

Graph Neural Networks for Modeling Reciprocity Relationships in Industrial Capacity Sharing Networks

Brace Frenklin

Department of Computer Science, University of North Texas, Denton, TX, USA.

brucework@unt.edu

Anil J. Mahajan

Department of Computer Science, University of Houston, Houston, TX, USA.

anil377@uh.edu

Abstract

Industrial capacity sharing has emerged as a strategic mechanism for firms to manage demand volatility, reduce capital expenditure, and enhance resource utilization across supply chains. The dynamics of such networks, however, depend critically on reciprocity relationships that extend beyond simple transactional exchanges and encompass trust, long-term mutual obligations, and strategic interdependence. This paper presents a system-level investigation into the use of graph neural networks for modeling these complex reciprocity patterns in industrial capacity sharing networks. We argue that graph neural networks offer a uniquely expressive computational substrate capable of capturing the multi-relational, temporally evolving, and context-dependent nature of reciprocal ties among firms. The discussion foregrounds the structural trade-offs involved in designing graph-based learning architectures for industrial settings, including the tension between expressive power and computational tractability, the need to incorporate heterogeneous firm attributes, and the challenge of aligning learned representations with operational governance mechanisms. Through detailed analysis, we explore how graph neural network models can be deployed to support fairness-aware resource allocation, early warning systems for reciprocity breakdown, and adaptive capacity pricing. We further examine critical deployment considerations, including data sovereignty, interoperability with legacy enterprise systems, and the long-term sustainability of data-driven reciprocity governance. Policy implications concerning algorithmic transparency, anti-trust considerations, and the reinforcement of cooperative norms are systematically addressed. The paper contributes a holistic framework that bridges machine learning research with the socio-technical realities of industrial ecosystems, offering a forward-looking perspective on the design of intelligent infrastructure for the circular and sharing economy.

Keywords

graph neural networks, reciprocity, capacity sharing, industrial networks, socio-technical systems, fairness, governance.

1. Introduction

The contemporary industrial landscape is increasingly characterized by the emergence of capacity sharing networks, wherein manufacturing firms, logistics providers, and service operators pool underutilized productive resources such as machinery, storage space,

transportation fleets, and specialized labor to achieve collective efficiency gains. These networks represent a departure from the traditional vertically integrated firm model and embody the principles of the sharing economy adapted to the high-stakes, capital-intensive domain of industrial operations. Unlike consumer-oriented sharing platforms, industrial capacity sharing involves substantial asset specificity, stringent quality requirements, regulatory constraints, and deeply embedded inter-organizational relationships that have often developed over decades of collaboration and competition. Within these dense relational fabrics, the concept of reciprocity emerges as a fundamental organizing principle. Reciprocity in this context refers not merely to the direct exchange of equivalent favors but to a multifaceted pattern of mutual responsiveness that encompasses preferential treatment in capacity allocation, flexibility during supply disruptions, joint investment in interoperability infrastructure, and the forbearance of short-term optimization in favor of long-term relational stability.

Understanding and modeling reciprocity relationships is essential for the design of effective governance mechanisms, dynamic pricing protocols, and risk management strategies in capacity sharing networks. Traditional operations management models have approached capacity sharing primarily through the lens of game theory and contract theory, often assuming rational agents with well-defined utility functions and complete information about partner reliability. While these analytical frameworks yield valuable insights into equilibrium behaviors, they struggle to capture the emergent, path-dependent, and high-dimensional nature of reciprocity as it unfolds in real-world industrial ecosystems. The rise of industrial digitization, manifested through the Industrial Internet of Things, digital twins, and enterprise blockchain systems, has generated a wealth of granular relational data that opens new possibilities for data-driven modeling. Graph neural networks (GNNs) have rapidly gained prominence as a powerful class of deep learning architectures specifically designed to operate on graph-structured data. By leveraging the inherent topology of industrial relationships, where firms constitute nodes and capacity-sharing transactions, contracts, and communication events form edges with multiple types of attributes, GNNs can learn latent representations that encode both nodal characteristics and the structural signatures of reciprocity.

This paper undertakes a comprehensive system-level examination of GNN-based modeling for reciprocity in industrial capacity sharing networks. Rather than focusing on algorithmic novelties or performance benchmarks, we situate the technical discussion within the broader socio-technical, governance, and infrastructural contexts that determine whether such models can be responsibly and effectively deployed in practice. The analysis addresses the architectural trade-offs inherent in designing GNNs for industrial relational data, the governance implications of encoding reciprocity norms into algorithmic systems, the challenges of ensuring fairness and preventing discriminatory outcomes in automated capacity allocation, and the policy frameworks required to sustain cooperative norms in increasingly algorithmic-mediated industrial ecosystems. By integrating perspectives from machine learning, operations management, organizational sociology, and information systems infrastructure, we aim to provide a roadmap for researchers and practitioners seeking to harness graph learning for the next generation of resilient and equitable industrial sharing networks.

2. Reciprocity in Industrial Capacity Sharing Ecosystems

The architecture of industrial capacity sharing networks diverges significantly from spot-market platforms in both structural complexity and temporal depth. Long-term supplier

relationships, joint ventures, industry consortia, and regional manufacturing clusters form dense, multiplex graphs where firms are connected by multiple simultaneous relations, including equity ties, shared certifications, collaborative research agreements, and historical exchange patterns. Reciprocity in such networks cannot be reduced to a scalar measure of balanced transactions. Instead, it manifests as a dynamic equilibrium that is maintained through the interplay of direct reciprocation, generalized exchange where a firm receiving capacity from one partner may compensate by offering preferential access to a third party, and what organizational theorists have termed serial reciprocity, where benefits are passed forward along a supply chain. The temporal dimension is equally critical. Reciprocity is often evaluated over extended time horizons that span multiple production cycles, meaning that a short-term deficit in balanced exchange may be tolerated if it serves to preserve a strategic relationship that yields benefits during future demand surges or component shortages.

The operationalization of reciprocity in capacity sharing involves several distinct but interrelated behavioral patterns. First, capacity providers often prioritize requests from partners with whom they have a history of balanced or favorable exchange, even when spot market prices might suggest more profitable alternative allocations. Second, firms may deliberately maintain a margin of slack capacity not as organizational inefficiency but as a relational resource that can be deployed to assist partners in distress, with the expectation of future reciprocation. Third, the sharing of sensitive demand forecasts or production plans, which constitutes a form of informational reciprocity, serves as a critical enabler of coordinated capacity planning but depends heavily on trust that the shared information will not be exploited opportunistically. These patterns collectively form a relational infrastructure that stabilizes industrial networks against the volatility of purely market-based coordination. Understanding how these patterns are encoded in transaction data, communication logs, and contractual structures is the central challenge that GNN-based modeling seeks to address.

Economic and sociological research on inter-firm relationships has long recognized that trust and reciprocity function as informal governance mechanisms that substitute for complete contracts in environments characterized by uncertainty and asset specificity. In the manufacturing sector, where capacity investments are lumpy and lead times for capacity expansion can span years, the ability to rely on reciprocal capacity sharing arrangements reduces the risk of both overcapacity and undercapacity. The financial crisis of 2008 and the supply chain disruptions caused by the COVID-19 pandemic vividly illustrated how firms embedded in robust reciprocal networks were better able to weather demand shocks by flexibly reallocating production loads among partners. However, the very informality that makes reciprocity adaptive also makes it fragile and opaque. When reciprocity expectations are violated, whether due to opportunistic behavior or to genuine constraints, the degradation of trust can propagate through the network, causing a cascade of defensive capacity hoarding that undermines system-wide efficiency. It is in this context that computational models of reciprocity, capable of detecting emergent shifts in relational patterns and providing early warnings of norm erosion, become not merely academic curiosities but operational necessities.

3. Graph Neural Networks as a Modeling Substrate

Graph neural networks provide a natural computational framework for learning from the relational data that embodies reciprocity in industrial networks. At their core, GNNs operate through a message-passing paradigm where each node iteratively aggregates information from its neighbors, updates its internal representation, and, over multiple layers, synthesizes information from increasingly larger receptive fields in the graph. This inductive bias toward

relational structure distinguishes GNNs from traditional feedforward or convolutional neural networks that assume independent and identically distributed data samples. In the context of capacity sharing, a firm node can be endowed with initial features derived from its production capabilities, financial health indicators, historical utilization rates, and geographic location, while edges can carry features representing transaction volumes, contract durations, frequency of communication, and qualitative assessments of collaboration quality. The GNN then learns to map these heterogeneous inputs into low-dimensional embedding spaces where firms with similar reciprocity profiles are positioned in proximity, and where the likelihood of future reciprocal behavior can be predicted.

Several architectural families within the broader GNN landscape offer distinct advantages for modeling industrial reciprocity. Graph attention networks allow the model to learn differential importance weights for different relational partners, which aligns intuitively with the observation that firms do not treat all capacity-sharing partners equally but accord greater weight to those with stronger historical reciprocity. Heterogeneous graph architectures, such as relational graph convolutional networks, are essential for handling industrial networks where multiple types of nodes, including manufacturers, logistics providers, and financial institutions, interact through diverse edge types like capacity transfers, credit guarantees, and knowledge exchanges. Temporally aware GNNs that incorporate recurrent or self-attention mechanisms over graph snapshots can capture the evolution of reciprocity over time, distinguishing between transient disruptions and systematic shifts in relational norms. The choice among these architectures entails fundamental trade-offs between expressive power, computational efficiency, and the interpretability of learned representations. Highly expressive models with multiple attention heads and heterogeneous message functions may capture nuanced reciprocity patterns but demand computational resources that can become prohibitive for networks spanning thousands of firms with millions of transactions, and their opacity can hinder the explanation of model decisions to human decision-makers.

The training of GNNs for reciprocity modeling raises critical questions about supervision signals and objective design. In many industrial settings, ground-truth labels for reciprocity are not directly observable. Instead, they must be inferred from proxy variables such as the persistence of partnerships, the symmetry of capacity flows over rolling time windows, or the speed with which firms respond to partner requests. These proxies are inherently noisy and may embed biases that the model could perpetuate. Self-supervised learning paradigms, where the model is trained to predict masked edges, reconstruct historical transaction sequences, or contrast positive pairs of reciprocating firms with negative samples, offer a promising avenue for learning reciprocity representations without explicit labeling. Yet, the design of pretext tasks must be carefully aligned with the operational definitions of reciprocity that matter for governance decisions, lest the model learn representations that are statistically sophisticated but operationally irrelevant. The interdisciplinary challenge lies in translating the rich theoretical vocabulary of reciprocity from organizational sociology and behavioral economics into computationally tractable objective functions that guide GNN training toward societally desirable outcomes.

4. Architectural Trade-offs and System Design

Designing a GNN-based modeling system for industrial capacity sharing networks requires navigating a complex landscape of architectural trade-offs that cannot be resolved through purely technical criteria. One central tension concerns the granularity of modeling. A fine-grained model that represents each individual transaction as an edge and updates node

embeddings after every capacity exchange can capture high-frequency reciprocity signals but introduces significant computational overhead and may be overly sensitive to operational noise. Conversely, a coarse-grained model that aggregates transactions over monthly or quarterly windows reduces volatility and computational cost but loses the temporal resolution needed to detect emergent reciprocity breakdowns. The choice of aggregation window must be informed by the natural business rhythms of the industry, including production cycle lengths, budgeting periods, and contract renewal timelines, thereby demanding close collaboration between system designers and domain experts.

Data heterogeneity presents a second major architectural challenge. Industrial networks integrate structured transactional data from enterprise resource planning systems, semi-structured data from electronic data interchange messages, and unstructured data from maintenance logs or engineer collaboration platforms. GNN architectures must incorporate modality-specific encoders that project these diverse data streams into a unified latent space while preserving the distinct informational content of each modality. The integration of these encoders into a cohesive message-passing framework raises questions about information bottleneck design and the risk of modality collapse, where one dominant data source, such as financial transaction volume, overwhelms subtler signals from communication patterns that may be more indicative of underlying trust and reciprocity. Regularization techniques and modality-specific attention mechanisms can mitigate such collapse, but they introduce additional hyperparameters that multiply the configuration space and complicate model maintenance in production environments.

A third trade-off concerns the geographic and organizational distribution of data. Capacity sharing networks often span multiple legal jurisdictions with divergent data protection regulations, and participating firms may be reluctant to share granular operational data with a centralized model repository due to concerns about competitive intelligence leakage. Federated GNN architectures, where each firm or regional cluster trains a local GNN on its own data and shares only encrypted gradient updates or model parameters with a central aggregation server, offer a privacy-preserving alternative. However, federated learning for graphs introduces challenges related to the non-IID nature of graph data, where the local subgraph of a single firm does not constitute an independent sample from the global distribution. The structural roles of firms in the capacity sharing network, such as hub versus peripheral positions, create statistical heterogeneity that can cause federated averaging to converge slowly or to erase the very relational patterns it seeks to learn. System designers must therefore weigh the privacy benefits of federation against the potential degradation in model quality and the increased complexity of orchestrating distributed training across heterogeneous industrial IT environments.

5. Governance, Fairness, and Policy Implications

The deployment of GNN-based reciprocity models in industrial capacity sharing networks is not a purely technical undertaking but a governance intervention that reshapes the information landscape within which firms make strategic decisions. When a GNN model informs capacity allocation by predicting which partners are likely to reciprocate, it effectively encodes a normative vision of desirable relational behavior into the operational fabric of the network. This raises fundamental questions about algorithmic fairness that go beyond standard definitions of demographic parity or equalized odds, because the entities affected are firms rather than individuals, and the harms of unfair allocation can cascade through interconnected supply chains, jeopardizing the viability of small and medium enterprises that lack the market

power to secure capacity through alternative means. Fairness in this context must be conceptualized in relational terms, considering whether the model systematically disadvantages firms based on their position in the network topology, their historical patterns of interaction that may reflect legacy structural inequalities, or the quality of data they can generate.

A particularly subtle governance challenge arises from the feedback loops between GNN predictions and firm behavior. If firms learn that the capacity allocation system favors those with high predicted reciprocity scores, they may alter their operational behaviors to game the scoring algorithm rather than to cultivate genuine cooperative relationships. This can lead to a statistical arms race where predictive features that were once informative become corrupted by strategic manipulation, degrading the model's ability to distinguish authentic reciprocity from performative cooperation. Countering such dynamics requires a governance framework that combines algorithmic transparency, so that firms can understand and contest the basis for allocation decisions, with an adversarial robustness regime that continuously monitors for distributional shifts indicative of gaming. The institutional mechanisms for implementing such governance—whether through industry associations, regulatory bodies, or multi-stakeholder consortia—must be designed to balance the efficiency gains from automated reciprocity assessment against the deliberative values of procedural fairness and stakeholder participation.

The policy implications of GNN-mediated reciprocity governance extend into competition law and industrial policy. When a group of firms uses a shared GNN model to coordinate capacity allocation based on reciprocity predictions, the line between legitimate operational collaboration and anti-competitive collusion can become blurred. Regulators must grapple with the question of whether algorithmic facilitation of reciprocal exchange constitutes a form of information sharing that reduces competitive intensity or whether it serves pro-competitive purposes by lowering barriers to entry for firms that lack the scale to self-provide all capacity. Policy frameworks that emphasize the design principles of algorithmic neutrality, mandated openness of reciprocity models for regulatory audit, and the preservation of firm autonomy to override algorithmic recommendations in favor of relational commitments that defy quantification, offer a potential path forward. Furthermore, as reciprocity models become embedded in the digital infrastructure of industrial sectors, they raise questions about infrastructure sovereignty and whether critical modeling capabilities should be provided by global technology platforms or maintained as shared digital public goods governed by industry collectives.

6. Infrastructure and Deployment Considerations

The translation of GNN-based reciprocity models from research prototypes to production-grade industrial systems necessitates a robust and resilient supporting infrastructure that spans data engineering, model serving, and continuous monitoring. The data pipelines that feed GNN models must ingest transaction records from a heterogeneous array of enterprise systems, many of which operate on legacy architectures with batch-oriented interfaces and inconsistent data quality. The construction of the dynamic graph upon which the GNN operates requires entity resolution to map transaction records to a consistent set of firm identifiers across systems, temporal alignment to ensure that edges reflect the actual time of capacity transfer rather than accounting post dates, and the handling of missing data that is endemic in industrial settings where not all capacity exchanges are formally recorded. The

engineering effort required to build and maintain these pipelines should not be underestimated and often constitutes the majority of the total project cost.

Model serving in industrial environments imposes stringent latency and reliability requirements that are not typical of academic GNN research. Capacity allocation decisions may need to be made in near-real-time during production line disruptions or logistics emergencies, and the GNN inference must complete within tight service-level agreements. Achieving low-latency inference on large-scale industrial graphs requires specialized model optimization techniques, including graph sampling strategies that limit the number of neighboring nodes considered during message passing, knowledge distillation from large teacher GNNs into compact student models suitable for edge deployment, and the pre-computation of node embeddings for stable portions of the network with incremental updates for dynamic regions. The trade-off between inference speed and model fidelity must be calibrated against the operational tolerance for suboptimal capacity allocations, which varies by industry and by the criticality of the specific resource being shared.

Continuous monitoring and model lifecycle management form the final pillar of deployment infrastructure. Reciprocity patterns in industrial networks can shift abruptly in response to macroeconomic shocks, technological discontinuities, or mergers and acquisitions among participating firms. A GNN model that was well-calibrated under one regime may become systematically biased or entirely non-representative under a new regime. Deployment architectures must therefore incorporate automated concept drift detection that compares the statistical properties of recent graph snapshots against the training distribution, triggering model retraining or human review when significant deviations are detected. The feedback loop from operational outcomes, such as the actual fulfillment of reciprocal commitments following capacity allocations, must be tightly integrated to enable the continuous refinement of the model's understanding of reciprocity. This infrastructure requires not only technical components but also organizational processes that define the roles and responsibilities of data engineers, model developers, domain experts, and business stakeholders in the model stewardship lifecycle.

7. Sustainability and Long-Term Robustness

The long-term sustainability of GNN-driven reciprocity modeling in industrial networks depends on aligning technical system design with the socio-economic incentives that sustain cooperative behavior. A purely efficiency-maximizing model that optimizes for short-term capacity utilization may, if left unchecked, erode the relational slack that underpins long-term reciprocity by always allocating capacity to the momentarily most efficient user at the expense of partners who have historically been reliable but are currently experiencing transient difficulties. System designers must therefore incorporate sustainability constraints that explicitly preserve relational diversity and prevent the network from collapsing into a brittle configuration where a small number of highly reciprocating firms dominate all capacity flows, creating single points of failure and stifling the entry of innovative newcomers. These constraints can be operationalized through regularization terms in the GNN training objective that penalize overly concentrated allocation patterns, or through post-processing fairness adjustments that ensure a minimum threshold of capacity access for firms that fall below certain reciprocity scores due to data sparsity rather than genuine uncooperativeness.

Robustness to adversarial threats presents another long-term concern. As the economic stakes of GNN-mediated capacity allocation grow, the incentive for malicious actors to manipulate the model through carefully crafted transaction patterns that inflate their reciprocity scores

increases correspondingly. Adversarial attacks on GNNs can take the form of structural perturbations, where attackers create fake firms and orchestrate sham transactions to build an artificial history of reciprocity, or feature-based attacks, where genuine firms misreport their capabilities or financial status. Defending against such attacks requires a multi-layered approach encompassing cryptographic identity management to ensure that firms in the network are who they claim to be, statistical anomaly detection that flags transaction patterns inconsistent with the physical constraints of industrial operations, and the development of GNN architectures with certified robustness guarantees against specific classes of graph perturbations. The intersection of robust machine learning and industrial security remains an underexplored frontier with significant practical importance.

From the perspective of environmental sustainability, the increasing computational footprint of training and serving large-scale GNN models for industrial applications must be weighed against the resource efficiency gains they enable. The carbon cost of training a large heterogeneous GNN on multi-year industrial transaction graphs can be substantial, particularly if hyperparameter optimization involves hundreds of trial runs. Architectural innovations that reduce this footprint, such as progressive model growth where capacity is added incrementally as data volume increases, and the use of sparsely activated mixture-of-experts layers that scale sub-linearly with graph size, represent important directions for future research. Moreover, the energy consumption of the model serving infrastructure must be accounted for in the overall sustainability calculus of industrial digitization initiatives. A holistic assessment that compares the computational energy cost of GNN-based reciprocity governance against the waste reduction, overcapacity avoidance, and logistics optimization it facilitates is essential to ensure that the cure is not worse than the disease.

8. Conclusion

This paper has presented a comprehensive system-level analysis of the application of graph neural networks to the modeling of reciprocity relationships in industrial capacity sharing networks. We have argued that GNNs provide a technically sophisticated and contextually appropriate framework for capturing the multi-relational, temporally dynamic, and structurally embedded nature of reciprocal industrial exchange. Yet, we have equally emphasized that the successful deployment of such models is contingent upon a much broader set of considerations than predictive accuracy alone. The architectural choices made in designing GNN systems for industrial settings embody trade-offs between granularity and scalability, heterogeneity and coherence, and centralization and privacy that must be navigated with careful attention to the specific operational and regulatory contexts of each deployment. The governance frameworks within which these models operate will determine whether they reinforce existing power asymmetries or contribute to a more equitable distribution of capacity access across industrial ecosystems. The policy implications, spanning competition law, algorithmic transparency, and digital infrastructure sovereignty, demand sustained dialogue between technologists, regulators, and industry stakeholders.

Looking forward, the convergence of GNN technology with other emerging computational paradigms, including physics-informed graph learning that can incorporate domain knowledge about production constraints, and the integration of large language model-based reasoning with graph-structured data for generating natural language explanations of reciprocity assessments, opens exciting avenues for making these systems more transparent and trustworthy. The path toward responsible adoption lies not in the blind pursuit of ever more complex models but in the deliberate co-design of technical, organizational, and

institutional mechanisms that together cultivate the kind of robust, adaptive reciprocity upon which resilient industrial networks depend. As the global economy confronts the imperatives of supply chain resilience, circular material flows, and distributed manufacturing, the intelligent modeling of cooperative capacity sharing will only grow in importance. Ensuring that the computational tools developed for this purpose are themselves designed with wisdom, humility, and a commitment to the long-term flourishing of industrial communities is the central challenge that this paper has sought to illuminate.

References

1. Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., & Vandergheynst, P. (2017). Geometric deep learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4), 18–42.
2. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *Proceedings of the International Conference on Learning Representations*.
3. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. *Proceedings of the International Conference on Learning Representations*.
4. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 1024–1034.
5. Schlichtkrull, M., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., & Welling, M. (2018). Modeling relational data with graph convolutional networks. *Proceedings of the European Semantic Web Conference*, 593–607.
6. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24.
7. Chen, J., Ma, T., & Xiao, C. (2018). FastGCN: Fast learning with graph convolutional networks via importance sampling. *Proceedings of the International Conference on Learning Representations*.
8. Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., ... & Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
9. Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437(7063), 1291–1298.
10. Hu, X., & Caldentey, R. (2023). Trust and reciprocity in firms' capacity sharing. *Manufacturing & Service Operations Management*, 25(4), 1436–1450.
11. Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91(3), 481–510.
12. Uzzi, B. (1997). Social structure and competition in interfirm networks: The paradox of embeddedness. *Administrative Science Quarterly*, 42(1), 35–67.
13. Gulati, R., & Gargiulo, M. (1999). Where do interorganizational networks come from? *American Journal of Sociology*, 104(5), 1439–1493.

14. Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2), 349–399.
15. Dyer, J. H., & Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4), 660–679.
16. Williamson, O. E. (1985). *The Economic Institutions of Capitalism*. Free Press.
17. Gereffi, G., Humphrey, J., & Sturgeon, T. (2005). The governance of global value chains. *Review of International Political Economy*, 12(1), 78–104.
18. Boudreau, K. J., & Lakhani, K. R. (2013). Using the crowd as an innovation partner. *Harvard Business Review*, 91(4), 60–69.
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. *Minds and Machines*, 28(4), 689–707.
20. Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
21. Shrivastava, P. (1995). The role of corporations in achieving ecological sustainability. *Academy of Management Review*, 20(4), 936–960.
22. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 115.
23. Pereira, A. G., & Stergioulas, L. K. (2023). Federated graph learning for privacy-preserving industrial analytics. *IEEE Transactions on Industrial Informatics*, 19(2), 1876–1885.