The Future of Generative AI in Architecture, Design, and Engineering

January 2024

This report explores how Generative AI (GAI) is set to disrupt fields such as architecture and construction, presenting both opportunities and challenges that demand strategic management.

JACOBS INSTITUTE AT CORNELL TECH

Prepared by :

Greg Lindsay

Anthony Townsend

Horizon Scan

Introduction

Generative AI (GAI) are a new class of tools enabling users to quickly generate digital content in response to written, spoken, or visual prompts. Large language models (LLMs) such as GPT-4 have been harnessed for such tasks as natural language processing, machine translation, and written content ranging from novels to software code. Launched in November 2022, its predecessor ChatGPT immediately surpassed more than 100 million users, becoming the most successful product launch of all time[1]. Meanwhile, image generators such as OpenAI's DALL-E, Stability AI's Stable Diffusion, and Midjourney have rapidly evolved to create renderings and animations virtually indistinguishable from photography and films.

GAI builds on the "deep learning" revolution of the last decade in which algorithms are trained on ever-larger datasets involving immense human- and machine effort. Their basic design leverages multiple recent advances in neural networks[2] — including diffusion models, generative adversarial networks (GANs), and transformers — to teach algorithms increasingly complex pattern-matching techniques through a process of "unsupervised learning[3]" using human-generated content, much of it scraped from public internet sites. The results are "foundation models" of exponentially growing size and sophistication — while ChatGPT's training dataset was ten times the size of its closest competitor at launch, GPT-4 is rumored to be vastly larger still.

The shock wave of GAI's seemingly-magical ability to produce thoughtful content and analysis at almost zero marginal cost to the end user has just begun to ripple through global politics, industry, and culture. But many firms and industries are already bracing for impact, as seen in the recent strikes strikes by both the Screen Actors Guild (SAG-AFTRA) and Writers Guild of America (WGA), with the former worried about replacement by AI-generated avatars modeled on their likenesses, and the latter winning concessions that forbid LLMs from doctoring scripts or using screenwriters' work as training data.

Another field grappling with the implications of GAI is architecture, engineering and construction (AEC). Curious about its capabilities and eager to collectively assess and mitigate the risks, several dozen AEC firms comprising the Innovation Design Consortium (IDC) approached the Jacobs Urban Tech Hub to scan the horizon out to 2030 for potential threats and opportunities. This briefing is drawn from interviews with experts and leading practitioners along with additional research, in an effort to understand how GAI might evolve as both a technology and industry generally, and disrupt AEC specifically.

"Technology is the answer, but what was the question?" This puzzle, famously raised by the British architect Cedric Price, offers a perfect starting point for how AEC firms should approach the rapid development of GAI. What exactly are the "answers" this technology will provide; why should they care; and what should they do about it? Our horizon scan reframes these questions into a set of three hypotheses reflecting strong-but-weakly-held views of what comes next in the

1 https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

2 Artificial neural networks are a machine learning technique modeled loosely on biological neurons, in which large clusters of densely interconnected processing nodes receive and relay date inputs. Each connection has an assigned numerical weight, which determined whether a given node "fires" in producing the final output.

3 Unsupervised machine learning techniques use algorithms to sort and associate data without the need for human input.

computing landscape. Building on these hypotheses, we describe three strategic dilemmas—problems that don't have easy solutions but must be continually managed over time.

Hypotheses: What Comes Next

The blistering rate of change in the development of GAI over the last year — which represents an acceleration in the two-and-a-half years between the release of GPT-3 in March 2020 and ChatGPT in November 2022 — creates a perception that forecasting its continued evolution is difficult at best and futile at worst. But while GAI has brought considerable volatility and speculation to the market for computation, several emerging trends will establish a baseline for future developments firms can use to develop their own hypotheses about the near-term future.

Hypothesis 1. This is a step-change disruption.

Clear limits on the size of LLMs are already beginning to emerge, with GPT-4 speculated to be an ensemble of smaller models rather than a single, super-massive model like its predecessors, to name just one example. As the Web is flooded with AI-generated content, training large models on publicly-scraped data will become increasingly less effective due to "model collapse" — an effect observed in foundation models only trained on their own output. Based on these and other signals, it would appear that bigger-is-always-better approaches to GAI may be approaching their limits, and thus the breakthroughs of the last year may be running their course. Rather than an exponential "singularity" breakout toward artificial general intelligence, what we are witnessing is a one-time step change setting a new baseline for strategy.

Hypothesis 2. The current boom will deliver abundant computational capacity.

Generative AI is "compute-bound," meaning adding more computing power produces better results. This has created enormous demand for limited GPU capacity, as seen in the arms race to procure or produce one's own silicon. Analysts at Andreesen-Horowitz speculate that demand currently outstrips supply as much as ten-fold[4]. Cornell Tech researchers routinely encounter throttling on AWS GPU clusters during peak hours when training workloads are running. The industry's response has been to spend an estimated $100 billion on new data centers in 2023 alone. As with past infrastructure investment bubbles, today's capacity crunch will soon give way to tomorrow's computing glut that will drive prices down for everyone, inviting entrepreneurial users to invest in new uses for this abundance.

Hypothesis 3. Pressure to adopt and adapt to AI will come from the bottom-up, and the outside-in.

Adopting and adapting workflows to GAI — much like prior waves of digital transformation — will happen more slowly and painfully in large, existing organizations than smaller, newer ones. The rate of change will also vary across industries, with AEC poised to lag behind others due to its structure — firms typically lack the time and resources to invest in developing new tools, skills, and workflows before their commercial value is proven elsewhere. Meanwhile, technology vendors are more focused on consumers than enterprises, relying on "bring your own device" policies to create organic demand (and subscriptions) rather than pitch CIOs on seat licenses. This strategy is playing out once again with GAI, as startups pivot to function-specific offerings — e.g. AI for copywriting, AI for programming — rather than industry verticals. Together, these trends mean the strongest pressures to harness GAI's potential and adapt to its current form will come from

4 https://a16z.com/navigating-the-high-cost-of-ai-compute/

outside AEC companies (e.g. clients, contractors) rather than within, and from the bottom-up (teams and individuals) rather than the top-down.

Dilemmas: Tradeoffs Shaping Strategy

Workflows: Human experience vs. Machine learning

The first strategic dilemma arises when organizations decide to develop and deliver generative AI use cases. Effectively dealing with this dilemma means being ambitious yet also pragmatic and flexible about the prospects of GAI to organize and exploit an organization's knowledge. Much as earlier waves of knowledge management systems (many of them based on earlier AI approaches) over-promised and under-delivered, so too, will GAI.

This isn't a question of identifying creative tasks with specific functions or steps to be handed over to AI— it's a matter of redesigning workflows to coordinate human and machine work in complementary ways that will inevitably change over time as individuals, algorithms, and organizations adapt. We foresee a number of open questions organizations must wrestle with.

1. How do we assess where GAI should augment human creativity and where it can replace it?

While much of the AI discussion and debate to date has focused on automation versus augmentation — i.e. in what areas should we enhance employees' capabilities rather than replace them? — a better starting point for firms might be tasks requiring little creativity at all.

For example, while architects are understandably obsessed with image generation, automating RFP responses using retrieval augmented generation to surface relevant expertise and projects while checking hundreds of pages of boxes has a vastly higher ceiling on ROI. The use of such systems has already begun transforming the RFP process from one requiring laborious manual effort to more of a quality assurance role in verifying GAI output, while raising more profound questions around how faster response times may lead to a corresponding increase in the number of responses — with downstream implications for firm strategy and operations. It also raises questions of how GAI might be used as a kludge to address workflow bottlenecks stemming from a lack in IT investment.

As seen in previous waves of enterprise software adoption, the ad hoc use of "freemium" tools by small teams and teams can be fruitful in discovering novel use cases. But a purely bottom-up approach leads to redundancies and a failure to scale. For now, firms may be content to experiment relying on self-identified early adopters and communities of practice, but they should start planning how best to assess which tasks and processes have reached the threshold of incorporating AI, and who is best positioned in firms leadership to make such decisions. Today's prompt engineers will give away to tomorrow's AI ombudsmen.

2. Which end of the workflow will GAI eat first — the front or the back?

While image generators may have captured architects' attention and imaginations, it's not a foregone conclusion that GAI will transform firms' workflows from front to back — in fact, initially prizing tools for conceptualization may lead to costly path dependencies further down the road.

For the moment, tools such as lookX and Gaia are leveraging GAI to create "Midjourney, but for architecture"-style tools used for initial conceptualization and rapid iteration with/on behalf of clients. But it remains to be seen — and is in fact doubtful — that this will be

> "It's a matter of redesigning workflows to coordinate human and machine work in complementary ways"

the primary driver of ROI by 2030. Rather, using GAI to translate and standardize output across various teams and functions — from conceptualization (rasterized images) to generating build drawings (prompt-to-point cloud) to refinement and revisions — is far likelier to create long-term value for firms.

This will in turn require significant and strategic investment in end-to-end data pipeline capable of supporting and exploiting new developments — means firms should be looking ahead to the end of pipeline and working backward rather than starting from the tools offered at the front and building out.

3. How can generative AI extract latent value from firms' IP?

While it's axiomatic clients own their IP, firms' body of work and the learning-by-doing it embodies represents an archival metadata of sorts that can be analyzed, internalized, and instrumentalized by firms or even a consortium of firms such as the IDC. More than an approximation of pure style, this tacit knowledge exists at both an individual- and firm level, and may comprise decades of experience. How might it be cost-effectively extracted and utilized with GAI?

Consider for a moment an image generator tuned to produce images drawn from a firm's work in a given typology (e.g. secondary schools) and timeframe (e.g. 1970-2010) — or even a foundation model trained from scratch on the collective output of a consortium's members. How would such tools alter how firms see their archives? What steps are necessary to digitize those archives for training input? How would the advent of such tools and models alter a firm's strategy and the competitive landscape? Would it accelerate M&A between firms and across adjacent industries in pursuit of vertically- or horizontally integrating past portfolios? (An obvious analogy would be music streaming, which has propelled massive asset prices and sales of artists' back catalogs in similar fashion.)

It's not difficult to see how in this instance firms' archives become their crown jewels of training data. How might a consortium like the IDC become an umbrella organization for federated sharing of data and models between firms? What cybersecurity and other risk management measures would be necessary to avoid ingestion and regurgitation by larger-scale models?

Tools: Off-the-rack v. Made-to-measure v. Bespoke

The second strategic dilemma revolves around acquiring GAI infrastructure, which includes machine learning pipelines, front- and back-end application components, and related data stores. But GAI is being packaged in many different forms — from browser-based SaaS to locally-hosted open source and everything in-between — to exploit the flexibility of cloud computing to decouple and distribute different parts of the supply chain. Where are firms to start and how are they to assess the tradeoffs?

We see three classes of GAI arising in the marketplace, distinguished by escalating resource commitments and increased control and customization, characterized by a familiar analogy from the world of tailoring:

1. Off-the-rack AI

Services such as ChatGPT and Bing Chat are designed for instant use without extensive tailoring. They cost less up front, but will never fit a user's needs precisely, even with slight modifications. In return, they provide the least control over how input data is processed and stored, as well as the terms and conditions for output. Use cases will tend to be tool-driven, which projects steered heavily (albeit implicitly) by the capabilities and limitations of whatever tool is under consideration. (Whether this selection process is ever revealed to the client is up to the firm.)

2. Made-to-measure AI

Garments tailored to individual wearers from a palette of pre-cut patterns are "made-to-measure." This approach attempts to balance the efficiencies of ready-to-wear mass production with the personalization of bespoke tailoring. Many GAI services offer an application programming interface (API) providing a standard approach to the core functionality of underlying foundation modals that can be integrated into an organization's own software built atop it. A made-to-

measure approach allows a firm's developers to focus on highly specific use cases and user experience along with potentially new capabilities, such as:

Customizing OpenAI models (e.g. ChatGPT) with proprietary data, using OpenAI's fine tuning model;

Constructing multi-step processes to develop more sophisticated GAI workflows, such as harnessing one LLM to verify another's output;

Integrating or comparing capabilities across multiple GAI platforms;

Developing applications capable of ingesting and outputting structured data for use in data analysis and visualization.

3. Bespoke AI

The most resource-intensive approach is building one's own model from the ground up. Befitting its namesake, this approach is costly and risky, but a surefire way to differentiate oneself in the marketplace — and one within reach for a growing number of organizations. Open source LLMs and image generators are already competitive with proprietary models, while relative data scarcity is pushing developers to create effective models on training sets several orders of magnitude smaller than current incumbents. And the surge of investment into GPUs is likely to drive a dramatic decrease in training costs due to a glut in capacity. (Energy costs may be another matter, however.)

The upsides of the bespoke approach are clear — full control over model design and training, as well as instant access to technological advances gleaned from open source projects. It also offers full control over inputs and outputs, all but eliminating risks arising from sensitive data and ownership of generated content.

Ethics: Reactive vs. Proactive

The final strategic dilemma concerns the use of GAI "responsibly," which is difficult to do when the scope of responsibility is so ill-defined. Whether the issue at hand is algorithmic bias, copyright (both input and output), data privacy violations, or even psychologically harmful content, the fundamental tension for firms is doing the right thing now versus doing it later, i.e. leading or following.

Leading can differentiate an organization from competitors, or collectively the industry from others. It involves considerable upfront resources and attention to course-correct ethical shortcuts by various players at various stages. Benefits include staying one step ahead of regulators in the EU and USA with regards to issues of copyright and model transparency. In rare cases, it may even create opportunities to address new markets that would otherwise be unable or unwilling to procure products and services produced with GAI.

Following involves different risks, and offers opportunities to learn from others' mistakes. A fast-follower strategy can often combine the best of both approaches, but requires developing a keen sensing system and the ability to act quickly to exploit clear opportunities. Opportunists may well benefit from lax regulation or overseas jurisdictions with regards to copyright violations, legal liability extending from GAI "hallucinations," and biases baked too deep into models to dislodge.

From a *realpolitik* perspective, the issue for the industry is one of leverage. Do architecture and engineering firms stand to gain from leading rather than following when it comes to attracting fresh talent, winning clients away from disruptive entrants, and winning the hearts and minds of their rank-and-file? If so, how might individual firms and the industry as a whole choose to lead? What might a GAI Hippocratic oath (or *AIA Code of Ethics*) look like? What role could they play in resolving tensions around data usage rights and copyrightable output? And how will firms reconcile their ESG and DEI goals with GAI's mounting energy footprint and algorithmic bias without leading the charge to stamp out both?

Insights: What It Means For Architectural Design And Engineering

What's to be done in light of these dilemmas, which will consume the remainder of this decade? Taking a wait-and-see approach won't suffice, given the wholehearted embrace by tech giants and rapid adoption in adjacent, potentially disruptive industries. So, where should architectural design and engineering firms begin?

1. Develop an open innovation strategy.

Engage with innovative practitioners, technologists, open source model builders, and startups. Look for opportunities to partner on projects and development, and find creative ways to bring them in-house for exposure and education. Make many small bets — while it's easy to speculate where disruptive entrants might emerge, it's impossible to predict. Accept they may come from any direction and be ready.

2. Experiment with custom tooling.

Rather than rely solely on off-the-rack AI for tasks such as image generation, firms should strongly consider the business case for developing made-to-measure or bespoke solutions offering greater control, differentiation from competitors, and strategic transformation rather than bolted-on features. With this increase in capacity comes a commitment to supporting application developers, end users, and both maintaining and constantly improving tooling. A major question for the IDC is whether and how to federate some of this commitment at a consortium- rather than firm level.

3. Define use cases.

Regardless of firms' commitment to the recommendation above, they should also invest in extensive, quasi-proprietary refinements to open source models where opportunities permit in order to both leverage their archive IP and hedge their bets if a combination of market forces and government regulation means only a handful of foundation models survive to dominate. Be ready for a future of open source abundance, and be prepared for one of proprietary oligopoly.

4. Develop trust requirements.

Signal to the market that the IDC and its members hold themselves to a higher standard when it comes to ethics, safety, and security. Draft a code of practice, require members commit to it, and market yourselves relentlessly to interested parties in order to overcome your structural disadvantage against potential disruptors unburdened by conscience.

Viewpoints

To provide context for both the industry- and horizon scans above, Jacobs Urban Tech Hub fellow **Greg Lindsay** interviewed multiple academics and practitioners (and an AIA executive) for their viewpoints on GAI's uses, abuses, and potential. Those interviews, edited for length and clarity, are presented below.

Phil Bernstein

Phil Bernstein is an architect and technologist who has taught at the School of Architecture since 1988 and where he received his B.A and M.Arch. He was a Vice President at Autodesk where he was responsible for setting the company's future vision and strategy for BIM technology. Prior to Autodesk Phil was a principal at Pelli Clarke and Partners Architects where he managed many of the firm's most complex commissions including projects 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) and co-editor of *Building (In) The Future: Recasting Labor in Architecture* (2010 with Peggy Deamer), and consults, speaks and writes extensively on technology, practice and project delivery. He is a Fellow of the American Institute of Architects, a Senior Fellow of the Design Futures Council, and former Chair of the AIA National Contract Documents Committee.

Greg Lindsay: Given you've literally written the book about the use of AI and machine learning in architecture, what stands out to you about generative AI? Is it a break from prior techniques, or merely an iteration? What features do you find particularly interesting, novel, and/or powerful?

Phil Bernstein: Let me put it this way — as I said to our Dean last week, it's theoretically possible — if not practically — for one of our students to design an entire project and never draw a line. One of the questions for us as educators and the profession at large is: What do we think of that?

It's clear these technologies provide a whole new avenue for generating ideas, and for doing things architects aren't good at, like prediction analysis and education. But as usual, there's an extreme fascination with form generation and image-making that's interesting but in danger of being a distraction.

Lindsay: And where do you fall on the question about drawing a line or not?

Bernstein: I think the healthiest attitude is the one we take about all tools, which is that it's just one in your toolbox. I like to say a generative image creator is like a bandsaw. We teach you how to use a band saw safely. We should teach you how to use a generative tool safely.

Lindsay: And you wouldn't want to use a band saw for everything, either. What should be off-limits when it comes to using generative tools?

Bernstein: When we don't understand the capability of the tool, it's a little early to decide what's off-limits. It's like asking what the speed limit should be when we haven't invented the internal combustion engine yet. I've been around long enough to see two "AI winters" in my own career, when everyone felt the world was about to change in a dramatic way before we hit the limits of computation. Nothing should be off-limits, but we need to proceed with caution because there some big red-and-yellow flashing lights around the use of data, intellectual property, ethics, hallucinations, and so on.

Lindsay: Focusing on the intellectual property issues for a moment, but what's the biggest flash red light there?

Bernstein: Well, it's tricky, because the legal theory behind the ownership of architectural intellectual property never anticipated technologies capable of ingesting huge amounts of that data and then drawing their own conclusions. We were experimenting yesterday with taking a rough hand sketch and running it through a generative environment we built in Comfy. "Make this look like a Louis Kahn building," we said. "Make this look like a Norman Foster. Make this look like a Zaha Hadid."

Under the law, as long as I can manipulate those images to be unique and not direct copies, there's nothing stopping me from doing that. Now expand that principle beyond formal strategies of image or composition — let's say HOK is really, really good at designing orthopedic operating rooms. Will all of that be generally available for ingestion by these algorithms? Right now, these tech companies are operating in an ethics-free zone where there's no legal principle stopping them from ingesting everything they can find on the Internet.

Lindsay: What should the industry do about that? Lock it down and use their IP crown jewels to train their models, either alone or in tandem? Will laws have to be passed to deal with this?

Bernstein: Look, I don't teach in the Law School. Well, actually I do teach in the Law School, but I'm not a law professor and don't have an informed opinion. The trajectory of legally protecting architectural ideas has evolved away from pure copyright toward protecting unique ideas. Some of that approach will be extrapolated into the use of this kind of data, but if the technology evolves in a way people can build their own training sets, then HOK can keep all of their orthopedic ORs off the Internet. The problem right now, of course, is that — at least at this moment — you can take every operating room HOK has designed since the beginning of time, and it's not nearly enough data to train anything.

Lindsay: Given such bottlenecks, how should firms incorporate such tools into their current workflows?

Bernstein: Because the architectural profession moves at glacial speed when it comes to adopting new technologies, I suspect the initial influence on the profession will be external, not internal, because firms don't have time, or money, or resources to work this out on their own. The fact that a group of very large firms you work for have banded together to figure this out because none of them have the resources to do it on their own is telling. If we were talking about any other industry, they would be competing against one another to find a market advantage. Instead, they're working on a certification label.

It's too early for anything other than wild speculation, but what I think will happen is that forces far larger from outside the profession — real estate development, manufacturing, supply chain management, construction management, finance — will likely create demands on architecture we'll have to respond to.

Lindsay: Which of those industries is likely to pose threats to the industry?

Bernstein: You used the word "threat." I didn't use "threat."

Lindsay: I did. Threatening incumbents in this case. From where is disruption likely to come from?

Bernstein: Let's speculate on both a positive and negative note. Positively speaking, those industries might create AI environments sufficiently data-rich that the insights they generate become accessible to the design profession. So, someone builds a fabulous set of generative AIs that make it much easier to generate a curtain wall design. Instead of keeping it to themselves, they make it widely available as a gateway drug to their platform. You visit their Website, give them a parameterized building and set some constraints, and they'll generate some options you can look at. That's the high road.

The low road follows The Innovators Dilemma, in which someone builds a tool that can only generate curtain walls, and that's not anywhere near a full building, so the industry shouldn't worry about it. But they get really good at curtain walls before moving onto the floor structures supporting them, and while architects are using these tools to make pretty pictures,

innovations happening outside the industry build up enough capability to supplant them. Terrible unresolved buildings are the result.

Lindsay: So where should AEC firms be focusing their energies and attention right now? On which piece of the tech stack?

Bernstein: I'll give the same advice I gave to firms during the early days of BIM, which is: We don't know where this is going. Firms with resources need to have someone on their staffs watching and building platforms for experimentation. I was just having lunch with a young architect — one of my former students — who's taken on this role for the 30-person firm he works for. He's watching. He's trying out things he thinks might be useful in practice, but it's way too early — certainly too early to be making intergalactic declarations about what's happening, because things are moving way too fast.

Lindsay: If Stable Diffusion is a tool the way a band saw is a tool, what are your classroom guidelines on when to use AI and when to use a band saw? How should firms — and educators — select their tools?

Bernstein: We can teach someone how to use a band saw because the band saw has been around for a hundred years. For this, it's too early to say. I'm going to give a presentation today where we're going to show three images from three different image generators with the exact same architectural prompt, straight out of the box. In these very early days, tweaking the carburetor or the fuel mix or the tire treads on these things, so to speak, creates very large changes in their output, so it's too early to be issuing detailed guidelines.

We're likely to tell our faculty here to follow the classroom guidelines the Cornell folks came up with, which is you have three options for generative tools. You can declare them off-limits and try to police it. You can tell students they can use them but have to carefully document it. Or you can make it a central theme of your teaching — "Take this prompt, put it in, and take the resulting image and show me how you manipulated it to get the final result" – which I think is clear and makes a lot of sense. It's a good way to think about how to use it within practices as well — ignore it, play around with it, or pick some part of your process you think would really benefit and apply it.

The last thought I want to leave you with is that these generative tools make increasingly convincing pictures of buildings. They do not make buildings — and they're nowhere close. They're barely two-dimension projections of complex, three-dimensional phenomena. In this classroom project I've been working on for the last two weeks, we've probably generated 300 or 400 images starting from the same sketch, and the image generators make up a lot of shit that makes no sense, right? They invent background buildings, don't understand roof lines, know nothing about how buildings meet the ground... The big open question is: what do machines for making evocative images mean in the long term?

Shelby Doyle

Shelby Elizabeth Doyle, AIA is an Associate Professor of Architecture at the Iowa State University College of Design and co-founder of the ISU Computation & Construction Lab (CCL) and director of the ISU Architectural Robotics Lab (ARL). The CCL and ARL the result of Doyle's ISU Presidential Impact Hire to rethink digital fabrication and design-build. The CCL works to connect developments in computation to the challenges of construction: through teaching, research, and outreach. She is the 2023 and 2024 President of the Association for Computer Aided Design in Architecture.

Greg Lindsay: The other interviews in this series have tended to focus on architecture and engineering. But how will generative AI impact construction and digital fabrication?

Shelby Doyle: I'm part of a research group at the Canadian Center for Architecture (CCA) called *The Digital Now: Architecture and Intersectionality* and my research looks at the history of computation and feminism in architecture. Specifically, I'm looking into the CCA's archive to see if there are stories that haven't been told, because there seems to be very few women

represented in the computational archive. And what I've found is a lot of absences — people and stories who weren't archived.

We can talk about archival politics and why this happens and what it means to formalize knowledge, but the reason this is connected to generative AI is I think that AI is really an archive project — it only exists because of archival records. And what I've found is that if you don't have those records to draw from, then the future you're pretending to create is equally, if not more problematic than the present, because you have a record missing any sort of knowledge that can't be translated into an understandable machine language. And there are also people, and histories, and forms of knowledge that don't want to be archived, because it's also a way of being found.

In the middle of this research project, Midjourney beta became available, and I started to ask Midjourney things like, "What does it look like for there to be an architecture involving computational feminism?" And every result was pink because the archive of data Midjourney includes millions of pink images tagged as "feminist" — an archive I'm now contributing to as I produce more images. That's a silly example, but it exemplifies the problems of using these tools to create a future that's based on an imperfect representation of the past.

Another project I'm working on involves 3D-printing concrete, and I've noticed there are a lot of folks working in that realm who come at it as a technology problem rather than a construction problem. Construction is local and relational, and about people knowing each other, and teaching forms of knowledge through practice — and many of the concrete 3D printing companies seem to be starting from the premise that any totalizing 3D model designed to solve construction will eliminate the human knowledge require to actually build things.

So, while I think the idea of mining every detail that's ever been drawn seems like it could be exciting in some ways, it also misses out on every form of construction that doesn't fall into a BIM model or a PDF, or is in some way machine readable. Anything that can't be taught through some sort of tradition or mentorship is lost because it doesn't count in the way it's archived.

Lindsay: What's your advice to AEC firms in light of these questions swirling around legibie- and illegible practices? Try to codify tacit knowledge as fast as they can? Forswear AI altogether as dangerously incomplete?

Doyle: I want to believe there are ways to use the technology well and create an ethical framework for it. But the whole point of these models is how they manage information at a scale beyond human capacity — and how quickly that gets away from us, exponentializing problems that already exist. Because now you can make a new image drawing from an incomplete dataset that becomes a form of fact, and then gets recursively worse. That's not just a semantic problem — it's a core issue of how the technology functions.

We have to be very careful we don't just magnify the biases we already know are the darker angels, if you will, of society and as architects and as firms. If you're looking at a very well-scoped data set — only elementary schools your firm developed from 1970 to 2010, for example — and using it for a very specific outcome, I think it can be very powerful. But if you don't have control over the data inputs — which none of us do, if you're using Midjourney for example — I think it's problematic.

Lindsay: In that case, is the industry's only solution to AI's "original sin" building tools atop their own data and models?

Doyle: That still misses out on the larger project of collective knowledge. We're really talking about rethinking everything from copyright to ownership models to what it even means to own an idea — and I don't know the industry's actually ready for that. If architecture as a field got together and declared, "We're going to create a shared dataset including every single drawing we've ever made," then perhaps... but I'm not sure we have the stomach for that. It's tough, because data can be so many things. It can be wall details, it can be renderings, it can be large language models, it can be how to assemble a brick wall — it's not like data is this one object.

During my time in practice here in the United States, the backsliding of social and legal protections for various types of people has been real and fast. If

you're collecting data on someone's gender identity, or sexuality, or race at a moment in time where we as a society have decide those are protected classes... well, in the next few years those laws could change. My real concerns about these models are the people tied to that data and what could happen to them if it's used for anything beyond its initial intent. What kind of protections can anyone promise that once data is ingested into a large model and it can't escape that model?

What if the knowledge you're trying to gather from someone who has done amazing construction work is technically not considered a legal resident of the United States? All of that knowledge will be lost if the only form we're willing to mine from are the machine legible models. But at the same time, you're asking people to trust you to take care of them if they're willing to share tacit knowledge in a way that it can be understood through these mechanisms.

Lindsay: Totally agreed, but again: given all this problematizing, what should practitioners and firms do? Roll their own? Hack it off at the root? Find some way to reconcile ultra-legible LLMs with illegible forms of knowledge? Or...?

Doyle: This is the entire crisis of higher education. AI is a humanities project, and I think any technology is doomed without that realization. Firms must find a way to treat that knowledge as valuable, for if architecture chooses to abandon its humanities project, then yes, this is a tool for evil. There's no way it can't be if it chooses not to embrace the humanities as something equally, if not far more valuable, because it's constantly challenging the technology. One of the problems is that there's no time for that in many firms. Things are moving too fast, and that ethos of "move fast and break things" is an ethos that has landed us in some bad places culturally.

Lindsay: How are you having this conversation with your students, given they are very likely using these tools already and haven't had enough experience to instill a similar perspective?

Doyle: We don't have an outright ban at Iowa State, and I haven't banned it, either. I try to ask questions in class that are hard for an image generator to answer — specific drawings and ways of working versus the random stuff you get from Midjourney. And then for the large language models, I ask them to show me their prompts and how they verified the output. We talk a lot about what it even means to verify knowledge, especially in the present day.

Again, it goes back to the questions of "should we?" and "how should we do this well?" But I appreciate firms asking these questions, because we've been living in this dream that technology would making architectural labor practices better. BIM was supposed to streamline workflows and people would work less egregious hours, and that hasn't played out. So, if you're going to adopt an LLM or whatever: how is that improving labor practices at the firm? Does it mean you're still working 80 hours a week doing different things?

Lindsay: My last question is: where would you start? How do we begin to leach the toxins out of these models?

Doyle: What would it mean to produce a counter dataset? For example, there are millions of images of hyper-sexualized women tagged as "women" online. I don't even know where we could start to produce images as a "counter dataset", or how many people would have to label other images in order to produce that counternarrative? That's what I mean about the problem of recursiveness — combatting the dominant narrative is probably something too hard and too big for most firms to do.

But my advice to them would be: if you're going to use imagery, it needs to be your own, and you need to have control over the dataset. Don't pull from these garbage piles of information, because that's what they are — rotting banana peels of data.

Tim Fu

Tim Fu is a renowned architectural designer specialised in advanced computation and artificial intelligence (AI). Emerging from *Zaha Hadid Architects*, he has founded *Studio Tim Fu*, a high-tech design practice pioneering the integration of AI into visionary design. As an active educator, he has run workshops at Harvard GSD, PA Academy, and lectured in various universities and conferences globally. Leveraging digital platforms, Tim has also built a notable online presence, sharing insights into the overlap of technology and design. His AI explorations have been featured by media worldwide and has been exhibited during the Venice Biennale.

Greg Lindsay: Would you please begin by telling me a bit about yourself, your background, and how you use generative AI in your practice?

Tim Fu: I'm a Canadian architectural designer who's worked at several offices in the UK, including Zaha Hadid Architects (ZHA) for more than two years. My specialization is using parametric algorithmic design and software for advanced computational design. I recently resigned to found my own architectural studio, partly due to the emergence of generative AI and algorithms. My work from ZHA to now has been more focused on the research side, using mostly-diffusion AI software to aid with the design and conceptual process of architecture, while using other tools to streamline ways of creating workflows from concept to construction. Everything is up for grabs now. Tech is moving very fast, we are all moving very fast as well, and my new office is fresh.

Lindsay: How does lookX differ from more general purpose generative AI tools such as Midjourney or Stable Diffusion?

Fu: If we're going to discuss lookX, we have to talk about Midjourney and other models, and why certain aspects of lookX are superior. First, it's trained specifically on architectural data sets, so the results are much more tuned — whether you're looking for an industrial quality rendering or photographic realism in a cityscape. It also has certain functionalities, such as its ability to use contours to generate new "hallucinations," if you will.

So, it's a bit like Stable Diffusion's ControlNet, which also takes an input image, keeps the contours of what you're looking at, and then generates new content. That's quite useful if you already have a model, a 3-D massing, or a hand-sketched image because the result will follow the positions of your original concepts. The accuracy of that visual placement is important, and it's what Midjourney is lacking because inputs are treated more like a suggestion — "here's the style I want you to replicate," and it's never a one-to-one match. You're always rolling a lot of dice in that sense.

We use lookX already. I find it especially useful with rendering massings — if you use computational software to calculate area, occupancy, and so on, you want your output and renders to follow as precisely as possible. That's what's aiding us right now in the process of early concept design — it expedites the process entirely.

What's lacking is the ability to go from a prompt directly to 3-D. There are some very rough models and point cloud generators, but nothing is viable enough for a practice at the moment. I think that will improve as we go forward. There's also the potential for depth mapping with 2-D to 3-D images by using a second AI to figure what a three-dimensional rendering of a 2-D image would look like. Again, there are rough variations, but it's not all there. Maybe a large language model and other tools can start tapping into this design flow and contribute to that side of the process.

Lindsay: What's the major technological impediment from crossing the chasm from 2-D to 3-D with point cloud generation?

Fu: I mean, it's a point cloud — it's not a beautiful render you can sell to clients. It also jogs a lot of creativity when you see images. Right now, Midjourney is quite useful in the creative aspects of concept designs. Point clouds pertain to mesh. You end up with mesh approximation operations and mesh geometry, which is probably fine for furniture design, but in architecture it would be chaos to rationalize for construction. Contrarily, mathematically-pure geometric repre-

sentations such as NURBS or SubD models are more flexible for geometric rationalization later down the pipeline, for processes such as algorithmic panelization and optimization of modularized units.

Lindsay: Now that you're starting your own studio from scratch, how do you intend to create a workflow optimized for AI? And how does that differ from your previous role at ZHA?

Fu: From what I've seen of other offices — including my own and a few friends' —right now it's all about adapting existing frameworks around these technologies and exploring where they can retrofit into the process. What I want to do is propose a fundamental shift in the whole process to make it as generative and as streamlined as possible.

I've worked with a lot of parametric algorithms in my past, and in the reality of the practice things often slow down because team members prefer to model it manually and solve localized issues. I think there are incentives for us to stop compartmentalizing these processes — solving them one-by-one — and start streamlining.

Typically, you might have the concept team, the massing study, the landscape team, the geometric rationalization team, and then you send it to the BIM team. Then there's facades, detailing, and all sort of other specialists expected of how current offices function. Our task is to tie these processes together so you're not pumping in data to one group and then have to pump it out. Instead of taking a bunch of base models and sending them to another team, I'd prefer to provide an algorithm that follows the base rules so what's inputted is automatically outputted, and my job would be ensuring the algorithms are functioning so the data pipeline is a continuous flow of information.

This is only possible recently because of the advent of diffusion AI and LLMs, which we can designate a lot of these tasks to. That's why I want to tie it altogether, putting machines where machines can be put to try and find a singular, streamlined workflow as much as possible.

Lindsay: How would you describe your relationship with lookX beyond simply being a user of the tool?

Fu: lookX is a Chinese company whose domestic brand is "Xkool," which is the name in reverse. They've been thriving for a couple years already in China, and they're one of the top dogs when it comes to this kind of AI generation and tools because their platform has a lot of crowdsourcing advantages. You have a database and a server, along with a forum of different users training their own AI models and sharing them. So, you have your own trained models you can share, and then you can attached those semi-independent models to main one you're running. We're working together to test the new lookX with my concepts and designs, using them to train data and some models and seeing how well they work.

Lindsay: How does lookX treat training data, user-created models, and other inputs from a privacy- and intellectual property perspective?

Fu: The small-scale models we train to attach to the main one are called LoRA — the low-ranked adaptation of large language models. Those models use clusters of even smaller models and the data image you can train on your own. The lookX Website hosts whatever the users train their own data on — you can't tell what the original images are trained from, but it's usually about 30 or 40 of them.

ArchiNet is lookX's proprietary architecture database used to train their main models. I don't know much more than that.

Lindsay: How do these issues factor into your decisions about which tools to use?

Fu: It matters a lot to me, albeit in a very different way. Most people are adhering to a certain trajectory of lawmaking to create a new environment in which everything has a digital footprint and a system is in place to monitor and accredit the source of each image. I don't think training a model is equal to copyrighting or stealing or plagiarizing it. I feel like I'm one of the few who believes we shouldn't be able to even opt out of training other people's computers because they're learning.

For example, if you're an artist and a student wanted to learn from you, they would study your work, pick up some patterns, and apply them to their own art. I

don't think you own any asset of that student's work just because you inspired them and they learned from you. The only difference here is that the student is a machine and this machine is learning a lot faster than a human. It's only because of that speed we forget that it's learning — which is fundamentally different from plagiarizing or taking patches of data and replicating it. We've never copyrighted learning in the past.

God forbid Zaha Hadid Architects one day copyrights its style through AI, because it's ridiculous — U.S. and international laws all state you cannot copyright "style," which is the only thing that's been replicated in these cases.

Lindsay: Speaking of which, will each firm create their own proprietary suite of sub-models based on open source foundation models, or what is your intermediate-term view of how these tools will evolve?

Fu: I'm sure everyone will learn to collate their own images and data and then produce their own models. But more than that, we've got to stop thinking with this capitalist, commodification mentality. The next era will be one of abundance where whatever you proprietize won't be as valuable because everyone will be able to do it as well. What's going to be of value are the people who know how to curate the results, be it images, music, or architecture. The Pandora's Box has been opened, and it will be too much work to try to regulate it. It would be rather dystopian, if you really think about it, to be forced to have a digital footprint for every image or piece of content ever generated. We should just let information flow freely, leaving everything to be open source and shared. Right now we're in an early phase of people recoiling, then patching together quick new laws when really we need a revolutionary change in lawmaking for this new paradigm.

We're in an early phase right now. Just organizing my office with a few collaborators and specialist consultants specializing in different fields, we're standing by as the technology develops, adapting what's available now and waiting for what will be available in the future. We're just trying to do our best to input human values into all these tools — because in the end, we are still the architects and designers, and these tools are here to express our intentions.

Nate Miller

Nate Miller is the founder and CEO of PROVING GROUND. Nate has developed trusted relationships with some of the most reputable organizations in the building industry to create strategies, workflows, and tools for enabling a data-driven business process. Nate strives to help his clients discover meaningful, ethical, and practical uses of new technology to support business growth the creation of better built environments.

Greg Lindsay: Let's start with your January blog post. Have the last eight months allayed or deepened your concerns about some of the questions flagged in your post?

Nate Miller: Some people have looked at that article and asked, "Shouldn't you be cheerleading technology rather than critiquing it?" Which is generally the tendency in a role like mine. I'm fundamentally optimistic about what technology can do — both for the world in general and especially in the building design and construction industry — but as I've grown more seasoned as a professional, I see myself less as a cheerleader than a critical implementer.

We saw a similar hype cycle last year around the "metaverse" and cryptocurrency, and it always felt suspicious to me — what problem is this trying to solve? — and if you can't answer that, then is it actually a solution looking for a problem? Maybe it was the similar level of hype combined with the addictiveness of prompt that gave some pause and caused me to delve into the terms of service. I wanted to look at the implications of what was running in the background, and suddenly I became far more skeptical of all this. Over the last eight months, my skepticism has grown quite a bit, and I think there's even more important issues to talk about than what I raised in the original post.

Lindsay: Well, then let's start there!

Miller: Let's start with a very quick recap of their terms of service first. They're the underlying basis for what the service provider can do with your information

— and what you can do with the output as well. There were a lot of concerning things then, and that's evolved quite a bit.

One thing is a growing legal precedent that users of generative AI do not retain copyright of the output. The image has to undergo significant post-processing by the human hand to be considered original, and the U.S. Copyright Office has issued guidance on that front to the effect of, "Yeah, this stuff is not yours." And that has massive implications for creative organizations.

I also think the terms of service has implications when it comes to placing liability on users — not only do you not own the stuff, but if your prompt generates something in violation of copyright due to the underlying training data, that may be on you as well. I'm not a lawyer, but as a user of technology, it's a potential risk I would be interested in.

There are also the ethical implications — how was the training data sourced, and did it involve using content without the original authors' permission? If you sample anything in the music industry, for example, you must provide credit — there's a specific trail of inherited attribution that's not quite present here. But that's starting to change with AIs generating valid references, which is a great step in the right direction, but when you dig deeper, more questions emerge — you can scrape the Internet for images, but this stuff still needs to be structured and set up in vector tables, and who's doing that and how?

So, earlier this year, around the same time I published my post, *Time* reported OpenAI was paying Kenyan workers just $2 an hour to train their models. In related reports, many of these workers reported mental health and trauma as a result from reviewing graphic images of sex and violence as part of flagging and labeling content. That doesn't seem right, either.

And then there's the environmental aspect. *The Wall Street Journal* just published an article titled, "Artificial Intelligence Can Make Companies Greener, but It Also Guzzles Energy." I got a good chuckle from this quote by Sasha Luccioni, a research scientist at Hugging Face: "The cost of all these shiny new toys is unsustainable if you're just switching out all these technologies that are working quite well to begin with, with these much more energy-intensive applications."

She indicates a practical step in limiting emissions in the context of AI would be to not integrated it into platforms that don't need it. And if you look at any new technology — especially when the hype is at a fever pitch — the tendency is to think, "I've got to put AI in this!" But it's not necessarily needed if there are other, more mature technologies that are better at solving the problem you're trying to solve with AI.

Lindsay: So, to recap: there are the copyright qualms, the ethical qualms, and the sustainability qualms. How should practitioners approach the three? Is generative like plutonium — dangerous, dirty, and useful only in narrow applications? Should firms build their own models to ensure they aren't tainted? Or should they only use 100% organic, officially-licensed models from Adobe, etc.?

Miller: My first piece of guidance is that there need to be more attention placed on these issues. One thing we talk to our clients about is how we don't want to get caught up in the hype of something new. We want to discover where you can meaningfully introduce technology that will benefit your organization and the clients you serve.

If we take that position, there may very well be aspects in any given AEC company that can greatly benefit from AI, but it needs to be managed and balanced against the question "What is this actually giving us?" versus a blind grab at technology simply because it's popular.

My other piece of advice is to ask how this tech can situate itself in a creative process. Very rarely have I ever been involved in a design process in which sending a prompt in, receiving an unexpected result back, and then using that is considered standard practice. Usually, when I come to a problem, I have a set of constraints I'm working within, and I want to balance those constraints to achieve some kind of result. In that context, you might think about AI — and machine learning as a subset of that — as sourcing data that's meaningful to the problem I want to solve; training algorithms and selecting ones best suited to the prob-

lems I'm trying to solve and might be reusable down the road, and how best to implement those solutions as firm grows or my problem set evolves.

That idea speaks to being the orchestrator of a data pipeline, with tools like AI and machine learning as part of the overall toolkit. In that context, you're being selective, deliberate, and purposeful with the tools you're putting forward, with the data you're sourcing, and the outputs you're given. That last piece is critical, too, because there's an aspect to any data-driven process where the output should match your intuition — if there's a big mismatch in expectations, maybe you should do a deep dive into why.

All of this is far more appealing to me in a professional design context than telling a prompt engine, "Give me my copy." Which might be useful for many, but for architects... not so much.

Lindsay: Who's in charge of building that data pipeline and defining guidance across firms? Is that the current chief information/innovation officer or some new role? And what does that investment roadmap look like?

Miller: I don't know. Maybe it won't necessarily be in the control of architects or professionals. Maybe you'll use AI whether you want to or not, because it's part of the Adobe program suite and the selection-and-fill function natively uses it to erase people and extend the ocean. There's sort of an inevitability to it. But there's also an aspect where we as professionals should critically guide this conversation to ensure we're not granting over rights to the products we're using. We want to be informed consumers of the tools we use, more than anything else.

Lindsay: Is there an alternative to make your own tools, however?

Miller: I've always been an advocate of architects making their own tools and composing systems that tie into their design process. That's a valid trajectory for this, and there are already tools out there that allow for it. Our freely available LunchBoxML sits right inside of Grasshopper, and we're seeing all sorts of really amazing stuff come out of that. They can also claim that AI/ML is being used in the design process, and they control the pipeline, which is a pretty appealing message. So, yes — make your own tools and find ways to tie your process together with them.

I think the worst thing you can do as a professional right now is sit on the sidelines and pretend nothing is happening. And unfortunately, I see that in some of the firms we've encountered. I think they have an opportunity to be very proactive in this conversation, if they'll just look beyond the popular media into the tech and not be afraid to try it out and see if they can glean value from it.

Rob Otani

Rob Otani is widely regarded as one of the industry's leading experts in automation and computational modeling. He established and oversees our CORE studio, which focuses on applications development, advanced computational modeling and research and development. Rob is also part of the leadership team that oversees our firm-wide research and development program, CORE Lab, and serves as an advisor to Thornton Tomasetti's innovation accelerator, TTWiiN. During his career, Rob has been responsible for the design and management of more than $2 billion of infrastructure, commercial, cultural and residential structures and infrastructure projects. He is highly proficient in seismic, wind, dynamic and thermal analysis, and has extensive experience detailing architecturally exposed structural steel and concrete structures. He manages and engineers many of the firm's special structures projects, including those involving innovative materials and delivery methods.

Greg Lindsay: What excites or interests you about generative AI from the perspective of an engineering firm, since I don't imagine you're using it to create images of buildings?

Rob Otani: What excites me about it is how it accelerates solving problems we wouldn't solve using manual methods. For instance, we've been pointing large language models at PDF documentation. We have a *lot* of PDFs, a lot of reports — and there's a lot of intelligence in them, from across the firm. And it's

effectively able to "read" those documents and answer questions within a minute or two. Which means that people in one office who wouldn't know where to start in asking about some topic within that office — much less someone in another office, because we have 50 or so — wouldn't even try. From a knowledge-sharing standpoint, it's extremely powerful. We're currently testing various methods and data sources to see how accurate it is, because it does hallucinate sometimes — that's an issue we'll have to address. We have to weigh the value versus the risk.

Lindsay: And have you already incorporated this into the firm's regular workflow?

Otani: Yeah. Our marketing department has thousands of what we call "project sheets" — one- or two-pagers with a short project description listing the design team, the architect, the contractor, the cost of construction, and so on. Our small marketing team used to ask people what they thought, based on their experience of past submissions, who would be the best expert for a particular project type. And then they would ask that person about which projects to include in the next submission as representative projects of that particular typology.

That was a very manual and slow process. Now we point a large language model at those thousands of project sheets, and it comes back with an answer. That's really powerful, and they're using it in production now. It's not perfect — it does make mistakes here and there. We're figuring out ways to QA/QC the answers. It's just like engineering — people make mistakes, you find them.

Lindsay: Where else do you feel confident enough to put LLMs into production?

Otanti: We have an intranet Website called Spark where Thornton Tomasetti employees can ask questions or just post fun stuff. It's a knowledge sharing and resource management platform. But it has a terrible search engine, and because of that, people are reluctant to sift through years and years of conversations and reference material. Again, we extracted all of that data and pointed large language models at it. Now you can get a clear, concise, time-based response from all of that information since the beginning of time — or at least since we've had the platform — in seconds. And it's Thornton Tomasetti information. It's not a Google search or OpenAI. It's our data.

Lindsay: Speaking of your data, how do you evaluate which LLMs or other tools to use in light of IP issues, including copyright and client data?

Otani: There's a risk of surfacing confidential information associated with some posts without explicitly stating which part is confidential. And some of the posts are just jokes — someone put an April Fool's up there, and the AI came back with a response. In many ways, data science is dumb — it's only as smart as the information you provide, which in that case was complete nonsense. And that's concerning, too — people using information without vetting it.

Lindsay: So, how have you gone about instituting fact-checking and quality control? Is there a formal, centralized group handling that?

Otani: We have an AI guiding principles policy document within the firm, which is going to change over time, obviously. But we thought it was worthwhile to alert people you have to be careful — that LLMs aren't necessarily going to be 100% right, and if there's any question about it, here's where you can find reference documents associated with that response. That's not perfect because it defeats the entire purpose — you don't want to go through 20, 30, 40 documents every time you receive a response.

Lindsay: In which case, is it left to individual users to verify? Is there any oversight? Who's in charge of vetting their use of AI?

Otani: Good question. We haven't figured it out. We're still researching how good large language models are and how might potentially do some level of post-processing to back-check a response. Or potentially build our own. It's worth looking into, let's just say that.

Lindsay: What would that entail, exactly?

Otani: Building a platform employing various logic models, pointing to various databases and sources, and having a UI within Spark very much like Bing's AI search engine. There are some downsides to that. One,

it's going to take a pretty long time. Two, we'd have to maintain the thing. Three, the technology changes all the time, and there are opportunity costs to consider — including people coming out of the woodwork with AI solutions. So, we'll start to experiment with our own tools and vet off the shelf solutions. For now, it's a low-cost effort of just doing it locally, on someone's computer, that we can test for other uses cases.

Lindsay: Beyond LLMs, are you using any generative AI for rasterized images?

Otani: We've experimented with that — it was one of the first things we did. The problem with rasterized documentation is that there's not much information in there. Okay, there's a wall and a there's a slab. At some point, we can extract vectors, we can extract text. Old, scanned blueprints are terrible quality, with a lot of noise in there. Clean PDFs exported directly from AutoCAD or Revit are pretty good, but there's still limited information in those documents. It's probably too much effort to try to go all the way with that, because there's going to be a lot of holes. And someone's going to have to fill those holes with manual documentation.

Although I will say one thing it does a very good job of is finding topics. We do a lot of litigation work, and one attorney set us thousands of image files of documents. Literally thousands. And not even PDFs! There are really good AI models that can extract text from that. We used it for querying gigabytes of information to pull out topics and related metadata — people, timelines, and things like that. It's not perfect, but if you're looking for a particular topic, you can start to mine a little deeper.

Lindsay: Are there any tools or any use cases for which you've ruled out AI completely? What are the red lines you won't cross?

Otani: There's nothing we've ruled out. As an engineering firm, we don't do much image generation. Whatever is created manually through a series of processes in architectural design is a result of many other factors that are hard to automate, or even glean, from an image. Unless someone can capture all those data points, images don't really mean much. It's hard to align that with the program requirements of a real project.

Lindsay: Returning to the idea of building your platform, what would have to change to make the ROI of that worthwhile? What would it take to make you reconsider?

Otani: It's going to be the cost of off-the-shelf software and the challenges of organizing certain data streams those platforms can access. It's just like any other software. It sounds great, but once you start actually using it, it's potentially a different thing.

The hardest thing is pointing these platforms to that data in a safe way. For example, we use something called Deltek that's a very common AEC platform that's also super old. It's gotten better, but we're going to move in a different direction. Regardless of where we go, there's a treasure trove of data in there that's hard to get out in a clean way. There will need to be teams to clean that data so we can meaningfully extract information from it. And we have the workings of a such a team, which for our industry is relatively new. But data-for-data's-sake wasn't very useful until we had these large language models — because everyone wants to visualize or access that data in a different way.

Lindsay: Different how?

Otani: I had a conversation the other day with some senior folks in our office about taking someone's emails and building a large language model out of them. It's tough — there's a lot of cleaning up to do before we can ask it questions. The person who volunteered is a steel expert with a 0.1% sort of expertise in welding and fabrication. He's probably 65 years old by now and he's stored a lot of his emails — just exchanges back and forth. We're going to ask it questions and see what comes back. Someone's going to have to vet that the answers are 100% right or only 50% right to see if this is really scalable, but he's the perfect test case. He's an irreproducible expert.

Praveen Patel

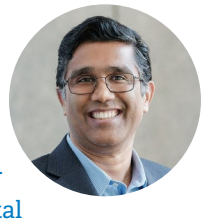

As senior vice president of technology, Praveen Patel leads The American Institute of Architects' (AIA) digital technology strategy, aligning it with the organization's business objectives while spearheading operational efficiency and fostering innovation. He directs the planning and implementation of enterprise systems in support of AIA's business operations to improve cost effectiveness, service quality, and business outcomes. He is responsible for all aspects of AIA's digital technology and systems. Prior to joining AIA, Praveen spent 18 years in technology consulting, enterprise architecture, and software consulting leadership, serving clients in professional and trade associations (nonprofit), the commercial sector, and state and federal government. Praveen holds a Master of Science in information systems and technology from Johns Hopkins University and a Bachelor of Science degree with a double major in chemistry and systems management.

Greg Lindsay: How does the AIA approach the growing flood of generative AI tools and how are you helping members understand and selectively deploy them?

Praveen Patel: What AIA has done as an organization is work with staff to identify opportunities where we can leverage artificial intelligence overall. We have "blue sky" sessions where we work with our staff to bring these innovative and creative ideas to the table, and then go through a brainstorming exercise to explore and prioritize few ideas for pilots. Executing on these new ideas are the focus in the short term and some of these ideas will fuel the long-term growth of the organization.

Lindsay: Speaking of focus, my other interview subjects raised issues ranging from IP to ethics to GAI's environmental footprint, etc. How does the AIA rank these issues and others in order of importance, and how are you advising members?

Patel: Currently, we haven't provided any guidance on this to our members, but what we have done internally is build up an AI usage policy and are working next on building our AI Strategy. While we want to encourage our staff to use AI, we want to put some guard rails around it to ensure everyone understands how best to deploy and use it responsibly. As part of that policy, we're providing general guidelines to staff on using AI tools while safeguarding organizational data and adhering to data privacy standards.

We're also promoting the use of some approved tools. When this buzz with generative AI started, our staff started looking into various tools and exploring them. Now, we're trying to put some governance in place towards which ones we want them to use, and the best practices using those tools.

A simple example is Grammarly. I think that's a tool that's done pretty well for many things tied to content creation and editing. That's something we don't see as a concern and will promote that for staff productivity. But there are other tools we're still assessing for staff use, — especially from a security and privacy standpoint — and given how many of these tools are duplicative – addressing the same business needs.

Lindsay: How does that evaluation process work, and what would be your advice for IDC firms in assessing the strengths and weaknesses of each? What's your diagnostic criteria, and which tools would you recommend for which tasks?

Patel: We assess tools and technologies based on gaps in our portfolio. There's a component of innovation, and assessment of how this tool will help us do our jobs better, like was the case with Grammarly. From there, we do an assessment of these tools from a security and privacy lens to protect organizational data and IP. I think that's hard nowadays, because most of these tools are changing their policies too fast — even Zoom, which had a recent change in its privacy policy that created immediate concern for us. [Editor's note: All of these interviews were conducted via Zoom.] Some questions we have to constantly ask ourselves - How does these changes impact us? What does it mean? Do we step back from using or continue leveraging it just like everybody else?

If the tool gets a clearance on security and privacy, then we enable the tool to be used to fill in the gap we have in our technology stack. Our goal while looking into new tools: Is it addressing something unique business need or solving a unique problem, or is it a productivity

tool? From a content creation standpoint, there are quite a number of tools that get you from 0% to 70-80% very, very fast. You don't have to start anything from zero anymore but you need to know how to use the tools to be able to get the best results. It's important to understand that it takes quite some effort to get to 100% completion and that's where our content creators play a huge role. Getting a structure in place using these AI tools is pretty straightforward and they do a magnificent job. But taking it from a template to a finished product is where we need experts to weigh in.

Lindsay: Beyond security and privacy, how do ethical and environmental issue factor into your evaluation? Given reports about underpaid and traumatized model trainers, coupled with the vast resource footprints associated with training models, how do these fit within the AIA's principles and guidelines?

Patel: That's a great question, and we're still trying to figure it out. The focus shouldn't be on using these tools to bypass any standards and codes that define our practice. Again, from a professional standpoint, leverage the tools to improve your productivity but not to do your job, because that's where you need to bring your expertise to the table to see whether the approach suggested is the right one for you as an organization or as a team. So, in that spirit, leverage the tools, but make sure you don't bypass any of the standard guidelines and processes we should all be following.

Lindsay: That's a good segue to discuss the labor implications. Given the focus on enhancing productivity versus using GAI to do the job, will the AIA issue any official guidelines about hiring and employment practices around the use of AI?

Patel: We haven't discussed any particular guidelines, but we have a subgroup currently focused on the AI within the design and architecture space. That team is mobilized to come up with many use cases that will help our profession going forward, and some of the use cases we're hearing are very promising.

We know these AI tools are at their best in areas where we have a lot of data and are able to make decisions faster than a human can. Those use cases are the ones we're currently exploring and developing pilot projects for.

For example, any task with a decision tree matrix can be easily leveraged with a private LLM. If there's a definition of something you can provide, and if your decision is based on that definition, does it comply or not comply? For those kinds of use cases, AI can do a relatively good job with a very high confidence factor, and you can continue to train the models to reach +90%. And for all the unique and edge cases, you will need a subject matter expert to weigh in and make decision.

We're trying to figure out these opportunities for automation where these AI tools can help supplement the job, and make it easier for us to deploy our forces elsewhere.

Lindsay: To your point earlier about many of the current tools being duplicative to the extent they focus on general purpose text- or image generation, what pieces of the AEC tech stack are the next to be affected? What do you have your eye on at the moment?

Patel: I attended a session recently where the presenter demonstrated how someone could design a small space by leveraging the design tools to arrive at and iterate through variations on design specifications quickly. It was interesting! And we see a lot more energy modeling tools for predicting which designs will perform optimally under various parameters and constraints.

Lindsay: Speaking as a representative of the AIA, what's your high-level message to the AEC firms and practitioners on whether and how they should be using GAI?

Patel: My message to the AEC firms and practitioners is to embrace this technology as a transformative tool that can significantly enhance creativity, productivity, and innovation in the industry.. We should be looking at Gen AI as a complement to human expertise, and not a replacement.