MIT Mobility Initiative 2025

Adam Grosser, Senior Fellow

Fall 2025

The Intelligent Workflow: AI's Role in the Renaissance of Engineering

---

Abstract

This paper is grounded in a year-long qualitative research effort aimed at understanding how artificial intelligence is likely to reshape the engineering profession across its full lifecycle. Rather than relying exclusively on published literature or theoretical extrapolation, the research draws primarily on in-depth interviews conducted with a diverse set of practitioners and thinkers, including academic researchers, practicing engineers, founders and entrepreneurs, and senior leaders within large industrial and technology organizations. These conversations were intentionally wide-ranging and exploratory in nature, focusing not on near-term product capabilities, but on how AI may alter the structure, pace, and locus of engineering work—from early research and conceptual design through simulation, validation, manufacturing, operation, and eventual system retirement. The methodology reflects a belief that meaningful insight into technological change emerges at the intersection of theory and practice: where formal research agendas, commercial constraints, and lived engineering experience meet. The resulting analysis seeks to synthesize these perspectives into a coherent framework that captures both the opportunities and the tensions introduced by AI, while remaining grounded in the realities of how engineering is actually performed today.

The engineering profession is approaching a tooling and workflow discontinuity comparable in magnitude to the introduction of Computer-Aided Design

(CAD). After decades of digitization across CAD, simulation, and manufacturing automation, productivity gains have proven incremental rather than transformative. This paper argues that Artificial Intelligence (AI) is reorganizing engineering around an *intelligent workflow*—one in which engineers increasingly supervise, validate, and orchestrate agentic systems that generate designs, run physics-based simulations, propose verification strategies, and continuously learn from operational data. The result is not merely faster tools, but a fundamental shift in the role of the engineer: from specialist operator of discrete software packages to polymath orchestrator of an end-to-end lifecycle. As physics-based foundation models, generative design systems, and predictive maintenance converge, the engineering cycle compresses from months to minutes, signaling a "ChatGPT moment" for the physical world.

---

1. Introduction: From Analog to Digital to Intelligent Engineering

For the past half-century, engineering practice has progressed through two dominant eras—Analog and Digital—and is now entering a third: the Intelligent era.

In the Analog era (pre-1980s), engineering was constrained by human computation, physical prototyping, and manual coordination. Work followed a *many-to-one* model: multiple specialists were required to produce a single validated artifact. Iteration was slow, expensive, and inherently limited by physical experimentation.

The Digital era (1980s–2020s) introduced CAD, CAE, PLM, and enterprise software, shifting engineering toward a *one-to-one* model. A single engineer could own a workflow—drawing, simulating, revising, and documenting designs independently. While this era dramatically improved precision, repeatability, and scalability, it preserved a largely serial workflow: design preceded simulation, which preceded testing, which preceded manufacturing. Engineers remained operators of tools rather than supervisors of systems.

Despite globalization, automation, and robotics, this digital paradigm has delivered diminishing productivity returns. Public productivity datasets show volatility rather than sustained step-change improvement over the past two decades. The bottleneck is no longer computation or tooling availability, but the *structure of the workflow itself*.

The emerging Intelligent era introduces a *one-to-many* relationship. A single engineer now orchestrates fleets of AI agents capable of generating designs, running physics surrogates, proposing tests, monitoring production, and updating operational models in real time. Engineering shifts from artifact creation to constraint definition, verification, and judgment. This reorganization—rather than any single model or algorithm—marks the true inflection point.

This shift matters because many classical productivity levers appear exhausted. U.S. manufacturing productivity and broader productivity measures have shown volatile and uneven performance in recent decades; public datasets (e.g., BLS/FRED series on manufacturing output per hour) provide a consistent reference frame for discussing the limits of digitization alone and the motivation for a new productivity regime.

Figure 1: Evolution of Engineering Workflows

Analog Era	**Digital Era**	**Intelligent Era**
Many → One	One → One	One → Many
Teams	Individual	Orchestrator
Manual	CAD/CAE	AI Agents
Physical Tests	Batch Simulation	Continuous Learning

Figure 1: The progression from manual coordination to digital ownership to AI-orchestrated engineering.

2. Phase I: Research and Discovery

Accelerating Science Through Inverse Design

Engineering begins with discovery—materials, architectures, and physical principles. Historically, this phase has been bottlenecked by laboratory throughput and the combinatorial explosion of possible configurations.

AI introduces inverse design. Rather than testing candidate materials forward ("what properties does this compound have?"), engineers specify target properties, constraints, and cost envelopes, and AI systems search the solution space backward. Virtual experimentation moves discovery from the wet lab to the server farm.

This shift does not eliminate physical experimentation; it reorganizes it. AI proposes hypotheses, automated systems execute focused experiments, results retrain the model, and the loop repeats. Discovery becomes a *closed-loop optimization problem* rather than a linear trial-and-error process.

2.1 Molecular intelligence and inverse design

In materials and chemistry, many performance properties emerge from interactions (mixtures, interfaces, microstructures) that are difficult to predict analytically. Platforms such as Aionics describe an approach combining physics-based simulation with machine learning to design new formulations, particularly for batteries and energy technologies. A specific example was the formulation of an electrolyte for General Motors that preserved full functionality at -90C versus the industry best -20C that was discovered in-silico, synthesized, and successfully tested in less than 3 months.

The workflow inversion is subtle but profound: rather than "try compound A, then B," the engineer specifies target properties (conductivity, stability, viscosity, flammability, manufacturability constraints), and the system searches for candidate formulations that satisfy them.

2.2 From literature review to model-grounded reasoning

Discovery also depends on knowing what has already been tried. AI literature systems can compress weeks of review into hours, but the real value is not summarization—it is hypothesis mapping: connecting methods, boundary conditions, failure modes, and datasets across disciplines.

2.3 The foundry model: automation as scientific scale

Ginkgo Bioworks' "foundry" framing illustrates how research becomes scalable when automation and standardized workflows are paired with algorithmic design and measurement—turning experimentation into something closer to "cloud compute" for biology.

The consequence is temporal compression. Breakthroughs that once required years of experimental iteration can now be explored in weeks, with physical testing reserved for validation rather than exploration.

---

3. Phase II: Design

From Geometry Creation to Requirements Curation

Traditional CAD workflows treat the computer as a passive canvas. In intelligent design systems, AI becomes an active collaborator.

Natural-language interfaces lower the barrier to entry for complex modeling operations, allowing engineers to describe intent rather than manually encode geometry. More importantly, generative design systems transform design from producing a single artifact to exploring entire *families* of solutions.

Engineers specify constraints—loads, materials, manufacturability, regulatory limits—and AI systems generate thousands of candidate geometries. Fast physics surrogates screen these candidates, narrowing the design space before high-fidelity validation.

The engineer's role shifts decisively:

- from drawing geometry
- to curating constraints
- to selecting among optimized alternatives

Design becomes an act of *judgment*, not drafting.

3.1 Natural language interfaces as a new layer of accessibility

Conversational interfaces can lower the barrier to entry for complex CAD operations, especially for junior engineers. The risk is that accessibility can mask the underlying constraints (datum structure, tolerances, manufacturing process capability). The best systems will likely pair "chat" with explicit constraint visibility—the engineer must be able to see *why* a design is valid.

3.2 Generative design: from "one design" to "a design distribution"

Tools like nTop and Infinitform emphasize reusable, automatable, field-driven workflows and lattice/implicit modeling that can generate complex structures quickly (particularly relevant in lightweighting, heat transfer, and additive manufacturing). The deeper shift is that the output is no longer a single geometry; it is a family of candidates conditioned on constraints.

3.3 Requirements curation becomes the engineer's bottleneck

As generation becomes cheap, *selection* becomes expensive. Engineers become curators and judges:

- Which objectives matter (weight vs. fatigue vs. cost vs. certification risk)?
- Which constraints are "hard" vs. negotiable?
- Which assumptions are fragile (boundary conditions, load cases)?

Figure 2: Constraint-Driven Design Funnel

Requirements
↓
Generative Candidate Set
↓
Fast Physics Screening
↓
High-Fidelity Validation
↓
Manufacturing Feasibility
↓
Final Design

4. Phase III: Simulation

From Numerical Solvers to Reasoning Systems

Simulation is where intelligent engineering becomes most disruptive.

Traditional solvers numerically approximate governing equations using mesh-based methods. They are accurate but slow and ill-suited for interactive exploration. AI introduces physics-based surrogate models, reduced-order models, and physics-informed neural networks (PINNs) that approximate system behavior at dramatically lower computational cost.

Once trained, these models enable *interactive physics*: designers can modify geometry and observe performance changes in real time. Simulation becomes conversational and exploratory rather than batch-oriented.

PINNs embed physical laws directly into the learning process, allowing models to respect conservation principles while solving inverse problems and interpolating sparse data. Hybrid pipelines—combining traditional solvers for anchor points with AI surrogates for exploration—are emerging as the dominant paradigm.

Most importantly, simulation becomes agentic. AI systems can autonomously generate a design, simulate it, assess constraint violations, and iterate—without human intervention. Engineers supervise the process, audit assumptions, and validate outcomes.

Simulation is where intelligent engineering can feel most discontinuous, because AI can turn simulation from a batch process into something approaching interaction.

4.1 Physics-based ML as a new simulation layer

Traditional solvers numerically approximate governing equations across meshes; they are accurate but expensive and slow to iterate. ML approaches increasingly act as surrogate models or hybrid accelerators—trained on simulation and/or experimental data to predict outputs quickly for nearby design points. An example is Neural Concept's public discussion of ML surrogates built from CAD/CAE data to predict performance in near real time for design exploration. BeyondMath similarly frames its offering as AI-enabled engineering simulation at industrial scales, reflecting the broader industry movement toward "generative physics" platforms.

4.2 Reduced-order models (ROMs) and "interactive physics"

ROMs compress high-dimensional physics into lower-dimensional representations that preserve the dynamics engineers care about. Paired with AI, ROMs can support rapid design-space exploration—where the engineer asks "what if?" and the system responds immediately with approximations and uncertainty bounds.

4.3 PINNs (Physics-Informed Neural Networks) as a bridge between data and PDEs

Physics-Informed Neural Networks embed the governing equations directly into the training objective, enabling the network to approximate solutions while satisfying physical constraints. The canonical reference describes PINNs as neural networks trained to solve supervised learning tasks while respecting physical laws expressed as nonlinear PDEs. In practice, PINNs are particularly attractive for:

- Inverse problems (estimating unknown parameters from sparse measurements)
- Data assimilation (reconstructing flow fields from limited sensor coverage)
- Hybrid pipelines (using conventional solvers for high-fidelity anchor points, and PINNs for interpolation/extrapolation)

4.4 Agentic orchestration: chaining generation → evaluation → refinement

The new stack is increasingly *compositional*: an agent proposes geometry, triggers a surrogate simulation, checks constraints, and iterates—then packages the results into an engineer-readable justification. The engineering challenge becomes governance: "What was assumed? What was tested? What uncertainty remains?"

Table 1: Simulation Paradigms

METHOD	SPEED	ACCURACY	BEST USE	RISK
Classical CFD/FEA	Slow	Very High	Final Validation	Cost
Reduced-Order Models	Fast	High	Design Exploration	Approximation
PINNs	Fast	Medium-High	Inverse Problems	Training Stability
AI Surrogates	Very Fast	Variable	Early Screening	Hallucination

5. Phase IV: Prototyping and Verification

Continuous Validation in a Probabilistic World

As simulation fidelity increases, prototyping shifts from proof to calibration. Physical tests validate models rather than designs.

However, AI introduces a new risk: *plausible but invalid designs*. Generative systems can produce geometries that satisfy visible constraints while violating hidden ones. This creates a verification gap.

The response is continuous, automated verification. Hardware-in-the-loop testing integrates physical components with virtual environments in real time. Verification is no longer a late-stage gate; it is embedded throughout the workflow.

As simulation improves, prototyping shifts from "prove the design works" to "calibrate the model and expose failure modes."

5.1 Virtual validation and hardware-in-the-loop

Hardware-in-the-loop (HIL) testing increasingly treats physical components as real-time counterparts to virtual environments, enabling continuous verification rather than a late-stage gate.

5.2 The verification gap: why AI makes V&V more important, not less

Generative models can produce plausible geometries that violate hidden constraints (manufacturability, fatigue hotspots, nonlinear couplings). This creates a verification gap: faster iteration increases the chance of shipping a false premise unless V&V is automated, traceable, and continuously applied.

Figure 3: Continuous Verification Loop

AI Design → Automated Checks → Targeted Tests
↑ ↓
← Model Update ←

Figure 3: Verification becomes continuous rather than episodic.

---

6. Phase V: Production and Manufacturing

From Rigid Automation to Adaptive Factories

AI transforms manufacturing from deterministic automation to adaptive systems.

Computer vision systems perform quality inspection with software-defined logic that evolves over time. Robots learn from data rather than fixed waypoints, adapting to variation in parts, lighting, and assembly conditions.

Manufacturing becomes a cyber-physical system: perception, planning, execution, and learning operate continuously. Production scheduling and supply chain coordination become inference problems rather than static plans.

Production is shifting from rigid automation toward adaptive manufacturing, where perception and planning are continuous rather than pre-programmed.

6.1 Computer vision as a quality system, not a camera

BMW has publicly described a pilot project ("GenAI4Q") aimed at tailored quality checks in vehicle assembly—illustrating how AI can adapt inspection logic to changing conditions rather than relying on static checklists.
The broader implication is that quality becomes a software-defined layer: retrainable, configurable, and increasingly predictive (detecting patterns earlier in the line, not only at end-of-line inspection).

6.2 Flexible robotics

Machine learning enables robots to generalize across variation—parts that don't arrive perfectly, lighting changes, small geometry deviations—reducing the need for brittle fixtures and exact positioning. UCR (Under Control Robotics) went from concept to a humanoid robot traversing complex terrain including stairs, and unpaved surfaces in under 4 months. The robot taught itself to walk using reinforcement learning in record time with no specific input from its creators. Similarly, UnitX has deployed over 1000 robots in the most complex manufacturing environments and is able to identify anomalies far more accurately than earlier systems - using as little as 30 samples to train across a wide spectrum of anomalies.

6.3 Supply chain and production scheduling as inference problems

Once production is instrumented, AI can detect bottlenecks, predict downtime risk, and propose schedule changes. The value is not prediction alone—it's closed-loop control across procurement, work-in-progress, and maintenance planning.

---

7. Phase VI & VII: Operation, Maintenance, and Disposition

Engineering Beyond Shipment

The intelligent workflow extends into operation and end-of-life.

Predictive maintenance models analyze sensor data to forecast failures before they occur, increasing availability and reducing cost. Products are accompanied by digital twins—living models that mirror physical systems throughout their lifecycle.

At end-of-life, these twins inform repair, remanufacture, recycling, and disposal decisions. Disposition becomes engineered, not improvised.

7.1 Predictive maintenance and OEE

Overall Equipment Effectiveness (OEE) is commonly framed as the product of Availability × Performance × Quality, providing a single metric for how much planned production time is truly productive. AI-driven predictive maintenance targets Availability by forecasting failures before they occur and by optimizing maintenance windows.

7.2 Digital twins as "living models"

A widely used definition describes a digital twin as a digital representation of a real-world entity or system that mirrors a physical object, process, or other abstraction. In the intelligent workflow, digital twins become:

- the container for lifecycle data (design → manufacturing → operation), and
- the mechanism for feedback into the next design revision.

7.3 Disposition and circularity

End-of-life decisions (repair, remanufacture, recycle) can be optimized when the twin contains a traceable history of materials, loading, degradation, and repairs—turning disposition from an afterthought into an engineered outcome.

---

8. Real-World Impact: Aerospace and Automotive

Aerospace illustrates the value of intelligent workflows under extreme complexity. Software content, system integration, and certification demands have grown exponentially, yet development timelines have increased only modestly. Digital threads and AI-assisted optimization enable teams to manage combinatorial explosion without proportional cost growth.

In automotive engineering, AI has collapsed the need for physical validation. High-fidelity simulation substitutes for crash tests, wind tunnels, and clay models.

Development cycles shrink, prototype counts fall, and validation accelerates—without sacrificing safety.

We can quantify at least one major driver: software complexity in aerospace systems. The Aviation Validation & Simulation Institute (AVSI) has described aerospace software complexity as increasing exponentially, with aircraft source lines of code doubling roughly every four years over decades. This matters because software complexity propagates: more code implies more interfaces, more test cases, more failure modes, and heavier certification burdens.

Aerospace case: digital threads + rapid design loops

In defense/aerospace, the industry push toward integrated digital ecosystems is often discussed as a way to accelerate innovation across the lifecycle. Northrop Grumman, for example, describes an "integrated digital ecosystem" designed to connect customers, suppliers, and internal teams across the program lifecycle. Once the thread exists, AI agents can operate across it—generating options, running analyses, and maintaining traceability.

Automotive case: simulation and "virtual substitution"

The automotive story is strongest when framed as virtual substitution: replacing physical prototypes and tests with high-fidelity simulation and data-driven modeling, while preserving traceability to requirements and failure modes. Recent industry commentary highlights the trend toward reducing physical prototyping through advanced simulation process/data management and faster results review cycles.

---

9. Challenges, Ethics, and Regulation

Governing Probabilistic Engineering

AI-mediated engineering raises new governance challenges:

- Provenance**:** Where did this design originate?
- Traceability**:** Which requirements does it satisfy?
- Auditability**:** Can it be reproduced and certified?

The primary risk in AI-mediated engineering is not that models are "black boxes" in the abstract; it is that the engineering organization loses the ability to explain why a design is safe.

9.1 Hallucinations become physical: plausibility is not validity

Generative systems can produce designs that *look* plausible but violate hidden constraints (fatigue life, thermal runaway edge cases, assembly tolerances). The mitigation is not "trust the model less," but to build workflows that automatically:

- enforce constraints,
- run independent checks, and
- document assumptions.

9.2 Data integrity and organizational readiness

AI systems are limited by the quality and coherence of organizational data: fragmented PLM, inconsistent naming, unstructured test logs, and missing ground truth. In practice, "AI readiness" often begins as a data architecture project.

9.3 Certification in a probabilistic world

Regulators and certification bodies need stable, explainable artifacts. The likely compromise is *not* explainability alone, but reproducible pipelines: deterministic seeds where possible, locked model versions, recorded prompts/constraints, and standardized test suites.

Certification in a probabilistic world will rely less on explainability alone and more on reproducible pipelines, locked models, documented assumptions, and standardized validation suites.

Culturally, organizations must frame AI as augmentation, not replacement. Ethical responsibility, judgment, and accountability remain human obligations.

---

10. The Future Workforce: The Polymath Engineer

Engineering is reversing its specialization curve.

The digital era demanded hyper-specialists—experts in CAD, CFD, or manufacturing systems. The intelligent era enables a return to *thin but powerful* generalists: engineers who understand the full lifecycle and orchestrate AI systems to supply depth on demand.

Education must adapt accordingly—training engineers to think in systems, constraints, verification, and governance rather than tool mastery alone.

This transition is not only technological; it is economic and demographic. Multiple sources referencing a Deloitte/Manufacturing Institute study report that the U.S. manufacturing skills gap could result in ~2.1 million unfilled manufacturing jobs by 2030, highlighting a structural shortage in technical labor. In this context, AI is not primarily about replacing engineers; it is about increasing the leverage of scarce talent. The "polymath" is not a return to shallow generalism—it is a return to lifecycle ownership enabled by AI systems that carry the depth on demand.

Practical implication:
Engineering education shifts from "master tool X" to:

- systems thinking across lifecycle,
- verification literacy,
- prompt/constraint specification,
- failure mode reasoning,
- and model governance.

---

11. Conclusion: The Bottleneck Has Moved

AI does not eliminate engineering constraints; it relocates them. When generation and simulation become cheap, clarity of requirements, rigor of verification, and organizational trust become the limiting factors. The intelligent workflow marks a structural reorganization of engineering—one in which a single modern engineer can accomplish what once required entire teams. Those who adapt their workflows, culture, and governance will define the next industrial renaissance.

For designers of hardware-based systems, this shift has immediate and consequential implications. As AI increasingly automates synthesis, optimization, and exploration across mechanical, electrical, and software domains, the primary design challenge moves upstream. System architecture, interface definition, and requirement decomposition become the dominant creative acts. Poorly specified constraints are no longer merely inefficient; they are multiplicative sources of error, amplified by automated generation. Conversely, well-posed problem formulations allow AI-driven tools to explore design spaces at scales previously inaccessible, surfacing non-obvious trade-offs and novel configurations. In this environment, engineering judgment is expressed less through manual calculation and more through the disciplined articulation of intent, assumptions, and boundaries within which automation operates.

Looking forward, the next domain of design tooling will center not on isolated point solutions, but on integrated, end-to-end systems that couple generative design, high-fidelity simulation, continuous verification, and lifecycle feedback. Effective leverage of AI will require engineers to think probabilistically about design outputs, to treat models as living artifacts, and to embed validation throughout the workflow rather than relegating it to late-stage gates. The most successful practitioners will be those who view automation not as a replacement for expertise, but as a force multiplier for it—augmenting human insight while demanding greater responsibility in oversight, ethics, and accountability. As the bottleneck moves from execution to intent, the defining skill of the engineer becomes the ability to design the process by which designs themselves are created.

References

[1]: CAD and AI - Google Sites

[2]: How is Artificial Intelligence Software Improving CAD Automation

[3]: How AI is going to revolutionize the CAD Industry - Australian Design & Drafting

[4]: Using AI libraries for Incompressible Computational Fluid Dynamics

[5]: Applying Machine Learning in CFD to Accelerate Simulation

[6]: AI-Enhanced Computational Fluid Dynamics (CFD) in Design

[7]: Pioneering the Future of CFD Simulation - Xitadel

[8]: Cfd Applications In Ai Simulations - Restackio

[9]: AI Use Cases in Manufacturing - JD Meier

[10]: 16 AI Use Cases Transforming Manufacturing in 2025

[11]: Manufacturing AI: Top 15 tools & 13 use cases & case studies - AIMultiple

[12]: How AI Integration Is Transforming AutoCAD in 2025

[13]: 3 practical ways industrial AI is reshaping manufacturing

[14]: The Role of Artificial Intelligence (AI) in the CAD Industry - Scan2CAD

[15]: How AI is Revolutionizing the CAD Industry

[16]: From Words to Creations: The Power of Conversational AI in CAD Design ...

[17]: 4 Myths about AI in CFD - Simcenter

[18]: AI in Manufacturing: Trends, Solutions, Use Cases | The Access Group

[19]: Integrating Deep Learning into CAD/CAE System

[20]: Navigating AI Integration In Manufacturing: Challenge and Benefits

[21]: Navigating AI Integration In Manufacturing - Forbes

[22]: Industrial transformation: Scaling AI across the manufacturing value ...

[23]: The Evolution of CAD Technology in 2024: Advances and Consequences

[24]: AI in Logistics: A 360-View on Challenges, Use Cases, Best ... - Trinetix

[25]: How can I use AI for Transportation Engineering?