

AI Architecture Studio Without Architects

Immanuel Koh

Then: Paperless Studios

The “paperless” studios of the mid-1990s stood out as a revolutionary experiment in the annals of architecture pedagogy, and their impact has remained exemplary in both academia and practice today. Initiated by Bernard Tschumi during his deanship at Columbia University’s Graduate School of Architecture, Planning and Preservation (GSAPP), the studios were taught by a team of then-young faculty who are now world-renowned architects, such as Greg Lynn, Hani Rashid, Jesse Reiser, Stan Allen, Alejandro Zaera-Polo, and Ben van Berkel. In fact, Lynn curated the 2013 exhibition “Archaeology of the Digital,”¹ held at Montréal Canadian Center for Architecture (CCA) and organized a two-week summer seminar called “Toolkit for Today” with lectures by Tschumi, Allen, and Rashid, retelling the history of the paperless studios. In Tschumi’s account, the story of the paperless studios was not simply the inclusion of computers in the studios, but more fundamentally, preceded by a deliberate shift in pedagogical vision, change in curriculum structure, sourcing of a large sum of funding (USD\$1.5 million), and renovation of existing studio spaces at the university to accommodate the newly purchased computing infrastructure. In Allen’s account, the interest in the computer was first and foremost intellectual, rather than practical, and was preceded by the theoretical transition from Derrida’s deconstruction to Deleuze’s difference which, in turn, was also a formal transition from the language of rupture to that of continuity. The paperless studio served predominantly to accelerate the pace of manifesting “the language and the operations ... theorized before the actual

1 Greg Lynn, *Archaeology of the Digital* (Sternberg Press, 2013).

use of the computer.”² It was not so much about the general use of computers for architectural designs, as was the case for most architecture schools that had introduced computers prior to the paperless studio at GSAPP such as at MIT or the AA, but very much about the specific use of computers for an a priori set of philosophically driven and materially visualized concepts to generate a new architectural language. For example, Lynn’s 1992 model of the Stranded Sears Tower was made of cut and twisted foam, expressing his philosophical concept of the supple with material suppleness prior to the use of computers.

In this light, it is thus not surprising that the studios were called neither “digital studios” nor “computer studios,” but “paperless studios”—a negation of what went before (i.e., “paper”) rather than a proposition of what was to come. Strictly speaking, “paperless” need not mean digital, nor even the computer. The “-less” in “paperless” should therefore be understood as a conceptual provocation rather than a literal pronouncement. In fact, Tschumi himself confessed that the use of paper in the studio projects did not end with the introduction of the computers but instead continued to support the digital exploration in parallel. In his text, “Building with Geometry, Drawing with Numbers,”³ when comparing the text–geometry approach of Vitruvius/Alberti before the age of printing and the drawings–measurements approach of Palladio/Vignola thereafter, Mario Carpo was in a way recasting the paperless studio as one characterized more by a shift in notation (e.g., spline modelling or animation keyframing). It is a shift in architectural language corroborated by new notational systems resulting in a new conception of architecture. Although Carpo has insisted that there will not be an AI-driven “third digital turn,”⁴ this might turn out to be premature. Increasingly, it seems inevitable that there could indeed be a “third digital turn” propelled by today’s rapid development and widespread infiltration of AI technology such as ChatGPT and Stable Diffusion. The notation is however different this time round, and its manifestation not necessarily straightforwardly oppositional as seen in the first two turns; namely, the “continuity” of the first digital turn and the “discreteness” of

2 Stan Allen, “The Paperless Studios in Context,” in *When is the Digital in Architecture?* ed. Andrew Goodhouse (Sternberg Press, 2017), 390.

3 Mario Carpo, “Building with Geometry, Drawing with Numbers”, in *When is the Digital in Architecture?* ed. Goodhouse, 35–44.

4 Mario Carpo, *Beyond Digital: Design and Automation at the End of Modernity* (MIT Press, 2023).

the second. In fact, for the very first time in the history of architecture, there may no longer be a need for any notation, whether it is notation as architectural drawings or programming codes. More concretely speaking, “natural language is all you need,”⁵ and text-prompting is now the dominant “AI notation.” A decade ago in his 2015 book, professor of computer science Pedro Domingos proposed the idea of a “master algorithm”—the ultimate machine-learning algorithm capable of solving all tasks.⁶ And, as early as 2017, computer scientist Andrej Karpathy (now director of AI at Tesla) wrote about a paradigm shift in software development with the emergence of deep neural networks from Software 1.0 to Software 2.0.⁷ The former refers to software traditionally written in programming languages that explicitly specify human-understandable instructions/algorithms to the computer. The latter instead refers to software written in the form of the weights of a neural network which learns an “algorithm” from its training dataset. As he puts it, “Software (1.0) is eating the world, and now AI (Software 2.0) is eating software.”⁸ AI is the master algorithm. In early 2024, Jensen Huang, CEO of NVIDIA, even went as far as to say “don’t learn to code” when asked what kids should learn in the age of generative AI. Might a “codeless studio” be indeed the “paperless studio” of the near future in architecture?

Now: Codeless Studios

There are a few important lessons to be learnt from the paperless studio of the 1990s as one begins to formulate a possible form of a codeless studio in the age of AI. First, the “-less” in “codeless” is to be understood in a similar manner whereby coding will not in effect end with the introduction of large language models (LLMs) like ChatGPT, but instead continue to support the AI exploration in parallel. Therefore, current computational design skills in visual programming (e.g., Grasshopper 3D) and text-based programming

5 Here, alluding to the title of the landmark AI research paper. Ashish Vaswani *et al.*, “Attention Is All You Need,” in *Advances in Neural Information Processing Systems*, vol. 30, ed. I. Guyon *et al.* (NIPS, 2017).

6 Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (Basic Books, 2015).

7 Andrej Karpathy, “Software 2.0,” *Medium* (blog), November 12, 2017, <https://karpathy.medium.com/software-2-0-a64152b37c35>.

8 Karpathy, “Software 2.0.”

(e.g., Python) will serve to support the AI exploration at the codeless studios, especially when integrating AI models with existing CAAD tools and creating specific AI models for the architecture discipline. Second, just like how setting up the paperless studio had incurred significant expenditure on software (e.g., Softimage at USD\$9,000) and hardware (e.g., Silicon Graphics IRIS Indigo Extreme R4400 machine at USD\$25,000),⁹ the codeless studio will need substantial funding for either paid subscriptions in accessing frontier closed-source AI models (e.g., OpenAI's GPT-4o or Midjourney) or purchasing local/cloud GPUs, computers, and data storage (e.g., A100 or AWS) for running open-source AI models (e.g., DeepSeek-R1 or Meta's Llama). Unlike the paperless studio, a codeless studio will also require datasets for training its AI models. Before the current paradigm of pretrained foundation AI models, it was still financially possible to train an AI model from scratch with one's own datasets and workstation. However, today's training cost for models like GPT-4 is in the tens, if not hundreds, of millions.¹⁰ Therefore, the training datasets for the codeless studio will mainly be used for the purposes of fine-tuning such pretrained foundation models or training smaller AI models, in order to adapt or customize the AI models for architecturally specific use cases. Third, it should be noted that the tutors (e.g., Lynn and Rashid) at the first paperless studios in 1994 were mostly non-experts in the technical operations of the computers and had digital assistants (DAs) who were computer-savvy students (e.g., Ed Keller) supporting them. This unusual pedagogical combination turned out to be synergetic in such an experimental studio context. However, unlike the predominantly platform-specific software application tooling the workflow of the paperless studios, AI workflows of the codeless studios can vary greatly in complexity and diversity, ranging from simply using paid online blackbox AI services to implementing one's own AI models. The latter will require a certain level of understanding and expertise in AI. AI models (especially agentic AI) of the codeless studio will also have greater design agency than the animation software of the paperless studio. It would therefore be more strategic to have faculty who are well-versed in AI, not only technically, but even more crucially, theoretically and conceptually, in

9 Endriana Audisho, "Screening Architecture: Architecture, Media, and Conflict since the 1990s" (PhD diss., University of Technology Sydney, 2024), <https://opus.lib.uts.edu.au/handle/10453/179447>.

10 Ben Cottier et al., "The Rising Costs of Training Frontier AI Models," *arXiv*, last modified February 7, 2025, <https://doi.org/10.48550/arXiv.2405.21015>.

order to steer studio projects into novel and unknown architectural territories. Instead of DAs, the faculty will ideally be supported by AI student assistants who might not necessarily be from the architecture department. Lastly, to broaden the conception of architecture, Tschumi wanted the paperless studio to borrow notations from other domains such as film and music and combine them with architecture's own axonometric mode. For the codeless studio, it will remain important to encourage relevant conceptual influences from domains outside of architecture. However, such "outside" borrowings will go beyond notations alone to include design briefs that are framed as more generally "architectural" (i.e., not just about making buildings), augmented by new AI processes. In short, coding skills in computational design should still be taught, funding for GPUs and AI-model access should be raised, teaching undertaken by faculty whose works conceptually resonate with the "yet-to-be-defined" language of AI, and AI-augmented concepts and processes should be borrowed from other domains outside architecture. In fact, as evidenced from the impactful career trajectories of the then-faculty and then-students at the paperless studio within academia and practice, the third and last lessons are the most crucial. The codeless studio is about putting in place a pedagogical platform that is theoretically critical, conceptually radical, and technically experimental. Only then could the germination of a new AI x Architecture culture emerge to instigate a third digital turn. In the next section of the text, we will look at how a similar "putting in place" of a codeless studio is beginning to take shape at the Singapore University of Technology and Design (SUTD).

Pedagogical Inversion: An Outside-In Architecture Studio

SUTD has recently announced itself as the world's first Design AI University,¹¹ with an SGD\$50 million investment in AI and design, focusing on "Human–AI interaction." Its vision of AI is to enhance and complement human capabilities through their mutual interactions. At first glance, SUTD might be mistakenly seen as yet another institution riding the current wave of AI hype and doing an AI-washing branding exercise to stay competitive. However, one should

11 Gabrielle Chan, "'Human–AI Interaction' Drives SUTD's \$50m Push for New Specialization in Design and AI," *Straits Times*, January 15, 2025, <https://www.straitstimes.com/singapore/human-ai-interaction-drives-sutds-50m-push-for-new-specialisation-in-design-and-ai>.

remember that SUTD was in fact the first university in the world to offer a four-year Bachelor of Science in Design and Artificial intelligence (DAI) degree back in 2020, predating even the late-2022 AI hype triggered by OpenAI's ChatGPT release. In fact, the first cohort of DAI students graduated in May 2024, exactly one and a half years after ChatGPT. In other words, the progression from the world's first Design AI degree to the world's first Design AI university is somewhat natural and convenient for a university of its size where a relatively agile administrative process is an advantage. Just as the paperless studio grew out of a convergence of theory, technology, media, and culture in the 1990s, SUTD's Design AI trajectories also grew out of an AI-First Nation¹² backdrop that consisted of Smart Nation 1.0 (2014), AI Singapore (2017), National AI Strategy (2019), and Smart Nation 2.0 (2024). These government-initiated strategic policies and multi-million-dollar fundings involving industry and academia to prepare an AI-trained competitive future workforce provided a very unique context which eventually propelled SUTD's rapidly focused adoption of AI in its university technology and design curriculum. SUTD's Architecture and Sustainable Design (ASD) degree program is no exception to this university-wide AI pivoting. Being the lead AI faculty jointly appointed at ASD and DAI degree programs, I have been responsible for designing and teaching all existing AI x Architecture courses offered at SUTD, namely, 20.318 Creative Machine Learning (2019–) and 20.224 Artificial & Architectural Intelligence in Design (2020–) at ASD, and 60.006 Spatial Design Studio (2023–) at DAI. Beyond the undergraduate level, through my own research laboratory, Artificial-Architecture, at SUTD, I have also been supervising AI x Architecture theses at both the MArch and PhD levels since 2020. In the following section, to shed light on the emergence of the so-called “codeless studio” at SUTD, the pedagogical structures developed for 60.006 Spatial Design Studio at DAI and the design research work produced at Artificial-Architecture shall be discussed.

A key pedagogical innovation of 60.006 Spatial Design Studio is the inversion of the standard architecture design studios. First, the students are not actually architecture students from ASD, but design and AI students from DAI. Second, there is no single architecture studio project brief enforced by a studio lead, but multiple problem statements that are inherently “architectural” and solicited from diverse industry partners curated by the studio lead (i.e., me). Third, specific AI approaches are not predetermined at the onset and can vary

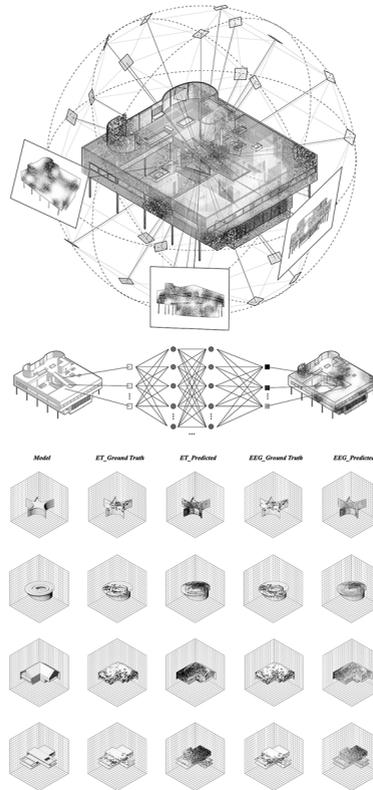
12 Laurence Liew, *AI-First Nation: A Blueprint for Policy Makers and Organisation Leaders* (AI-First Nation, 2024).

greatly from project to project throughout the duration of the studio. Thus, the creative exploration and critical evaluation of the appropriate AI techniques play an important role in the success of each project. Last, the final project outcome is not the design of a building, but the design of an AI prototype addressing existing gaps in architecture's conception and use of AI. In short, the pedagogical inversion here is the generation of a new AI x Architecture discourse from outside-in, rather than inside-out. That is, renewing existing architectural concepts with AI's notations through the production of the "architectural," instead of simply "adding" AI as a tool into the business-as-usual architecture design briefs and studios. In fact, the DAI Spatial Design Studio is designed to be a codeless studio prototype potentially paving the way to a future School of Architecture that must accept a transdisciplinary stance in the age of generative AI, especially when the architecture practice is at risk of de-professionalizing¹³ and the discipline losing its intellectual autonomy.¹⁴

The Brain-Perception Encoding project (or "Waves") at the DAI Spatial Design Studio in Fall 2024 built on the existing research work done at Artificial-Architecture to explore 3D architectural perception with neuroscientific and AI techniques.¹⁵

-
- 13 Richard E. Susskind and Daniel Susskind, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Oxford University Press, 2022).
 - 14 Patrik Schumacher, "The End of Architecture," *Khōrein: Journal for Architecture and Philosophy* 2, no. 2 (2024): 3–40.
 - 15 Immanuel Koh and Ashley Chen, "Your Memory Palace in the Metaverse with AI," *Proceedings of the AAAI Symposium Series* 1, no. 1 (October 3, 2023): 19–22, <https://doi.org/10.1609/aaaiss.v1i1.27469>; Elissa Hartanto, "Empirical Insights into Architectural Aesthetics: A Neuroscientific Perspective," in *Accelerated Design—Proceedings of the 29th CAADRIA Conference, Singapore, 20–26 April 2024*, vol. 3, ed. Nicole Gardner et al. (CUMINCAD, 2024), 69–78, https://papers.cumincad.org/cgi-bin/works/paper/caadria_2024_486.

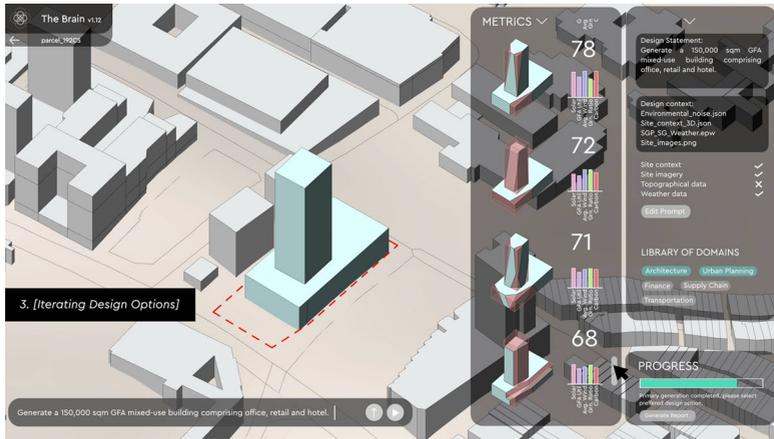
Fig. 42: Immanuel Koh, Brain-Perception Encoding, 2024. Electroencephalogram (EEG) and eye-tracking (ET) signals triggered when observing an architectural model in 3D via VR/AR are computed in parallel to generate a corresponding 3D saliency point cloud, where black-to-grey = high-to-low saliency and transparent = no saliency; A DGCNN AI model is trained to predict (with limited accuracy for now) the likely personal preference and attention distributions of any given unseen 3D CAD models.



The industry partner was HKS, Inc. which mainly supported the project by providing an initial dataset of 3D CAD models. In light of the increasingly prevalent, yet often conceptually superficial and technically simplistic, use of text-to-image generation models (e.g., Midjourney) among architects and architecture students, the project explores the potential of non-linguistic perceptions and generative designs. Six AI models were trained from scratch on electroencephalogram (EEG) and eye-tracking datasets recorded from human subjects (i.e., architecture and non-architecture students) interacting with 3D CAD architectural models in an AR environment. The trained AI models (specifically, a modified version of the dynamic graph convolutional neural network / DGCNN) were able to predict implicit attention and preference maps in the form of 3D saliency point cloud outputs given any new 3D model inputs. The project critically questions the affordances of today's generic text-to-image generation models while proposing an alternative form of the "non-generic" through the deep learning of brainwave and eye-tracking signals of individuals, thus potentially turning AI perception-predictive models into AI generative models. In short, a brain-to-architecture generation model.

The earliest conception of the ReGen City Design Brain project was seeded at the DAI Spatial Design Studio in Fall 2023 as the ArchitectMind.ai project, and is now a much larger ongoing research project at Artificial-Architecture. The then-industry partners were Autodesk and SAA Architects (a member of the Surbana Jurong Group). The former provided support and access to their Autodesk Forma API for the studio's development of custom add-ons, and the latter supported the project by granting the studio a glimpse of their existing non-AI workflows when conducting feasibility massing studies. The project leveraged LLMs to ingest design development planning guidelines from the online portal of the Urban Redevelopment Authority (URA) and trained a deep reinforcement learning model to generate feasible massing design solutions that ensure compliances such as building setbacks, while optimizing the maximum allowable building height and gross floor area (GFA). This project illustrates how a term-long studio could function as a prototyping platform for a subsequent in-depth lab research project, especially when the brief does not concern the design of a building, but the design of a system.

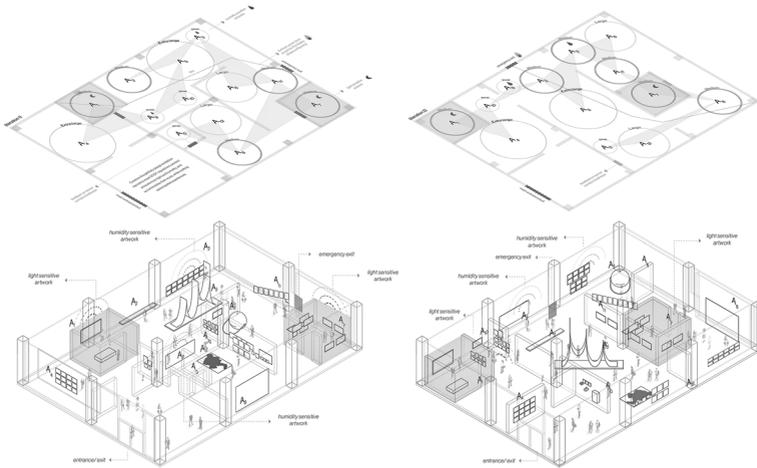
Fig. 43: Immanuel Koh, *ReGen City Design Brain*, 2024. An agentic multimodal large language model (MLLMs) prototype that generates 3D design massing models compliant with urban development planning guidelines and other relevant regenerative city design principles, while reasoning its own generative design processes in yielding optimal and explainable solutions.



The Curatorial AI Triangulator project was developed with the curators at the Singapore Art Museum (SAM) during Fall 2024 of the DAI Spatial Design Studio to explore ways in which artworks' semantics learnt by AI models could inform curators' own existing techniques of narrative spatialization. SAM's curators often use a linear circulation path embedded with a successive series of artworks' spatial triangulation to unfold an intended narrative/itinerary. In view of this curatorial practice, the project used an LLM to first generate a taxonomy for the artworks that included their physical attributes (e.g., dimensions and viewing distances), conservation requirements (e.g., lighting and humidity guidelines), and formal similarities (e.g., colors, materials, and figurations). The Curatorial AI Triangulator then computes the visual and narrative scores among all triads of artworks in order to construct a spatial network that adaptively adjusts itself with the generative placement of wall partitions, according to the given constraints of a linear path and gallery space. Using the 3D layout of an existing ongoing exhibition, "Everyday Practices,"¹⁶ the iterative outputs of the newly generated layouts were then evaluated as design options.

16 "Everyday Practices," Singapore Art Museum, accessed February 10, 2025, <https://www.singaporeartmuseum.sg/en/art-events/exhibitions/everyday-practices>.

Fig. 44: Immanuel Koh, Curatorial AI Triangulator, 2024. (LEFT) AI-encoding the individual artworks being displayed at the “Everyday Practices” exhibition held at the Singapore Art Museum; (RIGHT) New spatial configuration iteratively generated using LLMs for visual-narrative triangulations within the given gallery space.



Despite having only conducted the Spatial Design Studio twice thus far (Fall 2023 and Fall 2024), much can already be learnt from the students themselves when evaluating the effects of such pedagogical inversion. In substituting a common architecture design brief with a multitude of vastly different industry-informed “architectural” problem statements, students had first to overcome the immediate unfamiliarity of their chosen spatial domains, alongside the domain-specificity of workflows and tools, before even being able to articulate and then formulate an AI design prototype. It was both a technical and conceptual challenge. However, it was precisely the messiness of such pedagogical reconfiguration that students found not only exhilarating, but also liberating, because none of them had to act like architects in delivering building designs. It was like doing architecture without architects.

Neural Tectonics: An AI-Native Architectural Language

If the success of the paperless studios was measured by the intellectual impact of their faculty in manifesting a new architectural conception, language, or

theory, it is only fair to reflect on my own architectural development in the concluding section of the text. Indeed, a new language has emerged alongside a new theory called “neural tectonics.” It is a search for an AI-native architectural language for a third digital turn. Similar to how Prensky coined the terms “digital-immigrants” and “digital-natives” to differentiate two generations of human learners—those that were born into the digital age and those born prior¹⁷—the term “AI-native” of the third digital turn is to signal its difference from the first (digital-immigrants) and second (digital-natives) digital turns. An AI-native language is the language of the codeless studios and Artificial-Architecture. Neural Artefact Black (2023) and Neural Monobloc Black (2024) are perhaps most illustrative of this new language of the artificial. It is not the language of any particular “AI-simulated styles,” but that of AI’s very own “style of simulating” any styles. Both projects deliberately express the “Janus problem” as a form of AI aesthetics embedded in today’s text-to-3D diffusion models. This commonly generated computational glitch refers to the presence of an object’s canonical view (typically the front view) in several other non-canonical views (e.g., the side and back views), thus resulting in a 3D-generated object with multiple fronts or faces, much like the two-faced Janus, Roman god of beginnings. This glitch is a direct consequence of the inherent non-3D-aware behavior found in such 3D-generative models. Neural Artefact Black is the world’s first built physical public art-bench that is generated directly in 3D with a custom fine-tuned text-to-3D model and fabricated in an artisanal way with 100 percent upcycled wood. Sited in front of the Asian Civilisations Museum and along the historic Singapore River, it formally blends the learnt features from the antique Peranakan wooden furniture collection in the former and the long-disappeared small wooden boats (sampan) on the latter. The Neural Monobloc Black is a series of eight furniture pieces generated and fabricated in a similar way, except that it uses the generic white plastic Monobloc chairs as the learnt features to critique notions of optimized machine production and human consumption. The first project concerns an archival AI-reading of history while the second concerns an AI-critique of everyday design.

17 Marc Prensky, “Digital Natives, Digital Immigrants Part 1,” *On the Horizon* 9, no. 5 (2001): 1–6, <https://doi.org/10.1108/10748120110424816>.

Fig. 45: Immanuel Koh, Neural Artefact Black, 2023. Medium: Partially charred teak. Dimensions: 220cm x 108cm x 121cm. The “Janus” effect inherent in current non-3D aware text-to-3D diffusion models can be observed from the formal ambiguity of the bench’s “two-faced” seating configuration, which resulted in a two-way bifurcation at the leftmost end of the seat. 1 Express Place, Singapore. Photograph courtesy of Arts House Limited.



While AI researchers would call neural tectonics an undesirable hallucination problem, it is here a creatively desirable AI weirdness¹⁸ and formal subversion. After all, the paperless studios were not using the animation software to make movies, but as a new language to subvert then-existing modes of form-making in architecture. On a pessimistic note, however, when all creative AI hallucinations are resolved, and when all architectural ambiguities cease to exist, the notational space between academic studio experimentation and professional architecture practice might also dissolve, signaling the end of architecture studio culture, even the codeless studios themselves.

18 Mark Fisher, *The Weird and the Eerie* (Watkins Media Limited, 2017).

Fig. 46: Immanuel Koh, Neural Monobloc Black, 2024. Medium: Fully charred teak. On exhibition display at the National Design Centre in Singapore. Each of the eight functional chairs was generated with a custom fine-tuned non-3D aware text-to-3D diffusion model. The AI-driven “Janus” effect is even more pronounced here as observed from the multiple doubling/repetition of smoothly synthesized monobloc-like fragments.

