

AI Time, Timing, and Timelessness

Phil Bernstein

POSITION PAPER

When considering the evolution of technology in design, I have found that procrastination is a virtue, particularly when asked to speculate on the implications of emergent artificial intelligence. Just as soon as I think I have a foothold in a theoretical position or other polemic, the tools and their capabilities lurch ahead, reframing the proposition in ways previously unconsidered. Such was my experience when formulating thoughts for this essay in the Fall of 2023, being written from a vantage point seemingly light years beyond those in my RIBA (Royal Institute of British Architects) piece of 2022 referenced in the provocation for this edition of *The Plan Journal*. As technical iterations appear with increasing speed, and with them the expanding capabilities of the tools, one can barely organize a thought before being hit with the “next big thing” in autonomous computing, machine learning, and non-human intelligence.

Each of the major AI competitors (Microsoft, Google, Facebook, and soon-to-be Amazon) announced another iteration of their generative offerings over the weekend I began this essay; so-called “multimodal”¹ tools that create text, images, audio, and video in what seems to be a highly fluid, cross-connected network of algorithms, training sets, and correlations. Even the emergent discipline of “prompt engineering”—using text to nudge an image generator to a usable result—must now somehow include the mastery of text extraction from image descriptions and back again. Can three-dimensional representation be far from joining this technology party where generative AI can already create images, text, music, and video?

TIMING TECHNOLOGY

Let us assume that is the case, and maybe we are mere months (weeks?) away from generating complete models of buildings. If large language

models were the first spatial dimension of these capabilities, and images perhaps the second, maybe video is two-and-a-half-D (as we used to call the early 3D of AutoCAD's "Z" axis) and multimodal design representation is the third, and with it, perhaps, the ability to truly represent and manipulate full building design. This is, of course, an extraordinarily optimistic view and a huge leap from today's AI, algorithms that generate output that is miles wide with endless possibilities, and only pixels deep. They yield stuff that is sometimes fascinating, almost always surprising, and periodically wrong or even dangerous, nonsensical, or just hallucinogenic. Generative systems will have to overcome these critical failures to get beyond today's computational parlor tricks and become useful tools available to architects to design complete, interesting, competent, and beautiful buildings.

GETTING TO SEMANTIC VALIDITY

At the core of this challenge is the concept of "semantic validity," a theory of logic that "[a]n inference is valid if all interpretations that validate the premises validate the conclusion" ² yielded by a given model. In the case of our hypothetical, 2.5D bot-generated design model of a building, this is a pretty tall order. A design is a very complex hierarchy of interconnected decisions that the best architects orchestrate by carefully balancing the myriad parameters of each well; and today's AI generators do not work by inference, only correlation. As an example, take a relatively simple building component: an interior column. The column's precise configuration results in part from the following list of considerations, combined into a "column solution," in no particular order:

- location relative to nearby columns and other loads,
- load capacity,
- connection to adjacent structures (like beams),
- material choice,
- cost of that material to source, fabricate, deliver, purchase, and install,
- enclosure,
- finish and proportion,
- relationship to other elements in the space, technically (coordination), spatially (circulation), and aesthetically (composition),
- building code,
- construction sequence for installation, and
- [...] this list could go on.

Just representing this column in computable form is a huge challenge; a building information model is just a start. Our hypothetical bot that can generate a semantically valid column in this manner would need a comprehensive "understanding" of not just the essence of "column," but also every other element within the project with which this column has a relationship. That chain of relationships, connections, and necessary inferences expands exponentially as it connects to the larger enterprise:

the structural system, the cost estimate, the permitting, the ceiling system, and so forth.

Today's generative algorithms are built on very complex correlation schema called neural networks, which accumulate vast numbers of connections between adjacent data points, be they words or pixels, so at least that's a start. "Validity," or at least the appearance thereof, results from extensive tweaking of the output of these systems, a combination of training (looking at millions of pictures of columns on the internet) and tuning ("yes, that is a column") usually by human intervention. As of this writing, more than 250 tech companies are attempting AI-based tools for the building industry,³ many of whom are tweaking image generators toward specific ends like rendering a sketch or decorating a room.⁴ Each of these companies must build context-specific tweaking strategies—training their systems, creating intermediate rules-based data structures—to make the generalized infrastructure of the bigger generation platforms do something specifically useful for an architect. No one, however, has a theory about how these disparate systems might combine to yield a coherent, useful tool to handle an entire building design.

PRODUCTIVITY

Thus, the various start-ups in the building industry space are attacking the problem from another well-trodden angle: making design more productive by choosing specific, limited tasks where AI could be helpful. Who would not want to make it easy to render a hand sketch, or quickly see what a redecorated room might look like? The productivity ills of the larger building industry are well documented, if not well understood,⁵ as so much of the recent digitization of the building industry process has focused not on coherent data representation, analysis, or generation, but small process targets like data management, rendering, small scale fabrication (think 3D printers, not printed buildings) and data collection. These efforts, whether AI-based or otherwise, generate a broad, disorganized corpus of digital information largely without a taxonomy, and there are orders of magnitude with less digital information on the design enterprise than, say, the billions of sentences on the internet used to train large language models. This will make training design AI generators even more challenging. That said, an unintended result of this strategy is that the potential of AI to improve our productivity as designers in the long run, by generating, analyzing, and evaluating large-scale, multimodal digital representations of buildings, is pushed even further off by short-term early AI solutions that are emerging today. The episodic advantage of an AI that could, say, choose the right sort of glass for your curtain wall will likely never yield the longer-term ability to design the entire system as a whole in relation to the other components of the building.

PRODUCTS OF DESIGN

This more limited vision of AI efficacy in design is not necessarily a bad thing. Unlike the broad ethical challenges of today's "general purpose" text and image generators, it's likely to be a while before such systems can create buildings out of whole cloth, giving the profession a terrific chance to be opportunistic about where the vendors—likely to be the usual suspects such as Autodesk, Trimble, and Nemetschek—should focus their attention. Productivity opportunities abound, on both the design and construction sides of the computational equation. In design, one could imagine assistive systems operating in parallel with building information models that could report on cost implications, constructability, climate impacts, acoustics, or any one of several "could they not see that coming?" questions that arrive after an innovative project is complete. AI-based systems that work in those limited contexts, on specific issues where data is available, would serve designers in two ways: building credibility in the efficacy of their designs, and producing new opportunities to create value for which clients might be willing to pay a premium; whether that results from better insights and results, or merely more time spent in design and less working out complex, analysis-driven, but machine-worthy challenges.

During construction, when reams of digital information can result from a single project (LIDAR [light detecting and ranging] scans, drone photos, bills of materials, RFI [request for information] logs, invoicing, the list is endless) generative multimodal AI might play a different role, examining patterns and results in a series of projects and using the resulting correlations to predict labor requirements, supply chain optimization, or execution problems on the job site before they occur. The resulting insights could, when paired with an analytical AI cousin during design, provide construction information to the architect as decisions are made in real time. In a recent presentation on the implications of technology for project delivery,⁶ architect Craig Webber described the challenges of managing complex projects as "networks of interconnected risks." The interaction of constraints, information, conditions on the ground, and money affect the ability of a supply chain to deliver, in a world where climate change, demand, material prices, and even fair labor conditions may all combine to cause project failure. Future systems could both understand the correlative relationships between these factors and potentially predict them based on past experience and relationships between factors too complex for human analysts to coherently compile and understand.

PROFESSIONALISM

Perhaps the output of today's systems—what Gary Marcus calls "unreliable mimicry"⁷—is another opportunity for the design profession, rather than a threat. If the most promising direction for AI-generated material is the creation of design componentry, analytical conclusions, and episodic

optimization of discrete portions of a design (rather than an entire project writ large) then the designer's role as decision-maker and integrator is still well intact. Gary Marcus notes that the predictions about replacing radiologists with algorithms, even in the limited realm of reading X-rays, have yet to come to pass.⁸ If anything, the additional "stream of insight" that algorithms may provide to the designer makes her role even more, rather than less, important: sorting the useful from the hallucinogenic, or even dangerous, potential outputs and deploying them appropriately. In fact, one might argue that the concept of "semantic validation" in the service of solving the wicked problems of design is something of an oxymoron. Making decisions that require judgment, synthetic thinking, an understanding of trade-offs, and resolution of ambiguity is a distinctly human skill not soon to be replaced by machines, particularly when it comes to the enormously complex challenge of designing even the simplest of buildings.

Moreover, society has a vested interest in not delegating design authority to algorithms. Just as the U.S. courts have recently concluded that AI output cannot be protected by copyright,⁹ it is impossible to assign technical or legal responsibility to the output of an AI generator, and it is not even a good idea. The legal standard for competent performance by an architect, the standard of care, holds a designer to a comparable performance to someone in similar circumstances who did their job well, and precedent behavior is a critical part of that judgment. What is the precedent for the digital design thinking provided to the client by a bot? The recent unfortunate experience of a New York attorney who unwittingly used ChatGPT to write a brief—with the expected completely hallucinated result—was sanctioned by the court for his reliance on generative technology that he clearly did not understand.¹⁰ His lapse is a convenient early warning to both professionals and their clients that these programs need constant adult supervision, or potential disaster awaits. It is in everyone's interest to make sure professional architects are still in charge.

TIMELESSNESS

Further, perhaps having architects continue to create the spatial world is also in the best interest of architecture. The most significant projects, those that advance our understanding of our world and culture and express our values, are both of their time and simultaneously timeless, speaking to us across the years from their date of origin. As the discipline of design has evolved in parallel with the societies that need it, those timeless qualities emerge not from formulaic solutions, slavish adherence to precedent, or wild image creation, but from fundamental rethinking of the propositions of design and the spatial and temporal insights that the best designers deploy. The sources of those insights—whether digitally generated or otherwise—will evolve, but the human ability to create will expand along with, rather than be replaced by, new technologies, including these newly capable

AI bots. And much like the tools that came before—be they precision-measuring instruments, CAD, or even the internet—they will be absorbed into the discipline and eventually put to their best use. Let us hope today’s designers can guide that trajectory well into the future.

Notes

1. Sophie Bushwick and Lauren Leffer, “The State of Large Language Models,” *Science Quickly*, podcast audio, October 2, 2023 – <https://www.scientificamerican.com/podcast/episode/the-state-of-large-language-models1/>.
2. [https://en.wikipedia.org/wiki/Validity_\(logic\)](https://en.wikipedia.org/wiki/Validity_(logic)), accessed October 10, 2023.
3. See “AI in AEC: Apps Database” – <https://stjepanmikulic.notion.site/4fbe033065dc4780b714e45fce57b852?v=db7f465f76804419852a13f28071e6b3>, accessed October 1, 2023.
4. See <https://corbu.ai/> and <https://www.remodelai.co/>.
5. Austen Goolsbee and Chad Syverson, “The Strange and Awful Path of Productivity in the U.S. Construction Sector” (working paper, National Bureau of Economic Research, 2023) – <https://www.nber.org/papers/w30845>.
6. Craig Webber, “Unleashing Construction Innovation with IPD” (presentation, IPDA 2023 Conference – Accelerating Adoption of Integrated Project Delivery, Montreal, October 3, 2023).
7. Gary Marcus, “What Was 60 Minutes Thinking, in That Interview with Geoff Hinton?” *Substack* (blog), October 10, 2023 – <https://garymarcus.substack.com/p/what-was-60-minutes-thinking-in-that>.
8. Ibid.
9. Jane Recker, “U.S. Copyright Office Rules A.I. Art Can’t Be Copyrighted,” *Smithsonian Magazine*, March 24, 2022 – <https://www.smithsonianmag.com/smart-news/us-copyright-office-rules-ai-art-cant-be-copyrighted-180979808/>, accessed October 10, 2023.
10. Josh Russell, “Sanctions Ordered for Lawyers Who Relied on ChatGPT Artificial Intelligence to Prepare Court Brief,” *Courthouse News Service*, June 22, 2023 – <https://www.courthousenews.com/sanctions-ordered-for-lawyers-who-relied-on-chatgpt-artificial-intelligence-to-prepare-court-brief/>, accessed October 10, 2023.

References

- Bushwick, Sophie and Lauren Leffer. “The State of Large Language Models.” *Science Quickly*. Podcast audio, October 2, 2023. <https://www.scientificamerican.com/podcast/episode/the-state-of-large-language-models1/>.
- Goolsbee, Austen and Chad Syverson. “The Strange and Awful Path of Productivity in the U.S. Construction Sector.” Working paper, National Bureau of Economic Research, 2023. <https://www.nber.org/papers/w30845>.
- Marcus, Gary. “What Was 60 Minutes Thinking, in That Interview with Geoff Hinton?” *Substack* (blog), October 10, 2023. <https://garymarcus.substack.com/p/what-was-60-minutes-thinking-in-that>.
- Recker, Jane. “U.S. Copyright Office Rules A.I. Art Can’t Be Copyrighted.” *Smithsonian Magazine*, March 24, 2022. <https://www.smithsonianmag.com/smart-news/us-copyright-office-rules-ai-art-cant-be-copyrighted-180979808/>.
- Russell, Josh. “Sanctions Ordered for Lawyers Who Relied on ChatGPT Artificial Intelligence to Prepare Court Brief.” *Courthouse News Service*, June 22, 2023. <https://www.courthousenews.com/sanctions-ordered-for-lawyers-who-relied-on-chatgpt-artificial-intelligence-to-prepare-court-brief/>.
- Webber, Craig. “Unleashing Construction Innovation with IPD.” Presentation at the IPDA 2023 Conference – Accelerating Adoption of Integrated Project Delivery, Montreal, Canada, October 3, 2023.

Phil Bernstein is an architect and technologist who earned a B.A. and an M.Arch at Yale University School of Architecture, where he has taught since 1988. He was a Vice President at Autodesk, responsible for setting the company's future vision and strategy for BIM technology. Phil was a principal at Pelli Clarke and Partners Architects where he managed complex commissions for the Mayo Clinic, Goldman Sachs, and Reagan Washington National Airport. He is the author of *Machine Learning: Architecture in the Age of Artificial Intelligence* (2022), *Architecture | Design | Data – Practice Competency in the Era of Computation* (2018); co-editor of *Building (In) The Future: Recasting Labor in Architecture* (2010 with Peggy Deamer), and a Fellow of the AIA, a Senior Fellow of the Design Futures Council, and former Chair of the AIA National Contract Documents Committee. E-mail: phillip.bernstein@yale.edu