

Blockchain-Enabled Trust Formation in AI-Driven Inter-Firm Capacity Sharing Ecosystems

Arthur Bernatt

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL,
USA.

arthur.work@uab.edu

Yachan Peng

Department of Computer Science, George Mason University, Fairfax, VA, USA.

yichen1982@gmu.edu

Abstract

Inter-firm capacity sharing has emerged as a strategic instrument for coping with demand volatility, asset underutilization, and supply chain disruptions, yet its widespread adoption remains constrained by deficits in trust, information asymmetry, and coordination complexity. This paper investigates the architectural convergence of artificial intelligence and blockchain technologies as a mechanism to engineer trust in decentralized capacity sharing ecosystems. By analyzing system-level interactions, we demonstrate that AI-driven predictive and prescriptive models can significantly enhance the efficiency of capacity allocation, while blockchain-based distributed ledger technology provides an immutable, transparent substrate for recording commitments, executing smart contracts, and maintaining verifiable reputation trails without central intermediaries. The discussion is framed around structural trade-offs that arise when integrating these two computational paradigms: the tension between algorithmic opacity and ledger transparency, the scalability limits of consensus protocols in high-frequency capacity markets, and the governance challenges of embedding algorithmic decision rights into autonomous smart contracts. We further interrogate fairness considerations, including how access to AI-derived optimization and blockchain-anchored identity can asymmetrically benefit incumbent firms, and propose a layered architecture that decouples trust computation from capacity execution. Policy implications related to cross-border capacity pooling, data sovereignty, and the regulatory treatment of algorithmically managed capacity consortia are examined. The paper concludes that a carefully orchestrated fusion of AI and blockchain can transform inter-firm capacity sharing from a relationship-intensive, fragile practice into a resilient, semi-automated industrial infrastructure capable of underpinning sustainable supply networks.

Keywords

capacity sharing, blockchain, artificial intelligence, trust, inter-firm collaboration, decentralized governance, smart contracts.

1. Introduction

Contemporary manufacturing and logistics networks are increasingly characterized by fluid demand patterns, shortened product lifecycles, and frequent supply shocks. In this environment, the rigid ownership of production capacity, warehousing space, or transportation assets becomes economically untenable. Inter-firm capacity sharing, wherein companies trade temporarily surplus resources with competitors or complementary partners,

has been advanced as a powerful resource pooling strategy that can simultaneously reduce idle costs, improve service levels, and curtail the environmental burden of overbuilding. Foundational work in operations management has formalized the incentive structures and competitive effects of such arrangements, demonstrating that under certain contractual forms capacity sharing can yield Pareto improvements across the participating firms [1]. Yet actual industry uptake remains modest. Beyond the well-recognized difficulties of scheduling and pricing, the principal barrier is the fragility of inter-organizational trust: firms must trust that partners will honor commitments, disclose capacity constraints truthfully, and refrain from opportunistic exploitation of shared operational visibility.

Parallel advances in two computational domains have recently generated the raw materials for a systemic solution to these trust deficits. First, artificial intelligence (AI) techniques, ranging from deep demand forecasting to reinforcement learning-based resource allocation, have matured to the point where highly granular and adaptive capacity coordination is technically feasible [3]. Second, blockchain technologies, introduced through decentralized cryptocurrency systems [2] and later extended to programmable smart contracts [4], offer an institutional infrastructure that can render commitments, state transitions, and reputation scores tamper-proof and algorithmically enforceable. The convergence of these two technologies is not trivial, and their naive juxtaposition frequently leads to contradictory design pressures. AI models, especially those reliant on opaque deep architectures, undermine the auditability that blockchain is meant to guarantee. Conversely, the sluggish finality and limited throughput of many consensus protocols become bottlenecks for the kind of real-time, high-frequency coordination that AI-based dynamic sharing requires [7]. Resolving these tensions calls for a system-level perspective that situates trust formation at the intersection of algorithmic coordination and decentralized governance.

This paper provides a comprehensive, structurally oriented analysis of blockchain-enabled trust formation in AI-driven inter-firm capacity sharing ecosystems. We examine the socio-technical architecture, the economic and algorithmic trade-offs, and the longer-term implications for industrial policy and sustainability. Our analysis does not propose a single monolithic platform design; rather, it identifies the critical dimensions that any successful integration must navigate and evaluates how different design choices distribute risks, fairness, and resilience across the ecosystem. Throughout the paper, we weave together insights from supply chain management, cryptoeconomics, organizational trust theory, and algorithmic governance to construct a multi-layered understanding of this emerging paradigm.

2. Inter-Firm Capacity Sharing Ecosystems: Structure and Challenges

Capacity sharing emerges from a fundamental tension in industrial organization: assets are indivisible and expensive, while demand is stochastic and fragmented. Traditional responses such as safety stock, redundancies, or vertical integration absorb significant capital and generate waste. Horizontal capacity sharing, by allowing firms to pool complementary demand streams, can reduce the aggregate need for excess capacity, much as financial intermediaries pool uncorrelated risks. However, operationalizing this pooling is challenging. Capacity is not a perfectly fungible commodity; it is entangled with proprietary processes, quality specifications, geographical constraints, and timing relationships. Furthermore, the firms that most need capacity are often direct competitors, making the sharing arrangement a delicate balance between cooperative efficiency and competitive strategy [1].

Three interconnected challenges dominate the landscape. The first is verifiability of capacity states and commitments. A firm offering spare production slots must credibly signal that

those slots are genuinely surplus and not a veil for extracting production intelligence, while a firm seeking capacity must convincingly demonstrate creditworthiness and the authenticity of its demand forecasts. The second challenge is dynamic coordination under uncertainty. Capacity sharing is not a one-shot transaction but a repeated game overlain on rapidly changing conditions; renegotiation and rebalancing are constant, requiring sophisticated matching engines that traditional enterprise resource planning systems cannot supply. The third challenge is relational stickiness: in the absence of robust institutional trust, firms revert to long-standing bilateral partnerships even when those are economically suboptimal, locking the ecosystem into high search costs and low asset fluidity.

Addressing these challenges through conventional centralized platforms raises its own concerns. A third-party capacity exchange would amass sensitive operational data, become a single point of failure, and potentially exercise market power through information asymmetry. A genuinely decentralized alternative must, therefore, embed mechanisms for verifying states, enforcing agreements, and building reputation without relying on a trusted human coordinator. It is precisely this requirement that invites a joint consideration of AI for coordination intelligence and blockchain for institutional automation.

3. The Role of Artificial Intelligence in Capacity Coordination

AI contributes to capacity sharing ecosystems along several dimensions that transform the quality and speed of coordination. On the forecasting front, deep learning models trained on high-frequency demand signals, social media sentiment, and macroeconomic indicators can generate probabilistic capacity requirements far more granular than classical time-series methods [3]. When forecasting outputs are shared via standardized data models, potential capacity providers can price their spare assets with a level of confidence approaching that of internal planning. Beyond prediction, AI enables prescriptive orchestration. Multi-agent reinforcement learning systems can learn near-optimal capacity matching policies by simulating repeated interactions among autonomous firm agents, discovering equilibrium strategies that balance self-interest with ecosystem stability.

AI also reduces the cognitive load of trust assessment. Traditional trust formation relies on cumulative relational experience and costly due diligence. Machine learning classifiers can instead aggregate heterogeneous behavioral signals, such as historical fulfillment rates, response latencies, and deviation patterns, into dynamic trustworthiness scores. These scores can be updated continuously and made available to potential partners as lightweight, anonymized indicators, lowering the entry barrier for newer, less established firms. Importantly, AI-driven capacity coordination does not assume perfect rationality or complete information; it thrives precisely on patterns that are too subtle for human analysts and on adaptivity that manual contracts cannot offer.

Nevertheless, the reliance on AI introduces profound accountability problems. When an AI agent recommends a capacity commitment that later proves suboptimal, attributing responsibility becomes ambiguous. The opacity of many high-performing models, particularly deep neural networks, conflicts with the expectation that business agreements should be explicable and contestable. This tension is not merely philosophical; it has direct implications for legal enforceability and organizational acceptance. An ecosystem that blindly follows AI prescriptions without auditable reasoning risks alienating participants and attracting regulatory scrutiny. Thus, the introduction of AI must be accompanied by an institutional layer that preserves transparency and ownership of algorithmic decisions, a role that blockchain infrastructure is uniquely positioned to fulfill.

4. Blockchain as a Trust Infrastructure for Decentralized Capacity Markets

Blockchain technology offers a distinctive institutional fabric: a collectively maintained, append-only ledger that achieves consensus on state transitions without a central authority [2]. By encoding capacity sharing agreements as self-executing smart contracts [4], the technology can automate the lifecycle of a capacity lease, from reservation and escrowed payment to post-completion settlement. Because the contract code and its execution trace are immutably recorded, the parties gain a tamper-proof audit trail that dramatically reduces the scope for post-hoc disputes. In environments where legal enforcement is slow, costly, or cross-jurisdictionally fragmented, this on-chain finality acts as a functional substitute for court-backed contract law.

Beyond automating bilateral contracts, blockchain enables the emergence of ecosystem-wide trust through decentralized reputation systems. Every completed transaction, on-time delivery, capacity underutilization report, or default event can be hashed and anchored on-chain, forming a cryptographically secure reputation record that moves with the identity of the firm. Because such records are not owned by any single platform, firms cannot unilaterally delete or manipulate poor performance history, solving a long-standing problem of centralized reputation systems in which platform operators may selectively censor or inflate ratings. The immutability, however, is not without cost. Erroneously recorded data, whether due to sensor malfunction or malicious injection, becomes difficult to rectify, highlighting the need for carefully designed arbitration and rollback mechanisms that balance finality with fairness.

Blockchain also restructures the power dynamics of capacity markets. In a centralized exchange, the platform operator can unilaterally change fee structures, delist participants, or monetize trade data. In a permissioned or consortium blockchain, governance is distributed among the participating firms, typically through voting protocols encoded in the chain itself [5]. This shifts the locus of control from a single intermediary to a collective, aligning governance authority with the very actors whose capacity constitutes the market. However, this shift raises complex questions of voting power distribution, collusion resistance, and the representation of smaller firms, topics that are deepened when AI agents acquire the ability to vote or propose parameter changes on behalf of their principals.

5. Architectural Integration: AI Oracles and Smart Contract Governance

The integration of AI and blockchain revolves around the oracle problem: how can a deterministic, self-contained smart contract reliably ingest off-chain data, and especially the outputs of complex, non-deterministic AI models, without compromising the security guarantees of the ledger? Blockchain platforms cannot natively execute computationally intensive deep learning inference within their limited virtual machines while maintaining consensus efficiency. Instead, AI components must operate off-chain and supply their results through blockchain oracles. A decentralized oracle network aggregates data from multiple independent node operators, reducing the risk of a single manipulated feed [11]. For capacity sharing, an oracle might relay AI-generated demand forecasts, dynamic pricing signals, or partner trustworthiness scores, which then trigger smart contract clauses such as automatic capacity reservation or collateral adjustment.

The architectural layering required here is non-trivial. The AI layer operates in a probabilistic, continuously learning regime, while the blockchain layer is engineered for deterministic, append-only logic. Bridging these two paradigms demands careful attention to data provenance, model versioning, and dispute resolution. One promising approach is to record

cryptographic hashes of AI model weights and training datasets on-chain, allowing third parties to later verify that the output produced at a specific block height was indeed generated by the agreed-upon model. This design introduces an audit trail for algorithmic decisions without requiring nodes to replicate the computation, thereby preserving throughput. The governance of such a system then involves deciding which AI models are authorized, how frequently they can be updated, and under what conditions their outputs can be overridden by human arbitration committees. These are inherently socio-technical questions that cannot be resolved through engineering alone.

Smart contract governance becomes even more intricate when AI agents are granted limited autonomy to initiate or accept capacity trades on behalf of their firms. The principle of algorithmic authority, in which an algorithm's output is treated as binding, must be tempered by escape hatches that allow human intervention during exceptional circumstances. A well-designed ecosystem therefore encodes a hierarchy of decision rights, distinguishing between routine, low-value trades fully automated by AI and smart contracts, and non-routine, high-value commitments that require multi-signature authorization and off-chain deliberation. This graduated approach reflects an understanding that trust in inter-firm sharing is not a binary variable but a continuum that can be calibrated through institutional design.

6. Tradeoffs in Decentralization, Efficiency, and Scalability

A perennial theme in blockchain systems design is the tension between decentralization, security, and scalability, often termed the scalability trilemma [7]. In the context of capacity sharing, this trilemma acquires a specific operational flavor. High decentralization, implemented through a public permissionless blockchain, maximizes censorship resistance and inclusivity but limits transaction throughput to levels insufficient for an ecosystem that may need to match hundreds of capacity requests per minute during demand peaks. Permissioned or consortium blockchains, in which a known set of industry participants operate validating nodes, achieve far higher throughput and lower latency but reintroduce a degree of centralization and gatekeeping. The choice of consensus mechanism, whether proof-of-work, practical Byzantine fault tolerance, or a delegated proof-of-stake variant, therefore shapes not only technical performance but also the social fabric of the trust community.

Scalability challenges are compounded by the data volume that AI-augmented capacity sharing generates. Detailed sensor readings from production lines, GPS traces from logistics assets, and fine-grained demand forecasts can overwhelm on-chain storage. A modular architecture that relies on layer-two solutions, such as state channels or rollups, becomes essential. In such a design, the bulk of data exchange and AI negotiation occurs off-chain among directly interacting parties, with only dispute resolutions, final settlements, and compressed cryptographic proofs anchored to the main chain. This approach preserves the auditability and immutability benefits while keeping transaction fees manageable and commit times short enough for near-real-time capacity allocation.

The tradeoff extends into the realm of privacy. Inter-firm capacity sharing requires firms to reveal some operational data to potential partners, yet full transparency can erode competitive advantage. Zero-knowledge proofs and trusted execution environments have been proposed as privacy-preserving techniques that allow a smart contract to verify, for instance, that a firm possesses a certain quantum of idle capacity without revealing its entire production schedule. These cryptographic methods, however, impose computational overhead and reduce the richness of data available for AI-driven learning, leading to a delicate balance between

information concealment and coordination efficiency. System architects must navigate these tradeoffs not in isolation but as a coherent optimization problem that spans the entire stack, from cryptographic primitives to organizational norms.

7. Fairness, Incentive Design, and Algorithmic Governance

Trust is not built on transparency alone; it requires perceptions of procedural and distributive fairness. In an AI-driven capacity sharing ecosystem, fairness concerns arise at multiple levels. The algorithms that prioritize capacity requests may favor firms that have historically generated higher on-chain volumes, creating a rich-get-richer dynamic that marginalizes smaller or newer participants. Reputation scoring models, if trained on biased historical data, can inadvertently encode systemic inequalities, locking certain firms into unfavorable terms irrespective of their current reliability. The literature on interorganizational trust emphasizes that trustworthiness is multidimensional, encompassing competence, integrity, and benevolence [10], yet automated scores often collapse these dimensions into a single scalar metric, losing vital nuance.

Blockchain-based token economics offer powerful levers for incentive alignment but also risk introducing financialization dynamics that undermine operational trust. If capacity sharing tokens become tradeable assets, speculative behavior can decouple token value from underlying service quality, incentivizing short-term rent-seeking over long-term reliability. Careful mechanism design is therefore required to ensure that the on-chain incentive structure rewards behaviors that genuinely contribute to ecosystem resilience, such as capacity sharing during supply disruptions, rather than merely rewarding transaction volume. Drawing on principles from cryptoeconomics and mechanism design [6], it is possible to construct bonding curves, staking requirements, and slashing conditions that align individual profit maximization with collective service level targets.

Governance of these incentive structures must itself be decentralized to maintain legitimacy. When a small consortium of founding firms controls the parameter updates of the AI matching engine and the blockchain protocol, the system risks degenerating into a cartel. Incorporating algorithmic governance mechanisms that allow all stakeholders to vote on changes, perhaps weighted by a combination of stake and past contribution, can enhance legitimacy. Yet these voting systems must be shielded from manipulation by Sybil attacks or whale dominance. The field of algorithmic governance has begun to explore liquid democracy and quadratic voting in blockchain contexts, but their applicability to industrial consortia with asymmetric stakes remains an open research area [9]. Balancing efficiency and inclusiveness in governance design is as much a political as a technical challenge.

Recent empirical work on trust and reciprocity in firms' capacity sharing [13] has demonstrated that informal, relationship-based trust plays a significant role even in settings where formal contracts exist, suggesting that automated systems must not eliminate space for discretionary reciprocity. An over-engineered, fully deterministic smart contract regime may crowd out the very relational behaviors that sustain cooperation during unforeseen crises. The ideal socio-technical configuration thus embeds hard contractual automation for routine transactions while preserving flexibility for exceptions, creating layered trust structures that combine algorithmic reliability with human judgment.

8. Policy, Regulation, and Cross-Industry Deployment

The deployment of blockchain-AI capacity sharing ecosystems does not occur in a regulatory vacuum. Cross-border capacity pooling, particularly in industries such as pharmaceuticals,

aerospace, and energy, triggers a complex web of trade compliance, data localization, and antitrust regulations. The immutable and often pseudonymous nature of blockchain records can conflict with data protection frameworks that grant individuals and firms the right to rectification or erasure. The European Union's General Data Protection Regulation, for instance, poses challenges for permanent on-chain storage of transaction meta-data that may indirectly identify a firm. Architectural responses, such as storing only hashes and encrypted references on-chain while keeping personal data off-chain under strict access controls, can partially resolve these tensions, but they require continuous legal-technical co-design [15].

Antitrust authorities are likely to scrutinize capacity sharing consortia that use shared AI pricing algorithms and blockchain-governed rules, especially if the participating firms collectively hold substantial market power. The ability of smart contracts to enforce uniform contract terms across all participants may be interpreted as a form of algorithmic collusion, even in the absence of explicit human agreement. System designers must proactively embed regulatory compliance modules that can demonstrate to auditors that pricing and allocation algorithms do not exploit structural market power, for example by implementing verifiable differential pricing constraints or ensuring that capacity is allocated through transparent, non-discriminatory auctions whose code is open to inspection.

The potential for cross-industry deployment amplifies these regulatory complexities. A blockchain designed for automotive parts capacity sharing may need to interoperate with a separate chain for chemical logistics, requiring standardized identity frameworks, interoperable smart contract templates, and consistent semantic data models for capacity description. The development of industry-neutral token standards and oracle interfaces, driven by organizations such as the Enterprise Ethereum Alliance and the Decentralized Identity Foundation, points toward a future in which capacity sharing becomes a composable service that firms can plug into across multiple supply chains. In that future, the trust infrastructure becomes a public good, and its governance must transition from private consortium control to multi-stakeholder stewardship, involving industry associations, regulators, and civil society.

9. Sustainability and Long-Term Ecosystem Resilience

Capacity sharing has a direct sustainability dimension: by raising the utilization rate of expensive capital equipment, it reduces the need for new construction and the associated material and energy footprints. However, the computational overhead of operating AI training pipelines and blockchain consensus mechanisms can, if not carefully managed, offset these physical savings. Early blockchain implementations relying on proof-of-work have drawn justified criticism for their energy intensity [14]. Modern capacity sharing ecosystems must adopt energy-efficient consensus protocols such as proof-of-stake or federated consensus and leverage edge-AI processing to keep data and computation local where possible, reducing both latency and carbon footprint.

Long-term resilience of the ecosystem also depends on its ability to withstand adversarial behavior from both external attackers and internal defectors. Adversarial machine learning poses a relatively underexplored threat in this domain. A malicious firm could craft input data to poison AI forecasting models, causing competitors to misprice capacity or misjudge trustworthiness. The combination of on-chain model provenance records and off-chain anomaly detection systems offers a layered defense, but the arms race between adversarial AI and defensive monitoring is dynamic [19]. The immutability of the blockchain, while protecting the integrity of past records, also means that poisoned reputation marks can persist indefinitely unless explicit governance processes for expungement are embedded in the

protocol. Decentralized organizations must therefore design upgrade mechanisms, such as on-chain governance proposals to invalidate tainted blocks in extreme cases, that balance finality with corrective justice.

The ecosystem's adaptability over decades will depend on modularity. As AI algorithms evolve from static discriminative models to continuously learning agents, and as blockchain platforms adopt new cryptographic primitives such as post-quantum signatures, the trust fabric must be upgradeable without fracturing. A layered architecture that separates trust identity, capacity tokenization, AI orchestration, and dispute resolution into loosely coupled modules, connected through well-defined application programming interfaces and standardized data schemas, is essential for future-proofing. Such modularity also permits the gradual entry of industrial metaverse and digital twin technologies [16] that can supply ever richer representations of physical capacity states to the AI layer, further tightening the loop between virtual coordination and physical execution.

10. Conclusion

The convergence of AI and blockchain inaugurates a new class of industrial infrastructure in which inter-firm capacity sharing can transition from an ad hoc, relationship-dependent practice into a structured, algorithmically governed market with institutional-grade trust guarantees. This paper has argued that the viability of such ecosystems hinges not on any single technical breakthrough but on the deliberate articulation of an architecture that harmonizes the probabilistic intelligence of AI with the deterministic transparency of distributed ledgers. The structural trade-offs encountered along the path—between scalability and decentralization, opacity and auditability, automation and human discretion—are integral features of the design space rather than impediments to be eliminated. By treating these trade-offs as governance design variables, industry consortia can craft capacity sharing systems that are simultaneously efficient, fair, and resilient.

Our analysis underscores that trust in such ecosystems is not merely a byproduct of cryptographic proofs or algorithmic accuracy. It is a multi-layered construct that emerges from the interplay of verifiable identity, transparent rule enforcement, fair incentive design, and institutionalized spaces for discretionary reciprocity. The integration must therefore be approached with a socio-technical lens that respects the complexity of inter-organizational relationships [10] and the empirical realities of how firms reciprocate [13]. As regulators develop frameworks for algorithmic markets, and as federated learning [17] and privacy-preserving computation mature, the capacity sharing horizon will likely expand from industry-specific consortia toward open, composable capacity fabrics that span sectors and geographies. Realizing that vision will require sustained interdisciplinary collaboration among operations researchers, computer scientists, legal scholars, and industry strategists, yielding a field of inquiry as rich as it is practically urgent.

References

1. Fang, X., & Zhang, P. (2021). Capacity sharing between competitors. *Management Science*, 67(7), 4121–4144.
2. Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. White paper.
3. Toorajipour, R., Sohrabpour, V., Nazarpour, A., Oghazi, P., & Fischl, M. (2021). Artificial intelligence in supply chain management: A systematic literature review. *Journal of Business Research*, 122, 502–517.

4. Szabo, N. (1997). Formalizing and securing relationships on public networks. *First Monday*, 2(9).
5. Beck, R., Müller-Bloch, C., & King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10), 1020–1034.
6. Huberman, G., Leshno, J. D., & Moallemi, C. C. (2021). Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *The Review of Economic Studies*, 88(4), 1463–1501.
7. Croman, K., Decker, C., Eyal, I., Gencer, A. E., Juels, A., Kosba, A., ... & Wattenhofer, R. (2016). On scaling decentralized blockchains. In *International Conference on Financial Cryptography and Data Security* (pp. 106–125). Springer.
8. Josang, A., Ismail, R., & Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2), 618–644.
9. Diakopoulos, N. (2015). Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, 3(3), 398–415.
10. Zaheer, A., McEvily, B., & Perrone, V. (1998). Does trust matter? Exploring the effects of interorganizational and interpersonal trust on performance. *Organization Science*, 9(2), 141–159.
11. Al-Breiki, H., Rehman, M. H. U., Salah, K., & Svetinovic, D. (2020). Trustworthy blockchain oracles: Review, comparison, and open research challenges. *IEEE Access*, 8, 85675–85691.
12. Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7), 2117–2135.
13. Hu, X., & Caldentey, R. (2023). Trust and reciprocity in firms' capacity sharing. *Manufacturing & Service Operations Management*, 25(4), 1436–1450.
14. De Vries, A. (2018). Bitcoin's growing energy problem. *Joule*, 2(5), 801–805.
15. Zetzsche, D. A., Buckley, R. P., Arner, D. W., & Föhr, L. (2020). The distributed liability of distributed ledgers: Legal risks of blockchain. *University of Illinois Law Review*, 2020(4), 1361–1406.
16. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
17. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
18. Luu, L., Chu, D. H., Olickel, H., Saxena, P., & Hobor, A. (2016). Making smart contracts smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 254–269).
19. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Schafer, B. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.

20. Wright, A., & De Filippi, P. (2015). Decentralized blockchain technology and the rise of lex cryptographia. SSRN Electronic Journal.