2025 NSF/ASME Student Design Essay Competition

Theme: Challenges in the Design of Complex Systems 2025

Essay Title: **Preparing a High-Tech Manufacturing Enterprise for 2040**

Author Details:

Vispi Nevile Karkaria

Northwestern University,

Evanston, IL: 60208

Mentors:

Professor Wei Chen

Northwestern University

Category: Graduate Level

Challenges in the Design of Complex Systems: Preparing a High-Tech Manufacturing Enterprise for 2040

I. Introduction

In an era of rapid technological change, manufacturing industry face unprecedented complexity in system design. This essay is written from the perspective of a consultant advising a high-tech global manufacturing enterprise in 2025 with ambitions to become a world leader in design and production by 2040. To reach this goal, the company must master the design of complex systems, leveraging emerging technologies to outpace global competition. The integrated use of artificial intelligence (AI), advanced optimization techniques, and digital twin frameworks will be central to meeting these challenges. Recent research emphasizes that combining AI and digital twins can transform manufacturing, enabling real-time monitoring, predictive decision-making, and dynamic optimization. However, realizing this vision involves overcoming significant hurdles in data management, model scalability, uncertainty, and integration of design and control.

This essay first defines the envisioned characteristics of a successful high-tech manufacturing enterprise in 2040. We then identify key challenges that must be addressed to achieve those characteristics. We critically analyze how integrated AI, optimization algorithms, and digital twin technology can provide solutions. In particular, we focus on: (1) real-time optimization and AI-based decision making, (2) scalability of models using foundation-model approaches, (3) federated learning and edge-based data sharing across global operations, (4) rigorous uncertainty quantification (epistemic and aleatoric), (5) hybrid offline/online optimization workflows, and (6) control co-design of processes and materials. We present a structured discussion with sections on background, challenges, and proposed solutions, and conclude with concrete steps the company should begin implementing today. Throughout, we cite and critique state-of-the-art research to support our recommendations.

II. Background: Vision of a 2040 Manufacturing Leader

The schematic below (Figure 1) captures the three key capabilities of a 2040 manufacturing leader as three tightly integrated pillars—Digital Model Training & Update, Uncertainty Quantification, and Decision Making & Optimization—all centered around an interconnected digital-twin ecosystem:

1. **Digital Model Training & Update:** Every critical process and product is first encoded into an AI-based surrogate model via multimodal training (physics simulations, process logs, thermography, high-speed video) and then continuously calibrated with live sensor data to reduce model-calibration error and sensor noise [1]. These updated surrogates form the foundation models that power all higher-level decisions (Characteristic 1: Digital Twin Framework; Characteristic 2: AI-Driven Decision Making; Characteristic 3: Federated Learning & Edge Computing) [2].
2. **Uncertainty Quantification:** The digital twin explicitly quantifies numerical, prediction, and model uncertainties—distinguishing aleatoric (sensor noise, environmental, material variability) from epistemic (model limitations)—and embeds these confidence intervals into every recommendation [3]. This resilience

layer ensures robust decision making under unknown conditions (Characteristic 4: Resilience through Uncertainty Quantification) [4].

3. **Decision Making & Optimization:** Offline optimizers (Bayesian and linear programming) leverage historical data and high-fidelity twins to co-design materials, geometry, and process parameters—creating optimal baselines—while online solvers (dynamic programming, reinforcement learning, MPC) execute real-time adjustments on the physical system [5]. This continuous offline–online loop shrinks development cycles and drives sustainable, multi-objective trade-offs (Characteristic 5: Integrated Offline Planning & Online Control; Characteristic 6: Co-Design of Products & Processes; Characteristic 7: Sustainability & Resource Efficiency).

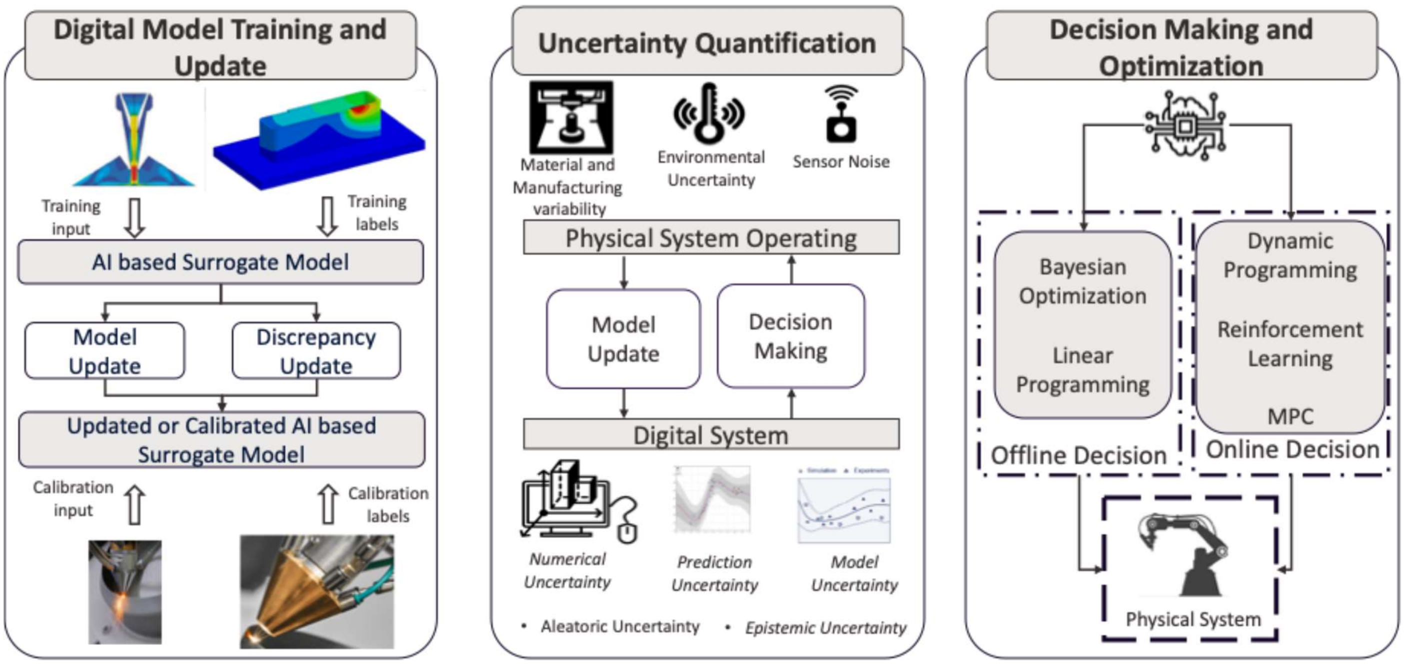

Figure 1. An integrated workflow in which AI-based surrogate models are trained and iteratively calibrated (left), uncertainty quantification distinguishes aleatoric and epistemic errors to inform robust model updates and decisions (center), and a dual offline–online optimization loop (Bayesian/linear programming offline, RL/MPC online) drives adaptive control of the physical system (right).

Together, these three pillars form a self-reinforcing, adaptive loop in which updated digital models inform quantified-uncertainty decisions, which in turn feed optimized actions back into both the physical system and the AI surrogates—empowering the enterprise to learn, adapt, and innovate at scale in 2040's complex global landscape.

Complexity in Modern Manufacturing Systems: The complexity of a high-tech manufacturing enterprise arises from multiple sources, including the heterogeneity of processes, the scale of data, and the coupling of decisions across domains (mechanical, computational, organizational) [6]. Digital twin frameworks, while promising, must integrate models at different scales – from detailed physics of materials to the macro-level production line and supply chain [7]. Ensuring consistency and accuracy across these scales is non-trivial, as models must be both detailed and computationally tractable. Data streams come from diverse Industrial Internet of Things (IIoT) sensors and databases, often in incompatible formats and varying quality [8]. Data quality and consistency are *foundational* challenges: effective optimization and AI rely on reliable, clean data, yet in practice data may be siloed or noisy across different factories and IT systems. In addition, manufacturing technologies evolve rapidly (e.g. new machines, materials, processes), so the enterprise's digital infrastructure must be flexible and updatable. Algorithms and models that were effective can become obsolete as new processes emerge. The company's complex system design must therefore accommodate continuous technology insertion and system reconfiguration [9].

In the following sections, we outline the most important challenges impeding the realization of the above vision and discuss how they can be addressed. While the focus is on technical solutions (AI, optimization, digital twins), we acknowledge that organizational and human factors (e.g. workforce training, change management) are also critical. Our analysis emphasizes research-backed strategies that this enterprise should adopt to systematically design and operate its complex system of the future.

III. Key Challenges in Designing a Complex Manufacturing System

Despite the clear potential of AI and digital twins, several challenges must be overcome to design an integrated manufacturing system that embodies the 2040 vision:

Challenge 1: AI-Driven Optimization with Digital Twins: Manufacturing lines must adjust instantly to variability (machine wear, material changes, failures). Digital twins need to do more than mirror operations—they must compute and enact optimal control moves (e.g., tuning parameters or rerouting workflows) within milliseconds to minutes, balancing safety and performance [10]. Achieving this demands a tight integration of offline planning (high-fidelity simulations) and online adaptation (model predictive or feedback control), plus AI that is both fast and interpretable so engineers can trust its on-the-fly recommendations.

Challenge 2: Scalability & Model Generalization (Foundation Models): Enterprises generate vast, heterogeneous data across products and sites; bespoke AI for each scenario is infeasible [11]. Foundation models—large, pretrained networks capturing general physics and manufacturing patterns—promise broad applicability but require massive, diverse training data (simulations, experiments, operations) and huge compute resources. Integrating and updating these monolithic models, while keeping them interpretable and respecting IP constraints, remains an open problem [12].

Challenge 3: Distributed Data & Collaboration (Federated Learning): Global factories produce siloed data that cannot be centralized for analysis [13]. Federated learning lets each site train local models on its own data and share only model updates, creating a global model without exposing raw data. Practical hurdles include network latency, heterogeneous equipment/processes, privacy/IP protection, and ensuring each site has sufficient edge compute capacity. Designing secure, robust protocols and adaptable aggregation algorithms is key to leveraging the enterprise's full global footprint [14].

Challenge 4: Uncertainty & Trustworthiness: Automated systems must quantify both epistemic uncertainty (model gaps) and aleatoric uncertainty (inherent randomness). Ignoring these leads to overconfident or unsafe decisions [15]. Embedding Bayesian methods, ensembles, and Monte Carlo simulations into digital twins provides confidence intervals, while real-time UQ dashboards and fallback strategies (e.g., human override) ensure decisions remain reliable. Building rigorous verification, validation, and uncertainty-aware controls is essential to earn user trust.

Challenge 5: Integrating Design & Production (Control Co-Design): Sequential product-then-process design yields suboptimal solutions. True co-design simultaneously optimizes materials, geometry, and factory parameters, but this creates a huge, nonlinear search space and heavy simulation loads. Advances in surrogate modeling and multi-fidelity optimization help, yet industry practice remains sparse. Overcoming siloed workflows and adopting shared digital-twin platforms will let cross-functional teams explore joint design–process trade-offs rapidly, slashing development cycles and unlocking novel, manufacturable solutions.

IV. Proposed Solutions through AI, Optimization, and Digital Twins

To meet the challenges above, we propose a multi-faceted strategy centered on integrating AI, advanced optimization, and digital twin frameworks into the company's operations and design processes. The solutions are deeply interrelated – a digital twin acts as the convergence point for many of these technologies – and together they form an ecosystem enabling the 2040 vision. Below, we outline the key solution approaches corresponding to each challenge, citing current research and practical considerations for implementation.

Solution 1: AI-Driven Optimization with Digital Twins:

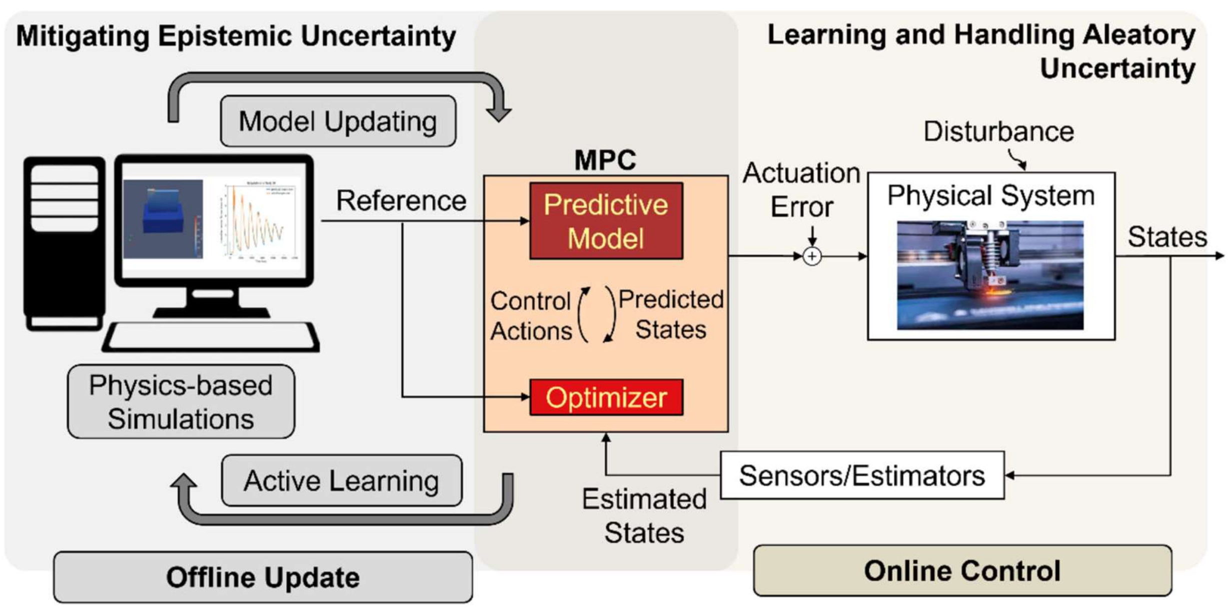

Figure 2. Offline physics-based simulations and active learning iteratively update the digital twin's predictive model to reduce epistemic uncertainty, while online model-predictive control uses real-time sensor feedback to counteract aleatoric disturbances and ensure robust system performance.

The Figure 2 illustrates a two-stage workflow in which a digital twin both refines its internal model and governs real-time operations to tackle epistemic and aleatoric uncertainties. In the offline update phase, high-fidelity, physics-based simulations generate reference trajectories and highlight areas of model inaccuracy [16]. These simulation outputs feed into an active-learning loop that periodically updates the twin's predictive model, thereby reducing epistemic uncertainty before deployment.

Once deployed, the online control phase employs model-predictive control (MPC) powered by the updated model: an optimizer computes control actions that drive the system along the desired reference [17]. These actions are executed on the physical system, where unpredictable disturbances and actuation errors inevitably introduce deviations. Real-time sensor measurements or state estimators capture the system's actual response, and this feedback closes the loop—allowing the MPC to adjust its commands instantly to counter aleatoric uncertainty. Any persistent discrepancies between predicted and measured behavior can then trigger the next offline update, completing a continuous cycle of model improvement and robust control [18].

Solution 2: Leveraging Foundation Models for Scalable AI. Addressing Challenge 2, the enterprise should build a unified foundation model trained on diverse manufacturing datasets—high-fidelity physics simulations, structured process logs, infrared thermography, high-speed melt-pool video, and microstructure imagery. Each modality is preprocessed via specialized encoders (e.g., autoencoder bottlenecks for simulation outputs, temporal convolutional layers for time-series signals, and convolutional backbones for images) [19]. The resulting embeddings feed into transformer or graph-neural-network layers that learn cross-modal correlations—such as how transient temperature spikes predict subsurface porosity and eventual surface defects. Training uses supervised losses (defect classification, wear-rate regression) alongside self-supervised objectives (masked reconstruction of sensor streams), enabling the model to internalize core manufacturing physics, from heat-transfer and phase changes to kinematic tolerances and vibration patterns [20].

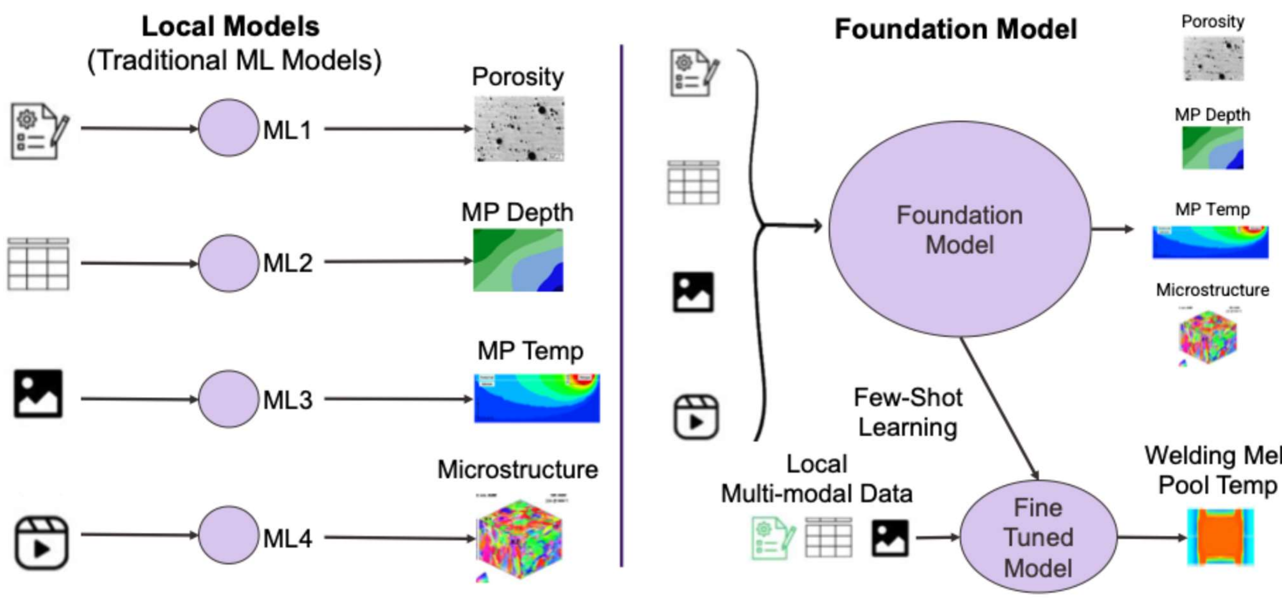

Figure 3. Diverse data sources—from physics-based simulation outputs and tabular process logs to thermal images and high-speed video—are encoded and fused into a central foundation model, which can then be fine-tuned on specific tasks to generate tailored predictions and control strategies.

Once pretrained, the foundation model can be fine-tuned on a new task—say, predictive maintenance for a freshly installed forging press—with as few as several dozen labeled examples. Techniques like low-rank adaptation (LoRA) and lightweight projection layers enable rapid convergence on limited local data [21]. This fine-tuned model is deployed within a robust MLOps pipeline featuring CI/CD for version control, automated retraining on edge-collected updates, and real-time dashboards tracking metrics (e.g., prediction error, falsealarm rates). Federated learning periodically aggregates encrypted local updates across plants, enhancing the global model without exposing raw data. To ensure safety and interpretability, rule-based verification layers enforce known material and process constraints, while explainability tools (integrated gradients, attention maps) highlight key features driving each AI recommendation—fostering engineer trust and enabling continuous, enterprise-wide learning and innovation [22].

Solution 3: Federated Learning and Edge-Based Collaboration. To tackle the distributed data challenge (Challenge 3), the company should implement a federated learning framework across its sites, combined with strong edge computing nodes (see Figure 4) [23]. In practice, this means establishing a central coordination server (cloud-based or at a main hub) and deploying federated learning software at each plant. Each site's digital twin runs AI model training or update routines locally on recent data (e.g., a week's worth of production logs), then periodically sends model parameters or gradients—encrypted to prevent eavesdropping—to the central server. The server aggregates these updates (for example, via FedAvg) into an improved global foundational model, which is then broadcast back to all sites so each benefits from others' experiences without ever sharing raw data [24].

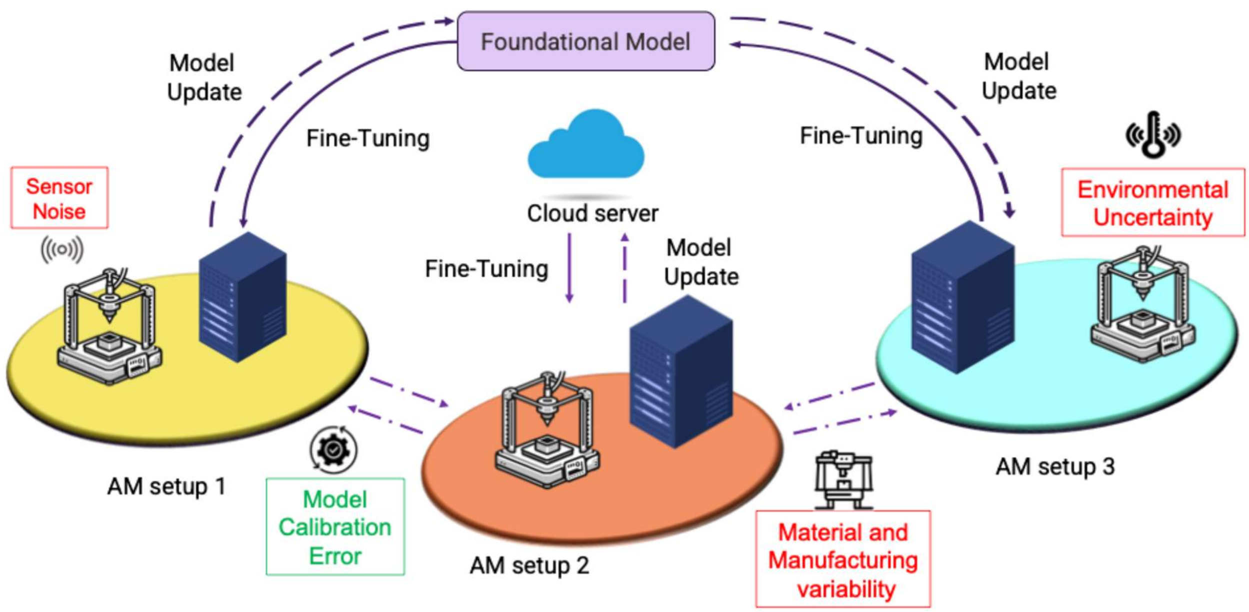

Figure 4. Edge-hosted digital twins at multiple additive-manufacturing sites fine-tune a shared foundation model via federated updates, enabling a global metatwin that learns from local sensor noise, model calibration errors, environmental uncertainty, and material variability without exchanging raw data.

The system must respect bandwidth and latency constraints: federated updates can occur during off-peak hours or at a manageable cadence (e.g., daily), while edge computing (Solution 1) handles latency-critical decisions in real time [25]. This synergy—local inference and immediate control at the edge, global intelligence improvements via asynchronous federated updates—ensures that even if connectivity is lost, each site functions autonomously on its last synced model, then seamlessly resynchronizes when online. Data privacy and security are inherently protected by keeping raw data on site; additional measures like encrypted updates and differential privacy can safeguard trade-secret production data. This approach maintains confidentiality and reduces network load while collectively optimizing processes across facilities.

A prime use case is predictive maintenance: similar machines across different plants can each train on local sensor noise and environmental variability, then contribute to a shared failure-prediction model that outperforms any single-site model. Early deployments in other sectors (for example, automotive fleets) have demonstrated that federated learning uncovers rare anomalies undetectable by isolated nodes. The enterprise should start by federating one or two critical models—such as quality prediction or tool-wear forecasting—gain operational experience, and then scale up. Over time, this network of site-level digital twins evolves into a “metatwin” for the entire enterprise: a higher-level digital twin that embodies learned knowledge from all facilities [26].

Solution 4: Robust Design via Uncertainty Quantification (UQ). To ensure the system handles uncertainty (Challenge 4), the enterprise must integrate UQ methods into both its design optimization and real-time operation. On the design side, this means performing stochastic optimization or robust optimization when planning processes and products [27]. Instead of optimizing for a single best-case scenario, the optimization should account for variability – yielding solutions that are optimal *on average* or minimize worst-case risk. Techniques like Monte Carlo simulation can evaluate how a candidate solution performs under many random trials (e.g., varying material properties, machine precision, environmental conditions). If a solution is too sensitive to variations, it would be penalized in the robust optimization objective [28]. This approach leads to choices that may not be extreme optima in any one scenario but deliver reliable performance across scenarios.

Digital twins provide an ideal platform for such simulations. The company’s digital twin models can be augmented with uncertainty models: for every input or parameter, an uncertainty distribution can be assigned (based on empirical data or expert judgment) [29]. Then large-scale simulations or even real-time Monte Carlo runs can be done to see outcome distributions. For instance, before implementing a new robotic assembly routine, the twin can simulate 1000 runs with slight random perturbations to gauge yield and identify any high-variance steps. If a certain joint in the robot often leads to misalignment only in certain conditions, that might be epistemic uncertainty (maybe the model of the robot isn't perfect in those conditions) which could be reduced by better calibration or adding a sensor; or it might be aleatoric (maybe random electrical noise) which you address by a more conservative speed setting.

For epistemic uncertainty reduction, the company should apply Bayesian updating and learning algorithms [30]. This can be as straightforward as continually comparing twin predictions with real outcomes and feeding back the error to adjust model parameters (a form of online learning for the twin). More formally, one can use Bayesian neural networks or Gaussian process models that naturally provide a measure of confidence with each prediction. As more data comes in, these models tighten their confidence intervals (reducing epistemic uncertainty). The enterprise could maintain uncertainty dashboards for critical metrics, which would visualize current confidence levels. If epistemic uncertainty for a key quality prediction is high, it flags that more exploration or data gathering is needed in that regime. Aleatoric uncertainty, on the other hand, informs the risk management strategies – if a particular process has inherent variability, the operations plan might include extra buffers or inspection steps to mitigate the risk of a bad outcome [31].

Importantly, when the AI system provides a recommendation, it should also output an *uncertainty estimate*. For example, an AI might suggest "Increase temperature by 5°C to improve yield, with an 85% confidence that this will not cause defects." This kind of information is invaluable for human experts to decide whether to trust the AI or not. Methods like quantile regression can directly predict certain quantiles of outcomes, giving insight into best-case and worst-case expectations. We recommend that by 2040 the company adopts a policy that no automated decision is made without an accompanying uncertainty analysis. This might involve embedding simulation-based certification: before a new AI control policy is uploaded to a production machine, it is tested in the twin under thousands of randomized scenarios to statistically validate its safety and efficacy (akin to virtual stress-testing).

From a cultural perspective, incorporating UQ will improve trust in AI among stakeholders (engineers, managers, regulators), because it demonstrates a cautious and scientific approach rather than a blind reliance on algorithms. Verification, validation, and uncertainty quantification (VVUQ) are *crucial components* in developing robust digital twins for complex manufacturing [32]. By systematically addressing VVUQ, the enterprise not only avoids costly surprises but can also gain competitive advantage by operating closer to performance limits safely, thanks to a better understanding of margins.

Solution 5: Integrated Co-Design of Products and Processes: To overcome the siloed design challenge (Challenge 5), the enterprise should adopt a co-design framework that brings product designers, materials scientists, and manufacturing engineers together around a shared digital model from day one [33]. Modern co-design platforms enable simultaneous, multi-domain simulation: for example, engineers can adjust alloy properties, part geometry, and machine parameters in one environment and immediately evaluate impacts on weight, strength, manufacturability, and cost. Extending the digital twin concept to a "design twin"—which simulates both product performance and its manufacturing—allows optimization algorithms to vary design and process variables in tandem, seeking global optima rather than settling for suboptimal sequential hand-offs [34].

Co-design merges offline and online optimization loops: offline simulations identify designs that perform well under ideal conditions, while online data from prototype builds feed back to refine the virtual design almost in real time [35]. This tight feedback loop, powered by the twin, drastically accelerates iteration and reduces reliance on costly physical prototypes. Multidisciplinary design optimization (MDO) techniques—already proven in aerospace—can be tailored for manufacturing to co-optimize part topology, material composition, and toolpaths, as demonstrated in additive-manufacturing pilots that achieved lighter, stronger parts with fewer build cycles. Implementing co-design also requires a unified data infrastructure—connecting CAD, CAM, and CAE tools via digital-thread standards—and cross-functional teams empowered to explore trade-offs rapidly. By embracing co-design, the company will simultaneously optimize materials, geometry, and processes, enabling agile responses to supply constraints or market shifts and securing a decisive competitive edge [36].

Integration of Solutions: While we have discussed each solution in isolation, it's important to stress that their real power comes when implemented together. For example, co-design (Solution 5) will benefit from uncertainty quantification (Solution 4) to ensure robust designs, and from foundation models (Solution 2) that provide quick estimates of performance for new combinations. Federated learning (Solution 3) will make the digital twins in each design center or factory smarter over time, which in turn improves real-time control (Solution 1) because the models those controls rely on are more accurate. In essence, we are advocating for a **holistic AI-driven system architecture** for the enterprise. Figure 1 (not included here) could be imagined to illustrate this synergy: a loop connecting design (co-design with offline optimization), operation (real-time control), global learning (federated updates), all underpinned by a digital twin network, with uncertainty analysis pervading each step. The more tightly these components are integrated, the closer the company gets to an autonomous, self-optimizing manufacturing ecosystem.

Crucially, human expertise remains in the loop at strategic points – AI is a tool to augment human decision-making and automate routine optimization, but domain experts will set goals, handle novel situations, and provide oversight. By implementing the above solutions, the company will effectively create a partnership between experienced engineers and powerful AI systems embodied in digital twins. This is how the most challenging aspects of complex system design can be managed.

V. Conclusion

Designing a leading-edge manufacturing system for 2040 is ambitious yet attainable by articulating a clear vision—an AI-empowered, agile, and resilient global enterprise—and then methodically addressing each obstacle. This essay identified five critical challenges—real-time AI optimization, scalable foundation models, federated learning across distributed sites, rigorous uncertainty quantification, and integrated product-process co-design—and offered concrete, research-backed strategies for each. Together, integrated AI, advanced optimization techniques, and digital twin frameworks provide the backbone for autonomous control, enterprise-wide knowledge sharing, and rapid virtual prototyping at a scale unprecedented just a decade ago.

To start this transformation today, the company should: initiate a Digital Twin pilot on a key production line to demonstrate AI-driven monitoring and control; invest in unified data infrastructure and governance, including edge/cloud platforms for federated learning; establish an AI Center of Excellence to develop tailored foundation models, explore reinforcement learning, and champion uncertainty quantification; form cross-functional design-manufacturing teams equipped with co-design tools to jointly optimize product and process; and mandate verification, validation, and uncertainty-quantification (V&V+UQ) for all AI models and simulations, training engineers to track epistemic and aleatoric uncertainty and prioritize data or sensor improvements accordingly. Executing these steps will build the foundational capabilities for a seamless digital-twin ecosystem, deeply integrated AI decision loops, and a co-design culture—positioning the enterprise to thrive as a world leader in complex system design and production by 2040.

References:

[1] Aydemir, H., Zengin, U., and Durak, U., "The Digital Twin Paradigm for Aircraft Review and Outlook," *AIAA Scitech 2020 Forum*, American Institute of Aeronautics and Astronautics.

[2] Ramu, S. P., Boopalan, P., Pham, Q.-V., Maddikunta, P. K. R., Huynh-The, T., Alazab, M., Nguyen, T. T., and Gadekallu, T. R., 2022, "Federated Learning Enabled Digital Twins for Smart Cities: Concepts, Recent Advances, and Future Directions," Sustain. Cities Soc., **79**, p. 103663.

[3] Battula, S., Alla, S. N., Ramana, E. V., Kiran Kumar, N., and Bhanu Murthy, S., 2024, "Uncertainty Quantification for Digital Twins in Smart Manufacturing and Robotics: A Review," J. Phys. Conf. Ser., **2837**(1), p. 012059.

[4] Karve, P. M., Guo, Y., Kapusuzoglu, B., Mahadevan, S., and Haile, M. A., 2020, "Digital Twin Approach for Damage-Tolerant Mission Planning under Uncertainty," Eng. Fract. Mech., **225**, p. 106766.

[5] Karkaria, V., Tsai, Y.-K., Chen, Y.-P., and Chen, W., "An Optimization-Centric Review on Integrating Artificial Intelligence and Digital Twin Technologies in Manufacturing," Eng. Optim., **0**(0), pp. 1–47.

[6] Wu, Z., and Liang, C., 2024, "A Review and Prospects of Manufacturing Process Knowledge Acquisition, Representation, and Application," Machines, **12**(6), p. 416.

[7] Böttjer, T., Tola, D., Kakavandi, F., Wewer, C. R., Ramanujan, D., Gomes, C., Larsen, P. G., and Iosifidis, A., 2023, "A Review of Unit Level Digital Twin Applications in the Manufacturing Industry," CIRP J. Manuf. Sci. Technol., **45**, pp. 162–189.
[8] Alabadi, M., Habbal, A., and Wei, X., 2022, "Industrial Internet of Things: Requirements, Architecture, Challenges, and Future Research Directions," IEEE Access, **10**, pp. 66374–66400.
[9] Leng, J., Guo, J., Xie, J., Zhou, X., Liu, A., Gu, X., Mourtzis, D., Qi, Q., Liu, Q., Shen, W., and Wang, L., 2024, "Review of Manufacturing System Design in the Interplay of Industry 4.0 and Industry 5.0 (Part I): Design Thinking and Modeling Methods," J. Manuf. Syst., **76**, pp. 158–187.
[10] Kumar, R., and Agrawal, N., 2024, "Shaping the Future of Industry: Understanding the Dynamics of Industrial Digital Twins," Comput. Ind. Eng., **191**, p. 110172.
[11] Singh, A., Jadhav, A., and Singh, P., 2024, "AI Applications in Production," *Industry 4.0, Smart Manufacturing, and Industrial Engineering*, CRC Press.
[12] Ren, L., Dong, J., Liu, S., Zhang, L., and Wang, L., 2024, "Embodied Intelligence Toward Future Smart Manufacturing in the Era of AI Foundation Model," IEEEASME Trans. Mechatron., pp. 1–11.
[13] Silbernagel, R., Wagner, C., Albers, A., Trapp, T.-U., and Lanza, G., 2021, "Data-Based Supply Chain Collaboration – Improving Product Quality in Global Production Networks by Sharing Information," Procedia CIRP, **104**, pp. 470–475.
[14] Tong, X., Hamzei, M., and Jafari, N., 2025, "Towards Secure and Efficient Data Aggregation in Blockchain-Driven IoT Environments: A Comprehensive and Systematic Study," Trans. Emerg. Telecommun. Technol., **36**(2), p. e70061.
[15] Ojha, J., Presacan, O., Goncalves Lind, P., Monteiro, E., and Yazidi, A., 2025, "Navigating Uncertainty: A User-Perspective Survey of Trustworthiness of AI in Healthcare," ACM Trans Comput Healthc.
[16] Wang, J., Li, Y., Gao, R. X., and Zhang, F., 2022, "Hybrid Physics-Based and Data-Driven Models for Smart Manufacturing: Modelling, Simulation, and Explainability," J. Manuf. Syst., **63**, pp. 381–391.
[17] Schwenzer, M., Ay, M., Bergs, T., and Abel, D., 2021, "Review on Model Predictive Control: An Engineering Perspective," Int. J. Adv. Manuf. Technol., **117**(5), pp. 1327–1349.
[18] Shang, C., Chen, W.-H., Stroock, A. D., and You, F., 2020, "Robust Model Predictive Control of Irrigation Systems With Active Uncertainty Learning and Data Analytics," IEEE Trans. Control Syst. Technol., **28**(4), pp. 1493–1504.
[19] Duan, J., Xiong, J., Li, Y., and Ding, W., 2024, "Deep Learning Based Multimodal Biomedical Data Fusion: An Overview and Comparative Review," Inf. Fusion, **112**, p. 102536.
[20] Joanes Agung', O., James, K., Samuel, K., and Evan, M., 2024, "Generative and Self-Supervised Ensemble Modeling for Multivariate Tool Wear Monitoring," Eng. Rep., **6**(6), p. e12788.
[21] Lin, Q., He, W., Zhang, qian, Peng, Z., Wu, Z., and Peng, L., 2024, "WGLora:Efficient Fine-Tuning Method Integrating Weights and Gradient Low-Rank Adaptation," *Proceedings of the 2024 12th International Conference on Communications and Broadband Networking*, Association for Computing Machinery, New York, NY, USA, pp. 108–114.
[22] Wang, J., Wu, Y., Li, M., Lin, X., Wu, J., and Li, C., 2020, "Interpretability Is a Kind of Safety: An Interpreter-Based Ensemble for Adversary Defense," *Proceedings of the 26th ACM SIGKDD*

International Conference on Knowledge Discovery & Data Mining, Association for Computing Machinery, New York, NY, USA, pp. 15–24.

[23] Nguyen, D. C., Ding, M., Pham, Q.-V., Pathirana, P. N., Le, L. B., Seneviratne, A., Li, J., Niyato, D., and Poor, H. V., 2021, "Federated Learning Meets Blockchain in Edge Computing: Opportunities and Challenges," IEEE Internet Things J., **8**(16), pp. 12806–12825.

[24] Zhuang, W., Chen, C., Li, J., Chen, C., Jin, Y., and Lyu, L., 2025, "When Foundation Model Meets Federated Learning: Motivations, Challenges, and Future Directions."

[25] Zhao, Z., Xia, J., Fan, L., Lei, X., Karagiannidis, G. K., and Nallanathan, A., 2022, "System Optimization of Federated Learning Networks With a Constrained Latency," IEEE Trans. Veh. Technol., **71**(1), pp. 1095–1100.

[26] Torres, J., San-Mateos, R., Lasarte, N., Mediavilla, A., Sagarna, M., and León, I., 2024, "Building Digital Twins to Overcome Digitalization Barriers for Automating Construction Site Management," Buildings, **14**(7), p. 2238.

[27] Lin, L., Bao, H., and Dinh, N., 2021, "Uncertainty Quantification and Software Risk Analysis for Digital Twins in the Nearly Autonomous Management and Control Systems: A Review," Ann. Nucl. Energy, **160**, p. 108362.

[28] Yazdani, D., Omidvar, M. N., Yazdani, D., Branke, J., Nguyen, T. T., Gandomi, A. H., Jin, Y., and Yao, X., 2024, "Robust Optimization Over Time: A Critical Review," IEEE Trans. Evol. Comput., **28**(5), pp. 1265–1285.

[29] Bárkányi, Á., Chován, T., Németh, S., and Abonyi, J., 2021, "Modelling for Digital Twins—Potential Role of Surrogate Models," Processes, **9**(3), p. 476.

[30] Hüllermeier, E., and Waegeman, W., 2021, "Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods," Mach. Learn., **110**(3), pp. 457–506.

[31] Li, Z., Jin, G., Yu, R., Chen, Z., Li, N., Han, W., Xiong, L., Leng, B., Hu, J., Kolmanovsky, I., and Filev, D., 2025, "A Survey of Reinforcement Learning-Based Motion Planning for Autonomous Driving: Lessons Learned from a Driving Task Perspective."

[32] Kaiblinger, A., and Woschank, M., 2022, "State of the Art and Future Directions of Digital Twins for Production Logistics: A Systematic Literature Review," Appl. Sci., **12**(2), p. 669.

[33] Meath, C., Karlovšek, J., Navarrete, C., Eales, M., and Hastings, P., 2022, "Co-Designing a Multi-Level Platform for Industry Level Transition to Circular Economy Principles: A Case Study of the Infrastructure CoLab," J. Clean. Prod., **347**, p. 131080.

[34] Khan, L. U., Han, Z., Saad, W., Hossain, E., Guizani, M., and Hong, C. S., 2022, "Digital Twin of Wireless Systems: Overview, Taxonomy, Challenges, and Opportunities," IEEE Commun. Surv. Tutor., **24**(4), pp. 2230–2254.

[35] Fails, J. A., Ratakonda, D. kumar, Koren, N., Elsayed-Ali, S., Bonsignore, E., and Yip, J., 2022, "Pushing Boundaries of Co-Design by Going Online: Lessons Learned and Reflections from Three Perspectives," Int. J. Child-Comput. Interact., **33**, p. 100476.

[36] Moenck, K., Koch, J., Rath, J.-E., Büsch, L., Gierecker, J., Kähler, F., Kalscheuer, F., Masuhr, C., Kipping, J., Prünte, P., Schoepflin, D., Eschen, H., Wulff, L. A., Rodeck, R., Wende, G., Gomse, M., and Schüppstuhl, T., 2025, "Industry 5.0 in Aircraft Production and MRO: Challenges and Opportunities," CEAS Aeronaut. J.

Appendix. Definitions of Key Technical Concepts

1. **Digital Twin**: A digital twin is a real-time virtual replica of a physical asset, process, or system that ingests live sensor data and simulation outputs to mirror its counterpart's behavior, enabling monitoring, analysis, and predictive control of manufacturing operations.
2. **Surrogate Model:** An AI-based surrogate model approximates complex physics-based simulations by learning input–output mappings (e.g., process parameters to product quality) from data, allowing rapid predictions and enabling iterative model updates and discrepancy corrections.
3. **Foundation Model**: A foundation model is a large-scale, pretrained neural network—often combining physics-based and data-driven components—that embeds general manufacturing knowledge (e.g., heat transfer, material behavior) and can be fine-tuned on limited new data for task-specific predictions.
4. **Federated Learning**: Federated learning is a decentralized training paradigm in which each site trains a local model on its own data and shares only encrypted model updates (e.g., gradients) with a central server for aggregation, preserving raw data privacy while building a global model.
5. **Edge Computing**: Edge computing refers to performing AI inference and model updates on hardware located physically close to the data source (e.g., on factory-floor controllers), minimizing latency and bandwidth use for real-time decision making.
6. **Aleatoric vs. Epistemic Uncertainty**
   - *Aleatoric uncertainty* arises from inherent randomness in processes or measurements (e.g., sensor noise, environmental fluctuations) and cannot be reduced by gathering more data.
   - *Epistemic uncertainty* reflects lack of knowledge or model inaccuracy (e.g., unmodeled physics, limited training data) and can be reduced via additional data collection or model refinement.
7. **Model Predictive Control (MPC)**: MPC is an optimization-based control strategy that uses a predictive model to compute a sequence of future control actions by solving a constrained optimization problem at each time step, then applies only the first action before repeating the process.
8. **Active Learning**: Active learning is a model-training technique in which the algorithm identifies and requests the most informative data points (e.g., boundary cases in simulations) to label or simulate, thereby improving model accuracy with fewer training samples.
9. **Multidisciplinary Design Optimization (MDO)**: MDO is an integrated optimization methodology that simultaneously considers multiple engineering disciplines (e.g., structural, thermal, manufacturing) and their interactions to identify globally optimal design solutions across competing objectives.
10. **Verification, Validation, and Uncertainty Quantification (VVUQ)**: VVUQ comprises:
    - *Verification*—ensuring the digital twin's numerical implementation is correct;
    - *Validation*—confirming the twin's outputs agree with real-world data;
    - *Uncertainty Quantification*—characterizing and propagating aleatoric and epistemic uncertainties to produce confidence bounds on predictions.