

AI Blueprint for the Future

A large, light gray background graphic. On the left, a stylized brain outline is formed by thick, flowing lines. On the right, a circuit board pattern with various lines and dots extends vertically, merging with the brain's structure.

Coalition for Innovation, supported by LG NOVA

Jami Diaz, Director Ecosystem Community & Startup Experience
William Barkis, Head of Grand Challenges & Ecosystem Development
Sokwoo Rhee, Executive Vice President, LG Electronics, Head, LG NOVA

Coalition for Innovation Co-Chairs

Alex Fang, CleanTech Chair
Sarah Ennis, AI Chair
Alfred Poor, HealthTech Chair

Authors

Adrien Abecassis, Johnny Aguirre, John Barton, Ann M. Marcus, Olivier Bacs, Taylor Black, Micah Boster, Mathilde Cerioli, Carolyn Eagen, Sarah Ennis, Annie Hanlon, Christina Lee Storm, Andrew Yongwoo Lim, Jess Loren, Refael Shamir, Svetlana Stotskaya

The views and opinions expressed in the chapters and case studies that follow are those of the authors and do not necessarily reflect the views or positions of any entities they represent.

Senior Editor, Alfred Poor
Editor, Jade Newton

October 2025



Preamble

The Coalition for Innovation is an initiative hosted by LG NOVA that creates the opportunity for innovators, entrepreneurs, and business leaders across sectors to come together to collaborate on important topics in technology to drive impact. The end goal: together we can leverage our collective knowledge to advance important work that drives positive impact in our communities and the world. The simple vision is that we can be stronger together and increase our individual and collective impact on the world through collaboration.

This “Blueprint for the Future” document (henceforth: “Blueprint”) defines a vision for the future through which technology innovation can improve the lives of people, their communities, and the planet. The goal is to lay out a vision and potentially provide the framework to start taking action in the areas of interest for the members of the Coalition. The chapters in this Blueprint are intended to be a “Big Tent” in which many diverse perspectives and interests and different approaches to impact can come together. Hence, the structure of the Blueprint is intended to be as inclusive as possible in which different chapters of the Blueprint focus on different topic areas, written by different authors with individual perspectives that may be less widely supported by the group.

Participation in the Coalition at large and authorship of the overall Blueprint document does not imply endorsement of the ideas of any specific chapter but rather acknowledges a contribution to the discussion and general engagement in the Coalition process that led to the publication of this Blueprint.

All contributors will be listed as “Authors” of the Blueprint in alphabetical order. The Co-Chairs for each Coalition will be listed as “Editors” also in alphabetical order. Authorship will include each individual author’s name along with optional title and optional organization at the author’s discretion.

Each chapter will list only the subset of participants that meaningfully contributed to that chapter. Authorship for chapters will be in rank order based on contribution: the first author(s) will have contributed the most, second author(s) second most, and so on. Equal contributions at each level will be listed as “Co-Authors”; if two or more authors contributed the most and contributed equally, they will be noted with an asterisk as “Co-First Authors”. If two authors contributed second-most and equally, they will be listed as “Co-Second Authors” and so on.

The Blueprint document itself, as the work of the group, is licensed under the Creative Commons Attribution 4.0 (aka “BY”) International License: <https://creativecommons.org/licenses/by/4.0/>. Because of our commitment to openness, you are free to share and adapt the Blueprint with attribution (as more fully described in the CC BY 4.0 license).

The Coalition is intended to be a community-driven activity and where possible governance will be by majority vote of each domain group. Specifically, each Coalition will decide which topics are included as chapters by majority vote of the group. The approach is intended to be inclusive so we will ask that topics be included unless they are considered by the majority to be significantly out of scope.

We intend for the document to reach a broad, international audience, including:

- People involved in the three technology domains: CleanTech, AI, and HealthTech
- Researchers from academic and private institutions
- Investors
- Students
- Policy creators at the corporate level and all levels of government



Chapter 5: Benefits and Drawbacks of Decentralized AI

Authors: Olivier Bacs, Carolyn Eagen



Overview

In an age where *ambient computing* – the seamless embedding of intelligent services into everyday environments – is gaining traction, decentralization is no longer an ideological ideal. It has become a commercial and infrastructural imperative. As inference (the act of “thinking” by AI models) increasingly needs to happen *offline-first*, privacy by default becomes not just a feature, but a requirement. This avoids transmitting sensitive tasks to remote servers, better aligning with legal, ethical, and user expectations ([Shi et al., 2016](#)).

This shift toward edge-based, privacy-preserving AI marks more than just a benevolent technical evolution; it reveals deeper structural tensions within the broader AI ecosystem. While decentralization is being driven by technical necessity at the edge, the artificial intelligence

landscape at large faces a critical juncture as the current centralized paradigm creates increasingly problematic bottlenecks in innovation, raises serious concerns about data privacy and algorithmic bias, and limits equitable access to AI capabilities across diverse organizations and communities ([Jobin et al., 2019](#)).

A handful of large technology companies dominate control over foundational models, training data, and computational infrastructure; this has resulted in concentration risk, data sovereignty issues, transparency deficits, access inequality, and compliance complexity that collectively threaten the democratic potential of AI development ([Barocas et al., 2019](#)). These dynamics raise structural concerns: Who decides what is permissible? Whose values get embedded into models? Who watches the watchers? These aren’t just ethical dilemmas; they’re market limitations. Governance – which is often the quietest element in environmental, social, and governance (ESG) debates – takes center stage



when decentralization is framed as a route to both resilience and self-determination.

In response to these systemic challenges, a paradigm shift toward decentralized AI systems has emerged, promising to distribute power and control more equitably while prioritizing transparency, user control, and community governance. This transition represents not merely a technical evolution but a fundamental reimagining of how AI systems are developed, deployed, and governed.

Decentralized AI envisions distributed infrastructure where computing, storage, and governance are spread across networks of participants rather than concentrated in centralized data centers; community ownership that provides stakeholders with meaningful participation in development and monetization; transparent operations through open-source models and auditable processes; consent-based data usage that maintains user control and fair compensation; and modular architecture that enables customization and innovation without platform lock-in ([Zuboff, 2019](#)).

List of Stakeholders (audience/readers)

The movement to decentralize AI is being shaped not only by community values but also by the emerging incentives of strategic players. From EleutherAI to Hugging Face, decentralization is now attracting both venture capital and developer mindshare. Even former insiders, such as Emad Mostaque (formerly of StabilityAI), have embraced open diffusion models, though critics note the ambiguity of such transitions, raising questions about whether decentralization is a narrative being co-opted or a movement being broadened.

To understand the real trajectory of this decentralization movement, it is essential to examine the diverse ecosystem of stakeholders actively involved in or impacted by this shift. Each group brings distinct priorities, challenges, and incentives that shape how decentralized AI

systems are being developed, adopted, and governed.

The technical community includes open-source developers and maintainers who build and sustain decentralized AI infrastructure; AI researchers and academics pursuing democratic access to computational resources; infrastructure providers and cloud services adapting to distributed architectures; and edge computing hardware manufacturers enabling local AI processing capabilities.

Commercial entities encompass AI startups seeking alternatives to big tech platforms and vendor lock-in; enterprise customers requiring compliance frameworks and auditability in their AI systems; SaaS companies building vertical AI solutions for specialized markets; and traditional software companies integrating AI capabilities into existing products and services.

The governance and policy sphere includes regulatory bodies developing AI compliance frameworks; government agencies implementing public sector AI initiatives; international organizations establishing AI standards and best practices; and digital rights advocates representing civil society interests.

Straddling both, startups such as Modular are making decentralized AI stack components commercially viable while still open-sourcing their research and runtime tools, illustrating that performance and profitability need not require enclosure. By lowering the barrier to sovereign infrastructure, these players are laying down the groundwork for sustainable decentralized ecosystems ([Modular, 2024](#)).

End users and communities represent perhaps the most critical stakeholder group, including data creators and content producers whose work trains AI systems; marginalized communities disproportionately affected by AI bias and discrimination; privacy-conscious individuals and organizations seeking greater control over their data; and emerging markets with limited access to centralized AI services due to cost or infrastructure constraints ([Benjamin, 2019](#)).



Challenges and Gaps

Current centralized AI systems exhibit several critical limitations that create urgent needs for alternative approaches. Concentration risk manifests as a small number of companies controlling the majority of AI capabilities, creating single points of failure that can disrupt entire sectors and limiting competitive dynamics that would otherwise drive innovation and reduce costs ([Parker, 2016](#)). This concentration enables these companies to set prices, determine access policies, and shape the direction of AI development according to their commercial interests rather than broader societal needs.

Data sovereignty represents another fundamental challenge, as users have minimal control over how their information is collected, processed, and monetized in AI training pipelines ([Lanier, 2013](#)). Personal data, creative works, and professional content are incorporated into training datasets without meaningful consent or compensation, creating extractive relationships that benefit centralized platforms while providing little value to data creators.

The transparency deficit inherent in proprietary models, which operate as "black boxes," makes it difficult to audit for bias, to understand decision-making processes, or to ensure compliance with evolving regulatory requirements ([Burrell, 2016](#)).

Access inequality creates significant barriers for smaller organizations, developing regions, and specialized use cases that cannot afford the high computational costs and platform restrictions imposed by centralized providers ([Birhane, 2021](#)). This digital divide threatens to exacerbate existing inequalities and limit innovation to well-funded entities in developed markets. Compliance complexity further compounds these challenges, as centralized systems struggle to meet diverse regulatory requirements across different jurisdictions and sectors, creating legal risks for organizations that depend on these platforms. This digital divide threatens to exacerbate existing inequalities and limits innovation. In addition, compliance complexity further compounds these challenges ([Aissaoui, 2021](#); [Marotta et al., 2021](#)).

A New Vision

We envision a decentralized AI ecosystem that fundamentally transforms how artificial intelligence systems are developed, deployed, and governed. This new paradigm prioritizes distributed infrastructure where computing power, data storage, and decision-making authority are spread across networks of voluntary participants rather than concentrated in corporate data centers controlled by a few powerful entities. Community ownership mechanisms ensure that stakeholders have meaningful participation in the development, governance, and monetization of AI systems, creating democratic processes for determining how these powerful technologies are used and who benefits from their value creation.

Transparent operations through open-source models and auditable processes enable scrutiny and accountability, allowing researchers, regulators, and affected communities to understand how AI systems make decisions and identify potential sources of bias or error. Consent-based data usage frameworks maintain user control over personal information while providing fair compensation for contributions to AI training datasets, addressing the extractive dynamics that characterize current data collection practices. Modular architecture designs enable interoperability and customization without vendor lock-in, allowing organizations to combine components from different providers and adapt systems to their specific needs without dependence on any single platform.

This vision extends beyond technical architecture to encompass new economic models that distribute value more equitably among all participants in the AI ecosystem. Rather than concentrating profits in a few large corporations, decentralized systems can provide direct compensation to data contributors, reward open-source developers for their contributions, and enable communities to capture value from AI systems that serve their needs. The goal is to create AI systems that are not only more technically robust and innovative but also more aligned with democratic values and social equity principles.



Driving Forces Behind AI Decentralization

The movement toward decentralized AI emerges from diverse actors with varying motivations and capabilities, each contributing unique perspectives and resources to this evolving ecosystem. Open-source communities have established themselves as fundamental drivers of democratization, with organizations such as Hugging Face, EleutherAI, and LAION working systematically to remove corporate gatekeeping mechanisms and to ensure that AI capabilities remain publicly accessible ([Osborne et al., 2024](#)). These communities have achieved remarkable success in producing competitive alternatives to proprietary models, including BLOOM, Falcon, and various fine-tuned variants that match or exceed the performance of closed systems in specific domains while maintaining full transparency about their development and capabilities.

The intersection of Web3 and blockchain ecosystems with AI development has introduced novel economic and technical frameworks for decentralized model training, governance, and monetization. Innovative startups including Ocean Protocol, Gensyn, Bittensor, and Fetch.ai leverage blockchain technology to create sophisticated incentive mechanisms for distributed computing, data sharing, and collaborative AI development ([Shi et al., 2016](#)). These platforms demonstrate how cryptoeconomic principles can align individual incentives with collective goals, enabling large-scale coordination without centralized control while ensuring fair compensation for all participants.

Infrastructure development provides the foundational layer for decentralized AI systems, with protocols like NEAR Protocol's Aurora, Ethereum, and Filecoin/IPFS delivering scalable, censorship-resistant capabilities for AI workloads ([Benet, 2014](#)). These protocols enable computing and storage solutions that operate independently of traditional cloud providers, creating new possibilities for autonomous AI development and deployment that cannot be controlled or shut down by any single entity.

Academic and research initiatives legitimize and advance decentralized AI through collaborative, multi-institutional efforts that prioritize scientific openness over proprietary advantages. Projects such as BigScience – which produced the BLOOM model – and OpenMined demonstrate how distributed research can achieve outcomes comparable to well-funded commercial projects while ensuring democratic access to results ([Scao et al., 2022](#)). These initiatives establish precedents for public-good AI development that serves broad community interests rather than narrow commercial objectives.

Beyond Ideology: Commercial Opportunities in Decentralized AI

While early decentralized AI efforts were often motivated by idealistic goals around democratization and transparency, the sector has increasingly attracted substantial commercial interest as viable business models have emerged and market opportunities have become apparent. Open-source AI innovators – including companies such as Hugging Face, LAION, BigScience, and Mistral.ai demonstrate that building and maintaining high-performing open models can create sustainable competitive advantages without relying on proprietary lock-in strategies ([Bommasani et al., 2021](#)). These organizations enable startups and enterprises to build applications on transparent, customizable foundations while generating revenue through ecosystem development, support services, and premium features rather than platform control.

Decentralized infrastructure builders represent a significant commercial opportunity, with projects such as Aurora (NEAR Protocol), Filecoin/IPFS, Gensyn, and Bittensor providing decentralized compute, storage, and smart contract capabilities that can support AI workloads at scale. These platforms enable cost-effective infrastructure for running and monetizing AI applications without dependence on traditional cloud providers, potentially disrupting established patterns of infrastructure ownership and creating new markets for distributed computing resources ([Keršič, V., et al., 2025](#)).



Vertical AI startups have found particular success leveraging modular open-source AI components to build specialized Software-as-a-Service (SaaS) products for underserved markets. Companies such as Kinstak (AI digital legacy vaults), Lex (AI for legal services), Phind (AI-powered coding search), Bendi (AI-powered supplier communications), and DoNotPay (legal automation) demonstrate how decentralized components enable rapid development and deployment while maintaining control over technology stacks and customer relationships (Chen et al., 2021). This approach allows smaller companies to compete with larger incumbents by focusing on domain expertise and customer service rather than foundational AI development.

The emergence of DAO-led data cooperatives introduces novel approaches to fair monetization and consent-based frameworks in AI development. Organizations such as Ocean Protocol, DataUnion.app, and Gitcoin enable communities to pool data resources, govern their use through democratic processes, and share revenue generated from AI training activities (Pentland et al., 2019). These models create new possibilities for equitable value distribution in data-driven AI systems while maintaining community control over how information is used and monetized.

Monetization Strategies for Decentralized AI

The transition to decentralized AI creates distinct opportunities and challenges for different stakeholder groups, fundamentally altering traditional patterns of value creation and distribution in the AI ecosystem. Creators and data contributors stand to benefit significantly through royalties, tokenized licensing, and consent-driven monetization mechanisms that provide direct compensation for their contributions to AI training datasets (Arrieta-Ibarra et al., 2018). This represents a fundamental shift from the current extractive model where personal data and creative works are incorporated into commercial AI systems without compensation or meaningful consent.

Open-source developers gain new opportunities to monetize fine-tuned models, plugins, and

specialized AI services, moving beyond volunteer contributions to sustainable careers in decentralized AI development. Emerging markets and underserved users benefit from access to low-cost, localized alternatives to expensive centralized services, enabling AI adoption in regions and sectors previously excluded from these capabilities. Decentralized autonomous organizations and cooperatives that govern AI systems democratically can share revenue among participants, creating new models of collective ownership and benefit distribution (Hakkarainen, 2021).

Edge hardware innovators benefit from increased demand for devices capable of supporting decentralized inference on consumer and IoT platforms, potentially shifting value from centralized data centers to distributed computing resources owned by end users. This creates opportunities for hardware manufacturers to develop specialized chips and devices optimized for local AI processing while enabling users to monetize their computational resources.

Revenue model innovations in decentralized AI span multiple approaches, each with distinct implications for different stakeholders. Pay-per-inference micropayments enable decentralized model usage tracking and billing through smart contracts, creating granular pricing mechanisms that better reflect actual usage patterns while enabling automated compensation for model providers (Catalini & Gans, 2020). Data royalty systems ensure that contributors earn ongoing compensation when their information is used to train or retrain AI models, addressing long-standing concerns about unpaid labor in AI development while creating sustainable income streams for content creators.

The Double-Edged Sword of Unregulated AI Generation

Decentralized AI presents a complex dual nature, offering significant benefits while simultaneously introducing new categories of risks that require careful management and mitigation strategies. As decentralized AI reduces dependence on hyperscalers and enhances privacy through local



inference, it also complicates governance and risk mitigation.

The positive aspects of decentralization include empowering user control and data sovereignty, which allows individuals and organizations to maintain greater autonomy over their information and its use in AI systems ([Winner, 1980](#)). Open models democratize innovation and access by removing barriers to entry and enabling developers worldwide to contribute to and build upon existing work without requiring permission from platform owners or paying licensing fees.

The acceleration of research, writing, and software development through widely accessible AI tools creates productivity gains across multiple domains, enabling smaller organizations and individual creators to accomplish tasks that previously required significant resources. Synthetic media capabilities support accessibility and creative expression for users with diverse needs and abilities, providing new forms of communication and artistic creation. Private inference capabilities preserve data sovereignty and privacy by enabling AI processing without exposing sensitive information to external parties, addressing fundamental concerns about surveillance and data misuse ([Bonawitz et al., 2017](#)).

However, these benefits come with corresponding risks that must be carefully managed. The absence of single data vendors ensuring accountability or content traceability can make it difficult to address harmful uses or assign responsibility for negative outcomes when decentralized systems are misused ([Jonas, 1984](#)). Lower barriers to abuse, including deepfake creation and disinformation campaigns, represent significant challenges for maintaining information integrity and social trust. The potential for AI tools to flood digital spaces with low-quality or misleading content poses risks to information ecosystems and public discourse more broadly ([Vosoughi et al., 2018](#)). Misaligned and malignant actors can exploit decentralization for surveillance, extremist mobilization, or even biomedical misuse through open-access model weights; this presents an ethical dilemma that is deeply tied to the lack of shared oversight. The accountability of high-flying corporate figures, liable for their actions and mismanagement, is now

replaced by thousands of faceless actors. The absence of platform-level chokepoints makes it difficult to track provenance, enforce moderation, or intervene in cases of misuse.

The continued erosion of trust in audio and video authenticity due to sophisticated synthetic media capabilities has implications for journalism, legal proceedings, and social communication. Additionally, the ability to conduct potentially harmful model training without oversight raises concerns about the development of AI systems that could be used for malicious purposes, including generating harmful content, conducting social engineering attacks, or developing capabilities that could be weaponized ([Chesney & Citron, 2019](#)).

Impact distribution across different populations reveals significant disparities in who benefits from and who bears the risks of unregulated AI generation. Marginalized communities face particular vulnerability to biased outputs, targeted misinformation campaigns, and synthetic identity attacks that can cause real harm to individuals and groups. Creators and intellectual property holders see their work scraped, replicated, or monetized without consent or compensation, undermining traditional models of creative economy and professional content creation.

Governance remains the critical “G” in ESG that is often overlooked. Yet without it, decentralization risks becoming an accelerant for harm, not a corrective. The illusion that decentralized systems are self-regulating is both a technical and political fallacy. Resilience and permissionless innovation must be matched with enforceable norms, trust-building tools, and protective standards.

Open Source as the Backbone of AI Decentralization

Open-source development serves as the fundamental infrastructure enabling AI decentralization, providing technical foundations, community governance models, and collaborative frameworks necessary for distributed AI systems to function effectively at scale. Foundational open-source communities – including Hugging Face, EleutherAI, LAION, Stability AI, Mistral, and BigScience – provide core models and tools that



enable independent AI development without reliance on proprietary platforms or corporate gatekeepers ([von Hippel, 2005](#)).

Projects such as llama.cpp and the ONNX Runtime are enabling a new class of fully local inference. These tools prove that open-source innovation can outpace closed ecosystems on accessibility, transparency, and performance efficiency, particularly for text generation and multimodal models ([Microsoft, 2023](#)). With Stable Diffusion now running on consumer laptops and TinyLlama operating with near-chatbot speeds on CPUs, the technical feasibility of decentralized AI has already arrived ([Mistral AI, 2023](#)).

Infrastructure layer contributors, including Filecoin, Aurora.dev, Gensyn, and Bittensor, supply computational and storage capabilities necessary for distributed AI systems to operate at scale while maintaining decentralized control and governance. Public sector and academic institutions prioritize open science principles and democratic access to AI capabilities, ensuring that research advances benefit broad communities rather than solely commercial interests ([Merton, 1973](#)). This institutional support provides legitimacy and resources for open-source AI development while establishing precedents for public-good technology development.

Grassroots developer ecosystems consisting of thousands of independent developers and small AI startups worldwide contribute to and build upon open-source foundations, creating diverse and resilient development communities that cannot be controlled by any single organization ([Raymond, 1999](#)). This distributed approach to innovation enables rapid experimentation and adaptation while maintaining collective ownership of core technologies, ensuring that fundamental AI capabilities remain accessible to all participants rather than controlled by commercial entities.

The strategic advantages of open-source development in AI include transparency – which allows for inspection, auditing, and verification of AI behavior – enabling trust and accountability mechanisms that are impossible with closed systems ([Lessig, 2001](#)). Reproducibility accelerates scientific progress by making research methods and datasets publicly available for verification and

extension by other researchers, creating cumulative knowledge development rather than duplicated proprietary efforts. Permissionless innovation allows developers to fork, modify, and extend tools without requiring approval from platform owners, removing gatekeeping mechanisms that can slow innovation and limit creativity.

Modular ecosystem development through tools such as LangChain, LlamaIndex, and open language models creates interoperable components that can be combined in novel ways, enabling rapid prototyping and system development without vendor lock-in ([Baldwin & Clark, 2000](#)). Open source removes platform control bottlenecks and enables truly distributed intelligence systems that no single entity can manipulate, providing fundamental infrastructure for democratic AI development that serves diverse community needs rather than narrow commercial interests.

Revenue Models and Competitive Advantages

Market participants in decentralized AI ecosystems employ diverse strategies to create sustainable business models while maintaining the openness and community control that define these systems. Decentralized AI startups building applications with open models and distributed infrastructure, such as Mistral, Gensyn, and Ocean Protocol, offer competitive alternatives to centralized services while maintaining transparency and user control that creates trust and reduces customer acquisition costs. These companies demonstrate that commercial success and open development can be aligned effectively when business models focus on value creation rather than platform control.

Data decentralized autonomous organizations (DAOs) and contributor communities monetize training datasets and participate in AI model governance through democratic decision-making processes that ensure fair compensation and community benefit. These organizations represent a fundamental shift from extractive data collection to collaborative value creation, where contributors maintain ownership and control over their



information while benefiting from its use in AI development. Specialized SaaS platforms use decentralized models to target niche verticals such as legal services, education, and healthcare with customized solutions that can be adapted to specific regulatory and professional requirements without platform restrictions.

Open-source maintainers earn revenue from fine-tuned models, plugins, commercial support, and wrapper services, creating sustainable careers in open AI development while maintaining community commitment to accessible technology. Developing markets create localized inference tools that operate independently of expensive cloud dependencies, enabling AI adoption in regions and sectors previously excluded from these capabilities due to cost or infrastructure limitations.

Emerging business models demonstrate the commercial viability of decentralized approaches across multiple revenue streams. Tokenized microtransactions enable pay-per-inference or storage costs tracked and settled on blockchain networks, creating granular pricing that better reflects actual usage while enabling automated compensation for providers. Consent-based royalties ensure data owners receive compensation when their contributions are used in training or inference, creating ongoing revenue streams that incentivize high-quality data contribution and maintain contributor engagement.

Vertical SaaS subscriptions provide specialized decentralized tools with recurring revenue models that can scale with customer success while maintaining competitive pricing compared to centralized alternatives. Freemium and open-core models offer basic functionality free with premium features or services requiring payment, enabling broad adoption while generating revenue from users who require advanced capabilities or commercial support. DAO and community governance fees allow users to participate in and pay for system upgrades, plugins, and computational resources while maintaining democratic control over development priorities and resource allocation.

Edge Computing vs. Centralized Performance

The architectural choice between edge computing and centralized systems in AI deployment presents fundamental trade-offs that affect performance, privacy, cost, and accessibility in complex ways that must be carefully evaluated for different use cases and stakeholder needs. Edge computing advocates – including IoT device manufacturers, privacy-focused startups, rural users with limited bandwidth, and companies like NVIDIA (Jetson) and Qualcomm – promote distributed processing solutions that bring computation closer to users and data sources while reducing dependence on network connectivity and centralized infrastructure ([Shi et al., 2016](#)).

The hardware shift enabling decentralized AI is already underway. Apple's Neural Engine, Qualcomm's Snapdragon X Elite, and AMD's Ryzen AI are offering 30 to 45 TOPS (Tera Operations Per Second) performance on-device, which is enough to run transformer models, image generators, and voice assistants locally. Microsoft's ONNX Runtime standardizes the deployment of these models across devices, ensuring that decentralized inference isn't just possible but broadly portable ([Microsoft, 2024](#)).

Centralization advocates, including hyperscale cloud providers such as Google, AWS, and Microsoft, along with AI laboratories OpenAI and Anthropic, among others, emphasize performance and scalability advantages that come from concentrating computational resources in optimized data centers with specialized hardware and efficient cooling systems ([Armbrust et al., 2010](#)). Each approach serves different stakeholder needs and use cases. For example, edge computing benefits end users who require privacy protection, offline functionality, or low-latency responses in applications such as healthcare monitoring, autonomous vehicles, robotics, and on-device AI assistants.

Centralized systems better serve enterprises demanding massive-scale training capabilities, real-time collaboration features, and centralized management of complex AI systems that require coordination across multiple users and



applications. The performance characteristics of these approaches differ significantly across multiple dimensions that affect user experience and system capabilities. Edge computing provides ultra-low latency through local processing, eliminating network delays that can be critical for real-time applications, while centralized systems experience higher latency due to network dependencies but can leverage connectivity for coordination and resource sharing across users and applications.

Privacy protection represents a significant advantage for edge computing, as data can remain on local devices without transmission to external servers, addressing concerns about surveillance, data breaches, and unauthorized access to sensitive information ([Bonawitz et al., 2017](#)). Centralized systems typically require data transmission and storage that creates privacy vulnerabilities and regulatory compliance challenges, particularly for applications involving personal, medical, or financial information.

Compute capacity differs dramatically between approaches, with edge computing limited by resource constraints on individual devices that may struggle with the most demanding AI tasks, while centralized systems can access massive graphic processing unit (GPU) and tensor processing unit (TPU) clusters with extensive scaling capabilities that enable training and running large models. Energy consumption patterns vary significantly between architectures, with edge computing potentially achieving lower overall system energy consumption by eliminating data transmission requirements and enabling more efficient local processing ([Strubell et al., 2019](#)).

However, decentralization may simply replace one form of dependency with another, from cloud monopolies to chip oligopolies. While the growing diversity of hardware providers introduces resilience, it does not eliminate lock-in risk entirely. What it does offer is lower latency, lower per-query cost, and better compliance with data sovereignty laws

These benefits, however, this must be balanced against potential inefficiencies. Distributed hardware environments can lead to

underutilization, and the environmental impact of manufacturing many smaller edge devices may outweigh that of maintaining fewer, more efficient centralized systems. Cost structures also differ substantially, with edge computing offering lower long-term operational costs for users who own their devices, while centralized systems typically operate on subscription-based or pay-per-use fee structures that can become expensive for high-volume usage but require lower upfront investment.

Balancing Performance with Responsible AI

The decentralized AI community must confront the reality that openness without stewardship often leads to abuse. While closed systems present ethical opacity, decentralized systems may enable unchecked experimentation or adversarial use. Tools such as Semantic Kernel are emerging to enable local, programmable ethical constraints and plug-in guardrails, embedding responsible AI principles into the toolkit of developers.

Still, decentralized AI governance remains underdeveloped compared to its centralized counterparts. It lacks the enforcement apparatus of major platforms, even as its reach grows. Building trust in decentralized models will depend on new forms of tooling, standardization, and community-led auditing to close the responsibility gap.

Yet embedding ethics at the infrastructure level is only one part of the equation. The intersection of performance optimization and responsible AI development presents one of the most complex challenges in contemporary AI systems, requiring careful navigation of competing objectives and stakeholder interests while maintaining both technical effectiveness and ethical standards. Model developers, including organizations such as OpenAI, Cohere, and Mistral, face the ongoing challenge of meeting both performance benchmarks and safety standards while remaining competitive in rapidly evolving markets where user expectations for capability and safety continue to increase ([Amodei et al., 2016](#)).



Deploying organizations – particularly startups and enterprises implementing language models in critical fields such as finance, healthcare, and legal services – must ensure reliability and compliance while maintaining the performance characteristics that make AI systems valuable for their use cases. This requires sophisticated understanding of both technical capabilities and regulatory requirements, as well as the ability to implement safety measures without compromising system effectiveness. Policymakers and regulators simultaneously develop accountability frameworks and safety standards that will shape the AI development landscape, creating new requirements that developers must integrate into their systems while maintaining innovation and competition.

Affected communities experience the real-world consequences of biased, incorrect, or unsafe AI outputs, making their perspectives crucial for understanding the true costs and benefits of different approaches to AI development ([Benjamin, 2019](#)). Their input is essential for identifying potential harms and developing mitigation strategies that address actual rather than theoretical risks. Standards organizations, including the Partnership on AI, OECD, IEEE, and UNESCO, provide frameworks for responsible AI development that attempt to balance innovation with safety and ethical considerations while creating industry-wide standards that enable interoperability and consistent expectations.

The fundamental tension between performance and responsibility manifests in multiple ways throughout AI system development and deployment. High-performance AI systems that prioritize speed, scale, and flexibility often sacrifice important qualities including fairness, explainability, data transparency, and comprehensive bias safeguards ([Barocas et al., 2017](#)). Conversely, responsible AI practices that ensure alignment with human values, legal compliance, and harm mitigation may reduce system performance and increase operational complexity, creating trade-offs that must be carefully managed.

Implementation strategies for balancing these concerns include fine-tuning with diverse datasets to improve representation and reduce bias across

demographic groups, ensuring that AI systems perform equitably for all users rather than optimizing for majority populations. Reinforcement Learning from Human Feedback (RLHF) aligns model behavior with human values and preferences, creating systems that are both capable and aligned with ethical standards ([Christiano et al., 2017](#)). Auditing and red-teaming practices help expose and mitigate risks before public release, while transparency protocols document model behavior, training sources, and known limitations for stakeholder review and ongoing monitoring.

Examples

Successful decentralized AI implementations provide concrete evidence for both the potential and practical challenges of alternative approaches to AI development and deployment. Hugging Face Model Hub represents a paradigmatic example of successful decentralized AI implementation, demonstrating how open-source model sharing can create thriving ecosystems where thousands of developers contribute improvements and specialized variants while maintaining quality and usability standards ([Wolf et al., 2020](#)). The platform's success illustrates how reducing barriers to participation and providing robust infrastructure can enable distributed innovation at scale while maintaining high standards for model quality and safety.

BigScience's BLOOM project demonstrates that collaborative, multi-institutional efforts can produce competitive large language models through coordinated open research, challenging assumptions that only well-funded commercial organizations can develop state-of-the-art AI systems ([Scao et al., 2022](#)). The project required sophisticated coordination mechanisms and shared governance structures that provide models for future collaborative efforts while maintaining scientific rigor and community accountability.

Ocean Protocol illustrates how blockchain-based data marketplaces can enable consent-driven data sharing and fair compensation for contributors, addressing fundamental concerns about data ownership and value distribution in AI systems



while maintaining data quality and utility for AI training ([Ocean Protocol Foundation, 2022](#)). The platform's implementation reveals both the potential and practical challenges of creating decentralized data economies that balance contributor rights with system functionality.

Open-source models can achieve commercial success while maintaining transparency and community engagement, demonstrating viable business models that do not rely on platform lock-in or proprietary advantages ([Mistral AI, 2023](#)). The company's approach shows how commercial and open-source objectives can be aligned effectively while creating sustainable competitive advantages through community building and ecosystem development.

However, implementation challenges and failures provide equally important insights for understanding the limitations and requirements of decentralized AI systems. Coordination difficulties have affected some decentralized projects, leading to fragmentation and reduced effectiveness compared to centralized alternatives that can make rapid decisions and implement consistent policies across their platforms ([Eghbal, 2020](#)). Performance gaps persist in certain distributed systems that cannot match the raw performance of well-resourced centralized systems, particularly for the most demanding AI tasks that require massive computational resources and specialized infrastructure.

Potential Benefits

Decentralized AI offers significant advantages that address fundamental limitations of centralized systems while creating new opportunities for innovation and equitable value distribution. Democratization and access represent perhaps the most significant potential benefits, as decentralized AI can provide broader access to advanced AI capabilities, particularly benefiting underserved communities, developing regions, and smaller organizations that cannot afford premium centralized services ([Birhane, 2021](#)). This increased access can level playing fields in education, healthcare, business development, and creative endeavors, enabling innovation and

economic development in previously excluded regions and sectors.

Innovation acceleration emerges from open-source development models that enable rapid experimentation and collaboration by removing barriers to entry and allowing developers to build upon existing work without restrictions or licensing fees. This permissionless innovation can lead to faster development cycles, more diverse applications, and creative solutions that might not emerge from centralized development processes focused on mass market applications. Privacy and data sovereignty provide users with greater control over their information and decision-making about how their data is used in AI training and inference, addressing growing concerns about surveillance capitalism and data exploitation.

Transparency and accountability through open models and auditable processes enable stakeholders to understand AI decision-making and identify potential biases or errors, creating trust and enabling continuous improvement through community oversight. This transparency is particularly important for applications in criminal justice, healthcare, education, and other high-stakes domains where AI decisions significantly impact people's lives. Economic opportunities emerge from new business models that distribute value more equitably among data contributors, developers, and users rather than concentrating profits in a few large corporations, creating sustainable income streams for a broader range of participants in the AI ecosystem.

Resilience and robustness result from distributed systems that are less vulnerable to single points of failure and can continue operating even if some nodes experience problems, creating more reliable AI services for critical applications. This distributed architecture also provides resistance to censorship and political control, enabling AI development and deployment that serves diverse community needs rather than narrow commercial or political interests.



Potential Risks & Mitigations

Decentralized AI systems face several categories of risks that require proactive mitigation strategies to ensure successful implementation and community benefit. Governance and coordination challenges represent significant risks, as decentralized systems may suffer from decision-making paralysis, conflicting objectives among stakeholders, and difficulty implementing consistent policies across distributed networks ([Eghbal, 2020](#)). Mitigation strategies include developing clear governance frameworks with defined decision-making processes, establishing dispute resolution mechanisms that can address conflicts efficiently, and creating incentive structures that align participant interests with collective goals through economic and social rewards.

Performance and reliability concerns pose risks that distributed systems might not match the performance, consistency, or reliability of well-managed centralized alternatives, particularly for mission-critical applications that require guaranteed uptime and response times. Mitigation approaches include investing in infrastructure optimization to improve distributed system performance, developing performance benchmarking standards that enable comparison and improvement across different implementations, and creating hybrid architectures that combine the benefits of both centralized and decentralized approaches for different use cases and requirements.

Security and safety vulnerabilities present risks with decentralized systems that may be more difficult to secure, update, and monitor for harmful usage, potentially enabling malicious actors to exploit AI capabilities for harmful purposes ([Jonas, 1984](#)). Mitigation strategies include implementing robust security protocols across all system components, creating distributed monitoring systems that can detect and respond to threats without central control, and developing rapid response mechanisms for addressing harmful usage while maintaining system openness and community control.

Quality control and standards represent risks that without centralized oversight, the quality and safety of AI models and applications may vary significantly, leading to unreliable or harmful outputs that damage user trust and community reputation. Mitigation approaches include establishing community-driven quality standards with clear criteria and enforcement mechanisms, creating reputation systems for contributors that incentivize high-quality work, and developing automated testing and validation tools that can assess model performance and safety without requiring centralized review.

Economic sustainability poses the risk that decentralized systems may struggle to generate sufficient revenue to fund initial launch, ongoing development, maintenance, and improvement, leading to degraded performance or system abandonment over time. Mitigation strategies include exploring diverse monetization approaches that can generate sustainable revenue streams, creating funding mechanisms through DAOs and cooperatives that enable community investment in system development, and developing partnerships with traditional organizations that can provide resources and market access while maintaining decentralized governance principles.

Next Steps

Successfully realizing the potential of decentralized AI requires coordinated action across multiple stakeholder groups, each contributing their unique capabilities and perspectives to build systems that serve broad community interests while maintaining technical excellence and ethical standards. For policymakers, the priority should be developing regulatory frameworks that support innovation while ensuring safety and accountability in decentralized AI systems, avoiding approaches that inadvertently favor centralized platforms or stifle beneficial innovation ([Calo, 2017](#)). This includes creating incentives for responsible AI development and deployment across both centralized and decentralized architectures, investing in public infrastructure and research that supports democratic access to AI capabilities,



and facilitating international cooperation on AI governance standards and best practices.

Policymakers should also focus on protecting data rights and ensuring fair compensation for data contributors while promoting transparency and accountability in AI systems regardless of their architectural approach. This may require new legal frameworks that recognize data ownership rights, establish mechanisms for consent-based data usage, and create enforcement mechanisms for holding AI developers accountable for system impacts on communities and individuals.

Technologists should prioritize developing tools and frameworks that make responsible AI practices easier to implement in decentralized systems, recognizing that technical solutions can often address governance challenges more efficiently than regulatory approaches ([Winner, 1980](#)). Creating interoperability standards that enable different decentralized AI components to work together effectively will be crucial for ecosystem development, while investment in research on hybrid architectures that combine the benefits of centralized and decentralized approaches may offer optimal solutions for many use cases.

Technical development should also focus on improving the performance and reliability of decentralized systems to ensure they can meet user expectations while maintaining transparency and community control that define these approaches. This includes developing better methods for measuring and comparing the performance, safety, and impact of different AI systems, creating tools for distributed governance and community coordination, and building security and safety mechanisms that protect users without compromising system openness.

Organizations should evaluate the potential benefits and risks of decentralized AI for their specific contexts and use cases, developing capabilities in open-source AI tools and decentralized infrastructure to reduce dependence on centralized providers while maintaining operational effectiveness ([Chesbrough, 2003](#)). Participation in community governance and standard-setting processes will help shape the development of decentralized AI ecosystems while

ensuring that organizational needs are represented in community decision-making. Organizations should also implement responsible AI practices regardless of underlying system architecture, ensuring ethical consistency and stakeholder trust across all AI implementations.

Communities and civil society groups should advocate for AI systems that serve community needs and values rather than just commercial interests, participate in the governance and oversight of AI systems that affect their members, demand transparency and accountability from both centralized and decentralized AI providers, and support education and capacity-building initiatives that enable broader participation in AI development and governance ([Winner, 1986](#)). Community engagement is essential for ensuring that decentralized AI systems truly serve diverse needs rather than simply replicating existing power structures in new technological forms.

The path forward requires recognizing that the future of AI will likely be characterized by hybrid ecosystems where different approaches serve different needs and contexts rather than complete dominance by either centralized or decentralized paradigms. Success will depend on ensuring that technological evolution serves broad human interests while maintaining the performance and safety standards that users and society require, viewing responsible AI development not as a constraint on innovation but as a prerequisite for building systems that can earn and maintain the trust necessary for beneficial long-term impact.

This chapter was developed collaboratively by the listed authors and reflects original analysis supported by properly cited academic and industry sources. AI tools, including OpenAI, were used to assist with editing and citation integration, with full transparency acknowledged in the document.

References

Aissaoui, N. (2021). The digital divide: A literature review and some directions for future research in light of COVID-19. Global Knowledge, Memory, and Communication. Advance online publication. <https://doi.org/10.1108/GKMC-06-2020-0075>



- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
<https://arxiv.org/abs/1606.06565>
- Armbrust, M., et al. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–5
<https://dl.acm.org/doi/10.1145/1721654.1721672>
- Arrieta-Ibarra, I., et al. (2018). Should we treat data as labor? *American Economic Association Papers & Proceedings*, 108, 38–42. Amazon:
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3093683
- Baldwin, C. Y., & Clark, K. B. (2000). *Design Rules: The Power of Modularity*. MIT Press.
<https://direct.mit.edu/books/monograph/1856/Design-Rules-Volume-1The-Power-of-Modularity>
- Barocas, S., Hardt, M., & Narayanan, A. (2017). *Fairness and Machine Learning*. fairmlbook.org.
https://fairmlbook.org/Algorithmic_injustice:_a_relational_ethics_approach_-_PubMed
- Barocas, S., Hood, S., & Ziewitz, M. (2019). Governing algorithms. *Science, Technology, & Human Values*, 44(1), 3–27.
<https://journals.sagepub.com/doi/abs/10.1177/0162243915608948>
- Bendi Software Ltd. (n.d.). *Bendi: AI-powered ESG risk management software (Prism)*. Retrieved September 10, 2025, from <https://www.bendi.ai>
- Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity.
<https://www.politybooks.com/bookdetail>
- Benet, J. (2014). IPFS - Content Addressed, Versioned, P2P File System. *arXiv preprint arXiv:1407.3561*.
<https://arxiv.org/abs/1407.3561>
- Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, 2(2), 100205.
<https://www.sciencedirect.com/science/article/pii/S2666389921000155>
- Bommasani, R., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
<https://arxiv.org/pdf/2108.07258>
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.
<https://doi.org/10.1145/3133956.3133982>
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1).
<https://journals.sagepub.com/doi/10.1177/2053951715622512>
- Calo, R. (2017). Artificial intelligence policy: A primer and roadmap. *University of California, Davis Law Review*, 51(2), 399–435.
<https://lawreview.law.ucdavis.edu/archives/51/2/artificial-intelligence-policy-primer-and-roadmap>
- Catalini, C., & Gans, J. S. (2020). Some simple economics of the blockchain. *Communications of the ACM*, 63(7), 80–90.
<https://doi.org/10.1145/3359552>
- Chesbrough, H. (2003). *Open Innovation*. Harvard Business School Press.
<https://store.hbr.org/product/open-innovation-the-new-imperative-for-creating-and-profiting-from-technology/8377>
- Chesney, R., & Citron, D. K. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6), 1753–1819.
https://proceedings.neurips.cc/paper_files/paper/2017/hash/d5e2c0adad503c91f91df240d0cd4e49-Abstract.html
- Christiano, P., et al. (2017). Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems*, 30.
https://proceedings.neurips.cc/paper_files/paper/2017/hash/d5e2c0adad503c91f91df240d0cd4e49-Abstract.html



Cointelegraph. (2025). AI agent tokens surge 322% in Q4 2024. *Cointelegraph Market Watch*. <https://cointelegraph.com/news/ai-agents-market-cap-surges-solana-leads-2024>

DataCamp. (2024). Federated learning and privacy-preserving AI. *DataCamp Whitepapers*. <https://www.datacamp.com/blog/federated-learning>

Eghbal, N. (2020). *Working in Public: The Making and Maintenance of Open-Source Software*. Stripe Press. <https://press.stripe.com/working-in-public>

Electronics MDPI. (2025). Blockchain-supported AI for legal automation. *Electronics*, 14(6), 1110. <https://www.mdpi.com/2078-2489/14/9/477>

Everstake. (2024). State of Web3 and blockchain adoption. *Everstake Insights Report*. <https://everstake.one/blog/blockchain-beyond-2024-trends-insights-and-predictions-for-2025>

Frontiers in Sustainable Cities. (2025). Blockchain-powered e-governance: Opportunities and tradeoffs. *Frontiers*, 3(9). <https://www.frontiersin.org/journals/sustainable-cities/articles/10.3389/frsc.2025.1623412/abstract>

Hakkarainen, J. M. (2021). *Naming something collective does not make it so: Algorithmic discrimination and access to justice*. *Internet Policy Review*, 10(4). <https://doi.org/10.14763/2021.4.1600>

Interexy. (2025). Decentralized AI market and growth projections. *Industry AI Report Q1 2025*. <https://interexy.com/>

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://www.nature.com/articles/s42256-019-0088-2>

Jonas, H. (1984). *The Imperative of Responsibility: In Search of an Ethics for the Technological Age*. University of Chicago

Press. <https://press.uchicago.edu/ucp/books/book/chicago/I/bo5953283.html>

Keršič, V., et al. (2025). *A review on building blocks of decentralized artificial intelligence*. *Journal of Information Security and Applications*, 77, 103779 <https://arxiv.org/pdf/2402.02885>

Kinstak. (n.d.). *Kinstak: AI-powered digital legacy vault*. Retrieved September 10, 2025, from <https://www.kinstak.com>

Lanier, J. (2013). *Who owns the future?* Simon & Schuster. Retrieved from <https://www.simonandschuster.com/books/Who-Owns-the-Future/Jaron-Lanier/9781451654974>

Lessig, L. (2001). *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House. https://www.researchgate.net/publication/43179722_The_Future_of_Ideas_The_Fate_of_the_Commons_in_a_Connected_World

Merton, R. K. (1973). *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press. https://archive.org/details/sociologyofscien0000mert_j7p1

Mistral AI. (2023). Open-weight innovation and scaling strategy. *Mistral Research Briefs*. <https://mistral.ai/>

Microsoft. (2023, November 17). *Empowering the AI generation: Microsoft's open-source initiative* [Blog post]. Microsoft Tech Community. <https://techcommunity.microsoft.com/blog/educatordeveloperblog/empowering-the-ai-generation-microsofts-open-source-initiative/3962888>

Microsoft. (2024, November 19). *Cross-platform Edge AI made easy with ONNX Runtime* [Blog post]. Microsoft Tech Community. <https://techcommunity.microsoft.com/blog/azure-ai-factory-blog/cross-platform-edge-ai-made-easy-with-onnx-runtime/4303521>

Modular Inc. (2025, June 20). *How is Modular democratizing AI compute ...* [Blog post]. Modular. <https://www.modular.com/blog/how-is-modular-democratizing-ai-compute>



- Marotta, A., & Madnick, S. (2021). Convergence and divergence of regulatory compliance and cybersecurity. *Issues in Information Systems*, 22(1), 10–50.
https://doi.org/10.48009/1_iis_2021_10-50
- Nardini, M., Helmer, S., El Ioini, N., & Pahl, C. (2020). A blockchain-based decentralized electronic marketplace for computing resources. *SN Computer Science*, 1, 251.
<https://doi.org/10.1007/s42979-020-00243-7>
- Ocean Protocol Foundation. (2022). Consent-based AI data marketplaces. *Ocean Protocol Whitepaper v4.0*.
<https://oceanprotocol.com/tech-whitepaper.pdf>
- OpenPR. (2025). AI governance market outlook 2024–2034. *OpenPR Research Series*.
<https://www.openpr.com/news/4059215/ai-governance-market-size-share-and-growth-report-2034>
- Osborne, C., Ding, J., & Kirk, H. R. (2024). *The AI community building the future? A quantitative analysis of development activity on the Hugging Face Hub*. *Journal of Computational Social Science*. <https://doi.org/10.1007/s42001-024-00300-8> [SpringerLink](#)
- Parker, G., Van Alstyne, M., & Choudary, S. P. (2016). *Platform Revolution*. W.W. Norton & Company.
https://www.google.com/books/edition/Platform_Revolution_How_Networked_Market/Bvd1CQAAQBAJ?hl=en&kptab=overview
- Pentland, A., Hardjono, T., & Lipton, A. (2019). *Digital asset marketplaces: Building digital asset marketplaces for the new economy*. MIT Connection Science.
<https://wip.mitpress.mit.edu/pub/aps1hrbe/download/pdf>
- Raymond, E. S. (1999). *The Cathedral and the Bazaar*. O'Reilly Media.
<http://www.catb.org/esr/writings/homesteading/cathedral-bazaar/>
- Scao, T. L., et al. (2022). BLOOM: A 176B-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
<https://arxiv.org/pdf/2211.05100>
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). *Edge computing: Vision and challenges*. *IEEE Internet of Things Journal*, 3(5), 637–646.
<https://doi.org/10.1109/JIOT.2016.2579198>
- Scientific Reports. (2025). Collaboration models in decentralized AI research. *Nature Scientific Reports*, 15(1). <https://www.mdpi.com/3482630>
- Sokolín, L. (2024). *Research: Decentralized AI Overview (2024)*. The Fintech Blueprint.
<https://lex.substack.com/p/research-decentralized-ai-overview>
- Splunk. (2025). U.S. 2025 Executive Order on AI innovation. *Splunk Policy Insights*.
https://www.splunk.com/en_us/blog/learn/ai-governance.html
- SpringerOpen. (2025). Federated learning and blockchain for secure AI. *Journal of Cloud Computing*, 14(3).
<https://www.springeropen.com/collections/abfltc>
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning. *arXiv preprint arXiv:1906.02243*.
<https://arxiv.org/abs/1906.02243>
- Technological Convergence Review. (2025). Smart contract monetization in AI ecosystems. *TechCon Review*, 11(2).
<https://www.forbes.com/sites/digital-assets/2024/11/12/watch-decentralized-ai-in-2025-the-convergence-of-ai-and-crypto/>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
<https://www.science.org/doi/10.1126/science.aaq9559>
- von Hippel, E. (2005). *Democratizing Innovation*. MIT Press.
<https://direct.mit.edu/books/book/2821/Democratizing-Innovation>



Winner, L. (1980). Do artifacts have politics? *Daedalus*, 109(1), 121–136.
<https://www.jstor.org/stable/20024652?seq=1>

Winner, L. (1986). *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. University of Chicago Press.
https://sciencepolicy.colorado.edu/students/envs_5110/whale_reactor.pdf

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Rush, A. M. (2020).

Transformers: State-of-the-art natural language processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38-45. Association for Computational Linguistics.
<https://doi.org/10.18653/v1/2020.emnlp-demos.6>

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.
<https://www.hachettebookgroup.com/titles/shoshana-zuboff/the-age-of-surveillance-capitalism/9781541758001/>

Author (In order of contribution)

Olivier Bacs, CTO and co-founder, Bendi

Olivier Bacs is the CTO and co-founder of Bendi, where he builds AI-powered tools that help companies gain visibility into their supply chains and collaborate more effectively with suppliers. His work combines geospatial analysis, automation, and natural language processing to uncover hidden risks while making complex compliance processes easier to navigate. Olivier is especially focused on decentralized and ethical approaches to AI, ensuring that technology enhances trust, equity, and resilience across global value chains.

Carolyn Eagen, MBA, Founder, Kinstak

Carolyn Eagen is the Founder and CEO of Kin Technologies and Kinstak, an AI-native platform pioneering private digital legacy management and decentralized digital asset manager for families and SMBs. She brings over 20 years of leadership in product strategy and innovation. Carolyn is passionate about building ethical, user-centered systems that unlock access, equity, and long-term resilience in the age of AI.





For more information about the Coalition for Innovation, including how you can get involved, please visit coalitionforinnovation.com.

[View the Next Chapter](#)

