

---

# FROM AUTOMATION TO AUGMENTATION: POLICY AND PRACTICE TO REDEFINE 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 contribution

## 1 Executive Summary

*“[T]he outstanding features of **the second machine age**: sustained exponential improvement in most aspects of computing, extraordinarily large amounts of digitized information, and recombinant innovation...have made it possible for humanity to create [one] of the most important one-time events in our history: **the emergence of real, useful artificial intelligence...**”*

— Erik Brynjolfsson and Andrew McAfee<sup>1</sup>

In the mid-2010s, as computing and other digital technologies matured, researchers began to speculate about a new era of innovation—with artificial intelligence as the standard-bearer of a “Fourth Industrial Revolution”.<sup>2</sup> The release of generative AI technologies (e.g. ChatGPT) in late 2022 reignited these discussions, prompting us to wonder: *what are the barriers, risks, and potential rewards to using generative AI (“Gen-AI”) for design and manufacturing?*

During the time that Gen-AI has entered the mainstream, geopolitics and business practices have shifted. Global supply chains were disrupted by the Covid-19 pandemic, tensions with import partners have risen, and military conflicts have introduced new uncertainties. As companies consider propositions like “reshoring,” or “nearshoring/friendshoring” production, we recognize several other hindrances, including: underutilization or misallocation of resources, labor market volatility and long-term talent trends toward an older and geographically mismatched labor force, and highly concentrated tech markets that foster anti-competitive and predatory business practices. As the US expands domestic manufacturing production capacity (e.g., for semiconductors and electric vehicles), we believe that the next generation of Gen-AI tools could help overcome these hindrances and accelerate the deployment of new production technologies. To investigate the current use and future potential of Gen-AI tools in design and manufacturing, we interviewed industry experts—including engineers, manufacturers, tech executives and entrepreneurs. These experts have identified many opportunities in their workflows where Gen-AI could be deployed:

1. Unifying the fragmented and iterative process of design, testing, prototyping, and manufacturing at scale;
2. Providing information to designers and engineers, including by identifying suitable design spaces and material formulations, and by incorporating idiosyncratic consumer preferences into the design process;
3. Improving the interpretation of test data to enable rapid validation and qualification;
4. Democratizing access and usage of manufacturing process data to provide real-time insights;
5. Empowering workers to become more productive—even enabling less-skilled workers to perform expert tasks.

Current Gen-AI solutions (e.g., ChatGPT, Bard, Claude) cannot accomplish these goals. We identify several issues that prevent the usefulness of current Gen-AI tools, including: inability to provide robust, reliable, and replicable output; lack of relevant domain knowledge; unawareness of industry-standards requirements for product quality; disjoint software

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

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

tools that do not integrate seamlessly with existing workflow; and inability to simultaneously interpret data from various sources and formats (e.g., text, images, video, audio, sensor data). To address these problems, we propose the following framework for the development of the next generation of generative artificial intelligence tools (“NextGen-AI”):

1. Augment creativity through NextGen-AI that integrates engineering tools, repositories, and search methods;
2. Integrate foundational adherence to first-principles when solving engineering problems across modalities;
3. Incorporate experience-gained knowledge from an aging workforce, to facilitate talent transitions and training;
4. Empower workers to be more productive, rather than lose productivity upside in pursuit of static automation;
5. Create a collaborative data ecosystem to train foundation models, but preserve data security and ownership;
6. Assess capabilities to continuously ensure that new tools are safe and effective.

These goals for NextGen-AI are extensive, and will require broad-based buy-in not only from business leaders, but also from researchers, engineers, policymakers, and operators. We recommend the following priorities to accelerate “the emergence of real, useful [generative] artificial intelligence” for design and manufacturing:

1. Improve systems integration to ethically collect real-time data in support of NextGen-AI software;
2. Establish clear data governance rules to ensure equal opportunity for all firms in development and ownership;
3. Leverage industry-wide opportunity, such as expanded collection of worker-safety data to assess AI usage;
4. Include engineers and operators in the development process, to maximize usefulness, buy-in, and uptake;
5. Focus on skills-complementary deployments of NextGen-AI to maximize productivity upside.

**Keywords** Generative AI, design, manufacturing, AI policy, Industry 4.0, future of work, automation, augmentation

## 2 Introduction

Manufacturing is a cornerstone of the world economy, accounting for approximately 16% of global GDP and 14% of employment in 2021.<sup>3</sup> In the US, manufacturing accounts for 12.5 million jobs and 35% of productivity growth.<sup>4</sup> A well-functioning manufacturing sector is crucial to many other industries, including consumer goods, information and communication technology, energy, food, healthcare, aerospace, and defense. While the US has long been a preeminent manufacturing power (see Figure 1),<sup>5</sup> 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

3. World Bank, *Manufacturing, value added (% of GDP)*, 2021, <https://data.worldbank.org/indicator/NV.IND.MANF.ZS>; United Nations, *Manufacturing employment as a proportion of total employment (%)*, 2021, <https://ourworldindata.org/grapher/manufacturing-share-of-total-employment?tab=table>.

4. U.S. Department of Defense, *U.S. Manufacturing Ecosystem Key to Economic Growth, Innovation, Competitiveness*, 2023, <https://www.defense.gov/News/News-Stories/Article/Article/3189049/us-manufacturing-ecosystem-key-to-economic-growth-innovation-competitiveness/>.

5. U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy, *Manufacturing Timeline: Advanced Manufacturing & Industrial Decarbonization*, 2019, <https://www.energy.gov/eere/amo/manufacturing-timeline>.

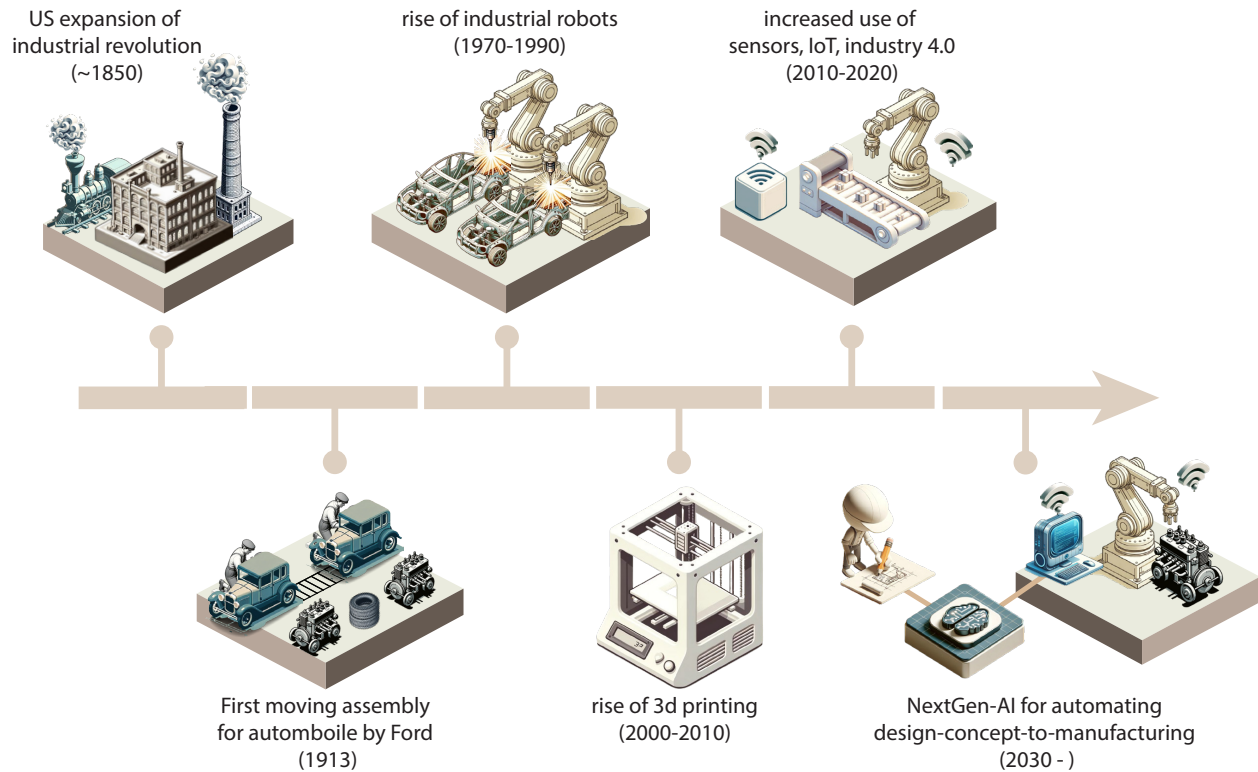


Figure 1: A brief history of manufacturing innovation in the US based on a timeline from the Office of Energy Efficiency and Renewable Energy and authors’ previous works (see footnotes 5–7). We anticipate that the next generation of generative AI (NextGen-AI) technologies will address several limitations of the current generation of generative AI technologies (Gen-AI) and could be the newest frontier for major productivity advancements in design and manufacturing.

domestic labor demand since the late 20<sup>th</sup> century.<sup>6</sup> This strategy has allowed companies to doubly boost their bottom lines by leveraging automation technologies to reduce the need for skilled labor domestically, while also offering the opportunity to squeeze less-skilled workers in emerging economies, paying lower wages and requiring less oversight for working conditions and safety than would be permissible domestically.<sup>7</sup>

In recent years, however, the priorities of firms and national security have shifted toward strategies including more reshoring or nearshoring/friendshoring of production. This shift has been catalyzed by recent supply chain disruptions (e.g., Covid-19) as well as increased tension and uncertainty about US-China relations and ongoing geopolitical conflicts

6. World Bank, *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>; Daron Acemoglu and Pascual Restrepo, “Robots and Jobs: Evidence from U.S. Labor Markets,” *Journal of Political Economy* 128, no. 6 (2020): 2188–2244; David Autor, David Dorn, and Gordon Hanson, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review* 103, no. 6 (2013): 2121–2168; Darrell M. West and Christian Lansang, “Global manufacturing scorecard: How the U.S. compares to 18 other nations,” *Brookings Institute*, 2018, <https://www.brookings.edu/articles/global-manufacturing-scorecard-how-the-us-compares-to-18-other-nations/>.

7. Acemoglu and Restrepo, “Robots and Jobs: Evidence from U.S. Labor Markets”; Christoph Boehm, Aaron Flaaen, and Nitya Pandalai-Nayar, “Multinationals, Offshoring, and the Decline of U.S. Manufacturing,” *Journal of International Economics* 127 (2020): 103391; Drusilla K. Brown, Alan V. Deardorff, and Robert M. Stern, “The Effects of Multinational Production on Wages and Working Conditions in Developing Countries,” in *Challenges to Globalization: Analyzing the Economics*, ed. Robert E. Baldwin and L. Alan Winters (University of Chicago Press, 2004).

(e.g., the Russo-Ukrainian and Israel-Hamas Wars).<sup>8</sup> This coincides with a technological shift in manufacturing, as the rapid advancements of new technologies—AI, advanced analytics, and frameworks like the Industrial Internet of Things (IIoT: increased integration and connectivity across devices, instruments, and sensors)—have convinced many industry experts that manufacturing is undergoing a fourth industrial revolution: “Industry 4.0”.<sup>9</sup>

However, the workflow in design and manufacturing remains a major bottleneck to delivering on the massive productivity promises of these new technologies.<sup>10</sup> 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 individualized workflows.<sup>11</sup>

In recent years, artificial intelligence (AI) has shown tremendous potential to overcome many of these limitations. AI implementations are quickly becoming pervasive in many manufacturing applications such as predictive maintenance,<sup>12</sup> quality control,<sup>13</sup> supply chain optimization,<sup>14</sup> process control<sup>15</sup> and risk management.<sup>16</sup> In order to make good on the promises of AI-based technologies, many areas of manufacturing will require rethinking and restructuring, including the processes of planning and conceptual design, system-level and detail design, fabrication, testing and scaling to

8. Laura Alfaro and Davin Chor, “Global Supply Chains: The Looming “Great Reallocation”,” NBER Working Paper No. 31661, *National Bureau of Economic Research*, 2023, <https://www.nber.org/papers/w31661>.

9. Jason Furman and Robert Seamans, “AI and the Economy,” *Innovation Policy and the Economy* 19, no. 1 (2019): 161–191; National Academies of Science, Engineering, and Medicine, *The Fourth Industrial Revolution: Proceedings of a Workshop in Brief* (The National Academies Press, 2017).

10. Marianne M. Francois et al., “Modeling of additive manufacturing processes for metals: Challenges and opportunities,” *Current Opinion in Solid State and Materials Science* 21, no. 4 (2017): 198–206; Carolyn Conner Seepersad, “Challenges and Opportunities in Design for Additive Manufacturing,” *3D Printing and Additive Manufacturing* (New Rochelle, NY) 1, no. 1 (2014): 10–13; The Manufacturing Institute, *Creating Pathways for Tomorrow’s Workforce Today: Beyond Reskilling in Manufacturing*, 2023, <https://themanufacturinginstitute.org/research/creating-pathways-for-tomorrows-workforce-today-beyond-reskilling-in-manufacturing/>; Bernard Marr, “Artificial Intelligence In Manufacturing: Four Use Cases You Need To Know In 2023,” *Forbes*, 2023, <https://www.forbes.com/sites/bernardmarr/2023/07/07/artificial-intelligence-in-manufacturing-four-use-cases-you-need-to-know-in-2023/>; S.K. Gupta, “The Importance Of Human-Centered Automation In Manufacturing,” *Forbes*, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/10/30/the-importance-of-human-centered-automation-in-manufacturing/>; S.K. Gupta, “How Generative AI Can Accelerate The Deployment Of Autonomous Robots,” *Forbes*, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/09/28/how-generative-ai-can-accelerate-the-deployment-of-autonomous-robots/>; S.K. Gupta, “Enabling High Manufacturing Quality And Trustworthiness With AI-Powered Robots,” *Forbes*, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/08/11/how-robot-use-in-manufacturing-can-impact-environmental-sustainability/>; Jorge F Arinez et al., “Artificial intelligence in advanced manufacturing: Current status and future outlook,” *Journal of Manufacturing Science and Engineering* 142, no. 11 (2020): 110804.

11. Scarlett R. Miller et al., “How Should We Measure Creativity in Engineering Design? A Comparison Between Social Science and Engineering Approaches,” *Journal of Mechanical Design* 143, no. 3 (2021): 031404.

12. Thyago Peres Carvalho et al., “A systematic literature review of machine learning methods applied to predictive maintenance,” *Computers & Industrial Engineering* 137 (2019): 106024, <http://dx.doi.org/10.1016/j.cie.2019.106024>; Gian Antonio Susto et al., “Machine Learning for Predictive Maintenance: A Multiple Classifier Approach,” *IEEE Transactions on Industrial Informatics* 11, no. 3 (2015): 812–820.

13. Juan Pablo Usuga Cadavid et al., “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 (2020): 1531–1558; Tianchi Deng et al., “Federated learning-based collaborative manufacturing for complex parts,” *Journal of Intelligent Manufacturing* 34, no. 7 (2023): 3025–3038.

14. Real Carbonneau, Kevin Laframboise, and Rustam Vahidov, “Application of machine learning techniques for supply chain demand forecasting,” *European Journal of Operational Research* 184, no. 3 (2008): 1140–1154.

15. Md Ferdous Alam et al., “Reinforcement Learning Enabled Autonomous Manufacturing Using Transfer Learning and Probabilistic Reward Modeling,” *IEEE Control Systems Letters* 7 (2022): 508–513.

16. Ian M. Cavalcante et al., “A supervised machine learning approach to data-driven simulation of resilient supplier selection in digital manufacturing,” *International Journal of Information Management* 49 (2019): 86–97.

production.<sup>17</sup> At the same time, with the rapid advancement of state-of-the-art AI technologies in many other domains, we must take a careful approach to assessing the opportunities and challenges of new AI across the manufacturing domain. 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 key opinion leaders and industry experts.

The current generation of generative artificial intelligence (i.e., “Gen-AI”) is dominated by text-based platforms,<sup>18</sup> image-based platforms,<sup>19</sup> and “multi-modal” platforms that incorporate both images and texts.<sup>20</sup> 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.<sup>21</sup> However, the potential of Gen-AI needs to be critically assessed for engineering applications such as design and manufacturing, which 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 still significantly struggle with complex tasks such as spatial reasoning and understanding complex design problems.<sup>22</sup> Therefore, some goals we identify for NextGen-AI development in design and manufacturing: enable augmented human creativity in (and assist the generation of) distinctive, optimized and actionable design concepts; evaluate designs at higher accuracy or lower cost than prohibitively expensive traditional methods; monitor quality control and preventative

17. Karl T. Ulrich and Steven D. Eppinger, *Product Design and Development*, 6th ed. (McGraw-Hill Education, 2015).

18. Ashish Vaswani et al., “Attention Is All You Need,” in *Advances in Neural Information Processing Systems*, vol. 30, 31st Conference on Neural Information Processing Systems (NIPS 2017) (Long Beach, CA, USA, 2017), 6000–6010, [https://papers.nips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://papers.nips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf); Hugo Touvron et al., “LLaMA: Open and Efficient Foundation Language Models,” arXiv Working Paper No. 2302.13971, *arXiv*, 2023, <https://arxiv.org/abs/2302.13971>; Jacob Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” arXiv Working Paper No. 1810.04805, *arXiv*, 2018, <https://arxiv.org/abs/1810.04805>; Tom Brown et al., “Language Models Are Few-Shot Learners,” in *Advances in Neural Information Processing Systems*, vol. 33, 34th Conference on Neural Information Processing Systems (NeurIPS 2020) (Vancouver, Canada, 2020), 1877–1901, [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf); Alec Radford et al., “Language Models are Unsupervised Multitask Learners” (2019), <https://api.semanticscholar.org/CorpusID:160025533>.

19. Robin Rombach et al., “High-Resolution Image Synthesis with Latent Diffusion Models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 10684–10695, [https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach\\_High-Resolution\\_Image\\_Synthesis\\_With\\_Latent\\_Diffusion\\_Models\\_CVPR\\_2022\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.pdf); Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising Diffusion Probabilistic Models,” in *Advances in Neural Information Processing Systems*, vol. 33, 34th International Conference on Neural Information Processing Systems (NeurIPS 2020) (Vancouver, Canada, 2020), 6840–6851, <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>; Jascha Sohl-Dickstein et al., “Deep Unsupervised Learning using Nonequilibrium Thermodynamics,” in *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37, Proceedings of Machine Learning Research (PMLR) (Lille, France, 2015), 2256–2265, <https://proceedings.mlr.press/v37/sohl-dickstein15.html>; Baptiste Rozière et al., “Code Llama: Open Foundation Models for Code,” arXiv Working Paper No. 2308.12950, *arXiv*, 2023, <https://arxiv.org/abs/2308.12950>.

20. Danny Driess et al., “PaLM-E: An Embodied Multimodal Language Model,” in *Proceedings of the 40th International Conference on Machine Learning*, vol. 202, Proceedings of Machine Learning Research (PMLR) (2023), 8469–8488, <https://proceedings.mlr.press/v202/driess23a.html>.

21. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, “Generative AI at Work,” NBER Working Paper No. 31161, *National Bureau of Economic Research*, 2023, <https://www.nber.org/papers/w31161>; Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381, no. 6654 (2023): 187–192; Sida Peng et al., “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” arXiv Working Paper No. 2302.06590, *arXiv*, 2023, <https://arxiv.org/abs/2302.06590>.

22. Cyril Picard et al., “From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design,” arXiv Working Paper No. 2311.12668, *arXiv*, 2023, <https://arxiv.org/abs/2311.12668>.

maintenance; provide personalized worker training and real-time problem-solving; optimize supply chain management practices. 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; disjointedness of software that precludes seamless integration with existing workflow; and inability to simultaneously interpret data from various sources and formats (e.g., text, images, video, audio, sensor data).

Based on the current trajectory of Gen-AI, we construct a strategic framework 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:

- **Concept Innovator:** Offering diverse and creative solutions for complex engineering design tasks with the possibility to automate design concept generation and evaluation, integrating multimodal information (e.g., materials, processes, geometries, carbon footprint). These models could assist in designing products tailored to individual customer preferences or specific market segments.
- **Task automation:** Accelerating detail-oriented tasks and workflows (e.g., certification of designs, conversion of data into different modes).
- **Expertise database:** Cloning unstructured pre-existing design and manufacturing domain expertise.
- **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 potential risks in the manufacturing process or product lifecycle and suggest decisions to mitigate risks.
- **Co-pilot for workers:** Improving information provision to empower workers with a variety of skillsets and educational backgrounds and boost worker productivity, including design engineers, and manufacturing operators. The co-pilot may also enhance human-robot collaboration in manufacturing settings.
- **Task planning:** Dynamically predicting optimal maintenance schedules by considering a multitude of factors. It will also help forecast demand, optimize inventory levels, and identify the most efficient logistics routes. This includes predictive analytics for supply and demand, and real-time adjustments to supply chain disruptions.
- **Digital twin:** Enable the creation of accurate digital twins using synthetic and real data. Using asset data, the models will enable automated quality checks, defect detection, and predictive maintenance.

Although integrating NextGen-AI tools in the design and manufacturing pipeline is attractive for the reasons mentioned above, these tools potentially also raise troubling issues, including the future of the workforce and competitive dynamics between firms. AI is a powerful technological platform and NextGen-AI tools could be designed to automate many

of the current engineering tasks and reduce the importance of skilled workers in the production process, or instead could (also) be used to augment workers' capabilities, increase expertise, and create more well-paid jobs in the coming decades. The direction of development for NextGen-AI is not predetermined: it will be shaped by the choices made by governments, executives, technologists, researchers, and workers.

The paper proceeds as follows: first, we provide a brief discussion on the underlying technical framework of Gen-AI tools, as well as the capabilities and limitations of current-generation Gen-AI in Section 3. Next, we discuss the societal impacts of Gen-AI in the design and manufacturing domain by focusing on its potential impacts on workers and firms, touching on how these impacts could affect the economy more broadly in Section 4. Then, we provide a strategic framework for designing suitable NextGen-AI tools to address current challenges in design and manufacturing in Section 5, with supporting evidence and an exploration of potential use-cases of NextGen-AI for the automotive and footwear industries in Section 6. Finally, we offer a policy guide for industry leaders and regulators to construct an ideal environment for productive and safeguarded development of NextGen-AI in Section 7 and present our concluding thoughts 8.

**Disclaimer:** This paper proposes some ideas and goals for the progression of AI as we transition from current Gen-AI to NextGen-AI tools that are suitable for wide-ranging applications in manufacturing and design. However, we also wish to emphasize that technological development often occurs in unexpected ways. We may overlook certain prohibitive complications of AI development or fail to predict all of the myriad possibilities for disruption in future AI advancements. Nonetheless, we believe that this is an important foray into the future of generative AI.

### 3 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,<sup>23</sup> improvements in computer hardware such as graphics processing units (GPUs)<sup>24</sup> and large-scale datasets.<sup>25</sup>

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,<sup>26</sup> BARD,<sup>27</sup> and Claude<sup>28</sup> are useful as general purpose question answering chatbots.<sup>29</sup> Preliminary evidence suggests that ChatGPT has even captured some web traffic from online Q&A sites like Stack Overflow.<sup>30</sup> 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,<sup>31</sup> writing assistance,<sup>32</sup> cooking instructions<sup>33</sup> and many others. We have seen similar trends in Gen-AI image-generation tools that create illustrations based on text, such as Midjourney,<sup>34</sup> DALL-E,<sup>35</sup> and Stable Diffusion.<sup>36</sup> 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.<sup>37</sup> Beyond these modalities, researchers are also exploring Gen-AI for building text- and image-based control of intelligent robots.<sup>38</sup>

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

---

23. Alex Krizhevsky, *Learning Multiple Layers of Features from Tiny Images*, technical report, Available at: <https://learning2hash.github.io/publications/cifar2009learning/> (University of Toronto, 2009); Vaswani et al., “Attention Is All You Need.”

24. Sara Hooker, “The Hardware Lottery,” *Communications of the ACM* 64, no. 12 (2021): 58–65.

25. Jia Deng et al., “Imagenet: A Large-Scale Hierarchical Image Database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2009), 248–255; Krizhevsky, *Learning Multiple Layers of Features from Tiny Images*; Han Xiao, Kashif Rasul, and Roland Vollgraf, “Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms,” arXiv Working Paper No. 1708.07747, *arXiv*, 2017, <https://arxiv.org/abs/1708.07747>; Leo Gao et al., “The Pile: An 800GB Dataset of Diverse Text for Language Modeling,” arXiv Working Paper No. 2101.00027, *arXiv*, 2020, <https://arxiv.org/abs/2101.00027>.

26. OpenAI, *ChatGPT*, 2023, <https://chat.openai.com>.

27. Google, *Bard*, 2023, <https://bard.google.com>.

28. Anthropic, *Claude*, 2023, <https://claude.ai>.

29. Ruiyun Xu, Yue Feng, and Hailiang Chen, “ChatGPT vs. Google: A Comparative Study of Search Performance and User Experience,” SSRN Working Paper No. 4498671, *SSRN*, 2023, [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4498671](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4498671).

30. Maria del Rio-Chanona, Nadzeya Laurentsyeva, and Johannes Wachs, “Are Large Language Models a Threat to Digital Public Goods? Evidence from Activity on Stack Overflow,” arXiv Working Paper No. 2307.07367, *arXiv*, 2023, <https://arxiv.org/abs/2307.07367>.

31. Rozière et al., “Code Llama: Open Foundation Models for Code.”

32. Touvron et al., “LLaMA: Open and Efficient Foundation Language Models”; OpenAI, *GPT-4 V(ision)*, 2023, <https://openai.com/research/gpt-4v-system-card>.

33. OpenAI, *GPT-4 V(ision)*.

34. Midjourney Inc., *Midjourney*, 2023, <https://www.midjourney.com>.

35. OpenAI, *DALL-E*, 2023, <https://openai.com/dall-e-3>.

36. Stability AI, *Stable Diffusion*, 2023, <https://stability.ai>.

37. Alec Radford et al., “Robust Speech Recognition via Large-Scale Weak Supervision,” in *Proceedings of the 40th International Conference on Machine Learning*, vol. 202, Proceedings of Machine Learning Research (PMLR) (2023), 28492–28518, <https://proceedings.mlr.press/v202/radford23a.html>; Matthew Le et al., “Voicebox: Text-Guided Multilingual Universal Speech Generation at Scale,” arXiv Working Paper No. 2306.15687, *arXiv*, 2023, <https://arxiv.org/abs/2306.15687>.

38. Brianna Zitkovich et al., “RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control,” in *Proceedings of The 7th Conference on Robot Learning*, vol. 229, Proceedings of Machine Learning Research (PMLR) (2023), 2165–2183, <https://proceedings.mlr.press/v229/zitkovich23a.html>; Siddharth Karamcheti et al., “Language-Driven Representation Learning for Robotics,” arXiv Working Paper No. 2302.12766, *Robotics: Science and Systems*, 2023, <https://arxiv.org/abs/2302.12766>.

computer-aided design (CAD).<sup>39</sup> For manufacturing applications, preliminary work has investigated the use of Gen-AI for computational design from text inputs<sup>40</sup> and for creating parts of spaceflight optical instrument structures.<sup>41</sup>

### 3.1 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. These qualities are essential for Gen-AI tools to be practical and trustworthy in engineering settings. Many studies have found that Gen-AI technologies often deviate from user inputs and self-contradict previously generated results.<sup>42</sup> Additionally, aligning Gen-AI to specific applications is a major challenge and an active research topic.<sup>43</sup> Although robustness issues might be less critical in certain applications like writing assistance, they become significant challenges in safety-sensitive areas of design and manufacturing. For example, 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. Therefore, it is crucial that NextGen-AI tools not only be robust and factually accurate but also capable of mitigating the limitations inherent in user inputs.

39. Rundi Wu, Chang Xiao, and Changxi Zheng, “DeepCAD: A Deep Generative Network for Computer-Aided Design Models,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2021), 6772–6782, [https://openaccess.thecvf.com/content/ICCV2021/papers/Wu\\_DeepCAD\\_A\\_Deep\\_Generative\\_Network\\_for\\_Computer-Aided\\_Design\\_Models\\_ICCV\\_2021\\_paper.pdf](https://openaccess.thecvf.com/content/ICCV2021/papers/Wu_DeepCAD_A_Deep_Generative_Network_for_Computer-Aided_Design_Models_ICCV_2021_paper.pdf); Aditya Sanghi et al., “CLIP-Forge: Towards Zero-Shot Text-to-Shape Generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 18603–18613, [https://openaccess.thecvf.com/content/CVPR2022/papers/Sanghi\\_CLIP-Forge\\_Towards\\_Zero-Shot\\_Text-To-Shape\\_Generation\\_CVPR\\_2022\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2022/papers/Sanghi_CLIP-Forge_Towards_Zero-Shot_Text-To-Shape_Generation_CVPR_2022_paper.pdf); Wamiq Para et al., “SketchGen: Generating Constrained CAD Sketches,” in *Advances in Neural Information Processing Systems*, vol. 34, 35th Conference on Neural Information Processing Systems (NeurIPS 2021) (2021), 5077–5088, [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/28891cb4ab421830acc36b1f5fd6c91e-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/28891cb4ab421830acc36b1f5fd6c91e-Paper.pdf).

40. Picard et al., “From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design”; Liane Makatura et al., “How Can Large Language Models Help Humans in Design and Manufacturing?,” arXiv Working Paper No. 2307.14377, *arXiv*, 2023, <https://arxiv.org/abs/2307.14377>.

41. Ryan McClelland, “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, 2022).

42. Percy Liang et al., “Holistic Evaluation of Language Models,” arXiv Working Paper No. 2211.09110, *arXiv*, 2022, <https://arxiv.org/abs/2211.09110>; Yuheng Zha et al., “AlignScore: Evaluating Factual Consistency with a Unified Alignment Function,” arXiv Working Paper No. 2305.16739, *arXiv*, 2023, <https://arxiv.org/abs/2305.16739>; Niels Münder et al., “Self-contradictory Hallucinations of Large Language Models: Evaluation, Detection and Mitigation,” arXiv Working Paper No. 2305.15852, *arXiv*, 2023, <https://arxiv.org/abs/2305.15852>.

43. Hunter Lightman et al., “Let’s Verify Step by Step,” arXiv Working Paper No. 2305.20050, *arXiv*, 2023, <https://arxiv.org/abs/2305.20050>; Long Ouyang et al., “Training language models to follow instructions with human feedback,” in *Advances in Neural Information Processing Systems*, vol. 35, 36th Conference on Neural Information Processing Systems (NeurIPS 2022) (New Orleans, LA, 2022), 27730–27744, [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf); Yotam Wolf et al., “Fundamental Limitations of Alignment in Large Language Models,” arXiv Working Paper No. 2304.11082, *arXiv*, 2023, <https://arxiv.org/abs/2304.11082>; Yang Wu et al., “An Early Evaluation of GPT-4V(ision),” arXiv Working Paper No. 2310.16534, *arXiv*, 2023, <https://arxiv.org/abs/2310.16534>.

**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<sup>44</sup> and Bookcorpus<sup>45</sup> 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,<sup>46</sup> with failures in basic tasks like grasping common manufacturing problems or recommending suitable materials. While these limitations might be addressed in fields with plentiful public data, this approach is not feasible for design and manufacturing, where most engineering design and manufacturing data are not only very industry-specific, but this information is also proprietary and subject to extensive rules (under both regulatory and private policy) 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 6.

**Does Gen-AI possess useful reasoning capabilities?** A major limitation of current Gen-AI technologies is the lack of reasoning capabilities.<sup>47</sup> 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.<sup>48</sup> Unfortunately, this poses a major challenge to adopting the current generation of Gen-AI tools.<sup>49</sup> 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

---

44. Jack W. Rae et al., "Compressive Transformers for Long-Range Sequence Modelling," in *8th International Conference on Learning Representations, ICLR 2020* (Addis Ababa, Ethiopia, 2020), [https://iclr.cc/virtual\\_2020/poster\\_SylKikSYDH.html](https://iclr.cc/virtual_2020/poster_SylKikSYDH.html).

45. Yukun Zhu et al., "Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books," in *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)* (2015), 19–27.

46. Makatura et al., "How Can Large Language Models Help Humans in Design and Manufacturing?"; Picard et al., "From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design."

47. Jie Huang et al., "Large Language Models Cannot Self-Correct Reasoning Yet," arXiv Working Paper No. 2310.01798, *arXiv*, 2023, <https://arxiv.org/abs/2310.01798>; Nouha Dziri et al., "Faith and Fate: Limits of Transformers on Compositionality," arXiv Working Paper No. 2305.18654, *arXiv*, 2023, <https://arxiv.org/abs/2305.18654>; Miles Turpin et al., "Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting," arXiv Working Paper No. 2305.04388, *arXiv*, 2023, <https://arxiv.org/abs/2305.04388>; Freda Shi et al., "Large Language Models Can Be Easily Distracted by Irrelevant Context," in *Proceedings of the 40th International Conference on Machine Learning*, vol. 202, Proceedings of Machine Learning Research (PMLR) (2023), 31210–31227, <https://proceedings.mlr.press/v202/shi23a.html>.

48. John R. Dixon and Clive L. Dym, "Artificial Intelligence and Geometric Reasoning in Manufacturing Technology," *Applied Mechanics Reviews* 39, no. 9 (1986): 1325–1330.

49. Makatura et al., "How Can Large Language Models Help Humans in Design and Manufacturing?"; Picard et al., "From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design."

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 a glimmer of hope for enhancing the reasoning abilities of Gen-AI systems, potentially making them more viable for complex engineering and manufacturing tasks in the future.

**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 organizations—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. To be trustworthy, it is also crucial for NextGen-AI to be adaptable 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. Integrating the tools required for adequate collection of data and rapid provision of data-driven insights for design and manufacturing workflow will require rethinking hardware and software architectures. 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 also 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. These efforts pose both capital and personnel constraints, as hardware and software are expensive investments—not only to purchase but also to train and encourage workers to use them. 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.

While the above-mentioned challenges pose difficulties in adopting Gen-AI technologies for the design and manufacturing domain, carefully developed NextGen-AI technologies have the potential to overcome many of these limitations. It is important to note that the above discussion provides a non-exhaustive list for brevity. Many other critical issues, such as integration with real-time multi-modal sensor data, tight integration with hardware and robotics, are also crucially important but beyond the scope of this brief overview. We provide potential strategic frameworks that can overcome these limitations in Section 5.

#### 4 Economic implications of Gen-AI for workers and firms

Although Gen-AI is a new frontier for manufacturing and design technology, the process of technological change and adaptation to that change is not unprecedented. In order to predict some of the challenges that workers and firms will face as they develop and incorporate NextGen-AI tools into their industries and workplaces, it is helpful to consider historical evidence and identify common patterns from earlier waves of new technology. From this investigation, several key questions emerge. Will NextGen-AI empower workers to become more productive, or will it instead aim to remove labor from the production process? Will NextGen-AI tools be designed and controlled by a few massive tech companies, further increasing their out-sized market power, stifling smaller firms and new ideas? Is it possible to shape the path of NextGen-AI development?

For more than a century, the United States was a global manufacturing superpower. The earliest innovations of US manufacturing included Oliver Evans' automatic flour mill in the late 1780s and the textile-flagship Slater Mill in 1793, the first American factory to implement Richard Arkwright's renowned water-powered cotton spinning frame technology. The successes of American industry have been driven in large part by clever innovators and their cutting-edge production machinery and mechanized processes.<sup>50</sup>

At this time, proliferative inventors benefited from the establishment of clear intellectual property rights (the US Patent system was founded in 1790), and a collaborative industrial environment that resulted in wide distribution of new ideas and uptake of new processes and machinery. In addition to obtaining the third-ever issued US Patent for his aforementioned flour-milling technology, Oliver Evans also self-published *The Young Mill-Wright and Millers' Guide* in 1795.<sup>51</sup> This opus contained chapters for grist millers—explaining how to implement Evans' improvements and how novice millers could emulate experts in their field—but also several chapters on general elements of mechanical and hydraulic engineering, a vital guide even for industries beyond the millwrighting trade.<sup>52</sup> The book's detailed explanations were expressly written with the next generation of industrialists and innovators in mind. They democratized

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

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

52. Oliver Evans, *The Young Mill-Wright and Miller's Guide* (Philadelphia, PA: Self-published by the author, 1795).

the technological frontier by providing cutting edge techniques and technology to an entire industry. This exemplifies how sharing “open science” and expanding the stock of existing knowledge can rise tides for entire industries and the economy overall.<sup>53</sup>

These nascent developments set the stage for American manufacturing productivity to increase dramatically, beginning in the latter half of the 19<sup>th</sup> 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.<sup>54</sup> What made this system so effective was its ability to make less-skilled or less-educated workers more productive and 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 20<sup>th</sup> 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.<sup>55</sup> Equally important to these gains were the contributions of workers, who were continually trained and upskilled to tackle new production tasks. Following landmark labor negotiations in the late 1930s, workers continued to push the automobile industry to 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.,<sup>56</sup>). These consequences have been driven by two commingled and depreciative forces. The first force has been 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.<sup>57</sup> The second force has been faltering creation of the new, productive tasks that are required to take full advantage of new technology.<sup>58</sup>

53. Paul M. Romer, “Increasing Returns and Long-Run Growth,” *Journal of Political Economy* 94, no. 5 (1986): 1002–1037.

54. Daron Acemoglu and Simon Johnson, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity* (New York, NY: Hachette PublicAffairs, 2023).

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

56. Acemoglu and Johnson, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*; Autor, Dorn, and Hanson, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States”; David Autor, “Work of the Past, Work of The Future,” *AEA Papers and Proceedings* 109, no. 2019 (2019): 1–32; Acemoglu and Restrepo, “Robots and Jobs: Evidence from U.S. Labor Markets”; Daron Acemoglu and Pascual Restrepo, “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica* 90, no. 5 (2022): 1973–2016.

57. Daron Acemoglu and Pascual Restrepo, “Artificial Intelligence, Automation, and Work,” in *The Economics of Artificial Intelligence: An Agenda*, ed. Ajay Agrawal, Joshua Gans, and Avi Goldfarb (Chicago, IL: University of Chicago Press, 2019), 197–236, <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/artificial-intelligence-automation-and-work>; Daron Acemoglu and Pascual Restrepo, “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives* 33, no. 2 (2019): 3–30; Daron Acemoglu, Andrea Manera, and Pascual Restrepo, “Does the U.S. Tax Code Favor Automation?,” *Brookings Papers on Economic Activity* 2020 (2020): 231–300.

58. David Autor et al., “New Frontiers: The Origins and Content of New Work, 1940–2018,” NBER Working Paper No. 30389, *National Bureau of Economic Research*, 2022, <https://www.nber.org/papers/w30389>.

The economic impacts of current and NextGen-AI innovations will depend on how innovators envision, develop, and deploy these tools. 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 as a whole.

But the frontier of artificial intelligence is as flexible as it is powerful. With careful development and guidance, NextGen-AI could offer enormous improvements to existing design and manufacturing workstreams, and also facilitate new capabilities that build and leverage the skills of manufacturing workers. For example, several companies (discussed further below) are attempting to design NextGen-AI tools that incorporate a wealth of data from machinery, instruments, and sensors on the manufacturing floor. Equipping these models with a well-crafted user interface that distills an overwhelming amount of operations data into easily interpreted metrics and accompanying machinery schematics could enable less-educated or less-technical workers to perform more-expert work by enabling them to interpret highly technical content with ease. This optimal outcome will require a collective vision from innovators, workers, management, policy, and civil society to best develop and deploy these technologies.

As Gen-AI tools become commercially available and more-tailored to the manufacturing industry, we have identified several key business challenges that these tools could address. The findings we discuss in the remainder of this section are partially derived from this historical analysis of how previous waves of new technologies have succeeded or failed in different domains, with corroboration and examples provided through conversations and interviews that we have conducted on these topics with various stakeholders in the future of NextGen-AI, including engineers and manufacturers, tech executives and entrepreneurs.

#### **4.1 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 metrics like Overall Equipment Effectiveness—a production factor’s interactional measure of availability, performance, and quality. This legacy approach is suboptimal for several reasons. Most obviously, potential productivity upside is forgone when such data is not systematically or durably captured, stored, and analyzed, as low-grade or system-wide issues may go undetected and unaddressed. This measurement process is also cumbersome and costly for workers who have many other (often urgent) responsibilities on the floor. When a worker’s job focus is on making parts, they may not have the time, interest, or ability to do the seemingly non-essential administrative function of data collection very effectively on top of their more-obviously essential functions. Lastly, if production planning and resource allocation relies on the

timeliness and accuracy of this data, static “whiteboard approaches” (manually tracking and updating key performance metrics on physical boards) will hinder optimal organization.

Frontier technologies can take several steps to address these issues. Companies can consider implementing automated data capture technologies, many of which are commercially available (e.g., Raven.ai,<sup>59</sup> Tulip,<sup>60</sup> Plex Production Monitoring,<sup>61</sup> etc.), to accurately capture downtime without requiring operators to log it manually. Capturing this data from machines in real-time also allows for more accurate and prescient assessments of equipment effectiveness and operability, allowing resources to be efficiently marshaled to production breakpoints. Many startups are also experimenting with automatic labeling, synthetic data generation, and unstructured data algorithms to produce historical insights for companies that have rich—but mostly unused—troves of operational data. These technologies seek to standardize data capture, as well as take advantage of information provided by legacy systems, to narrow the performance gap between “high-tech” and “low-tech” firms. 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 with less voluminous or systematically captured data, so the deployment of commercial solutions by startups and established industry affiliates (e.g., MontBlancAI,<sup>62</sup> Spencermetrics,<sup>63</sup> Lumafield,<sup>64</sup> Rockwell Automation<sup>65</sup>) could also help level the playing field for smaller competitors.

Another point of misallocation can happen at the disconnect between design innovation and shop-floor manufacturability, a pervasive source of frustration and waste in the manufacturing process discussed further in Section 6. NextGen-AI could be used to bridge these complications and improve the iterative process of converting groundbreaking design ideas into real products. A current area of exploration involves better integration between the tools used by design and manufacturing teams, to improve synchrony and availability of technical information. These attempts to unify and streamline the design-to-delivery pipeline align with the concepts of “Domain-Driven Design,” which suggest that business software should be built from the top-down, defining a domain-wide “ubiquitous language” with the entire business process in mind, instead of allowing the software of different organizational functions to fracture into many sub-domains (as is the case in many organizations today, including design and manufacturing firms) that each have their own vernacular, procedures, and isolated software ecosystems.<sup>66</sup> Improved NextGen-AI tools that incorporate multi-modal data could also accelerate the initial design process by providing relevant research and digitally simulating whether proposed designs can withstand expected use cases by incorporating data on material and fabrication specifications before prototyping the product. By more seamlessly linking the tools used by design and manufacturing

---

59. Raven Telemetry, *Raven.ai*, 2023, <https://raven.ai/>.

60. Tulip, *Tulip*, 2023, <https://tulip.co/>.

61. Rockwell Automation, *Plex*, 2023, <https://www.plex.com/products/production-monitoring>.

62. Mont Blanc AI, *MontBlancAI*, 2023, <https://montblanc.ai/>.

63. SpencerMetrics, *SpencerMetrics*, 2023, <https://www.spencermetrics.com/>.

64. Lumafield, *Lumafield*, 2023, <https://www.lumafield.com/>.

65. Rockwell Automation, *Rockwell Automation*, 2023, <https://www.rockwellautomation.com/en-us.html>.

66. Martin Fowler, *Patterns of Enterprise Application Architecture* (Addison-Wesley Professional, 2002); Eric Evans, *Domain-Driven Design: Tackling Complexity in the Heart of Software* (Addison-Wesley Professional, 2003); Vaughn Vernon, *Implementing Domain-Driven Design* (Addison-Wesley Professional, 2013); Vaughn Vernon, *Domain-Driven Design Distilled* (Addison-Wesley Professional, 2016).



teams, in addition to expanding the range of data that can be utilized, time-to-market could be significantly reduced via NextGen-AI driven feasibility testing.

#### 4.2 Variability and risk in supply chains and forecasting

Supply chains are predictably unpredictable, though the exogenous shocks posed in recent years by the COVID-19 pandemic, as well as political turmoil and military conflict, are salient examples of how quickly these systems can break down. Even when transacting under more typical circumstances, delays and interruptions are common in the production and transportation of raw materials or finished products. In addition to the arcanelly defined, multiplex environment of dynamic production processes, the related priority of accurately forecasting demand is perhaps even more complex. Our conversations with tech entrepreneurs and industry executives corroborate that the combination of these two interjoined and highly variable systems poses constant headaches for executives and other business planners.

NextGen-AI can provide a new frontier for manufacturing companies to adeptly adjust to shifting timelines with vastly improved predictive modeling that can parameterize both common problems (breakdowns, production delays, missed deliveries or mis-deliveries) and tail-risks to demand and delivery. In particular, predictive and Gen-AI systems have comparative strengths in modeling large and ill-defined problems by incorporating an incomprehensibly large number of parameters. As these models become better calibrated to understand spatial data and construct representative schema of real-world processes, their capabilities will be well-suited to address inventory optimization, production schedules, supply chain planning, forecast accuracy, and operations management. One drawback of high-dimensional models, however, is that they are often difficult for humans to interpret, which can hinder sensible intervention to override and course-correct when the model is off-track. There are also trade-offs in selecting simple causal models vs. complex data-driven models. This may preclude the use of these tools from safety-risk applications, which may be better performed by well-informed human operators.

#### 4.3 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 and manufacturing expert, Fortune 500 company

This sentiment indicates that there is a future for workers in manufacturing, but several challenges impede the ability of manufacturers to do business and threaten career security for workers:<sup>67</sup>

- A faster-aging (and already older) workforce than the overall US average, driven by older worker retention and 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;

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

- 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 and cogently present vast amounts of information in a relatively digestible format—even for workers who may have limited background knowledge or experience (as mentioned above). 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.<sup>68</sup> 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 a wealth of prior cases and equipped with advanced computer vision technology.

Research on past waves of technological innovation in manufacturing, discussed in the introduction to this section, suggests that these improvements are most effective when they augment human capabilities, rather than aim to remove humans from the process. Many new innovations and systems improvements are produced by expert workers who have a close familiarity with production processes—companies that can take advantage of these complementarities will gain an advantage in the market over their competitors. Worker empowerment drives buy-in for new technologies, and meteoric productivity improvements are necessarily worker-mediated; thus, technologies must be designed with workers in mind.

#### 4.4 Market power, predatory business, 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.<sup>69</sup> 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 technological 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. In our conversations with industry leaders, for example, several cited interest in implementing new, commercially available tools that simulate Amazon’s pricing strategies,<sup>70</sup>

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

69. Acemoglu and Johnson, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*; Daron Acemoglu and Simon Johnson, “Big Tech Is Bad. Big AI Will Be Worse.,” *New York Times*, 2023, <https://www.nytimes.com/2023/06/09/opinion/ai-big-tech-microsoft-google-duopoly.html>.

70. Daisuke Wakabayashi, “Does Anyone Know What Paper Towels Should Cost?,” *New York Times*, 2022, <https://www.nytimes.com/2022/02/26/technology/amazon-price-swings-shopping.html>.

which have the capability to force out competition with dynamic price-out algorithms and extract consumer surplus via taste-specific price discrimination. Even if the increasing availability of commercially developed replicas allow all businesses to have access to the same toolkit to compete on even footing, it will still be at the expense of consumers and could lead to long-term decline.

If new capabilities are 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 extract value from smaller competitors who rely on licensing the NextGen-AI products (e.g., OpenAI) or using the platforms of those largest companies (e.g., Amazon) in order to stay afloat. It could also result in many smaller firms exiting the market, or being acquired by the largest companies (if they appear to be mounting 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 shaped within a market mediated by healthy competition.

## **5 A 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 3, 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. In addition to general suggestions for the entire domain of manufacturing and design (summarized in Figure 2), we propose several domain-specific ideas that may be beneficial to improve robustness and efficiency in some subfields.

### **5.1 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

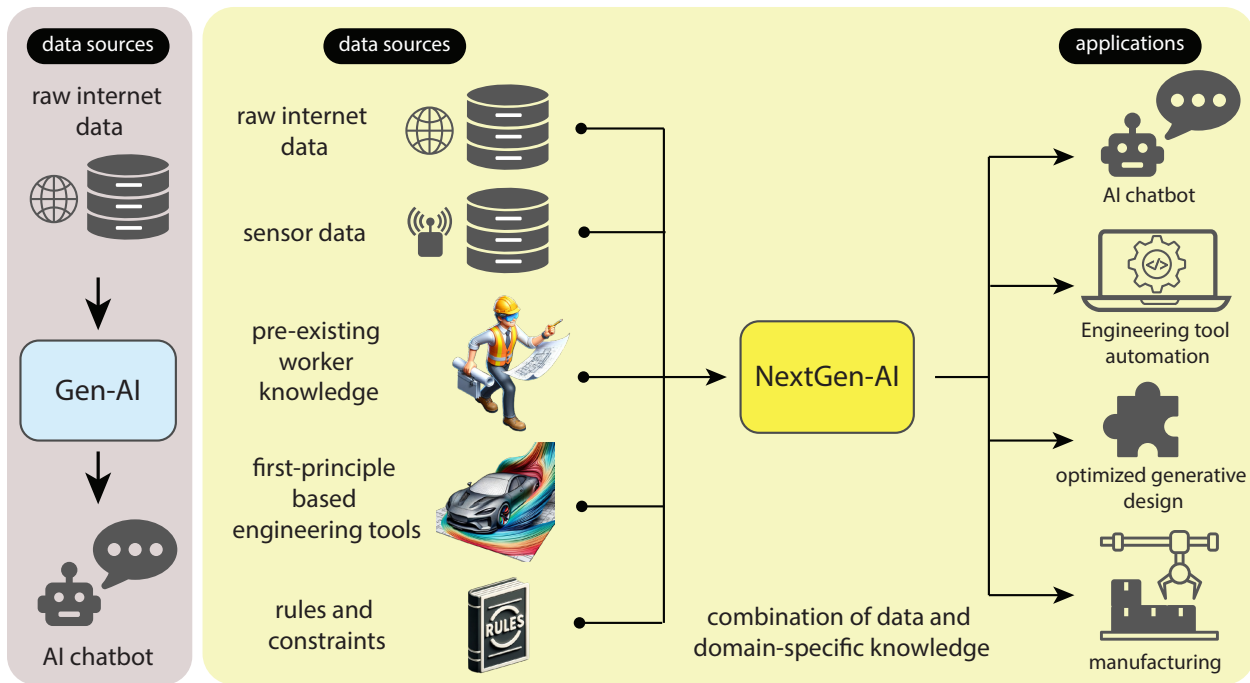


Figure 2: While current Gen-AI technologies are mostly developed using internet data (shown on left), we envision a combination of raw data, sensor data, first-principles-based engineering tools, pre-existing worker/expert knowledge and constraints for developing the NextGen-AI technologies (shown on 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.

in engineering design, where the diversity of ideas is crucial to fostering human creativity and innovation.<sup>71</sup> 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.

<sup>71</sup> Lyle Regenwetter, Amin Heyrani Nobari, and Faez Ahmed, “Deep Generative Models in Engineering Design: A Review,” *Journal of Mechanical Design* 144, no. 7 (2022): 071704; Yukari Nagai and John Gero, “Design creativity,” *Journal of Engineering Design* 23, no. 4 (2012): 237–239.

## 5.2 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.<sup>72</sup> This synergy between Gen-AI and first-principle models could lead to more robust and nuanced engineering solutions, where AI complements and extends the capabilities of traditional methods.

## 5.3 Enhancing Workforce Training and Knowledge Transfer

As discussed earlier in Section 4.3, 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.

---

<sup>72</sup> Timo Schick et al., “Toolformer: Language Models Can Teach Themselves to Use Tools,” arXiv Working Paper No. 2302.04761, *arXiv*, 2023, <https://arxiv.org/abs/2302.04761>; Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman, “Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models,” Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems 332 (New York, NY: Association for Computing Machinery, 2022), 1–7, <https://doi.org/10.1145/3491101.3519665>; Qiao Jin et al., “GeneGPT: Augmenting Large Language Models with Domain Tools for Improved Access to Biomedical Information,” arXiv Working Paper No. 2304.09667, *arXiv*, 2023, <https://arxiv.org/abs/2304.09667>; Wenlong Huang et al., “Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents,” in *Proceedings of the 39th International Conference on Machine Learning*, Proceedings of Machine Learning Research (PMLR) (Baltimore, MD, 2022), 9118–9147, <https://proceedings.mlr.press/v162/huang22a/huang22a.pdf>.

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.

#### 5.4 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.<sup>73</sup> In contrast, traditional robots and machines used in manufacturing require manual programming and struggle to function in highly variable 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%.<sup>74</sup>

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.<sup>75</sup> Additionally, integrating advanced sensory systems in robots, powered by Gen-AI, 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.

73. Allen Z. Ren et al., “Robots That Ask For Help: Uncertainty Alignment for Large Language Model Planners,” arXiv Working Paper No. 2307.01928, *arXiv*, 2023, <https://arxiv.org/abs/2307.01928>; Jacky Liang et al., “Code as Policies: Language Model Programs for Embodied Control,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2023), 9493–9500; Boyi Li et al., “Interactive Task Planning with Language Models,” arXiv Working Paper No. 2310.10645, *arXiv*, 2023, <https://arxiv.org/abs/2310.10645>.

74. Peng et al., “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.”

75. Matthias Guertler et al., “When is a Robot a Cobot? Moving Beyond Manufacturing and Arm-Based Cobot Manipulators,” *Proceedings of the Design Society* 3 (2023): 3889–3898; Maurizio Faccio et al., “Human Factors in Cobot Era: A Review of Modern Production Systems Features,” *Journal of Intelligent Manufacturing* 34, no. 1 (2023): 85–106.

## 5.5 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 the creation of a shared data repository, which could benefit a multitude of stakeholders by aiding in the training of *NextGen-AI foundation 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<sup>76</sup> and differential privacy<sup>77</sup> 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.<sup>78</sup> 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.

## 5.6 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

76. H. Brendan McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017*, vol. 54, Proceedings of Machine Learning Research (PMLR) (Fort Lauderdale, FL, 2017), 1273–1282, <https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf>.

77. Cynthia Dwork, "Differential privacy," in *International Colloquium on Automata, Languages, and Programming, ICALP 2006*, ed. Michele Bugliesi et al., vol. 4052, Lecture Notes in Computer Science (Venice, Italy: Springer, 2006), 1–12.

78. Muneeb Ul Hassan, Mubashir Husain Rehmani, and Jinjun Chen, "Differential Privacy Techniques for Cyber Physical Systems: A Survey," *IEEE Communications Surveys & Tutorials* 22, no. 1 (2019): 746–789; Kevin I-Kai Wang et al., "Federated Transfer Learning Based Cross-Domain Prediction for Smart Manufacturing," *IEEE Transactions on Industrial Informatics* 18, no. 6 (2021): 4088–4096.

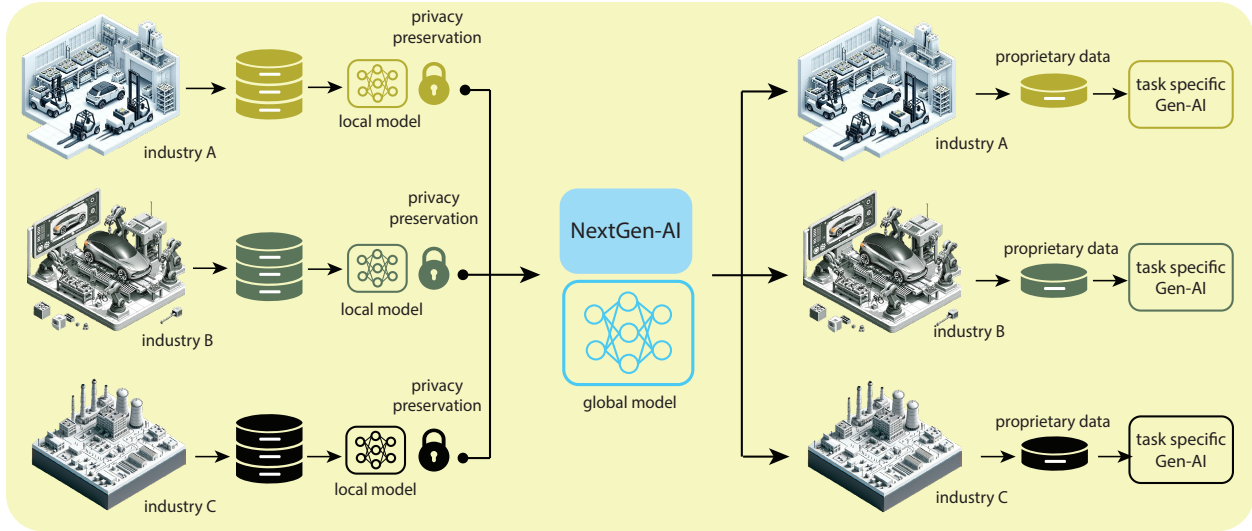


Figure 3: A two-step framework for training NextGen-AI models for manufacturing industries: A large repository of data is compiled from various cross-industry stakeholders to create a foundational Gen-AI model which can be further fine-tuned with industry-specific information and tasks to produce highly tailored NextGen-AI implementations. To protect proprietary data and competitive advantages, industries may opt for privacy-preserving machine learning techniques such as federated learning<sup>81</sup> and differential privacy.<sup>82</sup> 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.

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. Second, 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.



## 6 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. We delve into the practical applications and transformative potential of AI technologies with a particular focus on generative AI. 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. From small-scale improvements to industry-wide shifts, these case studies highlight the tangible benefits and challenges of integrating AI into the automotive manufacturing process.

### 6.1 Automotive industry

The automotive industry has been a pivotal force in shaping the American economy throughout the last century, as mentioned in Section 4. The interconnected complexities of the design and manufacturing workflow and the dense network of regulations with which they must comply introduce unique challenges and opportunities for the integration of AI technologies. As global automotive manufacturing is undergoing 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 their perspective on the trajectory of NextGen-AI for generative design and engineering assistance. At large companies like this, the design process is typically split into many specialized teams across design, engineering, and manufacturing.

*“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.”*

— *Manufacturing expert on the split specialities of parts-manufacturing*

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<sup>83</sup> of long design cycles that may span months to years.<sup>84</sup> Yet when designs are produced via this iterative process, costly late-stage design changes are often required. The likelihood of modifications being necessary increases as a design grows more complex. Centralizing information across multiple domains thus is essential to improving process efficiency. NextGen-AI design models could simultaneously capture and adhere to the constraints of multiple domains. The model could then produce

83. Mohammad Saravi et al., “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*, ed. Richard Curran, Shou-Yan Chou, and Amy Trappey, Proceedings of the 15th ISPE International Conference on Concurrent Engineering (CE2008) (London, U.K.: Advanced Concurrent Engineering, Springer, 2008), 123–130.

84. Naresh S. Iyer et al., “PATO: Producibility-Aware Topology Optimization using Deep Learning for Metal Additive Manufacturing,” arXiv Working Paper No. 2112.04552, *arXiv*, 2021, <https://arxiv.org/abs/2112.04552>.

designs that satisfy these constraints earlier in the design lifecycle, reducing the costs of a more-iterative process and disruption-prone process. This capability would reduce design costs and expedite the development of new products.

*“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.”*

— Manufacturing expert on useful NextGen-AI to support the design process

NextGen-AI tools could be developed either for expert-reference purposes or as a direct design-generation tool. Direct design generation is not currently used in the design workflow, as current Gen-AI tools are not powerful enough to accomplish the complex sub-tasks involved. Design problems may involve reiteration upon older designs or designing something entirely new. Generation is less applicable to reiteration and design of simple parts, where engineers may prefer precise control of design variables, or where setting up the generation problem may be as laborious as manually designing by hand. Generation is thus more likely to be advantageous when designing a complex assembly from the ground-up. It is also essential for generative tools to be capable of designing for specific manufacturing processes, such as metal stamping and injection molding. Instead of directly generating the design, NextGen-AI could also support designers by providing improved access to the resources (e.g., documentation) needed to produce a new design. 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]...”*

— Manufacturing expert on simplifying the information retrieval process

We highlight that developing either of these AI methods requires vast quantities of data to be collected and that the capabilities of our 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 will be an important challenge in the coming decades.

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

— Manufacturing expert on the difficulty of data sharing

## 6.2 Footwear industry

NextGen-AI could augment the creative process of design and facilitate functional aspects of manufacturability for complex, performance-optimized footwear. With NextGen-AI driven design, manufacturers could explore a vast array of

design possibilities, develop more innovative and functional footwear, and cater to the diverse preferences of consumers. In addition, the advent of smart footwear opens up exciting new avenues for integrating technology into everyday wearables.<sup>85</sup> 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 significant area where NextGen-AI can contribute to the footwear industry.<sup>86</sup> 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.<sup>87</sup> This rapid prototyping not only reduces the time-to-market for new designs but also allows for more iterative and responsive design processes. NextGen-AI may prove useful to footwear manufacturers that must quickly adapt to changing markets and consumer trends to ensure that their products remain relevant, resilient and desirable.

*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 practicably 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. There’s certainly a lot of repetitive work. . .”*

*— New Balance team describing the current, iterative design process*

Generative modeling capabilities are currently limited to the conceptual stage of design, mainly due to their inability to predict the manufacturability and ultimate performance of a product. Generative design models were adapted to the engineering domain on the basis of advancements in language processing and computer vision; however, while the design process does include image- and text-based information, it also requires data imputation from many other sources, as well as an array of technical specifications. 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 that are built specifically for solving engineering problems. For example, engineers often work with 3D geometries. Scaling generative models from 2D to 3D is a significant challenge, as the complexity represented by the model increases exponentially when adding full dimensionality.

85. Matthew A. Rhoades, “Smart Footwear: A Designer’s Perspective,” in *Smart Clothes and Wearable Technology (Second Edition)*, The Textile Institute Book Series (Woodhead Publishing, Elsevier, 2023), 509–527, <https://www.sciencedirect.com/science/article/pii/B9780128195260000114>.

86. Marco Binelli et al., “Digital manufacturing of personalised footwear with embedded sensors,” *Scientific Reports* 13, no. 1962 (2023).

87. Jochen Suessmuth, Florian Fick, and Stan Van Der Vossen, “Generative AI for Concept Creation in Footwear Design,” in *ACM SIGGRAPH 2023 Talks*, 17 (New York, NY: Association for Computing Machinery, 2023), 1–2, <https://doi.org/10.1145/3587421.3595416>.

*“...not everyone realizes that [a] 2D drawing is not necessarily something we can go off and build right away ... if [NextGen-AI models] don’t know how to realize something and make it work really well, I think having these super powerful generative tools is not that interesting ... [also,] the visual data doesn’t necessarily correlate to manufacturability, to comfort, to performance...”*

*— New Balance team on the shortcomings of current Gen-AI visual tools*

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.

*“Oh, goodness, 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 and try to understand them more holistically, there’s the biomechanics and physiology of the athletes that we collect ... The fit and durability of our shoes gets collected.”*

*— New Balance team on the wide variety of available data*

NextGen-AI models can leverage these data resources to provide advanced problem-solving capabilities, in order to keep pace with the ever-shifting focus of the footwear industry and the challenges posed by new goals.

*“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 [to maximize resilience]? ...Fit is super important. How do you take lots of feet and feed them [into a model] with preference data and all of the things that you collect to make [shoes] fit optimally?”*

*— New Balance team on evolving industry goals and challenges*

Even if models can be trained to accurately predict test results on the basis of collected data, the results are only as good as the test itself. NextGen-AI tools may ultimately be limited by the tests that are designed and implemented by human operators. This suggests that NextGen-AI tools will require (at a minimum) the flexibility to adapt to new testing strategies—or, more advantageously, perhaps even be deployed to help design new and better testing strategies.

*“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. We’re constantly looking to improve, but...it’s more about changing the mechanical tests than about scrubbing the data.”*

*— New Balance team on the limits of testing 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.”*

— *New Balance team on measuring users’ product experiences*

The utility of future AI tools is in their ability to be integrated into a productive workflow. Tools do not need to encapsulate the entire design cycle to be useful, nor is automating the entire design cycle inherently useful or 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.”*

— *New Balance team on exemplary implementations of AI-augmentation*

While NextGen-AI tools may automate some parts of the design process, it is essential that these tools are focused on complementing preexisting skills, keeping humans in the loop. Instead of attempting to encode and replace the experience and competency of an engineer, 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 that run in the background of CAD software 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.”*

— *New Balance team on what NextGen-AI cannot accomplish*

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 during the hiring process.

*“The rate at which [Gen-AI has] progressed is definitely intimidating. 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 criteria you can use to assess a designer’s potential or talent anymore. There’s definitely some tricky disruption.”*

— *New Balance team on Gen-AI changing hiring practice*

## 7 A policy guide for NextGen-AI development and adoption

The integration of Gen-AI and NextGen-AI tools in manufacturing does pose risks, though these can be largely avoided through careful design and deployment of these tools, as well as sensible industry norms and institutional safeguards.<sup>88</sup> These risks motivate the following guidelines for organizational and public policy, which we derive by considering how to best apply the framework in Section 5 to address the technical challenges facing Gen-AI tools from Section 3 while carefully considering the economic implications of the trajectory of NextGen-AI in Section 4.

**Systems integration for Gen-AI software:** Unlike most traditional software systems, the development of an AI model requires a lot of data. While large neural network-based AI software offer several potentials over traditional software architectures, they come with new challenges.<sup>89</sup> New models must be trained on very large datasets to learn patterns and correlations; then they must be tested on new, untouched datasets to ensure the tool’s accuracy and reliability; furthermore, adding new features or refining the model typically requires even more data. After deployment to real-world scenarios, AI models can be continually improved by taking advantage of *ethically* collected data from user interactions. Additionally, continuous monitoring is also crucial post-deployment to ensure that the AI models correctly adapt to new use-cases and contexts, while also maintaining their effectiveness over time. Unlocking the promise of NextGen-AI, equipped to handle multiple modalities of input and output, will require correct implementation of engineering first principles, seamless integration with preexisting processes and procedures, and a frictionless operator experience. To accomplish these goals, NextGen-AI systems must have access to training and real-time data from a wide range of sources—materials characteristics, spatial design specifications, time-dependent steps of operation, production batch records, inspection certification, demand forecasts, and NextGen-AI system user data, just to name a few. Because of these unprecedented demands for enormous volumes of data, Gen-AI also presents unique challenges for defining and enforcing data ownership, access, and governance.

**Data governance:** The preponderance of artificial intelligence has also raised further concerns about intellectual property—especially related to data ownership, access, and compensation.<sup>90</sup> As the use of current Gen-AI programs has become more widespread, there has been a rash of lawsuits—including many prominent authors and artists—over intellectual property matters, regarding whose data was used to train these models. Usage and reproduction rights for journalistic or literary publications and artistic works are well-defined and protected under intellectual property law, so these issues are likely to be sorted out soonest. However, similar ownership and compensation concerns arise for user-data collected by online platforms (e.g., social media, HRM/CRM software), but the rules and regulations for this sort of data are less clear. In addition to the risk that these data are being collected and leveraged unlawfully,

---

88. Daron Acemoglu, David Autor, and Simon Johnson, “Can We Have Pro-Worker AI? Choosing a Path for Machines in Service of Minds,” *MIT, Shaping the Future of Work Initiative*, 2023, <https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>; Valerio Capraro et al., “The Impact of Generative AI on Socioeconomic Inequalities and Policymaking,” SSRN Working Paper No. 4666103, *SSRN*, 2023, [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4666103](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4666103); Daron Acemoglu and Todd Lensman, “Regulating Transformative Technologies,” NBER Working Paper No. 31461, *NBER*, 2023, <https://www.nber.org/papers/w31461>; Nell Watson, *Taming the Machine* (London, UK: Kogan Page, 2024).

89. Andrej Karpathy, *Software 2.0*, 2017, <https://karpathy.medium.com/software-2-0-a64152b37c35>.

90. Daron Acemoglu, “Harms of AI,” in *The Oxford Handbook of AI Governance*, ed. Justin Bullock et al., Pre-print available at: <https://www.nber.org/papers/w29247> (New York: Oxford University Press, 2024).

there are acute risks to companies and their internal or statutory responsibilities for corporate and client data security. While many companies have provisionally banned such platforms as ChatGPT from corporate computers, to prevent security compromises and avoid releasing trade secrets or client information, this is not a long-term solution. During our conversations with industry experts, we heard concerns that prolonged abstention from implementing new generative AI tools will leave their companies behind the competition. Especially for smaller companies that do not have the capital or expertise to develop their own proprietary implementations of generative AI technologies, this presents a dilemma: use the enterprise tools provided by dominant market players to keep pace with the most technically advanced operations—at the risk of sharing, leaking, or otherwise losing control over their data—or stagnate and lose ground in the market. While the anti-competitive capabilities of companies to shut out their competition by leveraging frontier AI tools is highly concerning, the monopolized control of information poses a far greater threat to long-term economic growth within and beyond the manufacturing sector. Addressing these concerns will require clear guidelines for the companies designing and marketing NextGen-AI solutions, to define and protect ownership, management, access, sale, and protection of data from employees, clients, and customers.

**Collaboration among smaller firms and industries:** Another strategy for smaller manufacturing industries or firms that lack the resources to independently develop sophisticated NextGen-AI models is to pool resources in order to develop a powerful base model, which can then be refined to the individual manufacturers' needs (*see* 3). In some cases, this may be feasible with industry-level collaboration through trade organizations. As discussed in Section 4, several start-up and established industrial technology developers also aim to position themselves as intermediaries between companies and their data—leveraging the insights from shared models while protecting proprietary information and client data. This data-sharing strategy has worked to accomplish similar shared goals in other industries, such as improving airline safety through data-sharing provisions with the Federal Airline Administration (FAA). A similar strategy could be undertaken by OSHA to promote the use of NextGen-AI to improve worker safety in manufacturing plants.

**Focus on personalization and user experience:** We encourage a specific focus on the end users (e.g., operators and engineers) while building these technologies. For instance, when these AI technologies are employed for training operators in a manufacturing facility, they should be more than just one-size-fits-most instructional tools. An effective NextGen-AI tool must possess the capability to adapt to each operator's unique skill set. Such personalization ensures that the training is more engaging for the operator, and thus more likely to be effective. By analyzing data from each operator's performance, the AI model's worker training methods and feedback provision can be further tailored. This adaptive approach will allow every operator to receive appropriate training that is best suited to their learning curve and most complementary to their preexisting skills and capabilities.

**Worker safety and worker-complementary NextGen-AI:** The highest priority for the design of NextGen-AI should ultimately be on worker safety. This includes restricting the deployment of untested or insufficiently tested AI tools that could put workers at risk, both on the factory floor and in the back-office personnel decision-making processes like staffing, hiring, disciplinary action and termination. The introduction of new tools necessitates new training and

sufficient workforce preparation to both maximize the usefulness of new technologies and to ensure seamless and secure integration with existing work procedures. In many cases, unionized workers may have a vested right to new production tasks that are within the scope of their preexisting jobs; the most successful companies will carefully consider how to best leverage their workforces and upskill existing workers with highly specialized expertise to derive the best competitive advantage and treat workers fairly in the process. In addition, it is crucial that reasonable safeguards be implemented to prevent excessive workplace monitoring and surveillance to preserve the well-being and noncoercion of workers. This will require comprehensive updates to health and safety rules and guidelines (e.g., OSHA). Further, NextGen-AI technologies must address ethical concerns with a specific focus on end users. The field of AI ethics is critical and evolving as rapidly as the technology.<sup>91</sup> There are already many documented failure cases of text- and image-based Gen-AI technologies perpetuating sex- and race-based biases.<sup>92</sup> It is likely that a combination of well-crafted AI governance strategies<sup>93</sup> and ethical reporting of Gen-AI model outcomes<sup>94</sup> is required. These ethical and safety concerns should be addressed before companies can confidently adopt NextGen-AI technology.

## 8 Concluding remarks

Our view on the development of NextGen-AI technologies for engineering design and manufacturing is guided by the limitations and trajectory of current-generation 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 more broadly. 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 underutilized or misallocated resources, supply chain risk and forecast variability, 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 NextGen-AI technologies for the engineering design and manufacturing domain, which must include embedding first-principles

91. Anna Jobin, Marcello Ienca, and Effy Vayena, “The global landscape of AI ethics guidelines,” *Nature Machine Intelligence* 1, no. 9 (2019): 389–399.

92. Emily M Bender et al., “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?,” in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (New York, NY: Association for Computing Machinery, 2021), 610–623.

93. Inioluwa Deborah Raji et al., “Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance,” in *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (New York, NY: Association for Computing Machinery, 2022), 557–571; Inioluwa Deborah Raji et al., “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing,” in *Proceedings of the 2020 ACM Conference on Fairness, Accountability, and Transparency* (New York, NY: Association for Computing Machinery, 2020), 33–44.

94. Margaret Mitchell et al., “Model Cards for Model Reporting,” in *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency* (New York, NY: Association for Computing Machinery, 2019), 220–229.



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 new 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 will depend on the priorities we set and the choices we make.

## References

- Acemoglu, Daron. “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, 2024.
- Acemoglu, Daron, David Autor, and Simon Johnson. “Can We Have Pro-Worker AI? Choosing a Path for Machines in Service of Minds.” *MIT, Shaping the Future of Work Initiative*, 2023. <https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>.
- Acemoglu, Daron, and Simon Johnson. “Big Tech Is Bad. Big AI Will Be Worse.” *New York Times*, 2023. <https://www.nytimes.com/2023/06/09/opinion/ai-big-tech-microsoft-google-duopoly.html>.
- . *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York, NY: Hachette PublicAffairs, 2023.
- Acemoglu, Daron, Simon Johnson, and Austin Lentsch. “The Hollywood Writers’ AI Fight is Everyone’s Fight.” *Project Syndicate*, 2023. <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. “Regulating Transformative Technologies.” NBER Working Paper No. 31461, *NBER*, 2023. <https://www.nber.org/papers/w31461>.
- Acemoglu, Daron, Andrea Manera, and Pascual Restrepo. “Does the U.S. Tax Code Favor Automation?” *Brookings Papers on Economic Activity* 2020 (2020): 231–300.
- Acemoglu, Daron, and Pascual Restrepo. “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, 2019. <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/artificial-intelligence-automation-and-work>.
- . “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33, no. 2 (2019): 3–30.
- . “Robots and Jobs: Evidence from U.S. Labor Markets.” *Journal of Political Economy* 128, no. 6 (2020): 2188–2244.
- . “Tasks, Automation, and the Rise in U.S. Wage Inequality.” *Econometrica* 90, no. 5 (2022): 1973–2016.
- Alam, Md Ferdous, Max Shtein, Kira Barton, and David Hoelzle. “Reinforcement Learning Enabled Autonomous Manufacturing Using Transfer Learning and Probabilistic Reward Modeling.” *IEEE Control Systems Letters* 7 (2022): 508–513.

- Alfaro, Laura, and Davin Chor. “Global Supply Chains: The Looming “Great Reallocation”.” NBER Working Paper No. 31661, *National Bureau of Economic Research*, 2023. <https://www.nber.org/papers/w31661>.
- Anthropic. *Claude*, 2023. <https://claude.ai>.
- Arinez, Jorge F, Qing Chang, Robert X Gao, Chengying Xu, and Jianjing Zhang. “Artificial intelligence in advanced manufacturing: Current status and future outlook.” *Journal of Manufacturing Science and Engineering* 142, no. 11 (2020): 110804.
- Autor, David. “Work of the Past, Work of The Future.” *AEA Papers and Proceedings* 109, no. 2019 (2019): 1–32.
- Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. “New Frontiers: The Origins and Content of New Work, 1940–2018.” NBER Working Paper No. 30389, *National Bureau of Economic Research*, 2022. <https://www.nber.org/papers/w30389>.
- Autor, David, David Dorn, and Gordon Hanson. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103, no. 6 (2013): 2121–2168.
- Bender, Emily M, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. New York, NY: Association for Computing Machinery, 2021.
- Binelli, Marco, Ryan van Dommelen, Yannick Nagel, Jaemin Kim, Rubaiyet Haque, Fergal Coulter, Gilberto Siqueira, André Studart, and Danick Briand. “Digital manufacturing of personalised footwear with embedded sensors.” *Scientific Reports* 13, no. 1962 (2023).
- Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai-Nayar. “Multinationals, Offshoring, and the Decline of U.S. Manufacturing.” *Journal of International Economics* 127 (2020): 103391.
- Brown, Drusilla K., Alan V. Deardorff, and Robert M. Stern. “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, 2004.
- Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. “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, 2020. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. “Generative AI at Work.” NBER Working Paper No. 31161, *National Bureau of Economic Research*, 2023. <https://www.nber.org/papers/w31161>.
- Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W.W. Norton, 2014.

- Capraro, Valerio, Austin Lentsch, Daron Acemoglu, Selin Akgun, Aisel Akhmedova, Ennio Bilancini, Jean-François Bonnefon, et al. “The Impact of Generative AI on Socioeconomic Inequalities and Policymaking.” SSRN Working Paper No. 4666103, *SSRN*, 2023. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4666103](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4666103).
- Carbonneau, Real, Kevin Laframboise, and Rustam Vahidov. “Application of machine learning techniques for supply chain demand forecasting.” *European Journal of Operational Research* 184, no. 3 (2008): 1140–1154.
- Carvalho, Thyago Peres, Fabrizzio A.A.M.N. Soares, Roberto Vita, Roberto da P. Francisco, João P. Basto, and Symone G.S. Alcalá. “A systematic literature review of machine learning methods applied to predictive maintenance.” *Computers & Industrial Engineering* 137 (2019): 106024. <http://dx.doi.org/10.1016/j.cie.2019.106024>.
- Cavalcante, Ian M., Enzo M. Frazzon, Fernando A. Forcellini, and Dmitry Ivanov. “A supervised machine learning approach to data-driven simulation of resilient supplier selection in digital manufacturing.” *International Journal of Information Management* 49 (2019): 86–97.
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. “Imagenet: A Large-Scale Hierarchical Image Database.” In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. IEEE, 2009.
- Deng, Tianchi, Yingguang Li, Xu Liu, and Lihui Wang. “Federated learning-based collaborative manufacturing for complex parts.” *Journal of Intelligent Manufacturing* 34, no. 7 (2023): 3025–3038.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding.” ArXiv Working Paper No. 1810.04805, *arXiv*, 2018. <https://arxiv.org/abs/1810.04805>.
- Dixon, John R., and Clive L. Dym. “Artificial Intelligence and Geometric Reasoning in Manufacturing Technology.” *Applied Mechanics Reviews* 39, no. 9 (1986): 1325–1330.
- Driess, Danny, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, et al. “PaLM-E: An Embodied Multimodal Language Model.” In *Proceedings of the 40th International Conference on Machine Learning*, 202:8469–8488. Proceedings of Machine Learning Research (PMLR). 2023. <https://proceedings.mlr.press/v202/driess23a.html>.
- Dwork, Cynthia. “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, 2006.
- Dziri, Nouha, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Peter West, et al. “Faith and Fate: Limits of Transformers on Compositionality.” ArXiv Working Paper No. 2305.18654, *arXiv*, 2023. <https://arxiv.org/abs/2305.18654>.
- Evans, Eric. *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Addison-Wesley Professional, 2003.

- Evans, Oliver. *The Young Mill-Wright and Miller's Guide*. Philadelphia, PA: Self-published by the author, 1795.
- Faccio, Maurizio, Irene Granata, Alberto Menini, Mattia Milanese, Chiara Rossato, Matteo Bottin, Riccardo Minto, et al. "Human Factors in Cobot Era: A Review of Modern Production Systems Features." *Journal of Intelligent Manufacturing* 34, no. 1 (2023): 85–106.
- Ferguson, Eugene. *Oliver Evans, Inventive Genius of the American Industrial Revolution*. Greenville, DE: The Hagley Museum / Library, 1980.
- Fowler, Martin. *Patterns of Enterprise Application Architecture*. Addison-Wesley Professional, 2002.
- Francois, Marianne M., Amy Sun, Wayne E. King, Neil Jon Henson, Damien Turret, Ccut Allan Bronkhorst, Neil N. Carlson, et al. "Modeling of additive manufacturing processes for metals: Challenges and opportunities." *Current Opinion in Solid State and Materials Science* 21, no. 4 (2017): 198–206.
- Furman, Jason, and Robert Seamans. "AI and the Economy." *Innovation Policy and the Economy* 19, no. 1 (2019): 161–191.
- Gao, Leo, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. "The Pile: An 800GB Dataset of Diverse Text for Language Modeling." ArXiv Working Paper No. 2101.00027, *arXiv*, 2020. <https://arxiv.org/abs/2101.00027>.
- Google. *Bard*, 2023. <https://bard.google.com>.
- Guertler, Matthias, Laura Tomidei, Nathalie Sick, Marc Carmichael, Gavin Paul, Annika Wambsganss, Victor Hernandez Moreno, and Sazzad Hussain. "When is a Robot a Cobot? Moving Beyond Manufacturing and Arm-Based Cobot Manipulators." *Proceedings of the Design Society* 3 (2023): 3889–3898.
- Gupta, S.K. "Enabling High Manufacturing Quality And Trustworthiness With AI-Powered Robots." *Forbes*, 2023. <https://www.forbes.com/sites/forbestechcouncil/2023/08/11/how-robot-use-in-manufacturing-can-impact-environmental-sustainability>.
- . "How Generative AI Can Accelerate The Deployment Of Autonomous Robots." *Forbes*, 2023. <https://www.forbes.com/sites/forbestechcouncil/2023/09/28/how-generative-ai-can-accelerate-the-deployment-of-autonomous-robots>.
- . "The Importance Of Human-Centered Automation In Manufacturing." *Forbes*, 2023. <https://www.forbes.com/sites/forbestechcouncil/2023/10/30/the-importance-of-human-centered-automation-in-manufacturing>.
- Hassan, Muneeb Ul, Mubashir Husain Rehmani, and Jinjun Chen. "Differential Privacy Techniques for Cyber Physical Systems: A Survey." *IEEE Communications Surveys & Tutorials* 22, no. 1 (2019): 746–789.
- Hindle, Brooke, and Steven Lubar. *Engines of Change: The American Industrial Revolution, 1790–1860*. Washington DC: Smithsonian Institution Press, 1986.

- Ho, Jonathan, Ajay Jain, and Pieter Abbeel. “Denoising Diffusion Probabilistic Models.” In *Advances in Neural Information Processing Systems*, 33:6840–6851. 34th International Conference on Neural Information Processing Systems (NeurIPS 2020). Vancouver, Canada, 2020. <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.
- Hooker, Sara. “The Hardware Lottery.” *Communications of the ACM* 64, no. 12 (2021): 58–65.
- Huang, Jie, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. “Large Language Models Cannot Self-Correct Reasoning Yet.” ArXiv Working Paper No. 2310.01798, *arXiv*, 2023. <https://arxiv.org/abs/2310.01798>.
- Huang, Wenlong, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. “Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents.” In *Proceedings of the 39th International Conference on Machine Learning*, 9118–9147. Proceedings of Machine Learning Research (PMLR). Baltimore, MD, 2022. <https://proceedings.mlr.press/v162/huang22a/huang22a.pdf>.
- Iyer, Naresh S., Amir M. Mirzendehtdel, Sathyanarayanan Raghavan, Yang Jiao, Erva Ulu, Morad Behandish, Saigopal Nelaturi, and Dean M. Robinson. “PATO: Producibility-Aware Topology Optimization using Deep Learning for Metal Additive Manufacturing.” ArXiv Working Paper No. 2112.04552, *arXiv*, 2021. <https://arxiv.org/abs/2112.04552>.
- Jin, Qiao, Yifan Yang, Qingyu Chen, and Zhiyong Lu. “GeneGPT: Augmenting Large Language Models with Domain Tools for Improved Access to Biomedical Information.” ArXiv Working Paper No. 2304.09667, *arXiv*, 2023. <https://arxiv.org/abs/2304.09667>.
- Jobin, Anna, Marcello Ienca, and Effy Vayena. “The global landscape of AI ethics guidelines.” *Nature Machine Intelligence* 1, no. 9 (2019): 389–399.
- Karamcheti, Siddharth, Suraj Nair, Annie S Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, and Percy Liang. “Language-Driven Representation Learning for Robotics.” ArXiv Working Paper No. 2302.12766, *Robotics: Science and Systems*, 2023. <https://arxiv.org/abs/2302.12766>.
- Karpathy, Andrej. *Software 2.0*, 2017. <https://karpathy.medium.com/software-2-0-a64152b37c35>.
- Krizhevsky, Alex. *Learning Multiple Layers of Features from Tiny Images*. Technical report. Available at: <https://learning2hash.github.io/publications/cifar2009learning/>. University of Toronto, 2009.
- Le, Matthew, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, et al. “Voicebox: Text-Guided Multilingual Universal Speech Generation at Scale.” ArXiv Working Paper No. 2306.15687, *arXiv*, 2023. <https://arxiv.org/abs/2306.15687>.
- Li, Boyi, Philipp Wu, Pieter Abbeel, and Jitendra Malik. “Interactive Task Planning with Language Models.” ArXiv Working Paper No. 2310.10645, *arXiv*, 2023. <https://arxiv.org/abs/2310.10645>.

- Liang, Jacky, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. “Code as Policies: Language Model Programs for Embodied Control.” In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 9493–9500. IEEE, 2023.
- Liang, Percy, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, et al. “Holistic Evaluation of Language Models.” ArXiv Working Paper No. 2211.09110, *arXiv*, 2022. <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. “Let’s Verify Step by Step.” ArXiv Working Paper No. 2305.20050, *arXiv*, 2023. <https://arxiv.org/abs/2305.20050>.
- Lumafield. *Lumafield*, 2023. <https://www.lumafield.com/>.
- Makatura, Liane, Michael Foshey, Bohan Wang, Felix Hähnlein, Pingchuan Ma, Bolei Deng, Megan Tjandrasuwita, et al. “How Can Large Language Models Help Humans in Design and Manufacturing?” ArXiv Working Paper No. 2307.14377, *arXiv*, 2023. <https://arxiv.org/abs/2307.14377>.
- Marr, Bernard. “Artificial Intelligence In Manufacturing: Four Use Cases You Need To Know In 2023.” *Forbes*, 2023. <https://www.forbes.com/sites/bernardmarr/2023/07/07/artificial-intelligence-in-manufacturing-four-use-cases-you-need-to-know-in-2023>.
- McClelland, Ryan. “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, 2022.
- McMahan, H. Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. “Communication-Efficient Learning of Deep Networks from Decentralized Data.” In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017*, 54:1273–1282. Proceedings of Machine Learning Research (PMLR). Fort Lauderdale, FL, 2017. <https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf>.
- Midjourney Inc. *Midjourney*, 2023. <https://www.midjourney.com>.
- Miller, Scarlett R., Samuel T. Hunter, Elizabeth Starkey, Sharath Ramachandran, Faez Ahmed, and Mark Fuge. “How Should We Measure Creativity in Engineering Design? A Comparison Between Social Science and Engineering Approaches.” *Journal of Mechanical Design* 143, no. 3 (2021): 031404.
- Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. “Model Cards for Model Reporting.” In *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency*, 220–229. New York, NY: Association for Computing Machinery, 2019.
- Mont Blanc AI. *MontBlancAI*, 2023. <https://montblanc.ai/>.

- Mündler, Niels, Jingxuan He, Slobodan Jenko, and Martin Vechev. “Self-contradictory Hallucinations of Large Language Models: Evaluation, Detection and Mitigation.” ArXiv Working Paper No. 2305.15852, *arXiv*, 2023. <https://arxiv.org/abs/2305.15852>.
- Nagai, Yukari, and John Gero. “Design creativity.” *Journal of Engineering Design* 23, no. 4 (2012): 237–239.
- National Academies of Science, Engineering, and Medicine. *The Fourth Industrial Revolution: Proceedings of a Workshop in Brief*. The National Academies Press, 2017.
- Noy, Shakked, and Whitney Zhang. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381, no. 6654 (2023): 187–192.
- OpenAI. *ChatGPT*, 2023. <https://chat.openai.com>.
- . *DALL-E*, 2023. <https://openai.com/dall-e-3>.
- . *GPT-4 V(ision)*, 2023. <https://openai.com/research/gpt-4v-system-card>.
- Ouyang, Long, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, et al. “Training language models to follow instructions with human feedback.” In *Advances in Neural Information Processing Systems*, 35:27730–27744. 36th Conference on Neural Information Processing Systems (NeurIPS 2022). New Orleans, LA, 2022. [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf).
- Para, Wamiq, Shariq Bhat, Paul Guerrero, Tom Kelly, Niloy Mitra, Leonidas J. Guibas, and Peter Wonka. “SketchGen: Generating Constrained CAD Sketches.” In *Advances in Neural Information Processing Systems*, 34:5077–5088. 35th Conference on Neural Information Processing Systems (NeurIPS 2021). 2021. [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/28891cb4ab421830acc36b1f5fd6c91e-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/28891cb4ab421830acc36b1f5fd6c91e-Paper.pdf).
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirel. “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.” ArXiv Working Paper No. 2302.06590, *arXiv*, 2023. <https://arxiv.org/abs/2302.06590>.
- Picard, Cyril, Kristen M Edwards, Anna C Doris, Brandon Man, Giorgio Giannone, Md Ferdous Alam, and Faez Ahmed. “From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design.” ArXiv Working Paper No. 2311.12668, *arXiv*, 2023. <https://arxiv.org/abs/2311.12668>.
- Radford, Alec, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. “Robust Speech Recognition via Large-Scale Weak Supervision.” In *Proceedings of the 40th International Conference on Machine Learning*, 202:28492–28518. Proceedings of Machine Learning Research (PMLR). 2023. <https://proceedings.mlr.press/v202/radford23a.html>.
- Radford, Alec, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. “Language Models are Unsupervised Multitask Learners.” 2019. <https://api.semanticscholar.org/CorpusID:160025533>.



- Rae, Jack W., Anna Potapenko, Siddhant M. Jayakumar, and Timothy P. Lillicrap. “Compressive Transformers for Long-Range Sequence Modelling.” In *8th International Conference on Learning Representations, ICLR 2020*. Addis Ababa, Ethiopia, 2020. [https://iclr.cc/virtual\\_2020/poster\\_SyIKikSYDH.html](https://iclr.cc/virtual_2020/poster_SyIKikSYDH.html).
- Raji, Inioluwa Deborah, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing.” In *Proceedings of the 2020 ACM Conference on Fairness, Accountability, and Transparency*, 33–44. New York, NY: Association for Computing Machinery, 2020.
- Raji, Inioluwa Deborah, Peggy Xu, Colleen Honigsberg, and Daniel Ho. “Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance.” In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 557–571. New York, NY: Association for Computing Machinery, 2022.
- Raven Telemetry. *Raven.ai*, 2023. <https://raven.ai/>.
- Regenwetter, Lyle, Amin Heyrani Nobari, and Faez Ahmed. “Deep Generative Models in Engineering Design: A Review.” *Journal of Mechanical Design* 144, no. 7 (2022): 071704.
- Ren, Allen Z., Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, et al. “Robots That Ask For Help: Uncertainty Alignment for Large Language Model Planners.” ArXiv Working Paper No. 2307.01928, *arXiv*, 2023. <https://arxiv.org/abs/2307.01928>.
- Rhoades, Matthew A. “Smart Footwear: A Designer’s Perspective.” In *Smart Clothes and Wearable Technology (Second Edition)*, 509–527. The Textile Institute Book Series. Woodhead Publishing, Elsevier, 2023. <https://www.sciencedirect.com/science/article/pii/B9780128195260000114>.
- Rio-Chanona, Maria del, Nadzeya Laurentsyeva, and Johannes Wachs. “Are Large Language Models a Threat to Digital Public Goods? Evidence from Activity on Stack Overflow.” ArXiv Working Paper No. 2307.07367, *arXiv*, 2023. <https://arxiv.org/abs/2307.07367>.
- Rockwell Automation. *FactoryTalk Analytics LogixAI: Machine Learning and Logix*, 2023. <https://www.rockwellautomation.com/en-us/support/documentation/overview/factorytalk-analytics-logixai--machine-learning-and-logix.html>.
- . *Plex*, 2023. <https://www.plex.com/products/production-monitoring>.
- . *Rockwell Automation*, 2023. <https://www.rockwellautomation.com/en-us.html>.
- Rombach, Robin, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. “High-Resolution Image Synthesis with Latent Diffusion Models.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695. 2022. [https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach\\_High-Resolution\\_Image\\_Synthesis\\_With\\_Latent\\_Diffusion\\_Models\\_CVPR\\_2022\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.pdf).
- Romer, Paul M. “Increasing Returns and Long-Run Growth.” *Journal of Political Economy* 94, no. 5 (1986): 1002–1037.

- Rozière, Baptiste, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, et al. “Code Llama: Open Foundation Models for Code.” ArXiv Working Paper No. 2308.12950, *arXiv*, 2023. <https://arxiv.org/abs/2308.12950>.
- Sanghi, Aditya, Hang Chu, Joseph G. Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshan. “CLIP-Forge: Towards Zero-Shot Text-to-Shape Generation.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18603–18613. 2022. [https://openaccess.thecvf.com/content/CVPR2022/papers/Sanghi\\_CLIP-Forge\\_Towards\\_Zero-Shot\\_Text-To-Shape\\_Generation\\_CVPR\\_2022\\_paper.pdf](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. “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*, edited by Richard Curran, Shou-Yan Chou, and Amy Trappey, 123–130. Proceedings of the 15th ISPE International Conference on Concurrent Engineering (CE2008). London, U.K.: Advanced Concurrent Engineering, Springer, 2008.
- Schick, Timo, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. “Toolformer: Language Models Can Teach Themselves to Use Tools.” ArXiv Working Paper No. 2302.04761, *arXiv*, 2023. <https://arxiv.org/abs/2302.04761>.
- Schwab, Klaus. “The Fourth Industrial Revolution.” *Foreign Affairs*, 2015. <https://www.foreignaffairs.com/world/fourth-industrial-revolution>.
- Seepersad, Carolyn Conner. “Challenges and Opportunities in Design for Additive Manufacturing.” *3D Printing and Additive Manufacturing* (New Rochelle, NY) 1, no. 1 (2014): 10–13.
- Shi, Freda, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, and Denny Zhou. “Large Language Models Can Be Easily Distracted by Irrelevant Context.” In *Proceedings of the 40th International Conference on Machine Learning*, 202:31210–31227. Proceedings of Machine Learning Research (PMLR). 2023. <https://proceedings.mlr.press/v202/shi23a.html>.
- Sohl-Dickstein, Jascha, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. “Deep Unsupervised Learning using Nonequilibrium Thermodynamics.” In *Proceedings of the 32nd International Conference on Machine Learning*, 37:2256–2265. Proceedings of Machine Learning Research (PMLR). Lille, France, 2015. <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- SpencerMetrics. *SpencerMetrics*, 2023. <https://www.spencermetrics.com/>.
- Stability AI. *Stable Diffusion*, 2023. <https://stability.ai>.
- Suessmuth, Jochen, Florian Fick, and Stan Van Der Vossen. “Generative AI for Concept Creation in Footwear Design.” In *ACM SIGGRAPH 2023 Talks*, 1–2. 17. New York, NY: Association for Computing Machinery, 2023. <https://doi.org/10.1145/3587421.3595416>.

- Susto, Gian Antonio, Andrea Schirru, Simone Pampuri, Seán McLoone, and Alessandro Beghi. “Machine Learning for Predictive Maintenance: A Multiple Classifier Approach.” *IEEE Transactions on Industrial Informatics* 11, no. 3 (2015): 812–820.
- The Manufacturing Institute. *Creating Pathways for Tomorrow’s Workforce Today: Beyond Reskilling in Manufacturing*, 2023. <https://themanufacturinginstitute.org/research/creating-pathways-for-tomorrows-workforce-today-beyond-reskilling-in-manufacturing/>.
- . “The Aging of the Manufacturing Workforce: Challenges and Best Practices,” 2019. <https://themanufacturinginstitute.org/research/the-aging-of-the-manufacturing-workforce/>.
- Touvron, Hugo, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, et al. “LLaMA: Open and Efficient Foundation Language Models.” ArXiv Working Paper No. 2302.13971, *arXiv*, 2023. <https://arxiv.org/abs/2302.13971>.
- Tulip. *Tulip*, 2023. <https://tulip.co/>.
- Turpin, Miles, Julian Michael, Ethan Perez, and Samuel R. Bowman. “Language Models Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting.” ArXiv Working Paper No. 2305.04388, *arXiv*, 2023. <https://arxiv.org/abs/2305.04388>.
- U.S. Department of Defense. *U.S. Manufacturing Ecosystem Key to Economic Growth, Innovation, Competitiveness*, 2023. <https://www.defense.gov/News/News-Stories/Article/Article/3189049/us-manufacturing-ecosystem-key-to-economic-growth-innovation-competitiveness/>.
- U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy. *Manufacturing Timeline: Advanced Manufacturing & Industrial Decarbonization*, 2019. <https://www.energy.gov/eere/amo/manufacturing-timeline>.
- Ulrich, Karl T., and Steven D. Eppinger. *Product Design and Development*. 6th ed. McGraw-Hill Education, 2015.
- United Nations. *Manufacturing employment as a proportion of total employment (%)*, 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. “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 (2020): 1531–1558.
- Vaithilingam, Priyan, Tianyi Zhang, and Elena L. Glassman. “Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models,” 1–7. Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems 332. New York, NY: Association for Computing Machinery, 2022. <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. “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, 2017. [https://papers.nips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://papers.nips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf).
- Vernon, Vaughn. *Domain-Driven Design Distilled*. Addison-Wesley Professional, 2016.
- . *Implementing Domain-Driven Design*. Addison-Wesley Professional, 2013.
- Wakabayashi, Daisuke. “Does Anyone Know What Paper Towels Should Cost?” *New York Times*, 2022. <https://www.nytimes.com/2022/02/26/technology/amazon-price-swings-shopping.html>.
- Wang, Kevin I-Kai, Xiaokang Zhou, Wei Liang, Zheng Yan, and Jinhua She. “Federated Transfer Learning Based Cross-Domain Prediction for Smart Manufacturing.” *IEEE Transactions on Industrial Informatics* 18, no. 6 (2021): 4088–4096.
- Watson, Nell. *Taming the Machine*. London, UK: Kogan Page, 2024.
- West, Darrell M., and Christian Lansang. “Global manufacturing scorecard: How the U.S. compares to 18 other nations.” *Brookings Institute*, 2018. <https://www.brookings.edu/articles/global-manufacturing-scorecard-how-the-us-compares-to-18-other-nations/>.
- Wolf, Yotam, Noam Wies, Oshri Avnery, Yoav Levine, and Amnon Shashua. “Fundamental Limitations of Alignment in Large Language Models.” ArXiv Working Paper No. 2304.11082, *arXiv*, 2023. <https://arxiv.org/abs/2304.11082>.
- World Bank. *Manufacturing, value added (% of GDP)*, 2021. <https://data.worldbank.org/indicator/NV.IND.MANF.ZS>.
- . *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. “DeepCAD: A Deep Generative Network for Computer-Aided Design Models.” In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 6772–6782. 2021. [https://openaccess.thecvf.com/content/ICCV2021/papers/Wu\\_DeepCAD\\_A\\_Deep\\_Generative\\_Network\\_for\\_Computer-Aided\\_Design\\_Models\\_ICCV\\_2021\\_paper.pdf](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. “An Early Evaluation of GPT-4V(ision).” ArXiv Working Paper No. 2310.16534, *arXiv*, 2023. <https://arxiv.org/abs/2310.16534>.
- Xiao, Han, Kashif Rasul, and Roland Vollgraf. “Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms.” ArXiv Working Paper No. 1708.07747, *arXiv*, 2017. <https://arxiv.org/abs/1708.07747>.
- Xu, Ruiyun, Yue Feng, and Hailiang Chen. “ChatGPT vs. Google: A Comparative Study of Search Performance and User Experience.” SSRN Working Paper No. 4498671, *SSRN*, 2023. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4498671](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4498671).

Zha, Yuheng, Yichi Yang, Ruichen Li, and Zhiting Hu. “AlignScore: Evaluating Factual Consistency with a Unified Alignment Function.” ArXiv Working Paper No. 2305.16739, *arXiv*, 2023. <https://arxiv.org/abs/2305.16739>.

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

Zitkovich, Brianna, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, et al. “RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control.” In *Proceedings of The 7th Conference on Robot Learning*, 229:2165–2183. Proceedings of Machine Learning Research (PMLR). 2023. <https://proceedings.mlr.press/v229/zitkovich23a.html>.