

To excel at engineering design, generative AI must first learn to innovate, study finds

October 19 2023, by Jennifer Chu



MIT engineers trained several AI models on thousands of bicycle frames, sourced from a dataset of full bicycle designs, shown here color-coded by bike style. Credit: Massachusetts Institute of Technology

ChatGPT and other deep generative models are proving to be uncanny

mimics. These AI supermodels can churn out poems, finish symphonies, and create new videos and images by automatically learning from millions of examples of previous works. These enormously powerful and versatile tools excel at generating new content that resembles everything they've seen before.

But as MIT engineers say in a new study, similarity isn't enough if you want to truly innovate in engineering tasks.

"Deep generative models (DGMs) are very promising, but also inherently flawed," says study author Lyle Regenwetter, a mechanical engineering graduate student at MIT. "The objective of these models is to mimic a dataset. But as engineers and designers, we often don't want to create a [design](#) that's already out there."

He and his colleagues make the case that if [mechanical engineers](#) want help from AI to generate novel ideas and designs, they will have to first refocus those models beyond "statistical similarity."

"The performance of a lot of these models is explicitly tied to how statistically similar a generated sample is to what the model has already seen," says co-author Faez Ahmed, assistant professor of mechanical engineering at MIT. "But in design, being different could be important if you want to innovate."

In their study, Ahmed and Regenwetter reveal the pitfalls of deep generative models when they are tasked with solving engineering design problems. In a case study of bicycle frame design, the team shows that these models end up generating new frames that mimic previous designs but falter on engineering performance and requirements.

When the researchers presented the same bicycle frame problem to DGMs that they specifically designed with engineering-focused

objectives, rather than only statistical similarity, these models produced more innovative, higher-performing frames.

The team's results show that similarity-focused AI models don't quite translate when applied to engineering problems. But, as the researchers also highlight in their study, with some careful planning of task-appropriate metrics, AI models could be an effective design "co-pilot."

"This is about how AI can help engineers be better and faster at creating innovative products," Ahmed says. "To do that, we have to first understand the requirements. This is one step in that direction."

The team's new study [appeared recently online](#), and will be published in the December print edition of the journal [Computer-Aided Design](#). The research is a collaboration between computer scientists at MIT-IBM Watson AI Lab and mechanical engineers in MIT's DeCoDe Lab. The study's co-authors include Akash Srivastava and Dan Gutreund at the MIT-IBM Watson AI Lab.

Framing a problem

As Ahmed and Regenwetter write, DGMs are "powerful learners, boasting unparalleled ability" to process huge amounts of data. DGM is a broad term for any machine-learning model that is trained to learn distribution of data and then use that to generate new, statistically similar content.

The enormously popular ChatGPT is one type of deep generative model known as a large language model, or LLM, which incorporates natural language processing capabilities into the model to enable the app to generate realistic imagery and speech in response to conversational queries. Other popular models for image generation include DALL-E and Stable Diffusion.

Because of their ability to learn from data and generate realistic samples, DGMs have been increasingly applied in multiple engineering domains. Designers have used deep generative models to draft new aircraft frames, metamaterial designs, and optimal geometries for bridges and cars. But for the most part, the models have mimicked existing designs, without improving the performance on existing designs.

"Designers who are working with DGMs are sort of missing this cherry on top, which is adjusting the model's training objective to focus on the design requirements," Regenwetter says. "So, people end up generating designs that are very similar to the dataset."

In the new study, he outlines the main pitfalls in applying DGMs to engineering tasks, and shows that the fundamental objective of standard DGMs does not take into account specific design requirements. To illustrate this, the team invokes a simple case of bicycle frame design and demonstrates that problems can crop up as early as the initial learning phase.

As a model learns from thousands of existing bike frames of various sizes and shapes, it might consider two frames of similar dimensions to have similar performance, when in fact a small disconnect in one frame—too small to register as a significant difference in statistical similarity metrics—makes the frame much weaker than the other, visually similar frame.

Beyond 'vanilla'

The researchers carried the bicycle example forward to see what designs a DGM would actually generate after having learned from existing designs. They first tested a conventional "vanilla" generative adversarial network, or GAN—a model that has widely been used in image and text synthesis, and is tuned simply to generate statistically similar content.

They trained the model on a dataset of thousands of bicycle frames, including commercially manufactured designs and less conventional, one-off frames designed by hobbyists.

Once the model learned from the data, the researchers asked it to generate hundreds of new bike frames. The model produced realistic designs that resembled existing frames. But none of the designs showed significant improvement in performance, and some were even a bit inferior, with heavier, less structurally sound frames.

The team then carried out the same test with two other DGMs that were specifically designed for engineering tasks. The first model is one that Ahmed previously developed to generate high-performing airfoil designs. He built this model to prioritize statistical similarity as well as functional performance.

When applied to the bike frame task, this model generated realistic designs that also were lighter and stronger than existing designs. But it also produced physically "invalid" frames, with components that didn't quite fit or overlapped in physically impossible ways.

"We saw designs that were significantly better than the dataset, but also designs that were geometrically incompatible because the model wasn't focused on meeting design constraints," Regenwetter says.

The last model the team tested was one that Regenwetter built to generate new geometric structures. This model was designed with the same priorities as the previous models, with the added ingredient of design constraints, and prioritizing physically viable frames, for instance, with no disconnections or overlapping bars. This last model produced the highest-performing designs, that were also physically feasible.

"We found that when a [model](#) goes beyond statistical similarity, it can

come up with designs that are better than the ones that are already out there," Ahmed says. "It's a proof of what AI can do, if it is explicitly trained on a design task."

For instance, if DGMs can be built with other priorities, such as performance, design constraints, and novelty, Ahmed foresees "numerous engineering fields, such as molecular design and civil infrastructure, would greatly benefit. By shedding light on the potential pitfalls of relying solely on statistical similarity, we hope to inspire new pathways and strategies in generative AI applications outside multimedia."

More information: Lyle Regenwetter et al, Beyond Statistical Similarity: Rethinking Metrics for Deep Generative Models in Engineering Design, *Computer-Aided Design* (2023). [DOI: 10.1016/j.cad.2023.103609](https://doi.org/10.1016/j.cad.2023.103609). hdl.handle.net/1721.1/152444

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: To excel at engineering design, generative AI must first learn to innovate, study finds (2023, October 19) retrieved 21 May 2024 from <https://techxplore.com/news/2023-10-excel-generative-ai.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
