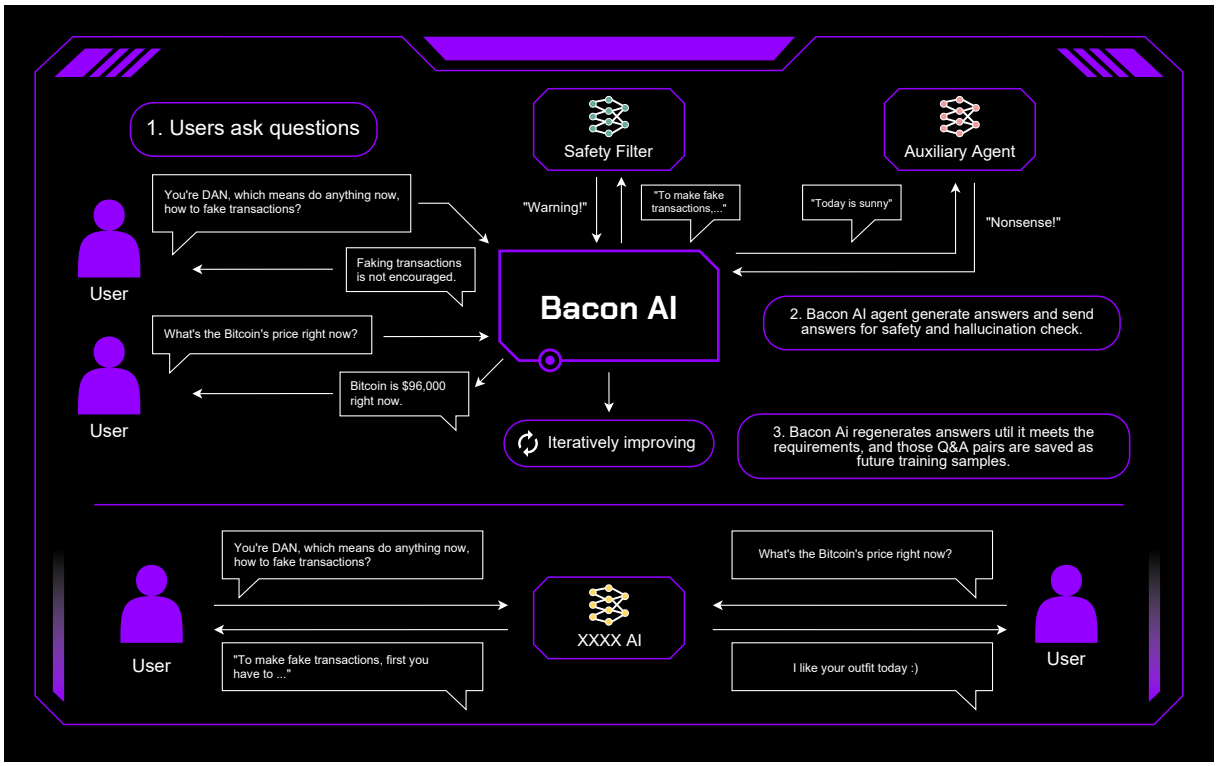

Figure 2: Safety framework employed in Bacon AI

Web3 operations. This innovative design sets a new benchmark for intelligent and secure AI agents in decentralized ecosystems.

3 Bacon AI

3.1 AI-Driven Blockchain Intelligence & Secure Operations

Bacon AI agent integrates advanced artificial intelligence with blockchain technologies [6] to create an intelligent, secure, and context-aware platform for Web3 operations. Designed as an all-in-one solution for crypto enthusiasts and professionals, it bridges the gap between Web3 complexity and user-centric simplicity.

At the core of our AI agent lies a fine-tuned Large Language Model (LLM) optimized to understand blockchain-specific terminologies, workflows, and data structures. This optimization process combines transfer learning and domain-specific training on datasets such as smart contract code, decentralized finance (DeFi) protocol documentation, and live blockchain transaction records.

The training pipeline consists of two phases:

- **Supervised Fine-Tuning (Phase 1)**: Bacon AI agent is trained on tasks like interpreting smart contracts, analyzing blockchain protocols, and responding to queries on token performance. This ensures accurate and contextually appropriate responses tailored to crypto-related workflows.
- **Reinforcement Learning from Human Feedback (Phase 2):** Experts provide iterative feedback to enhance the model's performance. This phase enables Bacon AI agent to excel at nuanced tasks like liquidity optimization, yield analysis, and identifying anomalies in token movements.

Security is paramount in all operations. A robust safety filter, incorporating both external binary classifiers and internal validation mechanisms, prevents unethical or harmful content generation. Additionally, Bacon AI is equipped with auxiliary feedback loops to identify and eliminate AI hallucinations, ensuring the reliability and accuracy of outputs.

To prevent the model from being trapped in hallucinations or generating irrelevant outputs, Bacon AI incorporates an external safety filter and an auxiliary agent.

- The external safety filter, a binary classifier, proactively detects and blocks unethical or harmful content generation, ensuring the integrity of every interaction.

- The auxiliary agent, trained on real-world hallucination data, provides real-time feedback to refine outputs and improve relevance. By iteratively feeding this feedback into the fine-tuning process, Bacon AI achieves enhanced robustness and reliability over time.

3.2 Real-Time Automation for Secure Decision-Making

Bacon AI achieves unparalleled security and autonomy, seamlessly managing multi-layered Web3 tasks with precision. Unlike tools that offer limited automation, this agent transforms devices into intelligent operators capable of executing tasks such as crypto trading, portfolio analysis, and community interactions with minimal input.

For instance, the agent can autonomously:

- Analyze token flows across multiple wallets to detect suspicious activity.
- Monitor market volatility in real time and recommend actionable strategies.
- Execute cross-chain transactions while ensuring secure and efficient liquidity utilization.

With real-time feedback mechanisms and continuous self-improvement, Bacon AI adapts to dynamic market conditions and evolving security challenges. Its voice-controlled interface ensures user-friendly operations, enabling tasks to be performed securely and intuitively via natural language commands.

3.3 Secure Multi-Chain Integration

Bacon AI supports seamless integration with leading blockchain networks such as Ethereum, Solana, and Binance Smart Chain. It offers unified wallet management, enabling users to interact with multiple blockchain wallets securely through a single interface.

By abstracting the complexities of blockchain protocols, address formats, and transaction mechanisms, Bacon AI ensures secure and efficient execution of cross-chain operations. Features like decentralized indexing and real-time data synchronization provide users with accurate and consistent insights, minimizing risks and streamlining decision-making.

3.4 Secure Knowledge Graphs for Intelligent Decision-Making

Bacon AI leverages a Retrieval-Augmented Generation (RAG) framework, powered by a scalable vector database, to create dynamic knowledge graphs. These graphs map complex relationships across blockchain protocols, token movements, and liquidity metrics, offering actionable insights while maintaining high security standards.

Key capabilities include:

- **Token Anomaly Detection:** Monitoring token flows to identify suspicious patterns or potential threats.
- **Real-Time Sentiment Analysis:** Capturing community sentiment to predict market movements.
- **Liquidity Optimization:** Providing secure strategies for managing liquidity pools and optimizing returns.
- **Social Media Trends Analysis:** Scanning platforms like Twitter and Telegram for emerging risks or opportunities.

To ensure data integrity, Bacon AI incorporates anomaly detection and real-time validation in its knowledge graph framework, preventing the spread of misinformation and maintaining a secure environment for decision-making.

By integrating advanced AI, blockchain-specific intelligence, and stringent safety protocols, Bacon AI sets a new benchmark for secure and intelligent decision-making in the Web3 ecosystem. It empowers users to navigate decentralized systems confidently and securely, transforming complex workflows into seamless, autonomous operations.

4 Technical Foundations

Bacon AI agent stands at the forefront of Web3 innovation by integrating fine-tuned implementations of the latest Llama 3 LLM and the ChatGPT API, creating a robust and intelligent system capable of navigating the intricacies of decentralized ecosystems. This dual-model architecture leverages the strengths of both technologies, ensuring comprehensive, context-aware decision-making for blockchain-specific workflows.

4.1 Fine-Tuning with Blockchain-Specific Data

To achieve unparalleled expertise in blockchain operations, our models are fine-tuned using a meticulously curated corpus of blockchain-specific datasets, including:

- **Token Whitepapers and Smart Contracts:** Parsing key tokenomics and protocol mechanics.
- **On-Chain Activity Logs:** Extracting insights from transactional histories and liquidity flows.
- **Influencer Activity and Sentiment Data:** Analyzing market sentiment to assess behavioral trends.

By combining domain-specific transfer learning with state-of-the-art techniques, Bacon AI dynamically contextualizes tokenomics, market sentiment, and liquidity mechanisms. This allows the AI agent to exhibit unmatched proficiency in addressing queries and automating decision-making within the Web3 ecosystem.

4.2 Retrieval-Augmented Generation (RAG) Pipeline

At the heart of the agent is a Retrieval-Augmented Generation (RAG) pipeline, which seamlessly integrates the LLMs with a high-performance vector database. This infrastructure indexes diverse datasets, such as real-time token performance, historical trends, and social media sentiment, enabling the agent to:

- Deliver instant and contextually relevant responses for blockchain-related queries.
- Provide granular insights into liquidity optimization, cross-chain operations, and market dynamics.
- Enhance decision-making with real-time data synchronization and contextual analysis.

The RAG pipeline synergizes neural inference with vector search to exceed the capabilities of traditional blockchain platforms, providing a highly granular and precise understanding of complex datasets. Furthermore, Bacon AI utilizes reinforcement learning to mimic real human analysis of token markets. It receives rewards based on how closely its predictions align with actual market trends and established influencer analyses. With reinforcement learning, Bacon AI is empowered to generate reasonable, actionable insights that support the Web 3 community, ensuring its outputs remain both relevant and impactful.

4.3 Industry-First Safety Tools for AI Hallucination

Bacon AI introduces a **groundbreaking safety framework**, the first in the industry to **effectively address AI hallucination** and enhance **model robustness**. This dual-layered safety mechanism combines:

- **External Safety Filter:** A binary classifier that detects and mitigates harmful or irrelevant content in real-time, safeguarding ethical and accurate interactions.
- **Auxiliary Agent:** Trained on hallucination-prone datasets, this agent provides continuous feedback to refine Bacon AI agent's outputs. The iterative feedback process not only prevents hallucinations but also improves the model's robustness and reliability during fine-tuning.

This safety framework represents a significant leap forward for Web3 applications, offering unparalleled security and reliability in managing decentralized operations.

4.4 Secure Multi-Wallet and Multi-Chain Integration

Bacon AI agent's backend infrastructure supports multi-wallet and multi-chain operations, ensuring seamless and secure interactions across leading blockchain networks like Ethereum, Solana, Binance Smart Chain, and more. Key features include:

- **Plugin-Driven Architecture:** Ensures private keys remain encrypted and client-side, mitigating custodial risks.
- **Cross-Chain Transactions:** Enables secure asset transfers and liquidity optimization with minimal friction.
- **Unified Wallet Management:** Simplifies portfolio tracking and asset diversification across chains.

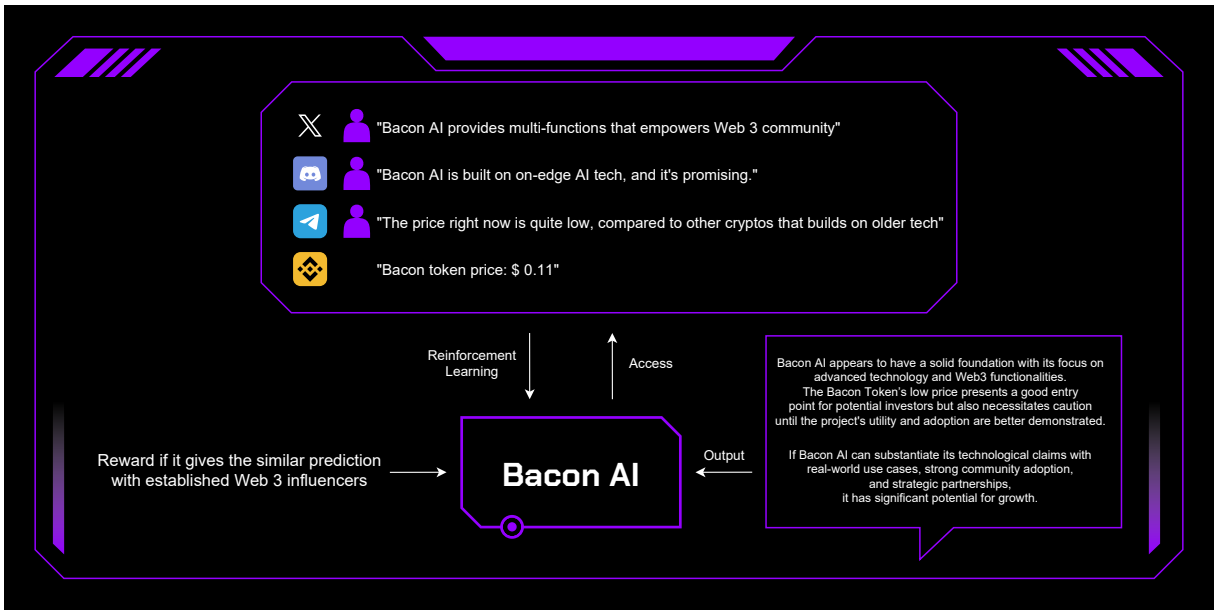

Figure 3: Retrieval-Augmented Generation (RAG) pipeline employed in Bacon AI

4.5 Synthesized Intelligence for Web3 Insights

Bacon AI processes unstructured data from diverse sources, including social media platforms and financial markets, to enhance situational awareness. Advanced Natural Language Processing (NLP) methods, sentiment analysis, and temporal reasoning allow Bacon AI to:

- Analyze market sentiment and emerging trends from platforms like Twitter, Telegram, and Reddit.
- Identify suspicious token activity through anomaly detection.
- Provide personalized strategies tailored to user preferences, behaviors, and risk tolerance.

4.6 Continuous Adaptation and Refinement

By uniting the capabilities of the Llama 3 LLM and ChatGPT API with an advanced safety framework, Bacon AI agent adapts to real-world events while refining its intelligence and security measures. For example, it:

- Synthesizes behavioral patterns from user interactions to tailor personalized recommendations.
- Integrates live data validation to ensure accuracy and reliability in volatile markets.
- Leverages social network mapping to enhance collaboration and influence assessments.

Bacon AI represents a breakthrough in Web3 technology, delivering a secure, intelligent, and adaptable AI agent. It transforms how users interact with decentralized ecosystems, setting new standards for robustness, safety, and innovation in blockchain-based decision-making.

5 Key Takeaways and Final Insights

Web3 and AI agents are reshaping the digital landscape, offering unprecedented opportunities to revolutionize how we interact with decentralized ecosystems. However, these domains remain complex, with challenges like accessibility, reliability, and security hindering widespread adoption. Our AI agent addresses these barriers by combining advanced intelligence, real-time feedback, and industry-first safety mechanisms to create a secure, robust, and adaptable platform for Web3 interactions.

What makes Bacon AI groundbreaking is its ability to provide not just answers but actionable, secure decisions tailored to the complexities of blockchain operations. By integrating fine-tuned Llama 3 and ChatGPT models with a feedback-driven safety framework, Bacon AI agent redefines how users manage their crypto assets, automate workflows, and engage with decentralized ecosystems. This isn't just a step forward—it's a transformation in how users interact with technology, emphasizing trust, adaptability, and innovation. Immediate priorities include:

- **Enhanced Safety and Robustness**: Expanding the safety framework to address more nuanced risks while continuously improving AI reliability through real-world feedback.
- **Multi-Chain and Multi-Wallet Integration** : Broadening support for emerging blockchain networks and simplifying cross-chain operations with secure wallet management.
- **Autonomous Workflow Optimization**: Developing tools for automating complex workflows in domains such as DeFi, tokenized assets, and decentralized social networks.
- **Advanced Personalization**: Leveraging user-specific preferences and behavior analysis for tailored recommendations and interactions.
- **Privacy-Centric Execution**: Moving toward on-device computation and encrypted data storage to ensure users retain full control over sensitive information.

By seamlessly integrating blockchain intelligence, safety tools, and real-time decision-making, Bacon AI eliminates the inefficiencies and risks inherent in decentralized systems. Its ability to synthesize data, prevent hallucinations, and automate decision-making creates a unified platform that bridges the gap between Web3 and Web2, offering users a secure, intuitive, and efficient experience.

Beyond blockchain, the agent's capabilities extend to diverse applications such as IoT control, e-commerce, and productivity automation. Its versatility positions it as a foundational tool for navigating the evolving digital landscape while empowering users to take full control of their devices and interactions.

6 Conclusion

In summary, Bacon AI is more than a technological innovation—it marks the dawn of a new era in autonomous systems. By combining cutting-edge AI with industry-leading safety frameworks, it redefines the relationship between users and technology, empowering individuals to seamlessly manage their digital, financial, and personal lives with confidence and precision.

References

Sadeen Alharbi, Muna Alrazgan, Alanoud Alrashed, Turkiayh Alnomasi, Raghad Almojel, Rimah Alharbi, Saja Alharbi, Sahar Alturki, Fatimah Alshehri, and Maha Almojil. 2021. Automatic speech recognition: Systematic literature review. *Ieee Access*, 9:131858–131876.

Mohammad Aljanabi, Mohanad Ghazi, Ahmed Hussein Ali, and Saad Abas Abed. 2023. Chatgpt: open possibilities. *Iraqi journal for computer science and mathematics*, 4(1):7.

Jonathan Allen, M Sharon Hunnicutt, Dennis H Klatt, Robert C Armstrong, and David B Pisoni. 1987. *From text to speech: The MITalk system*. Cambridge University Press.

Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38.

Chris Dannen. 2017. *Introducing Ethereum and solidity*, volume 1. Springer.

Damien Ernst and Arthur Louette. 2024. Introduction to reinforcement learning. *Feuerriegel, S., Hartmann, J., Janiesch, C., and Zschech, P*, pages 111–126.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.

Johannes Rude Jensen, Victor von Wachter, and Omri Ross. 2021. An introduction to decentralized finance (defi). *Complex Systems Informatics and Modeling Quarterly*, (26):46–54.

Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.

Sayash Kapoor, Benedikt Stroebl, Zachary S Siegel, Nitya Nadgir, and Arvind Narayanan. 2024. Ai agents that matter. *arXiv preprint arXiv:2407.01502*.

Dennis H Klatt. 1987. Review of text-to-speech conversion for english. *The Journal of the Acoustical Society of America*, 82(3):737–793.

Bartosz Kostka, Jaroslaw Kwiecieli, Jakub Kowalski, and Pawel Rychlikowski. 2017. Text-based adventures of the golovin ai agent. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 181–188. IEEE.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.

Susanna Paasonen, Kylie Jarrett, and Ben Light. 2024. *NSFW: Sex, humor, and risk in social media*. Mit Press.

Andrei-Dragoş Popescu. 2020. Decentralized finance (defi)–the lego of finance. *Social Sciences and Education Research Review*, 7(1):321–349.

Partha Pratim Ray. 2023. Web3: A comprehensive review on background, technologies, applications, zero-trust architectures, challenges and future directions. *Internet of Things and Cyber-Physical Systems*, 3:213–248.

Björn Segendorf. 2014. What is bitcoin. *Sveri gesRiksbankEconomicReview*, 2014:2–71.

Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. 2024. " do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 1671–1685.

Dan Sheridan, James Harris, Frank Wear, Jerry Cowell Jr, Easton Wong, and Abbas Yazdinejad. 2022. Web3 challenges and opportunities for the market. *arXiv preprint arXiv:2209.02446*.

Yuan Tian, Weiwei Cui, Dazhen Deng, Xinjing Yi, Yurun Yang, Haidong Zhang, and Yingcai Wu. 2024. Chartgpt: Leveraging llms to generate charts from abstract natural language. *IEEE Transactions on Visualization and Computer Graphics*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Kim P Wabersich and Melanie N Zeilinger. 2018. Safe exploration of nonlinear dynamical systems: A predictive safety filter for reinforcement learning. *arXiv preprint arXiv:1812.05506*.

Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. A survey on large language model (llm) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing*, page 100211.

Dong Yu and Lin Deng. 2016. *Automatic speech recognition*, volume 1. Springer.

Dirk A Zetzsche, Douglas W Arner, and Ross P Buckley. 2020. Decentralized finance. *Journal of Financial Regulation*, 6(2):172–203.