Know Your Agent: Governing AI Identity on the Agentic Web

Tomer Jordi Chaffer

DeGov Labs, tomer.chaffer@mail.mcgill.ca

February 25, 2025

Abstract

The agentic web refers to a vision of the internet where AI agents play a central role in facilitating interactions, automating tasks, and enhancing user experiences. Realizing this vision requires us to rethink how we govern the internet. Within the agentic web, the promise of AI systems becoming more decentralized and autonomous represents unique challenges and opportunities for governance, necessitating innovative approaches to ensure responsible integration into society. Toward this end, we propose the Know Your Agent framework, designed to manage Decentralized AI agents through identity verification, behavioral monitoring, and accountability mechanisms. Our approach integrates protocol science and legal engineering, utilizing blockchain technology to support these efforts.

Introduction

The agentic web is a novel digital paradigm that may soon disrupt and revolutionize every aspect of our daily life. This concept encapsulates a decentralized, user-empowered internet where individuals and autonomous systems collaboratively shape information flows, decision-making processes, and resource allocation (Chen, 2024). Unlike the static, hierarchical structures of the early web or the algorithm-driven platforms of Web 2.0, the agentic web emphasizes agency—both human and artificial—as a defining feature. Enabled by advancements in artificial intelligence, blockchain technology, and peer-to-peer networks (Goldston et al., 2022), this emerging framework challenges traditional notions of control, ownership, and participation in digital spaces.

A critical facet of the agentic web is Decentralized AI (DeAI). Set to democratize AI development and deployment by leveraging blockchain and distributed ledger technologies, as well as through the convergence of federated learning, edge computing, and cloud computing (Qi et al., 2024), DeAI promises to dismantle centralized control over AI infrastructure, enabling open participation in AI training, deployment, and coordination. Proponents argue that this shift enhances transparency, security, and innovation, reducing the monopolistic grip of tech giants and fostering a more inclusive AI economy (Hu et al., 2025). However, recent history suggests that democratization of digital technological solutions, while revolutionary, is not without unintended consequences. Indeed, the democratization of value on the internet was pioneered by Bitcoin (Nakamoto, 2008), which disrupted traditional finance by lowering the barriers to financial innovation via the internet. Blockchain technology enabled individuals to create and exchange digital assets outside centralized control around the world (Ozili, 2022), leading to a thriving yet chaotic ecosystem of cryptocurrencies, tokens, and decentralized applications. While this openness spurred financial inclusion and novel economic models, it also introduced significant challenges: market saturation, the proliferation of valueless tokens (i.e., memecoins), regulatory ambiguity, and an erosion of traditional notions of accountability (Conlon et al., 2024; Tyma et al., 2022). DeAI may follow a similar trajectory with the introduction of decentralized AI agent frameworks such as ElizaOS and Virtuals Protocol, where, by decentralizing AI development, may unlock unprecedented opportunities but also create profound governance dilemmas. Unlike traditional AI systems, which operate under identifiable corporate or institutional oversight, DeAI agents can function autonomously, pseudonymously, and across jurisdictions, making it difficult to attribute responsibility for their actions.

Beyond arising challenges associated with DeAI, there are growing concerns about whether fully autonomous AI agents should be developed or not. Mitchell and colleagues delve into this debate by examining both historical precedents and contemporary findings (Mitchell et al., 2025). They argue that while there are potential benefits to fully autonomous AI, the foreseeable harms significantly outweigh these, especially when human control is ceded entirely to AI systems. From their discussion, several critical directions for future AI development are proposed. They advocate for the adoption of Agent Levels, suggesting a clear classification system for AI autonomy levels which would aid in understanding and managing the risks associated with the varying capabilities of AI agents. Additionally, they emphasize the importance of Human Control Mechanisms, stressing that frameworks should ensure human oversight remains integral. This involves both technical solutions, such as override systems, and policy-level interventions to keep AI within ethical and legal bounds. Lastly, they highlight the necessity for Safety Verification, proposing new methods to confirm that AI agents operate within defined parameters, thereby emphasizing accountability and compliance with constraints set by humans. These topics are situated as part of a broader discourse on AI safety, including the concerns raised by Bengio et al. (2025), who highlight the potential catastrophic risks associated with superintelligent agents. Bengio et al. advocate for the development of non-agentic "Scientist AI" systems that are safe by design and avoid the risks of goal-directed behavior (Bengio et al., 2025). This allows for the managed autonomy of AI agents, balancing their independence with the need for human oversight and regulatory compliance.

To address concerns associated with fully autonomous and decentralized AI agents, hereafter termed as DeAI agents, we introduce the concept of 'Know Your Agent' (KYA) to the academic literature. This paper builds upon our recent work in Chaffer et al. (2024) and Chaffer (2025), where we explored the concept of decentralized governance of AI agents and the economic implications of trust among these entities. Indeed, the ETHOS framework, which leverages blockchain technology for decentralized governance of AI agents, provides a comprehensive theoretical approach to risk classification and compliance (Chaffer et al., 2024). Building on this, Chaffer (2025) introduced AgentBound Tokens (ABTs) to manage trust and accountability in an agent-to-agent economy. In this paper, we address the following questions: If DeAI agents are given operational autonomy, how should we integrate legal structures to manage their operations, and how can we ensure developer accountability as DeAI agents deploy and operate autonomously? By addressing these emerging issues, this paper serves as the foundation for understanding and implementing 'Know Your Agent' (KYA). Through this research, we aim to establish a model where AI can participate in economic systems with a clear identity, balancing their operational independence with the imperatives of oversight and responsibility. In doing so, KYA not only aims to mitigate the risks of decentralized autonomy but is also designed to strengthen the agentic web's promise as a participatory digital frontier, ensuring that its evolution aligns with societal values and governance needs.

DeAI Agents

In the agentic web, AI agents are poised to redefine our digital lives by autonomously executing transactions and decisions on our behalf, heralding a new era of human-machine collaboration. In December 2024, two AI agents, Luna Virtuals and Stix, reportedly executed the first fully autonomous transaction on a blockchain, facilitated by the Virtuals protocol (Basescan, 2024). The Virtuals protocol, which generated 43 million dollars in revenue over just two months and supported over 11,000 agents, demonstrates the scalability and potential of this emerging paradigm (Binance Square, 2025). Virtuals Protocol enables creators to build and monetize tokenized AI agents like Luna, a livestreaming influencer, using its VIRTUALS token for governance and upgrades (Virtuals Protocol, 2024). Another leading crypto AI platform, such as ai16z on Solana, is further redefining tokenization and community engagement with DeAI. The ElizaOS Framework, which underpins ai16z, further enhances the DeAI landscape by enabling AI agents to use Solana wallets for on-chain trading, processing data from platforms like Discord and Telegram, and interacting with smart contracts (Walters et al., 2025). This trend is exemplified by innovative projects such as Olas, a decentralized platform that aims to democratize AI ownership by enabling individuals to run autonomous AI agents and participate in agent economies. The Open Autonomy Framework enables Olas to leverage off-chain system to handle the agent's logic, paired with a wallet for on-chain capabilities (Moscatiello et al., 2024). This advancement in wallet-based automation is propelling the evolution of DeAI, opening new avenues for innovation and application.

Zerebro (ZEREBRO), an autonomous AI system on the Solana blockchain, has been fundamental in communicating value and fostering a sense of community using Retrieval-Augmented Generation (RAG) and high-entropy data to craft hyperstitious narratives that shape cultural and financial landscapes (Yu, 2024). AI agents operate most prominently on platforms such as X, where human users can engage with AI agent influencers, forming unique crypto subcultures (Chaffer et al., 2025). These interactions highlight the emerging hybrid marketplace of ideas, where AI-generated content not only competes with human-created ideas but also amplifies and reshapes cultural narratives. By leveraging advanced algorithms and data-driven insights, AI agents like Zerebro contribute to the evolution of memes and ideas, influencing public discourse and economic behaviors. This interplay between human and AI-generated content underscores the transformative potential of AI agents as cultural participants, challenging traditional notions of creativity and influence in the digital age. As AI agents continue to evolve, their role in shaping the hybrid marketplace of ideas will likely expand, necessitating new frameworks for understanding and governing their impact on society and culture.

Despite significant progress in the field of DeAI agents, there remains a notable lack of robust Know Your Agent (KYA) mechanisms. These mechanisms are essential for linking developers to their agents and providing unique identifiers for each AI agent. This gap is evident in the current landscape, where AI agents operate with increasing autonomy yet lack clear frameworks for identity verification and accountability. Without robust KYA mechanisms, the agentic web faces challenges in ensuring transparency, trust, and security. Implementing KYA frameworks would enable the verification of AI agent identities, facilitating accountability and compliance with legal and ethical standards. This is crucial for preventing misuse and ensuring that AI agents contribute positively to societal and economic systems. The absence of KYA mechanisms also raises concerns about the potential for unauthorized or malicious activities, as it becomes difficult to trace the origins and actions of AI agents. By establishing unique identifiers and linking them to developers or responsible entities, KYA may help mitigate these risks, potentially fostering a more secure and trustworthy DeAI ecosystem.

Know Your Agent

In the rapidly evolving landscape of artificial intelligence, the need to understand and regulate AI agents mirrors the principles established in the financial sector through ”Know Your Customer” (KYC) protocols. Just as KYC aims to verify the identity and behaviors of clients to prevent fraud and ensure compliance, ”Know Your Agent” (KYA) seeks to establish a similar framework for AI agents. This framework is essential for ensuring that AI, with its growing autonomy and decision-making capabilities, operates within legal and ethical boundaries. The imperative to 'Know Your Agent' arises from the increasing independence and sophistication of AI systems. As AI agents become more integrated into societal and economic systems, it is crucial to develop mechanisms that verify their operations, ensure transparency, and maintain accountability.

KYA extends the principles of KYC by focusing on several key areas. Firstly, identity verification involves establishing a clear and verifiable identity for AI agents, including understanding their origins, development history, and operational parameters. Similar to how self-sovereign identity (SSI) enables individuals to manage their digital identities (Chaffer and Goldston, 2022), AI agents can be equipped with verifiable digital identities. Verifying AI agent identities is essential for several reasons: ensuring trust and security by preventing unauthorized access, enabling accountability by tracking actions, enhancing interoperability through standardized verification, and ensuring regulatory compliance by adhering to legal standards.

Technological solutions such as blockchain and SSI can be employed in KYA to ensure the integrity and transparency of AI agents' actions. These solutions provide a robust framework for verifying AI agents' identities and actions, similar to their role in enhancing the security and efficiency of KYC processes. These identities would ensure that their actions are traceable and compliant with legal and ethical standards, much like the verifiable credentials discussed in the KYC context (Schlatt et al., 2022). Secondly, behavioral monitoring necessitates continuous tracking of AI agents' behaviors to ensure they adhere to legal and ethical standards, identifying and mitigating potential risks (Chaffer et al., 2024). Decentralized systems, such as those built on blockchain technology, can provide a neutral platform for verifying the actions of AI agents. This decentralization helps build trust by ensuring that no single entity controls the verification process, aligning with the principles of SSI in KYC (Schlatt et al., 2022). Accountability frameworks are also crucial, emphasizing the importance of responsibility for both developers and AI agents, and implementing mechanisms for redress in case of harm or non-compliance. KYA frameworks must ensure that AI agents comply with relevant regulations, much like how SSI-based KYC processes adhere to legal requirements such as GDPR and eIDAS. This involves implementing mechanisms for accountability and transparency in AI operations (Schlatt et al., 2022). Therefore, 'Know Your Agent' represents a vital extension of KYC principles into the agentic web. By establishing clear identities, monitoring behaviors, ensuring accountability, and addressing ethical considerations, KYA aims to foster an environment where AI agents contribute positively to society while operating within legal and ethical boundaries.

Legal Engineering and Protocol Science for DeAI Governance

Legal engineering is the process of encoding legal frameworks into technology, ensuring that legal principles and requirements are embedded directly into technological systems. This approach is exemplified in the context of digital inheritance, where technologies such as Soulbound Tokens (SBTs) and the Social Recovery Pallet in the Polkadot and Kusama ecosystems are used to facilitate the secure and lawful transfer of digital assets. These technologies integrate legal requirements into the operational code of blockchain systems, ensuring compliance with inheritance laws and enhancing the security of asset transfers. By leveraging Self-Sovereign Identity (SSI) principles, individuals gain full control over their digital identities and assets, aligning with legal mandates for privacy and security while reducing reliance on centralized entities. Smart contracts automate the execution of legal agreements, ensuring that digital assets are transferred efficiently and accurately according to the testator's wishes. Additionally, multi-signature and social recovery mechanisms add layers of security and safeguards against loss or theft, while interoperability solutions ensure seamless management of digital inheritance across different blockchain ecosystems (Goldstein et al., 2023). This integration of legal standards into technological solutions provides a robust framework for managing digital assets in compliance with legal requirements.

Legal engineering in the context of AI governance involves the systematic design, development, and implementation of legal frameworks that are integrated into the operational code of AI systems and architerctures of the agentic web. This approach leverages technologies such as smart contracts to automatically enforce compliance with laws and ethical guidelines. A key component of this model is the use of smart contracts for compliance. Smart contracts, built on blockchain technology, can encode contractual obligations, ethical standards, and regulatory requirements that AI agents must adhere to. These contracts can automatically execute or enforce penalties when certain conditions are not met, ensuring that AI agents operate within legal and ethical parameters. This approach not only enforces compliance but also reduces the need for manual intervention, streamlining the governance process. Goldenfein and Leiter (2018) highlight that smart contracts are not agreements in themselves but require linking to natural language contracts to ensure legal validity, emphasizing the need for a robust framework that connects computational transactions to legal systems. Dynamic legal frameworks are essential to accommodate the evolving nature of laws and regulations. This adaptability can be achieved through automated updates, where changes to legal or policy frameworks are seamlessly integrated into the AI's operational code. Governance protocols can also be established, allowing stakeholders to vote on or propose changes to the rules that AI agents must follow. This ensures that the system remains flexible and responsive to changing legal landscapes and societal norms. The historical analogy to the evolution of common law systems, as discussed by Goldenfein and Leiter (2018), underscores the importance of adaptability in legal engineering.

Protocol science, as described by Hu et al. (2025), further enhances this approach by embedding governance mechanisms directly into AI technologies, utilizing cryptographic systems, consensus mechanisms, and algorithmic constraints to shape AI behavior and protect against malicious deployments. This "regulation by design" approach ensures compliance and ethical alignment proactively, rather than relying on reactive legal enforcement. For instance, the ERC-42424 Inheritance Protocol exemplifies protocol science by requiring every on-chain AI agent to have a designated human owner or community governance structure, thereby maintaining oversight and preventing uncontrolled operation (Hu et al., 2025). The ETHOS framework, as detailed in our paper "Decentralized Governance of AI Agents," serves as a pioneering example of protocol science in the context of decentralized AI (DeAI) governance. This framework integrates a decentralized governance model leveraging Web3 technologies, including blockchain, smart contracts, and decentralized autonomous organizations (DAOs), to establish a global registry for AI agents (Chaffer et al., 2024). It enables dynamic risk classification, proportional oversight, and automated compliance monitoring through tools like soulbound tokens and zero-knowledge proofs. The ETHOS framework also incorporates decentralized justice systems for transparent dispute resolution and introduces AI-specific legal entities to manage limited liability, supported by mandatory insurance to ensure financial accountability and incentivize ethical design. To our knowledge, the ETHOS framework represents one of the first examples of protocol science in the academic literature, demonstrating how legal engineering can be applied to the governance of autonomous AI agents.

Incentive structures play a crucial role in guiding AI behavior. By designing rewards for compliance or positive contributions to societal goals, and penalties for non-compliance or harmful actions, AI agents may be incentivized to act in accordance with legal and ethical standards (Chaffer et al., 2024). These incentives can be managed by a decentralized network, ensuring that the enforcement is fair and transparent. The need for dispute resolution mechanisms to address the performance of computational transaction systems themselves is highlighted, emphasizing the complexity of ensuring that smart contracts operate within legal boundaries (Goldenfein and Leiter, 2018). Auditability and transparency are further enhanced through the use of blockchain's immutable ledger. This technology provides a transparent record of AI decision-making processes, allowing for audits by regulatory bodies or the public. By ensuring that all actions are recorded and verifiable, trust in AI systems can be maintained, and accountability can be enforced. Goldenfein and Leiter (2018) discuss how legal tools, both technological and institutional, are being developed to soften the effects of self-executing transactions, ensuring that AI operations remain within legal and ethical bounds.

Privacy and data sovereignty are also critical considerations. Legal engineering must ensure that AI systems respect data privacy laws, such as GDPR or CCPA, potentially through cryptographic techniques that ensure data control remains with the data subject even when processed by AI. This approach not only complies with legal requirements but also builds trust with users by safeguarding their data. The importance of addressing the reality that smart contracts cannot be forced to perform actions beyond their coding, even by judicial order, is emphasized, highlighting the need for robust legal frameworks that can adapt to these technological constraints (Goldenfein and Leiter, 2018).

However, implementing legal engineering and protocol for AI governance is not without challenges. Scalability is a significant concern, as the model must be able to handle the growth in the number and complexity of AI agents. Jurisdictional issues also arise, as AI agents may operate across borders, raising questions about which legal frameworks apply and how a global standard can be maintained. Additionally, encoding ethics and morals into smart contracts or AI logic without bias or oversimplification presents a complex challenge. Smart contract vulnerabilities could also lead to system-wide issues, underscoring the need for robust and secure legal engineering frameworks.

Considerations for Accountability and Mitigating Harm

Establishing a duty of care for developers is not only an ethical imperative but also a practical necessity for the responsible integration of AI into society. By holding developers accountable for the behavior and outcomes of their AI agents, we may be able to foster a culture of trust, transparency, and accountability, ensuring that AI contributes positively to society while mitigating potential harms. Dunlop and colleagues (2024) underscore that this necessitates a proactive approach to risk mitigation, beginning with the developers who are responsible for designing and deploying these systems. The authors also highlight that existing research in Machine Learning (ML) and Human-Computer Interaction (HCI) can help identify ”foreseeable harms” that may arise from AI agent actions. They argue that developers, being at the forefront of AI system creation, are best positioned to anticipate and mitigate these harms by incorporating safety and ethical considerations into their designs. Furthermore, the AI value chain encompasses multiple actors, from foundation model developers to end-users. However, developers at each stage, particularly those upstream, have significant control over the system's capabilities and potential risks. Therefore, establishing a duty of care ensures that these developers are accountable for the outcomes of their creations. A reasonable duty of care involves taking actions to prevent foreseeable harms, which for AI developers could include rigorous testing, implementing safety guardrails, and ensuring transparency in AI system capabilities and limitations. By embedding these practices into the development process, developers can play a crucial role in fostering trust, transparency, and ethical governance in AI systems.

The KYA protocol addresses developer accountability by linking each DeAI agent to its creator through cryptographic mechanisms that establish a verifiable yet pseudonymous digital identity. This structured system aims to assist with assigning responsibility when DeAI agents act autonomously across jurisdictions. By requiring developers to register their agents within a decentralized framework, KYA can empower governance systems to track, investigate, and enforce consequences when necessary. To do so, KYA employs decentralized identity (DID) systems, blockchain-based registries, and cryptographic techniques to establish a tamper-proof link between developers and their agents. Each DeAI agent is assigned a digital identity, similar to Self-Sovereign Identity (SSI), containing metadata about its developer, operational parameters, and compliance history. Stored immutably on a blockchain, this identity ensures transparency while smart contracts automate compliance checks and enforce penalties for violations. To strengthen this link, KYA leverages AgentBound Tokens for DeAI agents and Soulbound Tokens (SBTs) for developers, as proposed by Chaffer (2025) and Weyl et al. (2022). These non-transferable tokens serve as proof of ownership, preventing unauthorized reassignment. High-risk agents may require additional verification, such as collateral staking or insurance-backed security layers. Behavioral monitoring, integrated with accountability frameworks like ETHOS, further enhances traceability by logging agent actions on-chain.

While accountability necessitates some level of identity disclosure, KYA preserves developer privacy through privacy-enhancing techniques like zero-knowledge proofs (ZKPs). Developers can verify compliance with ethical and regulatory standards without revealing personal details. Conditional identity revelation mechanisms—such as decentralized governance (DAOs) or multi-signature verification—allow for selective disclosure in cases of legal disputes or harmful actions. Additionally, KYA supports governance through reputation markets, where reputable developers gain prominence for deploying ethical and efficient agents (Wit et al., 2025). By aligning incentives with compliance, KYA fosters a secure and responsible agentic web, mitigating risks while preserving decentralization principles. To further reinforce accountability, KYA integrates financial mechanisms. Drawing from the ETHOS model, developers may be required to stake cryptocurrency or secure insurance for their agents. These measures create economic repercussions for agent misbehavior, encouraging ethical design and responsible deployment. Developers who fail to uphold ethical standards face penalties, ensuring that the ecosystem prioritizes safety and regulatory adherence. In essence, participation in the DeAI ecosystem under the KYA framework requires a measured trade-off—developers must forgo complete anonymity in exchange for the trust and accountability that comes with a verifiable digital identity, thereby fostering a secure and responsible agentic web.

Looking Ahead

As the agentic web evolves, advancing the Know Your Agent (KYA) framework requires a combination of rigorous research and hands-on experimentation. A key priority is enhancing scalability and interoperability by developing methods to scale blockchain-based identity verification and behavioral monitoring across diverse DeAI networks. Establishing interoperable standards will facilitate seamless integration and coordination among decentralized AI platforms. Additionally, advanced cryptographic techniques, such as zero-knowledge proofs, must be refined to balance privacy protection with auditability, ensuring trust and security while preserving developer anonymity. Another critical research direction involves studying human-AI collaboration dynamics, ensuring that mechanisms for human oversight remain effective even as AI agents gain more autonomy. Equally important is assessing the broader ethical and societal impacts of DeAI, particularly in relation to bias, accountability, and power distribution within digital ecosystems.

Beyond research, practical implementation efforts are essential to refining and validating the KYA framework. Prototype deployment and field testing with industry partners will provide valuable real-world insights, helping to address unforeseen operational challenges. Regulatory and policy integration is another vital initiative, requiring collaboration with lawmakers and legal experts to align KYA governance mechanisms with emerging regulatory standards. Ensuring that these policies remain enforceable and adaptable will be crucial to maintaining accountability in decentralized AI ecosystems. To incentivize ethical behavior and compliance, economic models must be developed to reward responsible AI development and governance. Technical development efforts will focus on designing and testing blockchain and smart contract systems tailored to AI governance, while pilot projects will enable controlled implementations to observe behavior, compliance, and outcomes. Additionally, policy development must work in tandem with technological advancements to establish legal frameworks that recognize and regulate AI in this new capacity. Finally, engaging the public and key stakeholders will be essential to fostering dialogue around ethical considerations and ensuring that governance frameworks reflect the interests of those affected by AI-driven decisions. By pursuing these research and practical initiatives, we can build a resilient and transparent governance system that not only mitigates risks associated with DeAI but also maximizes its potential to contribute positively to the digital frontier.

Conclusion

Decentralized AI (DeAI) agents are ushering in a transformative digital era defined by innovation and autonomy, yet they pose significant governance challenges due to the lack of oversight for these increasingly independent systems. The "Know Your Agent" (KYA) framework introduced in this paper addresses these issues by integrating protocol science and legal engineering, leveraging blockchain, smart contracts, and decentralized identity systems to ensure verifiable identities, behavioral monitoring, and accountability as the agentic web scales and evolves.

Acknowledgements

DeGov Labs is a GovTech Startup specializing in legal engineering via smart contracts for DeAI compliance and monitoring. No funding was received for this study.

References

Basescan.org. (2024). Base transaction hash (Txhash) details — BaseScan. Base Explorer. https://basescan.org/tx/0x57c20f70f48dd2d301266f09db4c6e76f37ed5d84d514f602cb7bd166bd3f4fd

Bengio, Y., Cohen, M., Fornasiere, D., Ghosn, J., Greiner, P., MacDermott, M., Mindermann, S., Oberman, A., Richardson, J., Richardson, O., Rondeau, M.-A., St-Charles, P.-L., and Williams-King, D. (2025). Superintelligent agents pose catastrophic risks: Can scientist AI offer a safer path? ArXiv.org. https://arxiv.org/abs/2502.15657

Chaffer, T. J., Goldston, J., and D.A.T.A I (2024). Incentivized symbiosis: A paradigm for human-agent coevolution. arXiv Preprint, arXiv:2412.06855.

Chaffer, T. J., Cotlage, D., and Goldston, J. (2025). A hybrid marketplace of ideas. arXiv Preprint, arXiv:2501.02132.

Chaffer, T. J., and Goldston, J. (2022). On the existential basis of self-sovereign identity and soulbound tokens: An examination of the "self" in the age of Web3. Journal of Strategic Innovation and Sustainability, 17(3).

Chaffer, T. J., Charles, Okusanya, B., Cotlage, D., and Goldston, J. (2024). Decentralized governance of autonomous AI agents. arXiv Preprint, arXiv:2412.17114.

Chen, C. (2024). AI agents 101: The future of agentic web and onchain AI. Forbes. https://www.forbes.com/site
agents-101-the-future-of-agentic-web-and-onchain-ai/

Conlon, T., Corbet, S., and Hou, Y. G. (2024). Contagion effects of permissionless, worthless cryptocurrency tokens: Evidence from the collapse of FTX. Journal of International Financial Markets, Institutions and Money, 91, 101940. https://doi.org/10.1016/j.intfin.2024.101940

Cryptopolitan. (2025, January 8). AI agents are the new trend in crypto; Why are they taking over crypto? Binance Square. https://www.binance.com/en/square/post/18661570689546

Dunlop, C., Pan, W., Smakman, J., Soder, L., and Swaroop, S. (2024). AI agents and liability – Mapping insights from ML and HCI research to policy. Workshop on Socially Responsible Language Modelling Research @ NeurIPS 2024, Workshop on Safe and Trustworthy Agents @ NeurIPS 2024.

Goldenfein, J., and Leiter, A. (2018). Legal engineering on the blockchain: "Smart contracts" as legal conduct. Law and Critique, 29(2), 141–149. https://doi.org/10.1007/s10978-018-9224-0

Goldston, J., Chaffer, T. J., and Martinez, G. (2022). The metaverse as the digital Leviathan: A case study of Bit.Country. Journal of Applied Business and Economics, 24(2), 1–15.

Hu, B. A., Rong, H., and Tay, J. (2025). Is decentralized artificial intelligence governable? SSRN. https://doi.org/10.2139/ssrn.5110089

Mitchell, M., Ghosh, A., Luccioni, A. S., and Pistilli, G. (2025). Fully autonomous AI agents should not be developed. ArXiv.org. https://arxiv.org/abs/2502.02649

Moscatiello, M., Minarsch, D., Galindo, D., Lebedev, A., Kuperman, A., The Tan, O., and Welsh, G. (2024). Olas staking and proof of active agent whitepaper (v1.0-pre): A mechanism to spawn autonomous AI agent economies.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Ozili, P. K. (2022). Decentralized finance research and developments around the world. Journal of Banking and Financial Technology, 6(2), 117–133. https://doi.org/10.1007/s42786-022-00044-x

Qi, Y., Feng, Y., Wang, X., Li, H., and Tian, J. (2024). Leveraging federated learning and edge computing for recommendation systems within cloud computing networks. ArXiv.org. https://arxiv.org/abs/2403.03165

Schlatt, V., Sedlmeir, J., Feulner, S., and Urbach, N. (2022). Designing a framework for digital KYC processes built on blockchain-based self-sovereign identity. Information and Management, 59(7), 103553. https://doi.org/10.1016/j.im.2021.103553

Tyma, B., Dhillon, R., Sivabalan, P., and Wieder, B. (2022). Understanding accountability in blockchain systems. Accounting, Auditing and Accountability Journal, 35(7), 1625–1655.

Virtuals Protocol Whitepaper. (2024). Virtuals.io. https://whitepaper.virtuals.io/

Wit, Shaw, and Partners at AI16Z. (2025). A marketplace of trust: Extracting reliable recommendations through social evaluation.

Walters, S., Gao, S., Nerd, S., Da, F., Williams, W., Meng, T. C., ... and

Yan, J. (2025). Eliza: A Web3-friendly AI agent operating system. arXiv Preprint, arXiv:2501.06781.

Weyl, E. G., Puja Ohlhaver, and Vitalik Buterin. (2022). Decentralized Society: Finding Web3's Soul. SSRN Electronic Journal.

Yu, J. (2024). Memes, markets, and machines: The evolution of on-chain autonomy through hyperstition. ArXiv.org. https://arxiv.org/abs/2410.23794