

From Automation to Augmentation: Redefining Engineering Design and Manufacturing in the Age of NextGen-AI

Md Ferdous Alam^{*}

Department of Mechanical Engineering
MIT
mfalam@mit.edu

Austin Lentsch^{*}

Shaping the Future of Work Initiative
MIT
alentsch@mit.edu

Nomi Yu

Department of Mechanical Engineering
MIT
yunomi@mit.edu

Sylvia Barmack

Shaping the Future of Work Initiative
MIT
sbarmack@mit.edu

Suhin Kim

Department of Mechanical Engineering
MIT
sikim@mit.edu

Daron Acemoglu

Institute Professor &
Department of Economics
MIT
daron@mit.edu

John Hart

Department of Mechanical Engineering
MIT
ajhart@mit.edu

Simon Johnson

Sloan School of Management
MIT
sjohnson@mit.edu

Faez Ahmed

Department of Mechanical Engineering
MIT
faez@mit.edu

^{*}Equal contributions

Executive Summary

In the mid-2010s, as computing and other digital technologies matured (Brynjolfsson and McAfee 2014), researchers began to speculate about a new era of innovation—with artificial intelligence as the standard-bearer of a “Fourth Industrial Revolution” (Schwab 2015). The release of generative AI technologies (e.g. ChatGPT) in late 2022 reignited the discussion, prompting us to wonder: *what are the barriers, risks, and potential rewards to using generative AI for design and manufacturing?*

As generative AI (“Gen-AI”) has entered the mainstream, geopolitics and business practices have shifted. Covid-19 disrupted global supply chains, tensions with import partners have risen, and military conflicts introduce new uncertainties. As companies consider propositions like “reshoring” or “nearshoring/friendshoring” production (Yellen 2023), we recognize other hindrances: suboptimal resource allocation, labor market volatility and trends toward an older and geographically mismatched workforce, and highly concentrated tech markets that foster anti-competitive business practices. As the US expands domestic production capacity (e.g., semiconductors and electric vehicles), Gen-AI could help us overcome those challenges. To investigate the current and potential usefulness of Gen-AI in design and manufacturing, we interviewed industry experts—including engineers, manufacturers, tech executives and entrepreneurs. They have identified many opportunities for deployment of Gen-AI: 1) Reducing the incidence of costly late-stage design changes when scaling production; 2) Providing information to designers and engineers, including by identifying suitable design spaces and material formulations, and by incorporating consumer preferences; 3) Improving test data interpretation to enable rapid validation and qualification; 4) Democratizing workers’ access and usage of data to enable real-time insights and process adjustment; 5) Empowering less-skilled workers to be more productive and do more-expert work.

Current Gen-AI solutions (e.g., ChatGPT, Claude) cannot accomplish these goals due to several key deficiencies, including: inability to provide robust, reliable, and replicable output; lack of relevant domain knowledge; unawareness of industry-standards requirements for product quality; failure to integrate seamlessly with existing workflow; and inability to simultaneously interpret data from different sources and formats. We propose a development framework for the next generation of generative AI tools for design and manufacturing (“NextGen-AI”): 1) Provide better information about engineering tools, repositories, search methods, and other resources to augment the creative process of design; 2) Integrate adherence to first principles when solving engineering problems; 3) Leverage employees’ experiential knowledge to improve training and performance; 4) Empower workers to perform new and more-expert productive tasks, rather than pursue static automation of workers’ current functions; 5) Create a collaborative and secure data ecosystem to train foundation models; 6) Ensure that new tools are safe and effective.

These goals are extensive and will require broad-based buy-in from business leaders, operators, researchers, engineers, and policymakers. We recommend the following priorities to enable useful artificial intelligence for design and manufacturing: 1) Improve systems integration to ethically collect real-time data; 2) Regulate data governance to ensure equal opportunity in development and ownership; 3) Expand the collection of worker-safety data to assess industry-wide AI usage; 4) Include engineers and operators in the development and uptake of new tools; 5) Focus on skills-complementary deployments to maximize productivity upside.

1. Introduction

Manufacturing is a cornerstone of the world economy, accounting for approximately 16% of global GDP and 14% of employment in 2021 (World Bank 2021; United Nations 2021). In the US, manufacturing accounts for 12.5 million jobs and 35% of productivity growth (Bistarkey 2022). While the US has long been a preeminent manufacturing power (see Figure 1), the last 50 years have shown a significant decline in manufacturing employment, due in part to outsourcing or offshoring, and in part to a rise in the use of industrial robots to displace domestic labor demand (World Bank 2023; Acemoglu and Restrepo 2020; Autor, Dorn, and Hanson 2013; West and Lansang 2018). This strategy has allowed companies to doubly boost their bottom lines by leveraging automation technologies to reduce the need for skilled labor, while also squeezing less-skilled workers in emerging economies—paying lower wages and requiring less oversight for working conditions and safety than would be permissible domestically (Acemoglu and Restrepo 2020; Boehm, Flaaen, and Pandalai-Nayar 2020; Brown, Deardorff, and Stern 2004).

In recent years, however, the priorities of firms and national security have shifted toward strategies including more reshoring or nearshoring/friendshoring of production (Yellen 2023). This shift has been catalyzed by recent supply chain disruptions (e.g., Covid-19) as well as increased uncertainty about US-China relations and geopolitical conflicts (e.g., the Russo-Ukrainian and Israel-Hamas Wars) (Alfaro and Chor 2023). This coincides with a technological shift in manufacturing, as recent advancements—AI, advanced analytics, and the Industrial Internet of Things (IIoT)—have convinced many that a fourth industrial revolution is underway: “Industry 4.0” (Furman and Seamans 2019; National Academies 2017).

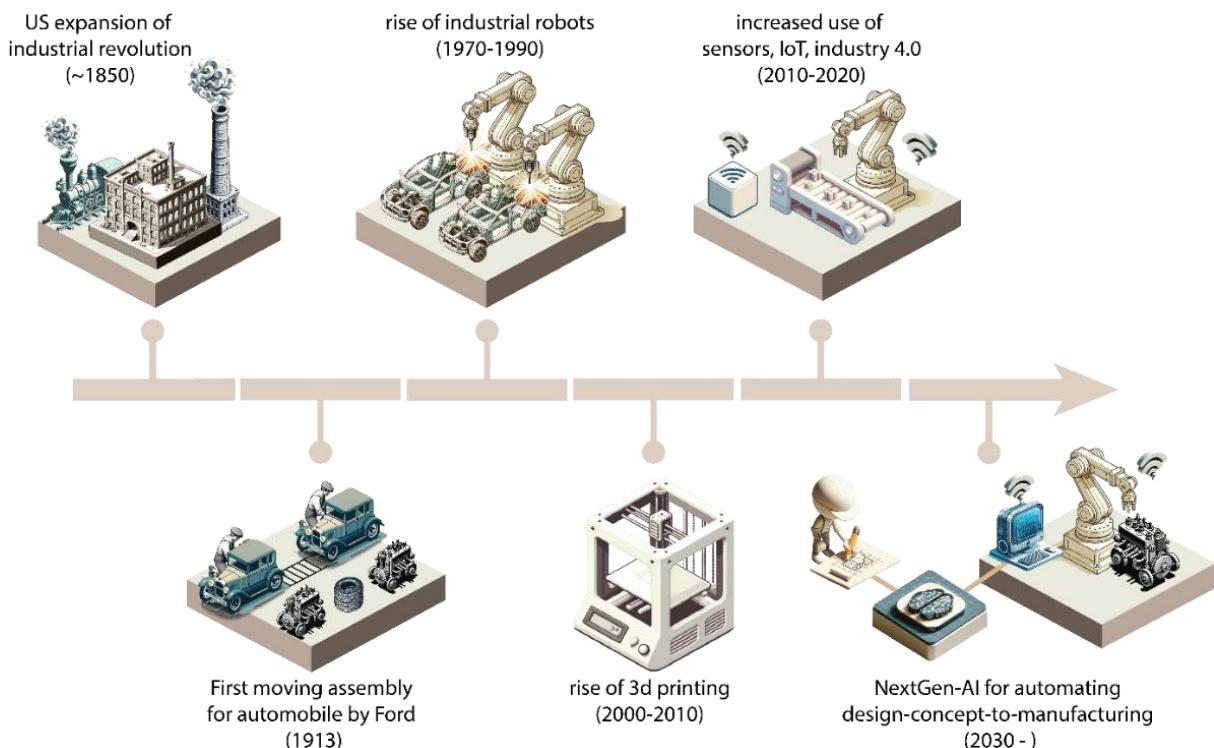


Figure 1: A brief history of manufacturing innovation in the US based on a timeline from the Office of Energy Efficiency and Renewable Energy (U.S. Department of Energy 2019) and authors’ previous works. We anticipate that the next generation of

generative AI (NextGen-AI) technologies will address several limitations of current generative AI technologies (Gen-AI) and could be the newest frontier for major productivity advancements in design and manufacturing.

However, the workflow in design and manufacturing remains a major bottleneck to delivering on the massive productivity promises of these new technologies (Francois et al. 2017; Seepersad 2014; Manufacturing Institute 2023; Bernard 2023; Gupta 2023c; 2023b; 2023a). For example, traditional iterative design and manufacturing workflows heavily rely on domain expertise, which requires substantial investments of time and effort by skilled engineers and operators who often rely on compartmentalized information to assess tradeoffs with a limited view of the complete design-manufacturing workflow. The full potential of human creativity for innovative product design faces headwinds of domain-specific knowledge and the manual nature of many workflows (Miller et al. 2021).

In recent years, artificial intelligence (AI) has shown tremendous potential to overcome many of these limitations and are quickly becoming pervasive in many manufacturing applications such as predictive maintenance (Carvalho et al. 2019; Susto et al. 2015), quality control (Usuga Cadavid et al. 2020; T. Deng et al. 2023), supply chain optimization (Carboneau, Laframboise, and Vahidov 2008), process control (Alam et al. 2022; 2021; 2020; Alam 2023) and risk management (Cavalcante et al. 2019). The goal of this article is to look beyond preexisting AI methods in design and manufacturing and focus on the potential of rapidly evolving generative AI (i.e., Gen-AI) tools and the possibilities of next-generation generative AI (i.e., NextGen-AI) that is designed to be more powerful and well-tailored to the needs of manufacturing and design. Our predictions and recommendations are guided by a careful analysis of historical facts and survey of current industry trends, including a series of interviews with industry experts.

The current generation of generative artificial intelligence (i.e., “Gen-AI”) is dominated by text-based platforms (Vaswani et al. 2017; Touvron et al. 2023; Devlin et al. 2018; Brown et al. 2020; Radford et al. 2019), image-based platforms (Rombach et al. 2022; Ho, Jain, and Abbeel 2020; Sohl-Dickstein et al. 2015), and “multi-modal” platforms that incorporate both images and texts (Driess et al. 2023). These tools, mainly trained on multimedia data from the internet for general-purpose use, have garnered much interest for their usefulness in applications like information processing, writing, and computer coding (Brynjolfsson, Li, and Raymond 2023; Noy and Zhang 2023; Peng et al. 2023). However, engineering applications such as design and manufacturing require foundational domain expertise and relevance to the creation and production of tangible goods. While preliminary studies have shown some promise of current Gen-AI for simple tasks such as material selection and manufacturability assessment, the models struggle with complex tasks such as spatial reasoning and complex design problems (Picard et al. 2023). Therefore, we have identified several goals for NextGen-AI development in design and manufacturing: enable augmented human creativity in distinctive, optimized and actionable design concepts; evaluate designs at higher accuracy or lower cost than prohibitively expensive traditional methods; monitor quality control and preventative maintenance; and provide personalized worker training and real-time problem-solving. Unfortunately, current Gen-AI tools are likely inadequate for a truly transformative impact in the design and manufacturing domain due to a variety of shortcomings: inability to provide robust, reliable, and replicable output; lack of relevant domain knowledge; unawareness of industry-standards requirements for product quality; failure to integrate with existing workflow; and inability to simultaneously interpret data from different sources and formats (e.g., text, images, video, audio, sensor data).

Based on the current trajectory of Gen-AI, we construct a strategic framework in Section 4 that can potentially guide the next generation of generative AI models (i.e., “NextGen-AI”) for use in design and manufacturing. We believe that this framework can align the shared interests of government, industry, workers, and research entities to ensure mutual benefits and prevent adverse consequences. We hypothesize that NextGen-AI has the potential to improve and accelerate the design and manufacturing pipeline in the following ways:

1. Concept Innovator: Offering multiple possible solutions for complex engineering design tasks and integrating multimodal information (e.g., materials, processes, geometries, carbon footprint)
2. Task automation: Supporting detail-oriented tasks and workflows by reducing the burden of rote, noncreative tasks (e.g., certification of designs, conversion of data into different modes)
3. Task creation: Leverage pre-existing domain expertise—captured by unstructured data used to train Gen-AI models—to enable less-skilled workers to perform more advanced work—or create new tasks that add value to the production process (e.g., predictive maintenance)
4. Decision support: Identifying design-manufacturability-cost optima through automation and combination of performance analysis (geometric optimization, performance simulation), discrete simulations (e.g., production systems), and a variety of data sources. The algorithms can predict risks in the manufacturing process or product lifecycle and suggest risk-mitigation tactics
5. Co-pilot for workers: Improving information provision to empower workers with a variety of skillsets and educational backgrounds and boost worker productivity by accelerating current design and manufacturing tasks, but also enabling new contributions from skilled workers
6. Digital twin: Enable the creation of accurate digital twins to facilitate quality checks, defect detection, and predictive maintenance

Although NextGen-AI tools hold great potential, they also raise troubling issues, including for the future of the workforce and competitive dynamics between firms. These tools could automate many engineering tasks and reduce the importance of skilled workers, or instead could (also) be used to augment workers’ capabilities, increase expertise, and create more well-paid jobs. The direction of development for NextGen-AI is not predetermined: it will be shaped by the choices of governments, executives, technologists, researchers, and workers.

2. Current Generation of Generative AI (Gen-AI)

Generative AI combines probabilistic methods and deep learning, rooted in artificial neural networks, to analyze and predict patterns within extensive datasets. This approach enables AI to generate new, complex data outputs. By feeding an AI model a wealth of examples in different formats (e.g., text, images, or speech), the model “learns” to discern complex patterns and relationships within this data. Once trained, the model acts like an algorithmic artist, creating unique outputs by sampling from its vast learned data. However, unlike human artists who blend technical skill with innate creativity, Gen-AI’s creations are entirely data-driven predictions. The AI does not “create” in the human sense but rather generates based on learned data. The Gen-AI model selects each ‘brush stroke’ based on a complex, multi-dimensional understanding of the generative process, drawing on its knowledge encoded in trained parameters. Thus, while the content generated by a Gen-AI model could be considered “original”, it is also a direct reflection

of the vast and diverse data it has been trained on. This data dependency is especially important in the context of design and manufacturing, where most high-quality data may not be freely available for training the models.

Recent advances in generative AI have primarily focused on two data modalities: text and image, though emerging areas like audio processing and 3D modeling are also gaining traction. Natural language processing (NLP) and computer vision (CV) have been important preceding developments, though the recent AI breakthroughs have been enabled by notable improvements in the design and usage of innovative deep neural network architectures (Krizhevsky 2009; Vaswani et al. 2017), improvements in computer hardware such as graphics processing units (GPUs) (Hooker 2021) and large-scale datasets (J. Deng et al. 2009; Krizhevsky 2009; Xiao, Rasul, and Vollgraf 2017; Gao et al. 2020). The recent progress in research and application of Gen-AI is striking, and has facilitated the rapid uptake of text- and image-based Gen-AI platforms for a number of applications. For example, Gen-AI personal assistant tools such as ChatGPT (OpenAI 2023a), BARD (Google 2023), and Claude (Anthropic 2023) are useful as general purpose question answering chatbots (Xu, Feng, and Chen 2023). Preliminary evidence suggests that ChatGPT has even captured some web traffic from online Q&A sites like Stack Overflow (del Rio-Chanona, Laurentsyeva, and Wachs 2023). Compared to the previous generation of chatbots, these Gen-AI chatbots are much more flexible. They can provide general-purpose information on a variety of topics, including computer programming (Rozière et al. 2023), writing assistance (Touvron et al. 2023; OpenAI 2023c), cooking instructions (OpenAI 2023c) and many others. We have seen similar trends in Gen-AI image-generation tools that create illustrations based on text, such as Midjourney (Midjourney Inc 2023), DALL·E (OpenAI 2023b), and Stable Diffusion (Stability AI 2023). In addition to image and text modalities, speech has also seen the application of Gen-AI with applications in text-to-speech synthesis, noise removal, content editing, style conversion, and diverse sample generation (Radford et al. 2023; Le et al. 2023). Beyond these modalities, researchers are also exploring Gen-AI for building text- and image-based control of intelligent robots (Zitkovich et al. 2023; Karamcheti et al. 2023). While Gen-AI has seen applications in text, image and speech, the applicability of these tools is somewhat limited in the design and manufacturing domain. In the design workflow, a few studies have explored the potential of Gen-AI for computer-aided design (CAD) (R. Wu, Xiao, and Zheng 2021; Sanghi et al. 2022; Para et al. 2021). Additionally, recent studies are exploring generative AI for the engineering design domain (Giannone et al. 2023; Giannone and Ahmed 2023; Bagazinski and Ahmed 2023; Mazé and Ahmed 2023). For manufacturing applications, preliminary work has investigated the use of Gen-AI for computational design from text inputs (Picard et al. 2023; Makatura et al. 2023) and for creating parts of spaceflight optical instrument structures (McClelland 2022).

2.1 Is Gen-AI suitable for design and manufacturing applications?

The current generation of Gen-AI technologies has sparked much interest in their mass adoption across many industries. Unfortunately, the promises of Gen-AI are underwhelming in design and manufacturing until certain limitations are addressed. Based on the literature and standard practices in design and manufacturing, we assess a few key challenges to adopting the current versions of Gen-AI technologies.

Are Gen-AI tools robust? In the context of design and manufacturing, the effectiveness of Gen-AI tools hinges on several key characteristics: they must be reliable, ensuring consistent performance; stable, to function predictably under varying conditions; accurate, providing

precise and correct outputs; adaptable, capable of adjusting to new or changing requirements; resilient, able to recover quickly from disruptions; and robust, strong enough to handle complex and diverse tasks effectively. Many studies have found that Gen-AI technologies often deviate from user inputs and self-contradict previously generated results (P. Liang et al. 2022; Zha et al. 2023; Mündler et al. 2023). Additionally, aligning Gen-AI to specific applications is a major challenge and an active research topic (Lightman et al. 2023; Ouyang et al. 2022; Wolf et al. 2023; Y. Wu et al. 2023). Gen-AI tools, such as intelligent chatbots, often fabricate information—commonly referred to as “hallucination”. Such inaccuracies present significant obstacles in developing Gen-AI tools for engineering design and manufacturing, as they can lead to unworkable design concepts or serious production disruptions.

Does Gen-AI possess useful domain knowledge? The majority of current Gen-AI technologies are trained on extensive internet data collections, which typically include limited content on design and manufacturing. For instance, widely used large textbook datasets like Project Gutenberg (Rae et al. 2019) and Bookcorpus (Zhu et al. 2015) contain minimal information on manufacturing topics. Consequently, these tools may completely lack pertinent domain knowledge or struggle to accurately and consistently extract relevant information from their training data, especially for highly specific user queries. Recent studies have identified significant limitations in Gen-AI models’ ability to comprehend and interpret complex instructions related to design and manufacturing tasks (Picard et al. 2023; Makatura et al. 2023), with failures in basic tasks like grasping common manufacturing problems or recommending suitable materials.

Most engineering design and manufacturing data are not only very industry-specific, but this information is also proprietary and subject to extensive rules of intellectual property and ownership. Furthermore, numerous industries employ unique domain-specific languages (DSLs) for data storage and maintenance, complicating the transfer and understanding of information across different tasks. Together, the lack of data sharing and localized data encoding creates a bottleneck to training large-scale Gen-AI models that could be useful across industries, as the current generation of Gen-AI technologies requires a tremendous amount of data to learn useful skills. While recent advancements in federated learning and differentially private techniques show promise in addressing these challenges, there remains a considerable journey ahead to fully harness the potential of Gen-AI in the specialized domains of design and manufacturing. We further discuss the importance of domain knowledge to design and incorporate NextGen-AI into the automotive and footwear industries in Section 5.

Does Gen-AI possess useful reasoning capabilities? A major limitation of current Gen-AI technologies is the lack of reasoning capabilities (J. Huang et al. 2023; Dziri et al. 2023; Turpin et al. 2023; Shi et al. 2023). Engineering design and manufacturing pipelines demand advanced reasoning capabilities to manufacture real-world products, as these workflows involve several iterative procedures to ensure that manufactured products meet user-specified qualifications (Dixon and Dym 1986). Unfortunately, this poses a major challenge to adopting the current generation of Gen-AI tools (Picard et al. 2023; Makatura et al. 2023). Although it may not be feasible (or even preferable) for Gen-AI tools to autonomously handle the entire product design process from beginning to end, it is important that the system can adhere to the pre-defined product specifications and design requirements during whichever steps *are* completed autonomously. For example, if an excessive amount of effort is required to simply police the Gen-AI tool’s accuracy and adherence to the pre-specified rules, it is unlikely to offer much

productivity benefit to the design process. Thus, the lack of consistency and feasibility of output make Gen-AI less useful for the design and manufacturing domain. However, ongoing developments in machine learning, such as advancements in neural network architectures and training methodologies, offer hope for enhancing the reasoning abilities of Gen-AI systems.

Are Gen-AI Tools Capable of Meeting Established Engineering Standards? Most engineering design and manufacturing tools currently rely on heuristic approaches and adhere to guidelines set by various standards organizations, such as the International Organization for Standardization (ISO) or ASME (American Society of Mechanical Engineers). For instance, original equipment manufacturers (OEMs) are required to comply not only with pertinent government and industry regulations but must also have a quality control system that meets guidelines from standards organization—for example, ISO and SAE International for engineering standards of manufactured equipment.

Furthermore, the outputs of Gen-AI tools must rigorously conform to specific resource and budget allocations. Current Gen-AI technologies struggle to consistently follow these sorts of strict guidelines. Therefore, NextGen-AI needs to be developed with the capability to fulfill standard protocols for standardization, scheduling, optimization, and efficient resource allocation. Incorporating mechanisms such as rule-based systems or integrations with existing compliance databases could be key in enabling Gen-AI to meet these standard requirements more effectively. NextGen-AI tools must be able to adapt to evolving standards, as engineering guidelines are continually updated to reflect new technologies and safety protocols.

Can Gen-AI software integrate seamlessly with existing software stack? Deployment of data-driven Gen-AI tools at scale will require seamless data collection to maximize usefulness. This may include evaluating the feasibility of product specifications or measuring the real-time efficiency of machinery on the manufacturing floor. Regrettably, much of the current infrastructure predates Gen-AI and often lacks compatibility with Gen-AI's principle of immediate data availability. The integration of Gen-AI with legacy systems could involve developing intermediary software layers or APIs that facilitate smooth data flow and compatibility. Success in this integration also hinges on the continuous involvement of end-users in the development process, ensuring that the tools developed are not only technically capable but also user-friendly and intuitive. For design engineers, whose specialization and earning potential heavily rely on software skills, the introduction of new software systems might face resistance; for machine operators, real-time data on equipment effectiveness is only useful if workers are able to easily interpret and react to the information being provided by the Gen-AI tool. Thus, a careful, worker-oriented implementation strategy will be crucial to ensure employee buy-in by assuring that i) NextGen-AI tools are developed to take advantage of preexisting worker skills and that ii) these tools are reliable and provide useful insights in practice. Moreover, ongoing training and support should be provided to help employees adapt to and embrace these new technologies, thus minimizing disruption and maximizing productivity.

We provide potential strategic frameworks that can overcome many of these limitations in Section 4. Ethical concerns are also critical for developing NextGen-AI as there are already failure cases of text- and image-based Gen-AI technologies perpetuating sex- and race-based biases (Bender et al. 2021). A combination of well-crafted AI governance strategies (Watson 2024; Raji et al. 2020; 2022) and ethical reporting of Gen-AI model outcomes (Mitchell et al. 2019) is required to address these concerns.

3. Economic implications of Gen-AI for workers and firms

For more than a century, the United States was a global manufacturing superpower. The successes of American industry have been driven in large part by clever innovators and their cutting-edge production machinery and mechanized processes (Hindle and Lubar 1986). In the earliest phase of the American Industrial Revolution, proliferative inventors benefited from the establishment of clear intellectual property rights (the US Patent system was founded in 1790), but also enjoyed a collaborative industrial environment that resulted in wide distribution of new ideas and uptake of new processes and machinery. For example, in 1795, Oliver Evans self-published *The Young Mill-Wright and Millers' Guide* (Ferguson 1980). This opus contained chapters for grist millers—explaining how to implement Evans' patented flour-milling technologies and offering advice to novice millers—but also several chapters on general elements of mechanical and hydraulic engineering, a vital guide even for industries beyond millwrighting (O. Evans 1795). The book helped to democratize the technological frontier for the next generation of industrialists and innovators.

These nascent developments set the stage for American manufacturing productivity to increase dramatically, beginning in the latter half of the 19th century, with the establishment of the American System of Manufacturing—noted for impactful productivity enhancing innovations like Henry Ford's moving assembly line, Eli Whitney's deployment of interchangeable parts, mechanized tools (e.g., drill presses), and the electrification of manufacturing facilities (Acemoglu and Johnson 2023b). What made this system so effective was its ability to make less-skilled or less-educated workers more productive while also enhancing the capabilities of skilled trade workers, enabling them to produce more consistent and high-quality output. This period of preeminence culminated in the fantastic success of automobile manufacturing during the 1920s–70s. At the turn of the 20th century, the entire industry produced only around 2,500 vehicles annually; however, by 1929, Ford and General Motors (GM) each produced around 1.5 million cars per year, with the entire industry employing nearly half a million Americans (Acemoglu, Johnson, and Lentsch 2023). Equally important to these gains were the contributions of workers, who were continually trained and upskilled to tackle new production tasks. Even after landmark labor negotiations in the late 1930s, the automobile industry continued to reach new heights: industry-wide employment *tripled* by the 1960s, while inflation-adjusted profits *quintupled*.

Sometime during the 1970s, however, the paradigm shifted. A combination of forces—including automation, offshoring, and failures to upskill and reskill workers for new technologies—have led to stagnated productivity for firms and a significantly worsened outlook for workers in the manufacturing sector (see, e.g., Acemoglu and Johnson 2023b; Autor, Dorn, and Hanson 2013; Acemoglu and Restrepo 2020; 2022; Autor 2019). These consequences have been driven by two commingled and depreciative forces. The first is a tendency to over-automate processes that workers can perform at least as effectively as (or better than) machines or algorithms, such that this automation offers underwhelming productivity benefits (Acemoglu and Restrepo 2019a; 2019b; Acemoglu, Manera, and Restrepo 2020). The second force has been faltering creation of the new, productive tasks that produce well-paid jobs and take full advantage of new technology (Autor et al. 2022).

The economic impacts of current and NextGen-AI innovations will depend on how innovators envision, develop, and deploy these tools (Acemoglu and Johnson 2023b; Capraro et al. 2023). If these technologies are conceived to rotely diminish the role of labor and chase mediocre

autonomous systems just for the sake of autonomy, the outcomes could be disappointing for firm productivity and profits, ruinous for highly skilled and specialized manufacturing workers who are displaced, and inimical to industry growth. Similarly, if the direction of development and control over these new tools is left to a small number of entrenched companies which can then further strengthen existing market power, these benefits are not likely to extend to smaller firms, to workers, or to the economy. But this is not inevitable: careful development and guidance of NextGen-AI could create new capabilities that leverage the skills of manufacturing workers.

3.1 Underutilization or misallocation of resources

One of the most important, yet often poorly measured, areas of manufacturing performance is resource allocation and utilization. Many firms manually record key performance indicators such as process “uptime” and “downtime”, often using a whiteboard on the factory floor, or a similarly analog approach, to estimate and track efficiency metrics. This legacy approach is suboptimal for several reasons. Productivity upside is forgone when data is not systematically or durably captured, stored, and analyzed, as low-grade or system-wide issues may go undetected and unaddressed. Current data capture processes are cumbersome and costly for workers who have other responsibilities. Many workers do not perform data collection very effectively when their priorities are on fulfilling immediate production functions. Lastly, if production planning and resource allocation relies on the timeliness and accuracy of this data, static “whiteboard approaches” will hinder optimal organization.

Companies can consider implementing automated data capture technologies—e.g., Raven.ai (Raven Telemetry 2023), Tulip (Tulip 2023), Plex Production Monitoring (Rockwell Automation 2023b)—to alleviate the need for operators to manually record performance data. Capturing this data real-time also allows for more accurate and rapid evaluations, efficiently marshaling resources to production breakpoints. The volume of data required to train high-performing Gen-AI models can be very large. This is a significant barrier to entry for small and mid-sized enterprises, so the deployment of commercial solutions by startups and established industry affiliates—e.g., MontBlancAI (Mont Blanc AI 2023), SpencerMetrics (SpencerMetrics 2023), Lumafield (Lumafield 2023), leela.ai (Leela AI 2023), Rockwell Automation (Rockwell Automation 2023c)—could also help level the playing field for smaller competitors. In some cases, intervention from industry trade organizations or federal regulators (e.g., OSHA) could also be helpful to facilitate industry-wide sharing safety data on the use of Gen-AI tools.

Another point of misallocation can happen at the disconnect between design innovation and shop-floor manufacturability, discussed further in Section 5. Better integration between the tools used by design and manufacturing teams, to improve synchrony and availability of technical information, is an important step to improve efficiency. Unifying and streamlining the design-to-delivery pipeline align with concepts of “Domain-Driven Design”, which suggest that business software should be built from the top-down with the entire business process in mind, instead of permitting different business areas to have their own isolated software ecosystems (Fowler 2002; E. Evans 2003; Vernon 2013; 2016).

3.2 Labor market volatility and long-term talent trends

“Long gone are the days of low-cost factory workers; we need highly skilled, knowledgeable workers.” —Automation expert, Fortune 500 company

There is a future for workers in manufacturing, but several challenges impede the ability of manufacturers to do business and threaten career security (Manufacturing Institute 2019):

- A faster-aging (and already older) workforce than the overall US average, driven by an inability to educate, attract, and retain younger workers;
- Geographic mismatch, driven by worker preferences to live near large urban centers (where few manufacturers do business) or arrange hybrid/remote work;
- A growing skills gap where less-educated or migrant workers may not be sufficiently equipped to manage increasing demand for technical and digital work.
- Emergent business model concepts like a top-down “delegated AI” workflow (i.e., only need one business manager to produce AI output and one lower-level employee to check and deploy results) or “autonomous companies” where even management and strategy decisions are delegated to an AI agent.

Gen-AI tools can provide massive opportunities for worker-augmentation in this setting, based on their demonstrated abilities to distill vast amounts of information—even for workers who may have limited background knowledge or experience. Many early generative and predictive AI tools are already being integrated with the automated data capture systems mentioned above to perform tasks like anomaly detection and monitoring machine deterioration to pinpoint crucial windows for routine or preventative maintenance (Rockwell Automation 2023a). NextGen-AI could offer even more use cases and larger productivity gains—for example, by enabling workers to remotely perform machine maintenance through interactive, 3D augmented reality interfaces with support from spatially informed generative systems. It could also expand the set of tasks that the average worker can perform, elevating middle-skilled trade workers to perform expert-level repair or production work with the guide of Gen-AI platforms trained on prior cases and equipped with advanced computer vision (Acemoglu, Autor, and Johnson 2023).

3.3 Market power and distribution

Already within the field of (generative) artificial intelligence, there are concerning trends toward a highly concentrated market with a small number of players who will wield control over the technological frontier (Acemoglu and Johnson 2023b; 2023a). This is concerning for several reasons, including unilateral ownership and control over who has access to the richest data and can attract and retain the most-capable technical talent, as well as the risk of allowing a singular, hegemonic vision to dominate the future of development for Gen-AI and NextGen-AI tools.

If the direction of NextGen-AI is controlled by a few major players, this will have deleterious distributional consequences for both firms and workers. If small and mid-sized firms do not have avenues to compete with the overwhelming capabilities of the largest operations, the opportunity for emergence of new businesses, business models, products, and innovation will all be diminished. This will leave large companies as the gatekeepers of new technology, enabling them to squeeze smaller competitors who rely on licensing the NextGen-AI products (e.g., OpenAI) or using the platforms of the largest companies (e.g., Amazon) to stay afloat. It could also result in smaller firms exiting the market, or being acquired by the largest companies (if they pose serious competition), leaving only the largest and most powerful companies to hire labor and pay whichever wage they wish for rank-and-file workers. Counteracting these concerns requires a shift in corporate norms of innovation and fair practice—or, if this is unsuccessful (which is likely), strict regulatory oversight from antitrust regulators to ensure that NextGen-AI is developed safely and shaped by healthy competition (Acemoglu and Lensman 2023).

4. Strategic framework for NextGen-AI in Design and Manufacturing

To overcome the shortcomings of current Gen-AI technologies, with regard to their applicability in design and manufacturing as discussed in Section 2, we offer a pathway for developing NextGen-AI technologies. Our recommendation is that new NextGen-AI tools must comport with domain knowledge and technical standards much more effectively than current Gen-AI. They should also align with demonstrated business needs and priorities, which we have identified through conversations with industry experts. NextGen-AI technologies should be capable of assessing a complex manufacturing problem with respect to technical specifications, outlining and explaining actionable steps toward a solution, and (with sufficient human oversight) even executing those steps. Targeting these discrete capabilities will produce powerful tools that can propose diverse and interesting solutions. Taking the full union of traditional design and manufacturing simulation tools, rules, constraints and expertise with NextGen-AI—instead of simply taking the intersection where automation is maximized—will provide more capable tools, empower engineers and operators, lead to exciting new innovations, and improve production and productivity across the fields of manufacturing and design (see Figure 2).

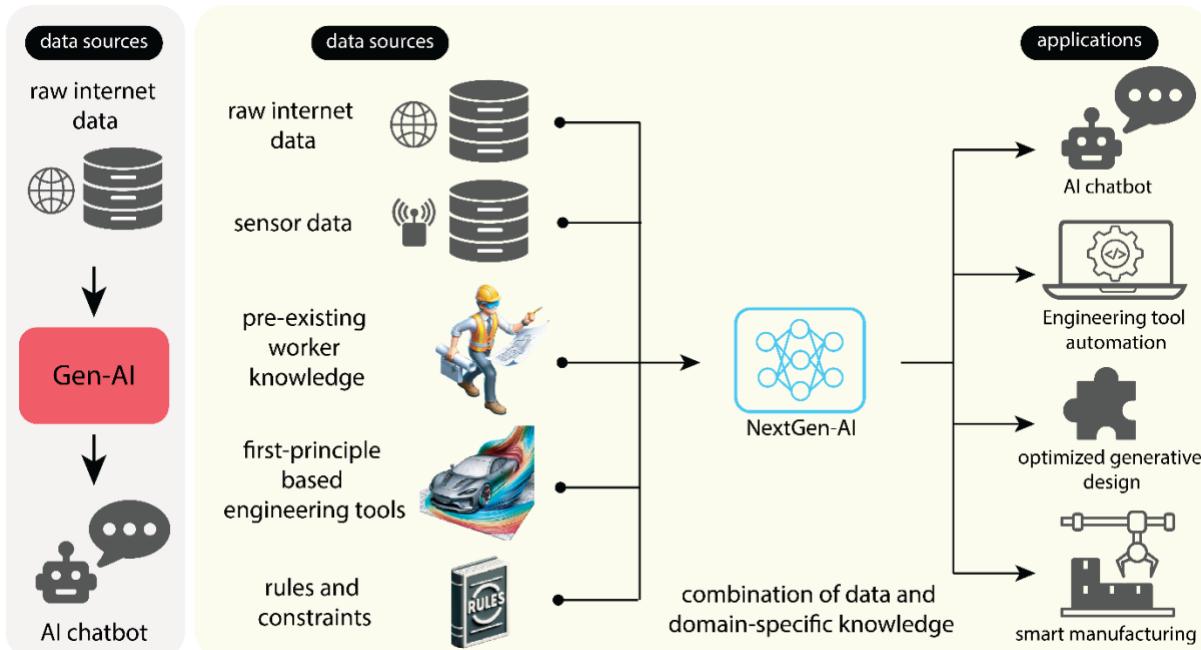


Figure 2: While current Gen-AI technologies are mostly developed using internet data (shown on left), we envision a combination of raw data, sensor data, first-principles-based engineering tools, pre-existing worker/expert knowledge and constraints for developing the NextGen-AI technologies (shown on right). The potential applications of NextGen-AI not only include intelligent AI assistant or ‘co-pilot’, but also seamless use in automation of engineering tools, optimized generative engineering designs and digital/smart manufacturing applications.

4.1 Enhancing Engineering Design with Human-Centric Gen-AI

Gen-AI is capable of rapidly generating a range of solutions for a single problem, contrasting with the traditional approach of many pre-Gen-AI systems that often sought a single, optimal solution. This feature is especially beneficial in engineering design, where the diversity of ideas is crucial to fostering human creativity and innovation (Regenwetter, Nobari, and Ahmed 2022; Nagai and Gero 2012). However, this capability of Gen-AI sometimes conflicts with the current

practices in manufacturing industries focused on design optimization, where the objective often revolves around achieving material efficiency under a singular definition, in adherence to stringent design and manufacturing constraints (requirements where current Gen-AI tools leave much to be desired). Aligning NextGen-AI tools' capabilities with the norms and needs of engineering design is a formidable challenge due to the complexity of high-dimensional design tasks: this underscores why most commercially available Gen-AI tools are currently not well-suited for engineering design. Therefore, the focus for NextGen-AI should be on enhancing human creativity and achieving design objectives by integrating existing design and optimization methods with Gen-AI. Furthermore, involving designers in the development process of NextGen-AI tools can ensure that these tools are intuitive and effectively complement the creative process. It is also vital for NextGen-AI to incorporate feedback loops, allowing continuous improvement and adaptation based on user interactions and real-world testing results.

4.2 Synergizing Traditional Engineering Methods and Gen-AI

In manufacturing, engineering analyses and solutions are heavily dependent on models based on first principles. These domain-specific simulation tools are invariably grounded in the fundamental physics underlying their respective processes. Widely used tools in this domain include finite element analysis (FEA), computational fluid dynamics (CFD), thermal analysis, and the creation of digital twins for complex systems. Although these tools are founded on comprehensive theoretical knowledge, they often struggle to accurately represent the complexities of real-world systems due to issues ranging from a lack of good theory, a lack of accurate data capture, noise, or changing environmental conditions. Consequently, engineers must rely not only on these fundamental scientific principles and processes but also heavily on their intuition, expertise, and judgment when utilizing these tools. The less quantifiable, experience-based skills that engineers bring to the table play a non-trivial role in the successful development of innovative new products.

The novelty of Gen-AI is in its ability to simulate some elements of those less-quantifiable abilities of intuition and judgment. The integration of Gen-AI with these traditional tools offers the potential to enhance their accuracy and predictive power, by incorporating complex, real-world data and learning from it. However, the current Gen-AI often lacks the ability to adhere to the first-principles-based models or real-world constraints required to construct valid and realistic engineering solutions. Encouragingly, recent research suggests that some Gen-AI systems are able to autonomously implement sophisticated tools to perform complex tasks (Schick et al. 2023; Vaithilingam, Zhang, and Glassman 2022; Jin et al. 2023; W. Huang et al. 2022). This synergy between Gen-AI and first-principles models could lead to more robust and nuanced engineering solutions, where AI complements and extends the capabilities of traditional methods.

4.3 Enhancing Workforce Training and Knowledge Transfer

As discussed earlier in Section 3, manufacturing is increasingly reliant on an aging workforce. While some firms have begun offering financial incentives to delay the retirement of long-tenured and experienced workers, this approach merely postpones the inevitable. When those workers do retire, the firm may lose the domain knowledge and fine-tuned production capabilities that the workers have acquired from many years of experience. This poses major risks and could lead to bottlenecks for productivity and efficient operations—historically, it has not been an easy task to replace this experience-based knowledge. Training new operators in heavy-asset manufacturing industries and original equipment manufacturing (OEM) industries is

a resource-intensive undertaking. NextGen-AI can help bridge this gap by offering interim solutions, such as knowledge stopgaps, or serving as a long-lasting, collective institutional memory resource—this will minimize the negative impacts of knowledge loss during talent transitions.

This application for NextGen-AI will also have implications for workers. NextGen-AI could be designed to learn from the expertise of experienced personnel in manufacturing industries. These NextGen-AI tools could also facilitate cross-generational knowledge transfer, ensuring that the invaluable tacit knowledge of veteran workers is passed on effectively to newer employees.

Alternatively, it could be designed as a resource for new employees in these industries by streamlining expensive and time-intensive training and information acquisition procedures that exist today. Incorporating AI-driven mentoring systems, which emulate the decision-making patterns and problem-solving strategies of experienced workers, could further mitigate the loss of expert knowledge. Further, these tools could even be used to upskill existing workers of varying skill levels by offering timely insights or “second opinions” to aid in solving current problems based on training data from previously solved workflow issues. Moreover, integrating NextGen-AI with virtual and augmented reality training modules can create immersive learning experiences that closely mimic real-world scenarios, accelerating the skill acquisition process. NextGen-AI tools will also support a high level of personalization in most aspects of the engineering workflow which may remove many inconvenient tools that exist today.

4.4 Enhancing Factory Floor Dynamics with Gen-AI and Autonomous Machines

Recently, there has been a growing interest in developing autonomous robots using Gen-AI techniques capable of processing unstructured data like language or text (Ren et al. 2023; J. Liang et al. 2023; Li et al. 2023). In contrast, traditional robots and machines used in manufacturing require manual programming and struggle to function in changing environments. Current advancements in Gen-AI hold the potential to simplify robot programming and make it more accessible to operators, utilizing text-to-code capabilities and programming “co-pilots” for experienced and inexperienced programmers alike. Preliminary evidence from tools like GitHub Copilot indicates that such augmentative Gen-AI applications can significantly enhance programmers’ speed (experienced programmers, no less) by more than 50% (Peng et al. 2023).

NextGen-AI should aim to enhance these productivity-boosting capabilities by incorporating advanced reasoning skills that reliably consider physical embodiment and environmental dynamics. In heavy-asset manufacturing industries, where safety risks are prominent, it’s essential that robots and machines are designed to interact safely with human operators on the factory floor. Training of NextGen-AI tools must prioritize human safety principles. We must focus on building NextGen-AI technologies with safety-sensitive features like collaborative robots (Guertler et al. 2023; Faccio et al. 2023). Additionally, integrating advanced sensory systems in robots can enhance their ability to understand and adapt to the physical workspace, thereby improving safety and collaboration. Moreover, focusing on user-friendly interfaces for these AI-enhanced robots and machines will ensure that they are not only technologically advanced but also accessible to workers with varying levels of technical expertise.

4.5 Building Collaborative Data Foundations for NextGen-AI in Manufacturing

To effectively train NextGen-AI with specialized manufacturing knowledge, it’s crucial to develop large-scale, domain-specific datasets. Many current Gen-AI tools are trained using various natural-language or image-based data sources that are available on (or have been

scraped/compiled from) the internet. While this approach yields a generally competent tool, it falls short in manufacturing and design domains that demand highly customized, sector-specific, and even organization-specific capabilities. Moreover, the likelihood of these models being effectively trained using publicly available information is low, as much of the knowledge in manufacturing industries is proprietary. Training robust NextGen-AI models will require careful collection of data. We envision a shared data repository, which could benefit all stakeholders by aiding in the training of *foundational NextGen-AI models*, as illustrated in Figure 3.

This data could originate from the design and manufacturing processes of discontinued products or be collected through comprehensive data aggregation and de-identification techniques to safeguard trade secrets and client privacy. We anticipate that manufacturing industries can benefit from privacy-preserving techniques, such as federated learning (McMahan et al. 2017) and differential privacy (Dwork 2006) for training large-scale Gen-AI models without sharing proprietary data. We seek motivation from recent studies that have explored the possibility of differential privacy and federated learning in cyber-physical and manufacturing systems (Hassan, Rehmani, and Chen 2019; Kevin et al. 2021). Furthermore, establishing industry-wide standards for data collection and sharing can streamline the process, ensuring that data from different sources is compatible and useful for training NextGen-AI models. Additionally, constructing meticulously designed digital twin representations of complex manufacturing systems presents an opportunity for cross-industry collaboration. Engaging in partnerships with academia and research institutions could also accelerate the development of these domain-specific datasets, leveraging their expertise in data science and AI. It is important to keep in mind that data governance should be an integral part of this framework as the preponderance of artificial intelligence has also raised further concerns about intellectual property—especially related to data ownership, access, and compensation (Acemgolu 2024).

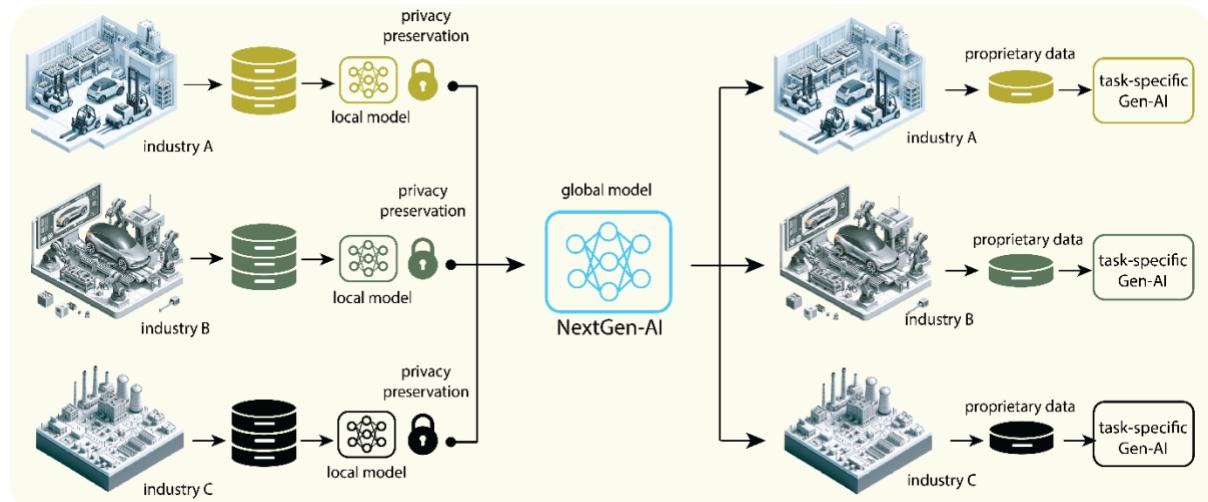


Figure 3: A two-step framework for training NextGen-AI models for manufacturing industries: A large repository of data is compiled from various cross-industry stakeholders to create a foundational Gen-AI model which can be further fine-tuned with industry-specific information and tasks to produce highly tailored NextGen-AI implementations. To protect proprietary data and competitive advantages, industries may opt for privacy-preserving machine learning techniques such as federated learning (McMahan et al. 2017) and differential privacy (Dwork 2006). As a motivating example, we show how federated learning can be employed to train a global Gen-AI model using several local Gen-AI models without sharing data amongst the industries.

4.6 Establishing Robust Evaluation and Feedback Mechanisms

To adopt NextGen-AI technologies in different aspects of manufacturing industries, we must ensure that the technologies are properly designed, incorporated, and continually evaluated to achieve successful and productive implementation. We believe that careful planning must include the following two steps, working backward from the end goal to ensure start-to-finish compatibility with this framework of safety and efficacy. First, we must focus on developing rigorous evaluation benchmarks for NextGen-AI technologies in the context of design and manufacturing. Evaluating current Gen-AI is already challenging, as these platforms are known to hallucinate information, and cannot self-regulate to identify and correct false information. The process of manual fact-checking is time consuming and may negate productivity upsides if it cannot be addressed. The challenge becomes even more difficult in design and manufacturing due to the high level of complexity inherent to engineering and production tasks. We argue the need for multiple evaluation benchmarks specific to each manufacturing industry. The idea of industry-specific Gen-AI evaluation benchmarks is motivated by the traditional use of rigorous benchmarks and standards in the engineering design and manufacturing of products. For example, to generate plausible designs for automotive industries, we must build evaluation benchmarks that take into consideration all design and manufacturing constraints for automotive parts. This evaluation must also include the constraints imposed by industrial standards organizations. Collaboration with industry experts during the benchmark development process will ensure that these standards are realistic, relevant, and effectively address industry-specific challenges.

Next, we advocate building an end-to-end pipeline for NextGen-AI that can provide continuous feedback on deployed AI tools. This is a particularly important step, in order to monitor for unexpected behavior and AI informational hallucinations. Additionally, several validation protocols must be embedded within NextGen-AI deployment, some of which should be mediated by highly skilled human engineers or operators. Implementing human oversight in this process will play a crucial role in ensuring the safe and reliable application of these tools in complex manufacturing settings. Developing continual learning inspired Gen-AI tools capable of real-time learning, self-assessment and error detection could further enhance the reliability and autonomy of these systems.

5. Potential of NextGen-AI in specific industries

To offer additional real-world insights on the possibilities and goals for NextGen-AI in manufacturing, we interviewed leading players in the automotive and footwear industries. These industry leaders—some already adopting this technology and others still assessing its viability—provide insights into how AI will (or will not) transform their operations.

5.1 Automotive industry

The automotive industry has been a pivotal force in shaping the American economy throughout the last century, as mentioned in Section 3. As global automotive manufacturing undergoes a non-AI related technological transformation (i.e., electric vehicles), NextGen-AI offers the opportunity to gain a competitive edge in a highly concentrated and competitive marketplace. We interviewed manufacturing experts at a top-10 automobile manufacturer to understand the opportunity of NextGen-AI in the industry; the quotes that follow are from those conversations.

The design process is fractured and often stilted, involving many specialized teams of designers, engineers, and manufacturers.

“For example, think about a bracket: there’s an engineer who manages the bracket, makes sure it goes through all the different design gates, [etc.]. Then, there is also a designer that they work with. So, if the engineer needs to make changes, they reference the designer and say hey, make ‘XYZ’ changes. The designer makes those, then goes back to the engineer. It’s a little bit less linked than what most people expect.”

This series of iterative hand-offs leads to inefficiencies, as design constraints need to be repeatedly communicated between each expert and the resulting changes need to be validated against those constraints. 70% of the life-cycle costs of a product are determined by decisions taken during the early design stages of long design cycles that may span months to years (Saravi et al. 2008; Iyer et al. 2021). Yet when designs are produced via this iterative process, costly late-stage design changes are often required. NextGen-AI models could help produce designs that satisfy engineering and manufacturing constraints earlier in the design lifecycle, reducing the costs of a more-iterative process and disruption-prone process.

“If you have [a NextGen-AI tool] you can ask: ‘I’m trying to make this part...what processes do you recommend and can you compile all of those?’ to give you some first-guesses, helping the first-time quality of that design. I think that’s where the value is. Not so much generating the design, it’s more so pointing them in the right direction and giving them the resources to help inform their decisions.”

As suggested in Section 3, NextGen-AI tools are likely to be most useful when designed for expert-reference purposes—supporting designers by providing improved access to resources—instead of as direct design-generation tools. At present, this is the preferred direction of development by many manufacturing design experts.

“We have [many] diagrams for a specific manufacturing process. Interacting with those resources could be easier. If you’re looking at this 60-page document and you’re trying to find what the radius is [of a specific part], then you can ask [the NextGen-AI tool]”

Given the substantial data demands for training Gen-AI models, the capabilities of Next-Gen AI tools will be directly dependent on the willingness of companies to collect and share their design and manufacturing data. Finding effective strategies to incentivize companies to share their data (discussed further in Section 3) will be an important challenge.

“Sharing of any data will be difficult...due to the level of confidentiality [companies uphold] for even the most uninteresting of parts.”

5.2 Footwear industry

NextGen-AI could augment the creative process of design and facilitate functional aspects of manufacturability for complex, performance-optimized footwear. The advent of such smart footwear opens up exciting new avenues for integrating technology into everyday wearables (Rhoades 2023). For example, AI-enhanced shoes can offer features such as activity tracking, gait analysis, and even adaptive cushioning, thus merging fashion with functionality and personal health monitoring. Personalization is also another area where NextGen-AI can contribute to the footwear industry (Binelli et al. 2023). A shift from “mass production” to “mass customization”

could set a new standard for consumer expectations. Further, new AI tools are already exploring ways to accelerate the process of moving from concept design to prototyping (Suessmuth, Fick, and Van Der Vossen 2023).

New Balance is an athletic footwear design company that has been developing Gen-AI tools for the past three years. These tools range from classical optimization techniques to modern diffusion models. We spoke to a team including CAD designers, computational design experts, and sports researchers to learn more about the development and integration of NextGen-AI tools for their workflow. Though eager to see the development and integration of AI tools, they acknowledge that current advancements in generative models have yet to practically influence the design process. They underscore a substantial frontier where NextGen-AI could support footwear design and manufacturing.

“There’s a lot of time in design hand-off. Like coming up with a 2D concept, and then moving it into 3D, and then moving that 3D into a set of technical drawings, then the factory is using [those drawings] to build their prototypes.”

Generative modeling capabilities are currently limited to the conceptual stage of design, mainly due to their inability to provide reliable manufacturability and product performance. Specifically, Gen-AI models are not designed with optimized engineering problem-solving as a core competency. The ideal path for NextGen-AI development requires deep-learning architectures built specifically for solving engineering problems. For example, scaling generative models from 2D design concepts to 3D geometries is a promising area of inquiry.

“...not everyone realizes that [a] 2D drawing is not necessarily something we can go off and build right away...[current Gen-AI tools] don’t know how to realize something and make it work really well...the visual data doesn’t necessarily correlate to manufacturability, to comfort, to performance...”

The representational power of deep learning can offer new findings from old data. As with many other advanced manufacturing industries, footwear companies already track and store data in vast quantities from many sources throughout design and production processes.

“...we have lots of data. So much data. We have material databases that all of our testing gets put into... We have multiple components of sports research... We do interviews with athletes...there’s the biomechanics and physiology of the athletes that we collect... The fit and durability of our shoes gets collected.”

NextGen-AI models can leverage these data resources to provide advanced problem-solving capabilities and keep pace with the ever-shifting focus of the footwear industry.

“The creation of our products—the chemistry that goes into them, creating optimal formulations. [For example, with] “super shoes”, you’re trying to get super-resilience out of the materials. How do you take varying formulations and pull them into a model and predict what the optimal one is? ...Fit is super important. How do you take lots of feet and feed them [into a model] with preference data...to make [shoes] fit optimally?”

Even if models can be trained to accurately predict test results, NextGen-AI tools may ultimately be limited by the tests designed and implemented by human operators. This suggests that

NextGen-AI tools will require flexibility to adapt to new testing strategies—or perhaps even help design new and better tests.

“As good as our mechanical tests are, we’re still constantly surprised [by] something broken in real life that we’ve not predicted through any of our mechanical tests... To improve...it’s more about changing the mechanical tests than about scrubbing the data.”

The utility of a generative model does not end in the engineering domain. If engineering is the process of achieving a target performance, the next step in the design cycle is determining whether that performance is useful to a consumer.

“We can take an external scan of someone’s foot and we can take the internal scan of someone’s shoe, but the preference element is...slightly unscientific. We don’t know where someone’s pain receptors are going to be...you’re trying to assign a number to human perception or human feelings of comfort, which are quite subjective.”

NextGen-AI tools do not need to encapsulate the entire design cycle to be useful, nor is automating the entire design cycle inherently desirable. Although some speculators contend that the future of AI should aim to achieve full automation of the design process, industry experts agree that AI tools will likely be much more useful as co-pilots for skilled workers.

“For instance, Photoshop has the diffusion model running in the background, and actually, I think it works better for a Photoshop user [as a background process]. I think Vizcom AI [also] shines because it aligns with the existing skill sets and experiences of designers, so they can use it [for] sketching, rendering these sketches and so on, instead of [just] ‘prompting’ it to generate an image.”

While NextGen-AI tools may automate some parts of the design process, they are best when designed to complement preexisting skills, keeping humans in the loop. Instead of attempting to encode and replace knowledge and competency, these tools should allow an engineer to better leverage their experience and intuition—in pursuit of generating greater insights than either a human or an AI-driven tool could accomplish alone. For example, design tools like CAD co-pilots could improve designers’ capabilities by providing recommendations, searching for information, or automating simple tasks.

“The challenge, from the AI side is this: not every experiential and perceptual thing can be quantified... I think [perception tests] are irreplaceable, to be honest. I don’t expect any fully perceptual, sensational understanding to be replaced by AI, anytime, ever.”

Nevertheless, the proliferation of Gen-AI based tools will inevitably change the skillsets that companies demand, as well as the criteria that firms use to determine the potential of new talent.

“I recently looked at a couple entry-level portfolios—and so much, historically, of what we’ve evaluated in those entry-level portfolios was a couple of really cool-looking drawings and sketches... That’s not a criterion you can use to assess a designer’s potential or talent anymore. There’s definitely some tricky disruption.”

Conclusion

Our view on the development of NextGen-AI technologies for engineering design and manufacturing is guided by the limitations and trajectory of current Gen-AI technologies, insights from industry experts and long-term trends, historical analysis of technological change, and opportunities for developments in NextGen-AI that can mutually benefit firms, workers, consumers, and the manufacturing sector. Our discussions with manufacturing experts and industry leaders motivate our concern that the current capabilities of Gen-AI tools are not sufficient to deliver significant productivity upsides for manufacturing and design. However, we believe that it is possible to overcome many of these challenges by considering domain-specific challenges while developing these powerful models.

We shed light on several issues that must be addressed by the next generation of Gen-AI tools, including improved robustness, accuracy, domain knowledge, reasoning, input data, standards adherence, and workflow integration capabilities. We also identify the economic challenges and upsides to the successful development of NextGen-AI, such as optimizing resource allocation, addressing labor market volatility and long-term talent trends, the increased market power of large firms, and the risks of unevenly distributed benefits from new NextGen-AI tools across firms and workers. We also offer a strategic framework for developing useful NextGen-AI, which must include embedding first-principles design techniques, supporting human creativity in the design process, digitizing expertise, identifying appropriate and inappropriate processes for automation, and broad-based collaboration among (especially smaller and mid-sized) manufacturing firms and industries. Finally, we encourage industry leaders and regulators to consider the importance of encouraging systems integration, clearly defining and protecting data ownership and governance guidelines, including design engineers and manufacturing workers in the development process, collaboration across firms and industries to promote shared interest in the development of powerful new tools, and a focus on the safety and usability of these tools by workers across education and skill levels. NextGen-AI is an exciting frontier for productivity, quality, and creativity in manufacturing and design. The trajectory of how these tools are designed, trained, and implemented will not be automatic—the outcomes depend on the priorities we set and the choices we make.

References

Acemoglu, Daron. 2024. “Harms of AI.” In *The Oxford Handbook of AI Governance*, edited by Justin Bullock, Yu-Che Chen, Johannes Himmelreich, Valerie Hudson, Anton Korinek, Matthew Young, and Baobao Zhang. Pre-print available at: <https://www.nber.org/papers/w29247>. New York: *Oxford University Press*.

Acemoglu, Daron, David Autor, and Simon Johnson. 2023. “Can We Have Pro-Worker AI? Choosing a Path for Machines in Service of Minds.” *MIT Shaping the Future of Work Initiative*, Policy Memo 2023. <https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>.

Acemoglu, Daron, and Simon Johnson. 2023a. “Big Tech Is Bad. Big AI Will Be Worse.” *New York Times*. <https://www.nytimes.com/2023/06/09/opinion/ai-big-tech-microsoft-google-duopoly.html>.

———. 2023b. *Power and Progress: Our 1000-Year Struggle Over Technology and Prosperity*. Hachette PublicAffairs.

Acemoglu, Daron, Simon Johnson, and Austin Lentsch. 2023. “The Hollywood Writers’ AI Fight Is Everyone’s Fight.” *Project Syndicate*. <https://www.project-syndicate.org/commentary/ai-wga-writers-strike-future-of-knowledge-work-by-daron-acemoglu-et-al-2023-08>.

Acemoglu, Daron, and Todd Lensman. 2023. “Regulating Transformative Technologies.” *NBER Working Papers*. <https://www.nber.org/papers/w31461>.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo. 2020. “Does the U.S. Tax Code Favor Automation?” *Brookings Papers on Economic Activity* 2020: 231–300.

Acemoglu, Daron, and Pascual Restrepo. 2019a. “Artificial Intelligence, Automation, and Work.” In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 197–236. Chicago, IL: *University of Chicago Press*. <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/artificial-intelligence-automation-and-work>.

———. 2019b. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.

———. 2020. “Robots and Jobs: Evidence from U.S. Labor Markets.” *Journal of Political Economy* 128 (6): 2188–2244.

———. 2022. “Tasks, Automation, and the Rise in U.S. Wage Inequality.” *Econometrica* 90 (5): 1973–2016.

Alam, Md Ferdous. 2023. *Efficient Sequential Decision Making in Design, Manufacturing and Robotics*. The Ohio State University.

Alam, Md Ferdous, Max Shtein, Kira Barton, and David Hoelzle. 2020. “Autonomous Manufacturing Using Machine Learning: A Computational Case Study with a Limited Manufacturing Budget.” *International Manufacturing Science and Engineering Conference* 84263.

———. 2021. “A Physics-Guided Reinforcement Learning Framework for an Autonomous Manufacturing System with Expensive Data,” 484–90.

———. 2022. “Reinforcement Learning Enabled Autonomous Manufacturing Using Transfer Learning and Probabilistic Reward Modeling.” *IEEE Control Systems Letters* 7: 508–13.

Alfaro, Laura, and Davin Chor. 2023. “Global Supply Chains: The Looming ‘Great Reallocation.’” *NBER Working Papers*. <https://www.nber.org/papers/w31661>.

Anthropic. 2023. “Claude.” <https://claude.ai>.

Autor, David. 2019. “Work of the Past, Work of The Future.” *AEA Papers and Proceedings* 109 (2019): 1–32.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2022. “New Frontiers: The Origins and Content of New Work, 1940–2018.” *NBER Working Papers*. <https://www.nber.org/papers/w30389>.

Autor, David, David Dorn, and Gordon Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103 (6): 2121–68.

Bagazinski, Noah J., and Faez Ahmed. 2023. “ShipGen: A Diffusion Model for Parametric Ship Hull Generation with Multiple Objectives and Constraints.” *arXiv*. <https://arxiv.org/abs/2311.06315v2>.

Bender, Emily M, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23.

Binelli, Marco, Ryan van Dommelen, Yannick Nagel, Jaemin Kim, Rubaiyet Haque, Fergal Coulter, Gilberto Siqueira, André Studart, and Danick Briand. 2023. “Digital Manufacturing of Personalised Footwear with Embedded Sensors.” *Scientific Reports* 13 (1962).

Bistarkey, Devon. 2022. “U.S. Manufacturing Ecosystem Key to Economic Growth, Innovation, Competitiveness.” *U.S. Department of Defense*. <https://www.defense.gov/News/News-Stories/Article/Article/3189049/us-manufacturing-ecosystem-key-to-economic-growth-innovation-competitiveness/>.

Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai-Nayar. 2020. “Multinationals, Offshoring, and the Decline of U.S. Manufacturing.” *Journal of International Economics* 127: 103391.

Brown, Drusilla, Alan Deardorff, and Robert Stern. 2004. “The Effects of Multinational Production on Wages and Working Conditions in Developing Countries.” In *Challenges to Globalization: Analyzing the Economics*, edited by Robert E Baldwin and L Alan Winters. *University of Chicago Press*.

Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, and others. 2020. “Language Models Are Few-Shot Learners.” In *Advances in Neural Information Processing Systems*, 33:1877–1901. 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Vancouver, Canada.
https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2023. “Generative AI at Work.” *NBER Working Papers*. <https://www.nber.org/papers/w31161>.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W.W. Norton.

Capraro, Valerio, Austin Lentsch, Daron Acemgolu, Selin Akgun, Aisel Akhmedova, Ennio Bilancini, Jean-François Bonnefon, et al. 2023. “The Impact of Generative AI on Socioeconomic Inequalities and Policymaking.” *SSRN Working Papers*.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4666103.

Carboneau, Real, Kevin Laframboise, and Rustam Vahidov. 2008. “Application of Machine Learning Techniques for Supply Chain Demand Forecasting.” *European Journal of Operational Research* 184 (3): 1140–54.

Carvalho, Thyago Peres, Fabrizzio A.A.M.N. Soares, Roberto Vita, Roberto da P. Francisco, João P. Basto, and Symone G.S. Alcalá. 2019. “A Systematic Literature Review of Machine Learning Methods Applied to Predictive Maintenance.” *Computers & Industrial Engineering* 137: 106024. <http://dx.doi.org/10.1016/j.cie.2019.106024>.

Cavalcante, Ian M., Enzo M. Frazzon, Fernando A. Forcellini, and Dmitry Ivanov. 2019. “A Supervised Machine Learning Approach to Data-Driven Simulation of Resilient Supplier Selection in Digital Manufacturing.” *International Journal of Information Management* 49: 86–97.

Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. “Imagenet: A Large-Scale Hierarchical Image Database.” In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–55. IEEE.

Deng, Tianchi, Yingguang Li, Xu Liu, and Lihui Wang. 2023. “Federated Learning-Based Collaborative Manufacturing for Complex Parts.” *Journal of Intelligent Manufacturing* 34 (7): 3025–38.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. “BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding.” *arXiv*. <https://arxiv.org/abs/1810.04805>.

Dixon, John R, and Clive L Dym. 1986. “Artificial Intelligence and Geometric Reasoning in Manufacturing Technology.” *Applied Mechanics Reviews* 39(9): 1325–30.

Driess, Danny, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, et al. 2023. “PaLM-E: An Embodied Multimodal Language Model.” In *Proceedings of the 40th International Conference on Machine Learning*, 202:8469–88. Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v202/driess23a.html>.

Dwork, Cynthia. 2006. “Differential Privacy.” In *International Colloquium on Automata, Languages, and Programming, ICALP 2006*, edited by Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, 4052: 1–12. Lecture Notes in Computer Science. Venice, Italy: Springer.

Dziri, Nouha, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Peter West, et al. 2023. “Faith and Fate: Limits of Transformers on Compositionality.” *arXiv*. <https://arxiv.org/abs/2305.18654>.

Evans, Eric. 2003. *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Addison-Wesley Professional.

Evans, Oliver. 1795. *The Young Mill-Wright and Miller’s Guide*. Philadelphia, PA: Self-published.

Faccio, Maurizio, Irene Granata, Alberto Menini, Mattia Milanese, Chiara Rossato, Matteo Bottin, Riccardo Minto, et al. 2023. “Human Factors in Cobot Era: A Review of Modern Production Systems Features.” *Journal of Intelligent Manufacturing* 34 (1): 85–106.

Ferguson, Eugene. 1980. *Oliver Evans, Inventive Genius of the American Industrial Revolution*. Greenville, DE: The Hagley Museum and Library.

Fowler, Martin. 2002. *Patterns of Enterprise Application Architecture*. Addison-Wesley Professional.

Francois, Marianne M, Amy Sun, Wayne E King, Neil Jon Henson, Damien Tourret, Ccut Allan Bronkhorst, Neil N Carlson, et al. 2017. “Modeling of Additive Manufacturing Processes for Metals: Challenges and Opportunities.” *Current Opinion in Solid State and Materials Science* 21 (4): 198–206.

Furman, Jason, and Robert Seamans. 2019. “AI and the Economy.” *Innovation Policy and the Economy* 19 (1): 161–91.

Gao, Leo, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, et al. 2020. “The Pile: An 800GB Dataset of Diverse Text for Language Modeling.” *arXiv*. <https://arxiv.org/abs/2101.00027>.

Giannone, Giorgio, and Faez Ahmed. 2023. “Diffusing the Optimal Topology: A Generative Optimization Approach.” *arXiv*. <https://arxiv.org/abs/2303.09760>.

Giannone, Giorgio, Akash Srivastava, Ole Winther, and Faez Ahmed. 2023. “Aligning Optimization Trajectories with Diffusion Models for Constrained Design Generation.” *arXiv*. <https://arxiv.org/abs/2305.18470>.

Google. 2023. “Bard.” <https://bard.google.com>.

Guertler, Matthias, Laura Tomidei, Nathalie Sick, Marc Carmichael, Gavin Paul, Annika Wambsganss, Victor Hernandez Moreno, and Sazzad Hussain. 2023. “When is a Robot a Cobot? Moving Beyond Manufacturing and Arm-Based Cobot Manipulators.” *Proceedings of the Design Society* 3: 3889–98.

Gupta, S. K. 2023a. “Enabling High Manufacturing Quality And Trustworthiness With AI-Powered Robots.” *Forbes*. <https://www.forbes.com/sites/forbestechcouncil/2023/08/11/how-robot-use-in-manufacturing-can-impact-environmental-sustainability/?sh=458d248262ad>.

———. 2023b. “How Generative AI Can Accelerate The Deployment Of Autonomous Robots.” *Forbes*. <https://www.forbes.com/sites/forbestechcouncil/2023/09/28/how-generative-ai-can-accelerate-the-deployment-of-autonomous-robots/?sh=702d13a7414c>.

———. 2023c. “The Importance Of Human-Centered Automation In Manufacturing.” *Forbes*. <https://www.forbes.com/sites/forbestechcouncil/2023/10/30/the-importance-of-human-centered-automation-in-manufacturing/?sh=e7d0d5549f58>.

Hassan, Muneeb Ul, Mubashir Husain Rehmani, and Jinjun Chen. 2019. “Differential Privacy Techniques for Cyber Physical Systems: A Survey.” *IEEE Communications Surveys & Tutorials* 22 (1): 746–89.

Hindle, Brooke, and Steven Lubar. 1986. *Engines of Change: The American Industrial Revolution, 1790–1860*. Smithsonian Institution Press, Washington DC.

Ho, Jonathan, Ajay Jain, and Pieter Abbeel. 2020. “Denoising Diffusion Probabilistic Models.” In *Advances in Neural Information Processing Systems*, 33:6840–51. 34th International Conference on Neural Information Processing Systems (NeurIPS 2020). Vancouver, Canada. <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.

Hooker, Sara. 2021. “The Hardware Lottery.” *Communications of the ACM* 64 (12): 58–65.

Huang, Jie, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. “Large Language Models Cannot Self-Correct Reasoning Yet.” *arXiv*. <https://arxiv.org/abs/2310.01798>.

Huang, Wenlong, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. “Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents.” In *International Conference on Machine Learning*, 9118–47. PMLR.

Iyer, Naresh S, Amir M Mirzendeht, Sathyaranayanan Raghavan, Yang Jiao, Erva Ulu, Morad Behandish, Saigopal Nelaturi, and Dean M Robinson. 2021. “PATO: Producibility-Aware Topology Optimization Using Deep Learning for Metal Additive Manufacturing.” *arXiv*. <https://arxiv.org/abs/2112.04552>.

Jin, Qiao, Yifan Yang, Qingyu Chen, and Zhiyong Lu. 2023. “GeneGPT: Augmenting Large Language Models with Domain Tools for Improved Access to Biomedical Information.” *arXiv*. <https://arxiv.org/abs/2304.09667>.

Karamcheti, Siddharth, Suraj Nair, Annie S Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, and Percy Liang. 2023. “Language-Driven Representation Learning for Robotics.” *arXiv*. <https://arxiv.org/abs/2302.12766>.

Kevin, I, Kai Wang, Xiaokang Zhou, Wei Liang, Zheng Yan, and Jinhua She. 2021. “Federated Transfer Learning Based Cross-Domain Prediction for Smart Manufacturing.” *IEEE Transactions on Industrial Informatics* 18 (6): 4088–96.

Krizhevsky, Alex. 2009. “Learning Multiple Layers of Features from Tiny Images.” University of Toronto. <https://learning2hash.github.io/publications/cifar2009learning/>.

Le, Matthew, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, et al. 2023. “Voicebox: Text-Guided Multilingual Universal Speech Generation at Scale.” *arXiv*. <https://arxiv.org/abs/2306.15687>.

Leela AI. 2023. “leela.ai.” <https://leela.ai/>.

Li, Boyi, Philipp Wu, Pieter Abbeel, and Jitendra Malik. 2023. “Interactive Task Planning with Language Models.” *arXiv*. <https://arxiv.org/abs/2310.10645>.

Liang, Jacky, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. 2023. “Code as Policies: Language Model Programs for Embodied Control.” In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 9493–9500. IEEE.

Liang, Percy, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, et al. 2022. “Holistic Evaluation of Language Models.” *arXiv*. <https://arxiv.org/abs/2211.09110>.

Lightman, Hunter, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. “Let’s Verify Step by Step.” *arXiv*. <https://arxiv.org/abs/2305.20050>.

Lumafield. 2023. “Lumafield.” <https://www.lumafield.com/>.

Makatura, Liane, Michael Foshey, Bohan Wang, Felix Hähnlein, Pingchuan Ma, Bolei Deng, Megan Tjandrasuwita, et al. 2023. “How Can Large Language Models Help Humans in Design and Manufacturing?” *arXiv*. <https://arxiv.org/abs/2307.14377>.

Manufacturing Institute. 2023. “Creating Pathways for Tomorrow’s Workforce Today: Beyond Reskilling in Manufacturing.” <https://themanufacturinginstitute.org/research/creating-pathways-for-tomorrows-workforce-today-beyond-reskilling-in-manufacturing/>.

———. 2019. “The Aging of the Manufacturing Workforce: Challenges and Best Practices.” <https://themanufacturinginstitute.org/research/the-aging-of-the-manufacturing-workforce/>.

Marr, Bernard. 2023. “Artificial Intelligence In Manufacturing: Four Use Cases You Need To Know In 2023.” *Forbes*. <https://www.forbes.com/sites/bernardmarr/2023/07/07/artificial-intelligence-in-manufacturing-four-use-cases-you-need-to-know-in-2023/?sh=71f5c9453bd8>.

Mazé, François, and Faez Ahmed. 2023. “Diffusion Models Beat Gans on Topology Optimization.” *arXiv*. <https://arxiv.org/abs/2208.09591>.

McClelland, Ryan. 2022. “Generative Design and Digital Manufacturing: Using AI and Robots to Build Lightweight Instruments.” In *SPIE Optics and Photonics*. International Society for Optics and Photonics. San Diego, CA.

McMahan, Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. “Communication-Efficient Learning of Deep Networks from Decentralized Data.” In *Artificial Intelligence and Statistics*, 1273–82. PMLR.

Midjourney Inc. 2023. “Midjourney.” 2023. <https://www.midjourney.com>.

Miller, Scarlett R, Samuel T Hunter, Elizabeth Starkey, Sharath Ramachandran, Faez Ahmed, and Mark Fuge. 2021. “How Should We Measure Creativity in Engineering Design? A Comparison between Social Science and Engineering Approaches.” *Journal of Mechanical Design* 143 (3): 031404.

Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. “Model Cards for Model Reporting.” In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–29.

Mont Blanc AI. 2023. “MontBlancAI.” <https://montblanc.ai/>.

Mündler, Niels, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2023. “Self-Contradictory Hallucinations of Large Language Models: Evaluation, Detection and Mitigation.” *arXiv*. <https://arxiv.org/abs/2305.15852>.

Nagai, Yukari, and John Gero. 2012. “Design Creativity.” *Journal of Engineering Design*. 23(4): 237–9.

National Academies of Science, Engineering, and Medicine. 2017. *The Fourth Industrial Revolution: Proceedings of a Workshop in Brief*. The National Academies Press.

Noy, Shakked, and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381 (6654): 187–92.

OpenAI. 2023a. “ChatGPT.” <https://chat.openai.com/>.

———. 2023b. “DALL·E.” <https://openai.com/dall-e-3>.

———. 2023c. “GPT-4 V(Ision).”

Ouyang, Long, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, et al. 2022. “Training Language Models to Follow Instructions with Human Feedback.” *Advances in Neural Information Processing Systems* 35: 27730–44.

Para, Wamiq, Shariq Bhat, Paul Guerrero, Tom Kelly, Niloy Mitra, Leonidas J. Guibas, and Peter Wonka. 2021. “SketchGen: Generating Constrained CAD Sketches.” In *Advances in Neural Information Processing Systems*, 34:5077–88. 35th Conference on Neural Information Processing Systems (NeurIPS 2021).
https://proceedings.neurips.cc/paper_files/paper/2021/file/28891cb4ab421830acc36b1f5fd6c91e-Paper.pdf.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.” *arXiv*.
<https://arxiv.org/abs/2302.06590>.

Picard, Cyril, Kristen M Edwards, Anna C Doris, Brandon Man, Giorgio Giannone, Md Ferdous Alam, and Faez Ahmed. 2023. “From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design.” *arXiv*. <https://arxiv.org/abs/2311.12668>.

Radford, Alec, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. “Robust Speech Recognition via Large-Scale Weak Supervision.” In Proceedings of the 40th International Conference on Machine Learning, 202:28492–518. *Proceedings of Machine Learning Research* (PMLR).
<https://proceedings.mlr.press/v202/radford23a.html>.

Radford, Alec, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. “Language Models Are Unsupervised Multitask Learners.”
<https://api.semanticscholar.org/CorpusID:160025533>.

Rae, Jack W, Anna Potapenko, Siddhant M Jayakumar, and Timothy P Lillicrap. 2019. “Compressive Transformers for Long-Range Sequence Modelling.” In 8th *International Conference on Learning Representations, ICLR 2020*. Addis Ababa, Ethiopia.
https://iclr.cc/virtual_2020/poster_SylKikSYDH.html.

Raji, Inioluwa Deborah, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic

Auditing.” In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44.

Raji, Inioluwa Deborah, Peggy Xu, Colleen Honigsberg, and Daniel Ho. 2022. “Outsider Oversight: Designing a Third Party Audit Ecosystem for Ai Governance.” In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 557–71.

Raven Telemetry. 2023. “Raven.Ai.” <https://raven.ai/>.

Regenwetter, Lyle, Amin Heyrani Nobari, and Faez Ahmed. 2022. “Deep Generative Models in Engineering Design: A Review.” *Journal of Mechanical Design* 144 (7): 071704.

Ren, Allen Z, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, et al. 2023. “Robots That Ask for Help: Uncertainty Alignment for Large Language Model Planners.” *arXiv*. <https://arxiv.org/abs/2307.01928>.

Rhoades, Matthew A. 2023. “Smart Footwear: A Designer’s Perspective.” In *Smart Clothes and Wearable Technology*, 509–27. Elsevier.

Rio-Chanona, Maria del, Nadzeya Laurentsyeva, and Johannes Wachs. 2023. “Are Large Language Models a Threat to Digital Public Goods? Evidence from Activity on Stack Overflow.” *arXiv*. <https://arxiv.org/abs/2307.07367>.

Rockwell Automation. 2023a. “FactoryTalk Analytics LogixAI: Machine Learning and Logix.” <https://www.rockwellautomation.com/en-us/support/documentation/overview/factorytalk-analytics-logixai--machine-learning-and-logix.html>.

———. 2023b. “Plex.” <https://www.plex.com/products/production-monitoring>.

———. 2023c. “Rockwell Automation.” Rockwell Automation. <https://www.rockwellautomation.com/en-us.html>.

Rombach, Robin, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. “High-Resolution Image Synthesis with Latent Diffusion Models.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–95. https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.pdf.

Rozière, Baptiste, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, et al. 2023. “Code Llama: Open Foundation Models for Code.” *arXiv*. <https://arxiv.org/abs/2308.12950>.

Sanghi, Aditya, Hang Chu, Joseph G. Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshan. 2022. “CLIP-Forge: Towards Zero-Shot Text-to-Shape Generation.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18603–13. https://openaccess.thecvf.com/content/CVPR2022/papers/Sanghi_CLIP-Forge_Towards_Zero-Shot_Text-To-Shape_Generation_CVPR_2022_paper.pdf.

Saravi, Mohammad, Linda Newnes, Antony Roy Mileham, and Yee Mey Goh. 2008. “Estimating Cost at the Conceptual Design Stage to Optimize Design in Terms of Performance and Cost.” In *Collaborative Product and Service Life Cycle Management for a Sustainable World*: Proceedings of the 15th ISPE International Conference on Concurrent Engineering (CE2008), 123–30. Springer.

Schick, Timo, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. “Toolformer: Language Models Can Teach Themselves to Use Tools.” *arXiv*. <https://arxiv.org/abs/2302.04761>.

Schwab, Klaus. 2015. “The Fourth Industrial Revolution.” *Foreign Affairs*. <https://www.foreignaffairs.com/world/fourth-industrial-revolution>.

Seepersad, Carolyn Conner. 2014. “Challenges and Opportunities in Design for Additive Manufacturing.” *3D Printing and Additive Manufacturing* 1(1): 10–13.

Shi, Freda, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, and Denny Zhou. 2023. “Large Language Models Can Be Easily Distracted by Irrelevant Context.” In *Proceedings of the 40th International Conference on Machine Learning*, 202:31210–27. Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v202/shi23a.html>.

Sohl-Dickstein, Jascha, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. “Deep Unsupervised Learning Using Nonequilibrium Thermodynamics.” In *Proceedings of the 32nd International Conference on Machine Learning*, 37:2256–65. Proceedings of Machine Learning Research (PMLR). Lille, France. <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.

SpencerMetrics. 2023. “SpencerMetrics.” <https://www.spencermetrics.com/>.

Stability AI. 2023. “Stable Diffusion.” <https://stability.ai/>.

Suessmuth, Jochen, Florian Fick, and Stan Van Der Vossen. 2023. “Generative AI for Concept Creation in Footwear Design.” In *ACM SIGGRAPH 2023 Talks*, 1–2.

Susto, Gian Antonio, Andrea Schirru, Simone Pampuri, Seán McLoone, and Alessandro Beghi. 2015. “Machine Learning for Predictive Maintenance: A Multiple Classifier Approach.” *IEEE Transactions on Industrial Informatics* 11 (3): 812–20.

Touvron, Hugo, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, et al. 2023. “Llama 2: Open Foundation and Fine-Tuned Chat Models.” *arXiv*. <https://arxiv.org/abs/2307.09288>.

Tulip. 2023. “Tulip.” <https://tulip.co/>.

Turpin, Miles, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. “Language Models Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting.” *arXiv*. <https://arxiv.org/abs/2305.04388>.

U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy. 2019. “Manufacturing Timeline.” <https://www.energy.gov/eere/amo/manufacturing-timeline>.

Ulrich, Karl T, and Steven D Eppinger. 2016. *Product Design and Development*. McGraw-Hill.

United Nations. 2021. “Manufacturing Employment as a Proportion of Total Employment (%), 2021.” 2021. <https://ourworldindata.org/grapher/manufacturing-share-of-total-employment?tab=table>.

Usuga Cadavid, Juan Pablo, Samir Lamouri, Bernard Grabot, Robert Pellerin, and Arnaud Fortin. 2020. “Machine Learning Applied in Production Planning and Control: A State-of-the-Art in the Era of Industry 4.0.” *Journal of Intelligent Manufacturing* 31: 1531–58.

Vaithilingam, Priyan, Tianyi Zhang, and Elena L. Glassman. 2022. “Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models.” In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–7. New York, NY: Association for Computing Machinery. <https://doi.org/10.1145/3491101.3519665>.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. “Attention Is All You Need.” In *Advances in Neural Information Processing Systems*, 30:6000–6010. 31st Conference on Neural Information Processing Systems (NIPS 2017). Long Beach, CA, USA. https://papers.nips.cc/paper_files/paper/2017/file/3f5ee243547dee91fdb053c1c4a845aa-Paper.pdf.

Vernon, Vaughn. 2013. *Implementing Domain-Driven Design*. Addison-Wesley Professional.

———. 2016. *Domain-Driven Design Distilled*. Addison-Wesley Professional.

Watson, Nell. 2024. *Taming the Machine*. Kogan Page, London.

West, Darrell, and Christian Lansang. 2018. “Global Manufacturing Scorecard: How the U.S. Compares to 18 Other Nations.” *Brookings Institute*.

Wolf, Yotam, Noam Wies, Yoav Levine, and Amnon Shashua. 2023. “Fundamental Limitations of Alignment in Large Language Models.” *arXiv*. <https://arxiv.org/abs/2304.11082>.

World Bank. 2021. “OECD: Manufacturing, Value Added, 2021 (% of GDP).” 2021. <https://data.worldbank.org/indicator/NV.IND.MANF.ZS>.

———. 2023. “Trouble in the Making? The Future of Manufacturing-Led Development.” 2023. <https://www.worldbank.org/en/topic/competitiveness/publication/trouble-in-the-making-the-future-of-manufacturing-led-development>.

Wu, Rundi, Chang Xiao, and Changxi Zheng. 2021. “DeepCAD: A Deep Generative Network for Computer-Aided Design Models.” In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 6772–82.

https://openaccess.thecvf.com/content/ICCV2021/papers/Wu_DeepCAD_A_Deep_Generative_Network_for_Computer-Aided_Design_Models_ICCV_2021_paper.pdf.

Wu, Yang, Shilong Wang, Hao Yang, Tian Zheng, Hongbo Zhang, Yanyan Zhao, and Bing Qin. 2023. “An Early Evaluation of GPT-4V(ision).” *arXiv*. <https://arxiv.org/abs/2310.16534>.

Xiao, Han, Kashif Rasul, and Roland Vollgraf. 2017. “Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms.” *arXiv*. <https://arxiv.org/abs/1708.07747>.

Xu, Ruiyun, Yue Feng, and Hailiang Chen. 2023. “ChatGPT vs. Google: A Comparative Study of Search Performance and User Experience.” *SSRN Working Papers*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4498671.

Yellen, Janet. 2023. “Remarks by Secretary of the Treasury Janet L. Yellen at Press Conference following Bilateral Meeting with Secretary of Finance and Public Credit Rogelio Ramírez de la O of Mexico.” *U.S. Department of the Treasury*. <https://home.treasury.gov/news/press-releases/jy1964>.

Zha, Yuheng, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. “AlignScore: Evaluating Factual Consistency with a Unified Alignment Function.” *arXiv*. <https://arxiv.org/abs/2305.16739>.

Zhu, Yukun, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. “Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books.” In *Proceedings of the IEEE International Conference on Computer Vision*, 19–27.

Zitkovich, Brianna, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, et al. 2023. “Rt-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control.” In *7th Annual Conference on Robot Learning*, 229: 2165–83. Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v229/zitkovich23a.html>.