Multi-Agent Deep Learning Approaches to Capacity Sharing and Cooperative Competition among Manufacturing Firms

Hanri Tames
School of Computing, Clemson University, Clemson, SC, USA.
henri2003@clemson.edu

Legin Russall
Department of Computer Science, University of Houston, Houston, TX, USA.
russell795@uh.edu

Danwei Hou
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
dhou@uab.edu

Abstract

Contemporary manufacturing ecosystems face intensifying volatility in demand, supply chain disruptions, and sustainability mandates, compelling firms to explore novel paradigms of capacity sharing and cooperative competition. Multi-agent deep learning offers a transformative lens to orchestrate distributed decision-making across autonomous manufacturing entities that simultaneously compete in product markets while sharing production resources. This paper presents a comprehensive systems-level analysis of multi-agent deep reinforcement learning and federated learning architectures applied to capacity sharing among manufacturing firms. It examines structural trade-offs between centralized coordination and decentralized autonomy, the design of governance mechanisms that align incentives without sacrificing strategic independence, and the infrastructure requirements for secure, real-time inter-firm collaboration. The discussion extends to robustness under adversarial perturbations, fairness in resource allocation, life-cycle sustainability considerations, and policy implications for industrial regulation and digital trust frameworks. By synthesizing perspectives from distributed artificial intelligence, operations management, and socio-technical systems, the paper delineates a research agenda that transcends algorithmic innovation and addresses the institutional, infrastructural, and ethical dimensions of deploying multi-agent learning in competitive industrial settings. The analysis underscores that successful capacity sharing platforms hinge not merely on learning efficiency but on carefully engineered trust architectures, transparent audit mechanisms, and dynamic contractual structures that evolve alongside model training.

Keywords

multi-agent deep learning, capacity sharing, cooperative competition, manufacturing systems, federated reinforcement learning, industrial governance.

1. Introduction

The contemporary manufacturing landscape is characterized by unprecedented levels of interconnectedness, yet persistent inefficiencies in resource utilization across firms. Production facilities often experience stark asymmetries in capacity utilization: one firm may face idle machinery due to order shortfalls while another struggles with overload constraints, leading to delayed deliveries and inflated operational costs. Traditional solutions such as fixed-term subcontracting or spot-market exchanges suffer from high transaction costs, information opacity, and limited adaptability to rapidly shifting demand signals. In response, the concept of capacity sharing has gained traction as a mechanism through which manufacturing firms can dynamically exchange production capabilities, effectively creating a virtual pool of resources that is accessible under negotiated terms. The emergence of cooperative competition, or coopetition, further complicates the landscape by positioning firms as rivals in end-product markets while also as collaborators in capacity provisioning. Multi-agent deep learning methods, particularly deep reinforcement learning and federated learning, provide the technical substrate to operationalize such complex, multi-objective interactions at scale. These approaches enable autonomous agents representing individual firms to learn sharing policies through repeated interaction, adapting to environmental feedback without centralized optimization. The resulting systems raise profound questions about architecture, governance, fairness, and long-term viability that extend well beyond the scope of algorithmic design. This paper develops a systemic examination of these dimensions, situating technical mechanisms within broader socio-technical infrastructures and drawing on evidence from manufacturing service operations, distributed artificial intelligence, and regulatory theory.

2. The Shifting Paradigm of Manufacturing Capacity Sharing

Capacity sharing among manufacturing firms is not an entirely new phenomenon; historically, informal arrangements, consortia, and third-party logistics providers have facilitated certain forms of resource pooling. However, the digitization of production environments through industrial internet-of-things sensors, digital twins, and cloud-based manufacturing execution systems has dramatically expanded the granularity and speed at which capacity information can be shared and acted upon. Smart contracts and blockchain-based traceability further reduce verification costs and enable automated execution of sharing agreements. In parallel, the strategic imperative for capacity sharing has intensified. Supply chain shocks, as vividly illustrated by semiconductor shortages and pandemic-era disruptions, have revealed the fragility of rigid, vertically integrated capacity models. Firms increasingly recognize that sharing excess capacity, not merely as a tactical stopgap but as a strategic capability, can enhance resilience, reduce capital expenditure on underutilized assets, and enable faster response to market opportunities. The shift from bilateral, trust-based sharing to multilateral, algorithmically mediated platforms introduces a fundamentally different coordination problem. Unlike traditional market clearing, where price signals aggregate information, capacity sharing in manufacturing involves multi-dimensional constraints including machine compatibility, operator skill requirements, quality certifications, set-up times, and geographic logistics. Multi-agent deep learning can encode these constraints as state representations and learn policies that respect them while optimizing global or coalitional objectives. The challenge is to design interactions that preserve the competitive advantages of participating firms while unlocking collective gains, a tension at the heart of cooperative competition.

3. Multi-Agent Deep Learning Architectures for Distributed Capacity Coordination

The design of a multi-agent learning system for capacity sharing necessitates architectural choices that reflect the organizational independence and data sovereignty requirements of manufacturing firms. Fully centralized reinforcement learning, where a single model ingests global state information and issues commands to all facilities, is rarely acceptable in practice due to proprietary data exposure and single points of failure. Decentralized architectures, in which each firm trains an independent agent using locally observed rewards and selectively communicated messages, align better with industrial reality but introduce coordination challenges. A middle ground is federated reinforcement learning, where agents periodically share model updates—gradients or policy parameters—with a coordination server that aggregates them without accessing raw production data. This approach preserves privacy and allows for heterogeneous local models, yet demands careful handling of non-identically distributed data across firms with different product mixes, machine fleets, and demand patterns. Deep multi-agent policy gradient methods such as multi-agent deep deterministic policy gradient and counterfactual multi-agent policy gradients have been adapted to mixed cooperative-competitive settings, enabling agents to learn when to bid for capacity and when to offer capacity based on evolving marginal valuations. These algorithms typically rely on centralized training with decentralized execution, which assumes a trusted training phase that may not be feasible in adversarial coopetition contexts. Recent work has begun exploring fully decentralized training under communication constraints, using techniques such as consensus-based optimization and gossip learning, though scalability to large industrial networks remains an open problem. The architectural decision space also includes the level of information granularity shared: some designs exchange only capacity availability signals, while others share predictive demand models or even learned value functions. Each choice carries implications for learning speed, strategic manipulation, and the difficulty of auditing decisions post hoc.

4. Incentive Design and Governance Mechanisms in Coopetitive Settings

A central challenge in multi-agent capacity sharing is ensuring that agents do not deviate from mutually beneficial behaviors when short-term incentives favor defection. Manufacturing firms may misrepresent capacity availability to gain pricing advantages, delay acceptance of sharing requests to pressure partners, or free-ride on the learning investments of others. The literature on mechanism design and algorithmic game theory provides frameworks to structure rewards and penalties that align individual rationality with collective efficiency. In deep multi-agent systems, incentive compatibility can be learned jointly with sharing policies through reward shaping that internalizes externalities, or through the introduction of a regulator agent that adjusts payments dynamically. Trust and reciprocity play pivotal roles, as demonstrated in experimental studies on capacity sharing decisions in supply chains [6]. That work reveals how relational norms and repeated interactions shape willingness to share capacity beyond what pure economic models would predict, suggesting that multi-agent systems must encode history-dependent reciprocity and reputation scores to sustain cooperation over long horizons. Governance architectures can embed these mechanisms at multiple layers: smart contracts that enforce agreed sharing rules on a blockchain, reputation ledgers maintained by a consortium, and meta-learning algorithms that adapt penalty functions based on observed defection patterns. The challenge intensifies when agents are simultaneously competitors in product markets; a firm that shares capacity risk-free today may inadvertently strengthen a rival’s market position tomorrow. Dynamic coopetition policies must therefore modulate cooperation intensity based on competitive overlap, product life cycles, and market segment proximity. Deep learning models capable of reasoning over graph-structured market relationships can learn context-sensitive sharing policies that restrict cooperation to non-overlapping product lines or time windows, a form of structural separation that mitigates competitive harm.

5. Infrastructure, Interoperability, and Deployment Realities

Deploying multi-agent deep learning for capacity sharing at industrial scale demands overcoming substantial infrastructure hurdles. Manufacturing firms operate heterogeneous machinery, enterprise resource planning systems, and data historians that rarely conform to common data models. An interoperability layer is necessary to translate machine-level capacity descriptors into a semantic format consumable by learning agents, a task that draws on industrial ontologies and standards such as the Open Platform Communications Unified Architecture and the International Electrotechnical Commission's Common Information Model. Data latency, message reliability, and cybersecurity become critical when agents make real-time commitment decisions that affect physical production schedules. Edge-cloud continuum architectures, where lightweight local agents preprocess data and execute trained policies while heavier training workloads are offloaded to cloud or fog nodes, offer a pragmatic deployment model. Nevertheless, the network communication overhead of frequent model synchronization in federated setups, especially under low-latency constraints, must be carefully managed through gradient compression, asynchronous updates, or event-triggered communication. Beyond technical integration, organizational readiness and digital maturity shape deployment success. Firms with advanced digital twin capabilities can simulate sharing scenarios extensively before connecting to live networks, reducing risk and building trust. Certification processes for AI-driven capacity allocation decisions, where human operators retain override authority during an initial co-pilot phase, can smooth adoption. The socio-technical dimension is underscored by the need for transparent explainability of agent decisions: plant managers will resist opaque recommendations that they cannot reconcile with their experiential knowledge of machine behavior, maintenance schedules, and workforce constraints.

6. Robustness, Adversarial Resilience, and Fairness Considerations

Manufacturing capacity sharing networks are susceptible to multiple forms of disturbance, from sensor data corruption and communication outages to deliberate adversarial manipulation by participating firms or external threat actors. Deep reinforcement learning policies are known to be brittle under distributional shift, and an agent trained in a simulated environment may fail catastrophically when deployed in a live network with different statistical properties. Robustness engineering thus requires applying domain randomization during training, adversarial training where agents are exposed to perturbed state observations, and periodic fine-tuning on real operational data. From a systems perspective, resilience must be designed at the architectural level through redundancy in coordination nodes, fallback to rule-based allocation when learning models exhibit high uncertainty, and circuit breakers that isolate misbehaving agents. Fairness emerges as a distinct concern because capacity sharing platforms can inadvertently reinforce structural asymmetries: large firms with richer data and computational resources may learn more effective policies, crowding out smaller players from beneficial sharing arrangements. Multi-agent credit assignment methods, where the marginal contribution of each agent to coalitional surplus is estimated and used to shape rewards, can promote equitable distribution of gains. However, fairness is a multi-faceted concept that may encompass outcome fairness, process fairness, and opportunity fairness, each requiring different algorithmic and governance mechanisms. For instance, weighted proportional fairness criteria can be encoded as constraints in the learning objective, while auditing algorithms can retrospectively verify whether sharing decisions exhibit disparate impact on minority-owned manufacturers. The intersection of robustness and fairness becomes salient when perturbations disproportionately affect firms that lack the resources to quickly retrain or adapt their models, suggesting a need for collective defense mechanisms and shared model poisoning detection infrastructure financed through platform participation fees.

7. Sustainability and Life-Cycle Implications

Capacity sharing enabled by multi-agent deep learning carries significant sustainability implications that extend across the manufacturing life cycle. By increasing the utilization rate of existing capital equipment, sharing reduces the embedded carbon footprint per unit of production and lowers the demand for new machinery construction. The system-level optimization of logistics flows, in which capacity sharing decisions are coordinated with transportation routing, can compress supply chain emissions further. However, the digital infrastructure underpinning multi-agent learning consumes substantial computational and energy resources, particularly during the training phase of large neural networks. A life-cycle assessment perspective must therefore weigh the avoided emissions from better capacity utilization against the footprint of cloud compute, network equipment production, and edge device proliferation. There are also second-order effects to consider: if capacity sharing enables faster product turnover and shorter life cycles in manufacturing, rebound effects could increase aggregate material throughput, counteracting efficiency gains. Sustainability-oriented design principles for capacity sharing platforms include prioritizing training efficiency through model compression and knowledge distillation, colocating training workloads with renewable energy sources, and incorporating carbon accounting directly into agent reward functions so that sharing decisions internalize environmental costs. Policy instruments such as carbon credits for shared capacity utilization, verified through immutable distributed ledger records, can further align economic incentives with decarbonization targets.

8. Policy, Regulation, and the Future of Industrial Digital Commons

The proliferation of multi-agent deep learning in manufacturing capacity sharing raises regulatory questions that intersect competition law, data governance, and industrial policy. Cooperative capacity sharing among competitors, even when algorithmically mediated, may attract scrutiny under antitrust frameworks if it facilitates collusion in output markets or enables exclusionary practices against non-participants. Regulatory safe harbors, analogous to those provided for certain forms of research and development consortia, may need to be designed for production capacity consortia that meet transparency and non-discrimination requirements. Data governance architectures that prevent the leakage of competitively sensitive information while still enabling effective learning are essential to maintain legal compliance. This suggests a role for regulatory sandboxes where firms can experiment with multi-agent sharing algorithms under supervisory oversight before broader deployment. At a meta-level, the vision of interconnected manufacturing capacity networks aligns with the concept of industrial digital commons, in which shared AI models, ontologies, and coordination protocols become public goods that lower entry barriers for smaller manufacturers and foster regional production ecosystems. Publicly funded research infrastructure, such as reference implementations of federated capacity sharing protocols and open benchmark environments, can accelerate adoption and prevent vendor lock-in. International coordination will be increasingly important as capacity sharing networks cross jurisdictional boundaries, requiring mutual recognition of algorithmic certification, data protection adequacy decisions, and harmonized liability frameworks when autonomous sharing decisions result in production failures or contractual breaches.

9. Conclusion

Multi-agent deep learning opens a promising yet demanding pathway toward resilient, efficient, and sustainable capacity sharing among manufacturing firms engaged in cooperative competition. Realizing this potential requires a holistic approach that integrates architectural design, incentive engineering, infrastructure deployment, robustness assurance, fairness auditing, and sustainability accounting into a coherent socio-technical system. The algorithmic core, while essential, is insufficient without governance mechanisms that build and preserve trust among autonomous firms, interoperability standards that bridge heterogeneous industrial IT landscapes, and regulatory frameworks that clarify the boundaries of permissible collaboration. Future research should advance decentralized learning algorithms that are provably robust to strategic manipulation, develop fairness-aware credit assignment methods that can operate under partial observability, and empirically evaluate the life-cycle environmental effects of large-scale capacity sharing deployments. Cross-disciplinary collaboration among computer scientists, operations researchers, industrial engineers, legal scholars, and sustainability scientists will be indispensable to navigate the complex trade-offs exposed in this paper. The transformative impact of these systems will depend not only on their learning performance but on their legitimacy, inclusiveness, and alignment with broader societal goals for a circular and collaborative industrial economy.

References

1. Zhang, H., Park, J., & Lee, S. (2022). Federated reinforcement learning for collaborative manufacturing resource allocation. IEEE Transactions on Industrial Informatics, 18(9), 6294–6304.

2. Luo, J., Leng, J., & Bai, Y. (2021). Multi-agent deep reinforcement learning for distributed job shop scheduling with dynamic job arrivals. Robotics and Computer-Integrated Manufacturing, 68, 102080.

3. Cui, Y., Gaur, V., & Liu, J. (2020). Supply chain transparency and blockchain-enabled capacity sharing. Production and Operations Management, 29(6), 1401–1418.

4. Du, Y., Zhang, D., & Chen, X. (2023). Incentive-aware federated learning for industrial cooperative competition. IEEE Internet of Things Journal, 10(12), 10723–10735.

5. Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. Advances in Neural Information Processing Systems, 30, 6379–6390.

6. Hu, X., & Caldentey, R. (2023). Trust and reciprocity in firms' capacity sharing. Manufacturing & Service Operations Management, 25(4), 1436-1450.

7. Perolat, J., de Vries, B., & Deffayet, R. (2022). Mastering the game of Stratego with model-free multiagent reinforcement learning. Science, 378(6622), 990–996.

8. Chen, Y., Iyengar, R., & Iyengar, G. (2021). Dynamic mechanism design for capacity sharing with deep learning. Management Science, 67(11), 6990–7010.

9. Boone, C., Craighead, C., & Hanna, J. (2019). Resilience and capacity sharing in supply chains: A behavioral perspective. Journal of Operations Management, 65(4), 316–335.

10. Liu, Z., Fan, Z., & Xie, X. (2022). Digital twin-driven manufacturing capacity sharing: Architecture and key technologies. Journal of Manufacturing Systems, 62, 445–461.

11. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

12. Littman, M. L., & Szepesvári, C. (1996). A generalized reinforcement-learning model: Convergence and applications. In Proceedings of the 13th International Conference on Machine Learning (pp. 310–318).

13. Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. IEEE Transactions on Systems, Man, and Cybernetics, Part C, 38(2), 156–172.

14. Sheng, M., Saif, L., & Choo, K.-K. R. (2022). Blockchain for secure and trustworthy capacity sharing in Industry 4.0. IEEE Transactions on Engineering Management, 69(4), 1230–1244.

15. Wang, F., Xiong, Z., Niyato, D., & Wang, P. (2021). Decentralized federated learning over mobile edge computing: A survey. IEEE Communications Surveys & Tutorials, 23(2), 1172–1205.

16. Ding, Y., Li, F., & Chen, X. (2022). Fairness-aware multi-agent reinforcement learning for distributed resource allocation. IEEE Transactions on Neural Networks and Learning Systems, 33(9), 4201–4215.

17. Feng, Q., Shanthikumar, J. G., & Wu, X. (2022). Coopetition and capacity sharing in manufacturing networks. Manufacturing & Service Operations Management, 24(2), 899–917.

18. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

19. Boutilier, C. (1996). Planning, learning and coordination in multiagent decision processes. In Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge (pp. 195–210).

20. Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. Science, 347(6221), 509–514.