

Multimodal Action-Conditioned Digital Twins for Adaptive Human–Robot Collaboration

Gaojiang Zhu

School of Computing, Clemson University, Clemson, SC, USA.
contactgao@clemson.edu

Derraen D. Park

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
park179@unh.edu

Pierra Lewrence

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
pierrework@unr.edu

Abstract

The increasing deployment of collaborative robots in manufacturing, logistics, and service sectors demands systems that can dynamically interpret human intent and adapt to fluid task environments. This paper proposes a conceptual architecture for multimodal action-conditioned digital twins as a unifying paradigm for adaptive human–robot collaboration. By integrating high-fidelity virtual representations with streaming multimodal sensory data and action-conditional prediction models, digital twins can continuously mirror the physical workspace, anticipate operator behaviors, and orchestrate safe, efficient joint activity. The discussion emphasizes system-level trade-offs including architectural decomposition between edge and cloud, latency-aware synchronization, cross-modal fusion, and the placement of learning algorithms within the twin hierarchy. A core argument is that action-conditioned world models, which learn to transition virtual states in response to candidate actions, offer a principled foundation for proactive adaptation. The paper further examines governance mechanisms needed to ensure robustness against sensor perturbation and adversarial inputs, fairness in adaptive allocation of robotic assistance, and the sustainability of large-scale twin infrastructures over operational lifetimes. Through an integrative lens that spans perception pipelines, control logic, deployment topologies, and regulatory frameworks, we delineate research trajectories toward trustworthy, scalable socio-technical systems. The analysis concludes that carefully designed multimodal action-conditioned digital twins can reconcile real-time responsiveness with long-term operational intelligence, but only if structural choices regarding data stewardship, model interpretability, and human agency are embedded from the earliest design stages. This perspective contributes to the emerging discourse on intelligent infrastructure for human-centric automation.

Keywords

Digital Twin, Human–Robot Collaboration, Multimodal Perception, Action-Conditioned Models, Adaptive Systems, World Models, Edge Computing, Fairness, Governance.

1. Introduction

Contemporary industrial and service environments are witnessing a rapid proliferation of robots operating in close physical proximity to human workers. Unlike caged automation that functions behind safety barriers, collaborative robots must engage in fluid interaction, interpreting subtle nonverbal cues, adjusting trajectories in response to operator motion, and dynamically sharing task responsibilities. This transition challenges conventional control architectures that rely on scripted sequences and rigid environment models. The digital twin paradigm, originally conceived for product lifecycle management and virtual commissioning, offers a compelling foundation for addressing these challenges. By maintaining a dynamic virtual counterpart of the physical workcell, a digital twin can integrate real-time sensor streams, track the evolving state of humans and machines, and simulate possible futures to guide joint action. Early conceptualizations of digital twins emphasized high-fidelity mirroring of physical assets to detect deviations and predict maintenance needs [1], while subsequent work extended the concept to full production systems, proposing frameworks where simulation, monitoring, and optimization are continuously coupled [2]. Despite these advances, the full potential of digital twins for adaptive human–robot collaboration remains underexplored, particularly regarding the incorporation of action-conditioned prediction and the governance of multimodal data flows.

A central limitation of many existing twin architectures is their predominantly descriptive nature: they reflect what is happening but lack the capacity to proactively generate and evaluate candidate courses of action in real time. Human–robot collaboration demands anticipatory reasoning that can answer the question, “What will happen if the robot moves left, or if the operator reaches toward a component?” Answering this question requires a world model that conditions its predictions on prospective actions. Concurrently, the richness of human behavior necessitates multimodal sensing that combines vision, depth, auditory signals, force-torque feedback, and possibly wearable biosignals. The integration of such multimodal streams into an action-conditioned digital twin raises profound system-level questions about architectural decomposition, synchronization latency, dataset bias, and algorithmic transparency. This paper develops a conceptual architecture for multimodal action-conditioned digital twins tailored to adaptive collaboration, systematically examining the structural trade-offs, infrastructure requirements, governance frameworks, and policy implications that arise when such systems are scaled beyond laboratory prototypes. Throughout, we emphasize that technical choices in module placement, fusion strategy, and adaptation logic carry direct consequences for safety, fairness, and human trust, and that these socio-technical dimensions must be treated as first-order design constraints.

2. Foundations and Related Work

The proposed architecture draws on several converging research streams. The field of human–robot interaction has long recognized that effective collaboration depends on mutual understanding and shared control. Foundational surveys established core principles of human–robot communication, highlighting the importance of transparency in robotic intent and the need for adaptive autonomy [3]. More recent reviews have catalogued progress in physical human–robot collaboration, underscoring key advances in compliant actuation, real-time motion planning, and intention estimation [4]. These works make clear that purely reactive strategies are insufficient for pro-locutive tasks where anticipation of partner actions can dramatically improve fluency and safety.

A parallel development concerns the use of machine learning to enable robots to acquire collaborative behaviors from experience and demonstration. Reinforcement learning has

emerged as a powerful paradigm for training policies that optimize long-term task performance under uncertainty [5]. When such policies are conditioned on inferred human goals, they can adapt robot behavior to changing task contexts. However, learning directly in physical environments is costly and risky, motivating the use of learned world models that can simulate environment dynamics and provide a safe sandbox for planning. This idea has gained significant traction through work on video prediction models that generate future frames conditioned on action sequences [7] and through the concept of recurrent world models that encode latent dynamics and enable policy search entirely within compact latent spaces [8]. The fusion of action-conditional prediction with multimodal sensing introduces additional complexity: the world model must not only predict visual appearances but also anticipate auditory events, haptic contacts, and shifts in operator posture that are registered through distinct sensor modalities.

Multimodal fusion has been extensively studied in the context of affective computing and human–computer interaction, with comprehensive reviews documenting techniques for combining heterogeneous data streams to improve recognition accuracy and robustness [6]. Extending these methods to the streaming, time-critical domain of human–robot collaboration requires architectural decisions about where and how fusion occurs: at the sensor level, at the feature level, or at the level of abstract semantic representations. More recently, the concept of interactive world models has been expanded through action-aware memory mechanisms that enable a model not only to explore a static environment but to actively simulate the consequences of interventions [9]. These advances provide a conceptual bridge from passive digital twins to active, action-conditioned counterparts that can propose and test candidate robot behaviors before execution. Our work situates these technical threads within a system-of-systems perspective, addressing how such components can be integrated into a coherent twin infrastructure that respects operational constraints and governance requirements.

3. Multimodal Action-Conditioned Digital Twin Architecture

We conceive the multimodal action-conditioned digital twin as a layered system spanning sensing, state estimation, world modeling, action generation, and execution monitoring. At the lowest layer, a heterogeneous sensor suite continuously captures visual imagery, depth maps, microphone arrays, force-torque measurements, and, in some deployments, wearable inertial or electro-physiological signals. These raw streams are time-stamped and transmitted to a spatial-temporal fusion engine that constructs a coherent representation of the workcell state, including rigid-body poses of robots and manipulable objects, skeletal poses of human operators, and bounding regions of occupancy and contact. The fusion engine must resolve temporal misalignments, occlusions, and varying sensor noise profiles, drawing on established multisensor data fusion techniques that operate across levels of abstraction [13].

Above the fusion layer resides the core action-conditioned world model. This module learns to predict future multimodal states given the current fused representation and a candidate sequence of robot actions. Architecturally, the world model may be implemented as a recurrent state-space model, a transformer over tokenized sensorimotor sequences, or a hybrid physics-informed neural network. The critical design choice concerns the granularity at which actions are represented. For fine-grained trajectory adaptation, continuous control actions must be conditioned; for task-level reasoning, discrete symbolic actions such as pick, place, or handover may suffice. The world model continuously receives updates from the fusion layer, enabling it to track changes in human behavior and environmental configuration. Crucially, it

also provides queries for counterfactual “what-if” simulations that can be evaluated in accelerated virtual time, generating a set of predicted outcomes that inform action selection.

The twin’s structural trade-off centers on the division of labor between edge and cloud resources. Edge nodes near the robotic cell can host lightweight, low-latency components for fast state estimation and reactive safety monitoring, while more computationally intensive world model training, long-horizon planning, and fleet-wide learning can be offloaded to cloud or fog infrastructure. This decoupling aligns with broader trends in edge computing [10] and industrial IoT architectures [11], but introduces tension between model synchronization and staleness. When a world model trained in the cloud is deployed at the edge for real-time inference, discrepancies between the training distribution and the live sensor stream can degrade prediction quality and lead to unsafe actions. Maintaining coherence demands continuous data pipelines, online fine-tuning strategies, and versioned model management that are themselves complex systems engineering challenges. The digital twin approach offers a natural framework for managing these multi-tier deployments by maintaining a unified logical state that is physically partitioned across compute tiers, with well-defined interfaces for state access and model updates.

The architectural configuration also shapes how fairness and safety properties are enforced. Embedding rule-based safety monitors at the edge layer that operate on raw sensor data ensures low-latency intervention independent of the world model’s predictions, providing a defense-in-depth safety strategy. Simultaneously, the world model’s future predictions can be audited for disparate performance across different operator demographics or task variations, but this requires logging and governance mechanisms that span edge and cloud, adding regulatory complexity.

4. Adaptivity and Collaboration Mechanisms

Adaptivity in human–robot collaboration goes beyond precomputed contingency plans; it requires the digital twin to dynamically re-rank potential robot actions based on evolving estimates of human intent, fatigue, and task priority. The action-conditioned world model serves as the engine for this continual replanning. By generating multiple rollouts conditioned on different robot action sequences, the twin can compare predicted outcomes against a multi-objective cost function that encodes task efficiency, human ergonomic risk, and social acceptability. The best action can then be selected through model-predictive control or policy networks trained in the world model’s simulated environment.

Learning collaborative behaviors from demonstration provides a rich source of prior knowledge that can bootstrap the world model and the action selection policy. Techniques that infer force, impedance, and trajectory profiles from human-robot dyad demonstrations enable the robot to internalize compliant interaction patterns that are difficult to specify analytically [16]. When such learned behaviors are embedded within the digital twin, they can be adapted online by conditioning on the current state and the inferred partner strategy. This creates a closed loop in which the twin not only predicts the human’s future motion but also anticipates how the human will react to the robot’s own actions, embodying a mutual adaptation that is characteristic of fluent collaboration.

Trust is a critical moderating variable in this adaptive loop. If the digital twin’s predictions are frequently inaccurate or if the robot behaves in ways that seem erratic to the human, trust erodes and collaboration breaks down. Research on automation trust has shown that transparency about the system’s reasoning, appropriate reliance, and gradual calibration

through experience are essential for maintaining operator confidence [17]. A multimodal action-conditioned twin can contribute to trust by offering explainable projections of intended actions, such as augmented reality overlays that show the robot’s planned trajectory and the twin’s confidence bounds. However, the very complexity of the multimodal world model may impede explainability; thus, architectural choices that favor modular, interpretable representations—such as explicit object-centric state encodings—can support both accuracy and human understandability.

An underappreciated dimension of adaptivity is the fair allocation of robotic assistance across workers. If the action selection policy is optimized solely for aggregate throughput, it may consistently prioritize faster or more experienced operators, leaving others with less support and reinforcing skill disparities. Embedding fairness constraints directly into the twin’s planning objectives requires operational definitions of equitable assistance that balance individualized adaptivity with group-level parity. This implicates both the training data used to build the world model, which must capture diverse operator behaviors, and the runtime optimization that must reconcile conflicting objectives in real time.

5. Infrastructure and Deployment Considerations

Deploying multimodal action-conditioned digital twins at production scale demands a robust infrastructure for data ingestion, model lifecycle management, and cross-site transfer. The continuous stream of high-bandwidth sensor data generated by even a single collaborative cell—combining multiple cameras, depth sensors, microphones, and haptic interfaces—can overwhelm local network links if not subjected to intelligent filtering and compression. Edge gateways must perform early sensor fusion and feature extraction to reduce volume before transmission to aggregating services. This requirement aligns with the broader vision of edge-centric architectures that push computation close to data sources to satisfy latency and bandwidth constraints [14].

Digital twin synchronization across geographically distributed sites raises further challenges. When a fleet of robotic workcells is operated by a single organization, there is a compelling economic incentive to aggregate de-identified behavioral data in a centralized cloud twin to improve world model accuracy through federated or centralized learning. However, data transfer regulations, intellectual property concerns, and the sheer cost of cloud egress impose limits on how much raw information can be centralized. A practical infrastructure must therefore support hybrid learning topologies in which local twins train on-site and periodically share model weights or gradient updates rather than raw sensor streams. This federated paradigm preserves data locality while enabling fleet-wide improvements, but it introduces non-trivial consistency management because local twins diverge due to site-specific drift in sensor calibration, lighting, and operator population. Dealing with such drift requires versioned model repositories, continual validation against reference tasks, and rollback mechanisms that are rare in today’s production robotic systems but are well-understood in large-scale software operations.

The integration of digital twin instances with manufacturing execution systems and enterprise resource planning platforms is a further deployment dimension that affects the granularity and timeliness of action-conditioned adaptation. If the twin accesses production schedules, inventory levels, and quality metrics, it can modulate collaboration strategies to prioritize throughput, quality inspection, or worker training depending on broader factory objectives. Implementing such integration demands standardized APIs and semantic interoperability layers that can translate high-level business goals into cost function weights understood by the

world model's planner. The absence of widely adopted standards for digital twin interfaces remains a significant barrier, although recent research on reference architectures for digital twin-driven manufacturing provides foundational guidance [12] and specific case studies demonstrate feasibility in robotic assembly cells [15].

6. Robustness, Fairness, and Operational Governance

Any system that conditions robot actions on multimodal predictions must be robust to sensor perturbations, adversarial inputs, and domain shifts. Multimodal pipelines based on deep neural networks are known to be vulnerable to carefully crafted perturbations that are imperceptible to humans but cause severe mispredictions. In a collaborative setting, an adversary could project a subtle pattern into the workspace or inject spoofed sensor readings that lead the world model to generate dangerous action plans. Evaluating and hardening such models against adversarial threats is an active area of research, and techniques originally developed for image classification are being extended to multimodal and sequential prediction tasks [20]. Robustness must be designed into the twin at multiple levels: through sensor redundancy and consistency checks at the fusion layer, through adversarial training of the prediction models, and through architectural safeguards that limit the authority of model outputs when confidence is low.

Fairness in adaptive collaboration is a multi-faceted governance concern. As noted, unequal distribution of robotic assistance can emerge inadvertently from optimization criteria that emphasize global efficiency. Beyond allocation fairness, the world model itself may exhibit performance disparities across demographic groups if its training data over-represents certain operator characteristics. The growing literature on fairness in machine learning provides formal frameworks for measuring and mitigating such disparities, including notions of individual fairness, demographic parity, and equality of opportunity [18, 19]. Translating these concepts from classification and ranking domains to sequential decision-making in physical human–robot interaction is non-trivial, because fairness must be evaluated over extended collaborative episodes and must account for the safety-critical nature of the domain. A promising approach is to integrate fairness constraints directly into the multi-objective optimization at the planning layer, ensuring that no operator is systematically disadvantaged. However, this raises the question of who defines the fair outcome and how trade-offs between fairness and productivity are negotiated—a socio-political decision that technical systems alone cannot resolve.

Operational governance frameworks must also address continuous monitoring, incident logging, and accountability. When a digital twin actively influences robot behavior in a shared workspace, any injury or property damage requires traceability to understand whether the root cause was sensor noise, a model error, a planner misconfiguration, or a human override. This necessitates tamper-proof logging of all sensory inputs, model predictions, and selected actions, akin to flight data recorders in aviation. In the European regulatory context, such requirements intersect with the proposed Artificial Intelligence Act, which classifies safety-critical AI applications in industrial settings as high-risk and imposes obligations for transparency, human oversight, and robustness. Established safety standards for collaborative robots, particularly ISO/TS 15066, provide a baseline for mechanical and control safety but do not yet fully address the risks introduced by learning-enabled components [21]. Bridging this gap requires new conformance assessment procedures that can evaluate the behavior of action-conditioned twins under a representative range of operational scenarios.

7. Sociotechnical Integration and Policy Implications

The introduction of multimodal action-conditioned digital twins into workplaces inevitably reshapes the balance of agency between human workers and automated systems. While the technical architecture aims to augment human capabilities and adapt to individual needs, the lived experience of operators may be one of surveillance and algorithmic management if the twin is perceived as a tool for performance monitoring rather than assistance. Addressing this requires deliberate sociotechnical design that involves workers in shaping the twin's behaviors, defines clear boundaries on what data are collected and how they are used, and ensures that workers can override or modify the twin's suggestions without penalty. An ethical framework for AI that emphasizes beneficence, non-maleficence, autonomy, justice, and explicability provides a useful lens for evaluating such systems [22]. Applying this framework, a well-governed digital twin should enhance worker autonomy by offering information and suggestions that the worker can accept or reject, rather than imposing unilateral robot actions. It should also be designed to avoid harmful physical and psychological consequences, such as excessive performance pressure induced by continuous visibility of efficiency metrics.

Policy implications extend to data governance, liability assignment, and workforce development. As digital twins accumulate granular data about human motor behavior, voice patterns, and possibly physiological states, they implicate data protection regulations such as the General Data Protection Regulation in Europe. The action-conditioned nature of these systems raises additional questions about whether predictive models of human behavior constitute personal data and what rights individuals have to access or contest such models. Liability frameworks must evolve to accommodate the distributed nature of twin-based robot control, where errors may arise from a combination of edge hardware, cloud-trained models, and third-party software components. Clarifying the responsibilities of system integrators, model developers, and deploying organizations is essential for industry adoption. Finally, workforce policies must support upskilling initiatives that enable operators to collaborate effectively with digital twin-enhanced robotic systems, transforming roles from manual execution to supervisory control and co-creation.

8. Sustainability and Long-Term Evolution

Beyond immediate deployment, the sustainability of multimodal action-conditioned digital twin infrastructures warrants careful attention. The computational demands of continuously training and updating world models on high-dimensional multimodal data consume significant energy, and the carbon footprint of large-scale cloud-based model training has become a concern across the AI community. System designers should therefore consider trade-offs between model accuracy and resource consumption, and explore techniques such as model compression, sparse architectures, and event-driven computation to reduce the operational energy profile of twins. Smart manufacturing principles that emphasize resource efficiency and circular economy goals provide a broader context for evaluating the life cycle impact of these digital infrastructures [23].

Long-term evolution of digital twins presents another sustainability challenge. Workspaces change, new products are introduced, and operator populations evolve over time, causing distributional drift that degrades world model performance unless the twin continuously adapts. This calls for lifelong learning mechanisms that can incorporate new data without catastrophically forgetting previously acquired collaborative skills. Architecturally, this may require a separation between stable low-level sensorimotor priors and rapidly adaptable high-level task models, together with mechanisms for safe exploration that prevent the twin from

suggesting unvalidated actions during learning phases. The governance implications of continuous learning are significant: if a twin's behavior changes gradually over weeks, it may drift out of its certified safety envelope without a discrete upgrade event that triggers reassessment. Regulatory frameworks must therefore move toward performance-based continuous assurance rather than static pre-deployment certification, a shift that is already under discussion in aviation and medical device regulation and could offer a model for collaborative robotics.

9. Conclusion

Multimodal action-conditioned digital twins represent a promising integrative framework for advancing adaptive human–robot collaboration beyond reactive control and pre-scripted interaction patterns. By fusing real-time multimodal sensor data with predictive world models that simulate the consequences of candidate actions, these systems can proactively shape robot behavior to match operator intent, task context, and ergonomic constraints. This paper has offered a system-level analysis of the architectural, infrastructural, governance, and policy dimensions that must be addressed to translate this vision into safe, fair, and sustainable practice. The key structural tensions—between edge autonomy and cloud intelligence, between model accuracy and computational sustainability, between individualized adaptation and group fairness—do not admit universal resolutions but must be negotiated within the specific constraints of each deployment context. What remains clear is that the technical design of digital twins cannot be decoupled from the human values they encode. As researchers and practitioners develop the next generation of collaborative systems, embedding transparency, worker agency, and rigorous safety assurance into the twin's architecture from the outset will be essential for earning trust and realizing the full societal benefits of human–robot collaboration.

References

1. Grieves, M., & Vickers, J. (2017). Digital twin: mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer.
2. 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.
3. Sheridan, T. B. (2016). Human–robot interaction: status and challenges. *Human Factors*, 58(4), 525–532.
4. Ajoudani, A., Zanchettin, A. M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., & Khatib, O. (2018). Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5), 957–975.
5. Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: a survey. *The International Journal of Robotics Research*, 32(11), 1238–1274.
6. Poria, S., Cambria, E., Bajpai, R., & Hussain, A. (2017). A review of affective computing: from unimodal analysis to multimodal fusion. *Information Fusion*, 37, 98–125.
7. Oh, J., Guo, X., Lee, H., Lewis, R., & Singh, S. (2015). Action-conditional video prediction using deep networks in Atari games. In *Advances in Neural Information Processing Systems* 28 (pp. 2863–2871).

8. Ha, D., & Schmidhuber, J. (2018). Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems* 31 (pp. 2450–2462).
9. Xiong, Z., Song, Y., Kang, H., Yan, Q., Jiang, L., Yang, J., ... & Jacobs, N. (2026). ActWorld: From Explorable to Interactive World Model via Action-Aware Memory. *arXiv preprint arXiv:2606.17730*.
10. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
11. Xu, L. D., Xu, E. L., & Li, L. (2018). Industry 4.0: state of the art and future trends. *International Journal of Production Research*, 56(8), 2941–2962.
12. Lu, Y., Liu, C., Wang, K. I.-K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61, 101837.
13. Lahat, D., Adali, T., & Jutten, C. (2015). Multimodal data fusion: an overview of methods, challenges, and prospects. *Proceedings of the IEEE*, 103(9), 1449–1477.
14. Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39.
15. Nikolakis, N., Alexopoulos, K., Xanthakis, E., & Chrysosolouris, G. (2019). A digital twin-based approach for designing and multi-objective optimization of a robotic manufacturing system. *Procedia CIRP*, 79, 295–300.
16. Rozo, L., Calinon, S., Caldwell, D. G., Jimenez, P., & Torras, C. (2016). Learning physical collaborative robot behaviors from human demonstrations. *IEEE Transactions on Robotics*, 32(3), 513–527.
17. Lee, J. D., & See, K. A. (2004). Trust in automation: designing for appropriate reliance. *Human Factors*, 46(1), 50–80.
18. 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), 1–35.
19. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (pp. 214–226). ACM.
20. Carlini, N., & Wagner, D. (2017). Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy* (pp. 39–57). IEEE.
21. Matthias, B. (2015). ISO/TS 15066 – The new technical specification for collaborative robots. *IEEE Robotics & Automation Magazine*, 22(1), 28–33.
22. 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.
23. Kusiak, A. (2018). Smart manufacturing. *International Journal of Production Research*, 56(1–2), 508–517.