The Plan Journal 8 (2): 217-236, 2023
doi: 10.15274/tpj.2023.08.02.3
Peer-Reviewed
THE PLAN Journal

Jarvis, Hal, or AlphaZero? Looking Beyond Conventional Narratives Concerning AI and Architecture

Cem S. Kayatekin, June Aoun,
Yusuf Sühan Bozkurt, Daphné Fournel

ABSTRACT - *Since early 2021, the discourse concerning the potential and impacts of artificial intelligence on architecture has radically expanded. Discussions have largely focused on the heightened levels of productivity or efficiency that can be achieved within the existing ecology of architectural production processes, as well as the potential disruptions that may arise through human–AI co-authorship of the built world. What this paper asserts is that these dominant narratives appear to be extensions of quite conventional storylines which either frame artificial intelligence as a hyper-computational prosthetic for the enhancement of the architect or architectural office or as a critically disruptive force that will trigger micro- to macro-scale reconfigurations of the domain of built- environmental authorship. The dilemma is that we appear to be thinking of AI on old models of brute-force computation (i.e., Deep Blue) or dystopian conceptions of AI systems that can readily cross-pollinate with and radically disrupt existing societal configurations and dynamics (i.e., HAL-9000). What we have not quite considered are the real capacities and limits exhibited by artificial neural networks anchored around self-play reinforcement learning models (i.e., AlphaZero).*

Keywords: AI in architecture; AI in design; neural networks; reinforcement learning models

While artificial intelligence platforms have been in development and at significant levels of functionality within academic and research circles for several decades, the recent past has witnessed rather critical breakthroughs in the field—e.g., Ian Goodfellow's overnight development of the concept of Generative Adversarial Networks in 2014.[1] In the past handful of years in turn, this domain has stretched into the realm of individual use and consumption. DALL-E meandered into the public realm in January of 2021; ChatGPT was opened up to individual use in November of the same year; Midjourney arrived a bit later on the scene, launching its open beta model in July of 2022, followed a month later by Stable Diffusion. The public consumption of these platforms grew at an anomalous exponential rate with ChatGPT, for instance, reaching 100 million users two months after its initial launch.[2]

The topic of artificial intelligence has subsequently exerted a critical weight over the discursive landscapes of a wide range of disciplines. As expected, not much literature has been produced on the subject that is of a nonchalant tone. Although there is a discursive gradient with some subtlety that can be observed, the literature has tended to be predominantly anchored around utopian and dystopian projections concerning the nature, role, and impacts of artificial intelligence in the decades to come. The discursive landscape of architecture has replicated many of these trends. On one end of the range, AI is speculated to take form as an artificial neural networked, deep-learning-oriented high-tech assistant—akin to a JARVIS (Just A Rather Very Intelligent System) to a Tony Stark.[3] On the more disruptive end of the range, artificial intelligence is speculated as a source of stochastic noise and displacement that will trigger micro- to macro-scale disarrangements within the discourse, profession, and educational models underpinning architecture—more akin to a Clarkian conception of HAL-9000.[4]

Around the JARVIS polarity, speculations situate the functionality of artificial intelligence within the confines of problems requiring extensive quantitative analysis. In this domain, AI takes on the character of a semi-controlled brute-force supercomputer helping to resolve the complex layers of the design and management of the built environment bound to large data sets and non-qualitative design methodologies. Within the HAL-9000 polarity, in turn, artificial intelligence is envisioned striding into the wrinkles of the design process itself. At the micro-scale of the projected potential disturbances, it is cast as a co-author of the built environment, working side by side with the architect, with slightly varying degrees of authorship potential. At the macro-scale of the disruptive range, AI takes shape as an active competitor within the profession and marketplace of architecture at large.

Whether it is the utopian or dystopian range of the discursive landscape, there is a seeming feasibility or reasonability which underpins either speculative polarity. Yet, despite how legitimately extrapolated either projection may seem, there is the inescapable haze of limited foresight which comes to the fore—evoking in the reader that oft-repeated

Figure 1. Human and A.I. systems engaged in the co-authorship of the built world.

Kierkegaardian reminder: "It is really true what philosophy tells us, that life must be understood backwards. But with this, one forgets the second proposition, that it must be lived forwards," [5] There is, however, an oversight of the recent past which, if examined, may help to more accurately weigh the legitimacy of the future for artificial intelligence being envisaged. Namely, an extant and robust story arc concerning AI that has already reached near full maturity over the past three to four decades, found within the contemporary developmental narrative of chess. In examining the storylines underpinning the developments of AI chess systems over the recent past, what emerges is that the nature of the upheaval triggered by artificial intelligence was rather different than what was expected. And if this cross-disciplinary history is indeed permitted to serve as a partial precedent in this ongoing discussion concerning artificial intelligence and the built environment, then it would seem that the projections about AI and the shaping of the built world are also potentially critically off the mark.

THE JARVIS MODEL

Under the JARVIS pole of AI speculation, artificial intelligence is cast in the role of a hyper-technological operator, bearing the brunt of the quantitative and analytical requirements of various built environmental project stages—ranging from design development, construction documentation, project management, through post-occupancy evaluation.[6] To support this outlook, it becomes a requisite that the architect-AI relationship is shaped around certain models. One potential scenario is that the artificial intelligence system is given a vast and detailed amount of information upfront, and is then mobilized to process these datasets toward certain goals. The interaction between human and AI is limited to the initial input and final output phase, with nothing in between—an interactive arrangement referred to as the "black-box" approach.[7] While this format allows for certain advantages in terms of human control, it falls short of leveraging the potential of more back-and-forth iterative communication cycles. It also requires the development of hyper-detailed analytical and computational programming in order to achieve the desired output results.

Potentially a more practical and productive scenario is one which allows for more flexible and intermittent interaction between the human and synthetic intelligence system to take place. Known as "gray-boxing," this format allows for more incremental input–output cycles as well as refinement of the machine-learning-oriented neural networks to take place, while still maintaining a degree of human control over the process as a whole.[8] A more dynamic relationship between human and artificial intelligence emerges as a result.

> The ability to intervene throughout the generating process is a fundamental dimension: as each step of the pipeline represents a distinct portion of architectural expertise, each model can be trained independently, opening the way to significant improvements and experimentation in the near future. Indeed, improving this entire pipeline end-to-end could be a long and cumbersome task, while amending it step by step remains a manageable process, within the reach of most architects and engineers in the industry.[9]

Regardless of the relationship dynamic adopted, the existing discourse operating around JARVIS model presuppositions tends to frame the constructive potentials of AI around certain, typically quantifiable, layers of built environmental design and analysis—specifically focused around, but not limited to the topics of:

1. project energy and carbon footprint modeling/tracking; [10]
2. environmental control systems design/refinement/evaluation; [11]
3. structural systems design/refinement/evaluation; [12]
4. critical analysis of building codes and urban planning regulations; [13]
5. design/refinement of urban infrastructures; [14]
6. post-occupancy analysis and workspace planning; [15] and
7. the amplification of participatory design frameworks and processes.[16]

The overarching theme to the speculated JARVIS storyline is the confinement, whether of a “black-box” or “gray-box” nature, of artificial intelligence systems to the quantitative domain, and either at a significant distance or simply slightly out of reach from the more critical qualitative portions of the design development process. While “collaboration” is a term that is used within this discursive landscape, what is envisaged is a rather neutered and lopsided type of collaboration wherein the artificial intelligence system is relegated to the role of a quantitative prosthetic to the human author.

> Through such collaborative intelligence, humans and A.I. actively enhance each other’s complementary strengths: the leadership, teamwork, creativity, and social skills of the former, and the speed, scalability, and quantitative capabilities of the latter. [...] humans need to perform three crucial roles. They must train machines to perform certain tasks; explain the outcomes of those tasks, especially when the results are counterintuitive or controversial; and sustain the responsible use of machines (by, for example, preventing robots from harming humans).[17]

THE HAL-9000 MODEL

Within the more Clarkian extrapolation of the future of AI and built environmental authorship, artificial intelligence is envisioned as breaking its quantitative confinements and delving at varying depths into the more qualitative, discretionary, intuitive, or intellectual domains of project development. While the literature shaped around the JARVIS model generally adopts a more technocratic and optimistic tone, at the HAL-9000 end of the spectrum an atmosphere of a loosening (or loss) of control pervades. On one end of this discursive niche, artificial intelligence is cast in the role of intra-office collaborator or co-author of the built world. A micro-scale example of this is found in firms engaged with the use of

AI image generation as a means of propelling early design development and ideation. Patrik Schumacher for instance, noted that "most" of Zaha Hadid Architects' projects now lean upon AI-generated imagery within early conceptual development phases, asserting that prompting a generative AI platform felt rather akin to prompting a human design team (although this anthropomorphization of artificial intelligence was also accompanied by Schumacher's claim of full authorship over the images generated by said AI platforms at Zaha Hadid Architects).[18]

Globally an increasing array of large, well-known firms are clearly in active internal discussions about how to incorporate artificial intelligence systems into the iterative design process in comparable ways, however, the idealized parameters for their incorporation have predominantly not been clearly articulated to the public.[19] At smaller scales of operation, some offices are quite transparent about the precise nature of how generative AI platforms are being mobilized—e.g., the leveraging of anomalies or output noise (and the maximizing of such potential output noise) within AI-generated media as instigators for design development: "In a way, I'm finding that this is like sketching, where there is a weird artifact that's serving as this kind of crack, [that] is exposing this entirely new opportunity that I hadn't previously considered." [20]
Herein, AI-generated imagery, or more specifically the anomalies, distortions, or "cracks" within such imagery are utilized as *objet trouvé* on which unexpected iterative design-process trajectories can be anchored. While the co-authorship role of artificial intelligence here is quite compartmentalized, this still overcomes a critical juncture wherein the outputs authored via artificial intelligence systems are validated as legitimate anchors for ideation.

A comparable, mostly theorized (and marginally observed within student bodies by the authors of this article), area for co-authorship emerges with AI taking on the role of sounding board—i.e., "a thought processor instead of a word processor" [21] that checks and corrects for deviations in the theoretical, ideological, or narrative structures being utilized to frame the developmental trajectory of a project. Particularly in dealing with complex multivariate built environmental problems, the narratives developed to underpin a design proposal can often develop theoretical incompatibilities and contradictions. These can be difficult to spot, or even if observed, difficult to disentangle. Within the AI-as-thought-processor scenario, an individual or design team can directly engage with the AI system (e.g., ChatGPT), updating it on the current state of affairs—i.e., the narratives being utilized within the design process and the decisions that have thus far been made. When queried to do so, the artificial intelligence platform can then check for theoretical, ideological, or intellectual deviations, and articulate an array of external viewpoints to help reroute the theoretical or narrative structure to better cohere with the choices that have been made, or inversely rewire the design decisions that have been made so that they remain in better accord with the existing narrative structure.

At the other, more macro end of the HAL-9000 landscape, AI breaks out from these micro-scales of co-authorship potential and is envisaged as a liberated operator within the free market of built environmental authorship. Rather than the competition between human and synthetic agents taking place on a somewhat level playing field, however, the chances for long-term monopolization are assumed to be weighted toward the non-human sector. This is not an entirely unfounded concern, as the growing impact of artificial intelligence systems on certain professional domains is observable in the current day. The exponential development of ChatGPT's capacity for programming, for instance, is reshaping not only the nature of software engineering in the twenty-first century, but by extension the labor ecosystem coinciding with said profession.[22] Within architecture, much of the literature predicts a comparable corollary on the horizon. While AI may make smaller offices more efficient and able to compete with larger-scale offices by increasing their levels of efficiency and productivity,[23] it is extrapolated that artificial intelligence will also shed light on, and take advantage of, the innate constraints of human output. "There are signs that AI is becoming not only good, but terrifyingly good, to the point that it is beginning to expose our own limitations as human beings, and putting our jobs as architects unquestionably at risk." [24]

Looking across this discursive arc, as opposed to the more technocratic and quantitative nature of the JARVIS spectrum, the dynamics of the HAL-9000 pole are anchored around the qualitative and less-controlled output potentials of AI platforms. Whether in terms of media or ideas, artificial intelligence is extrapolated as an external contributing voice and agent, potentially altering the design process and the landscape of built environmental authorship, in a much more stochastic and less predictable manner. An oddity to note is that while the micro and macro extremes of this HAL-9000 landscape are fairly well defined, the middle sector is quite nebulous. On the minimum disruption end, artificial intelligence is granted some co-authorship or collaborative role, however, it tends to be inherently confined in terms of potential corporate growth—e.g., with AI essentially hitting an impassable glass ceiling at the role of junior architect, whether as sounding board or ideation source. On the maximum disruption end, AI is allowed to sever ties completely with its human counterpart and engage in free market competition as an independent contractor. It would be expected that what is left in between these poles, would be a collaborative interaction in the form of a legitimate and rigorous discourse between independent and discretion-exercising human and synthetic agents.

In reality, however, this middle ground remains predominantly populated with texts authored as:

1. philosophical treatises concerning the radically changing trajectory of the relationship between homo sapiens and technology—e.g., the incorporation of smartphones and the varying structures of the internet into our socio-psychological conception of "self"; [25]

2. as reiterations of utopian-Marxist societal projections wherein technological advancements can foster the liberation of homo sapiens from industrial toil; [26] and
3. as politico-economic analyses that offer valid critiques of utopian and dystopian extrapolations for the future, without quite clarifying what manner of middle ground could be expected in their stead.[27]

ALPHAZERO: A THIRD MODEL?

Taking both the JARVIS and HAL-9000 discursive geographies under survey, the overarching storylines appear to be dominated by propositions of confinement and projections of fear or inadequacy—specifically, propositions to confine AI within certain built environmental roles, whether they be strictly quantitative technocratic positions or qualitative discretionary ones within the lower rungs of an architectural office hierarchy; and the fear that if this confinement is broken, the Clarkian narrative may reach full maturity and expose the innate vulnerabilities and inadequacies of human built environmental authorship. Within the story arc of artificial intelligence observed within the domain of chess, two important findings emerge that are of use to this discussion:

1. the speculated brute-force supercomputer trajectory of AI (which the JARVIS model leans upon), reaching its apex with IBM's Deep Blue in the mid-1990s, proved to be a tangential aside to the final conceptual framework which overtook the domain of artificial intelligence development in chess; and
2. the ability of synthetic intelligence to communicate with, overlap with, relate to, or have a direct impact upon, human intelligence (which much of the HAL-9000 oriented projections of the future lean upon) also proved to be less consequential than expected.

While the broader literature stretches much further into the past, it was pointedly in the 1980s and 90s that discussions concerning artificial intelligence and chess started to clearly enter into the domain of public and interdisciplinary academic discourse. Within this time frame, AI system development was primarily centered around a brute-force approach to the game, wherein increasing processing power and quantitative computational potential were the main objectives.[28] With each new generation of synthetic innovation, more powerful computer systems emerged and were subsequently placed in competition with human players of varying chess caliber across the board.

This storyline reached its apogee in the mid-1990s with the widely publicized matches between IBM's Deep Blue and then-world chess champion and grandmaster Gary Kasparov. In the first 1996 match, Kasparov defeated Deep Blue 4 games to 2, winning games 2, 5, and 6, and drawing games 3 and 4. A general sigh of relief echoed across the public discourse, as the dominance of homo sapiens in this simultaneously creative and quantitative game was once more underscored. The period of calm, however, did not last, with the then-updated Deep Blue coming back to the table in 1997. It could evaluate 200 million positions per second, weighed around 1.4 tons, had 418 custom processors, and was approximately the size of a refrigerator.[29] Deep Blue won games 2 and 6, and drew games 3, 4, and 5. Kasparov was defeated with a result of 2.5 games to 3.5. While the grandmaster demanded a rematch, IBM rejected the offer, dismantling and shelving Deep Blue quickly after.[30] This was the epitome of the brute-force trajectory for chess-oriented artificial intelligence.

In the current day, the pinnacle of the game is occupied by a system called AlphaZero, developed by Google subsidiary DeepMind in 2017. Compared to Deep Blue's potential evaluation of 200 million positions/moves per second, AlphaZero operates at a mere fraction of that capacity—specifically, 60,000 positions/moves per second. Yet, across the hundred matches played between AlphaZero and Stockfish (the then-reigning computer chess champion) in 2017–18, AlphaZero won 28 games and drew 72. Not one victory for Stockfish; not one loss for AlphaZero.[31] "Most unnerving was that AlphaZero seemed to express insight. It played like no computer ever has, intuitively and beautifully, with a romantic, attacking style. It played gambits and took risks. In some games, it paralyzed Stockfish and toyed with it." [32]

As opposed to the Deep Blue brute-force hyper-processing approach, AlphaZero relies on the use of self-play deep reinforcement learning and artificial neural networks for the purposes of pattern recognition and generation. It does this better, but not necessarily quantifiably faster, when compared to its synthetic predecessors such as Deep Blue. A most critical distinction is that AlphaZero was not fed vast archives of human-versus-human chess games, compiled over the past centuries, in order to understand the game. It was simply given the rules underpinning chess and proceeded to play against itself through a simulator. Over the course of four hours, it refined its gameplay to the level at which it currently stands.[33] Due to the nature of how Elo ratings are calculated, it is difficult to quantify the exact level of AlphaZero's gameplay, but it is clearly beyond Stockfish, which is estimated to sit around an Elo rating of 3,546. MuZero, AlphaZero's successor, has seemingly surpassed this level, although a precise number is, to date, not known. By comparison, the highest-ranked human player of the current day, grandmaster Magnus Carlsen of Norway, is situated around an Elo rating of 2,839.

In studying AlphaZero's games against Stockfish, two significant points arise. First, AlphaZero's strategic understanding of the art of sacrifice is significantly anomalous. When sacrificing pieces, human players will tend to expect a return on investment, whether in terms of positional advantage or an exchange of pieces down the line, in the relatively immediate future. AlphaZero plays with a radically different rhythm, sacrificing pieces with seemingly no immediate strategic return on the horizon. One layer underpinning this behavior is that AlphaZero is likely able to visualize a much longer gameplay, and hence extrapolate a much more delayed return on investment via multiple unfolding storylines when compared to a human player.[34] A rather more eerie explanation is that AlphaZero may be pursuing "some sort of positional domination [that is] hard to understand." [35] Fellow grandmaster Hikaru Nakamura echoes Magnus Carlsen's above point, noting that AlphaZero's assessment of positional advantage appears to be quite fundamentally different from what a high-level human assessment of the structural state of a game will yield.[36]

Second, while AlphaZero clearly pursues and unfolds a type of gameplay that is "alien," [37] there is a surprising limit as to how much this new approach to the game has been able to alter or refine the human approach to chess. Nakamura, for instance, attributes this lack of cross-pollination to the difficulty, or perhaps impossibility, for AlphaZero (if it were even equipped to do so) to translate its positional and strategic understanding of the game into terms that homo sapiens could comprehend, let alone apply,[38] echoing Wittgenstein's famous observation that "If a lion could talk, we would not understand him." [39] In 2019, Magnus Carlsen, deeply inspired by AlphaZero's radically different approach to chess, adopted certain nuances of its gameplay. For a short period of time, he was sacrificing pawns with "great joy" and with some degree of success.[40] However, other players soon developed offensive and defensive lines that took advantage of these new unfolding structures, causing Carlsen to reel his newly adopted gameplay back toward more human lines.[41]

> In terms of man versus machine, that battle was lost for humans even before I entered top-level chess. So that's never been an issue for me. I never liked playing against computers much anyway. So that's completely fine. But it was amazing to see how they 'thought' about chess in such a different way, in a way that you could mistake for creativity.[42]

In the late 1980s through the 1990s, with the rise of Deep Thought, Deep Blue, and the latter's matches against grandmaster Gary Kasparov, the overarching speculation for the future of AI chess systems was the emergence of a platform that would effectively play chess like a human, but better. The brute-force concept underpinning the development and refinement of such systems relied on the synthetic player's access to immense knowledge sets concerning defensive and offensive openings, positional understandings, strategic narratives, etc., as fed into the system by human programmers. The AI system's primary advantage was its capacity to quantitatively out-read the human player. Anchoring this was the presupposition that both sides would be reading the same discursive landscape; or perhaps, that there was only one landscape to read.

AlphaZero upended this prediction quite significantly. Cognitively, analytically, and (if it dare be said) creatively, it is another species. It has significantly different capacities for foresight, positional understanding, and strategic narrative development. AlphaZero has access to and plays within a landscape of the game that was effectively not known to exist. Arriving hand in hand with this discursive disarrangement is the reality that AlphaZero either won't be able to vocalize the nature of this landscape within a terminological framework that homo sapiens will be able to understand, or that the fundamentals of this newly uncovered geography are so beyond the limits of homo sapiens that even if AlphaZero could describe how it navigates this realm with clarity, the human chess player would still be unable to work within this terrain due to our species' cognitive, analytical, and by extension creative, limits. In other words, AlphaZero's Elo rating isn't a new level of chess mastery for the human player to aim for. AlphaZero simply reads the game in a way that, fundamentally, homo sapiens cannot.

AI AND THE BUILT WORLD

Is there a corollary to this third pole, this AlphaZero domain, that is applicable to the architectural discourse's conception of AI? The extant bounds of the JARVIS and HAL-9000 domains rest upon assessments of weak AI platforms. The main difference between the two poles appears to be how these structures are accepted and incorporated into the underlying process of built environmental authorship — e.g., as prosthetic or as co-author/author. Both extremes utilize systems designed to satisfy rather narrow tasks, often but not necessarily using some form of artificial neural networks, trained via various learning models, and oftentimes relying on large external data sets to operate. Some systems (e.g., Midjourney and ChatGPT) are anchored around Generative Adversarial Networks and appear to be closer to leveraging the stochastic potentials of artificial intelligence. However, an architectural AI system trained via an internalized self-play reinforcement learning model such as in the case of AlphaZero, does not currently exist. This self-play reinforcement learning model incidentally is considered "crucial for achieving successful AI." [43]

Given the absence of such a system within the landscape of architecture, an example of a comparably "alien" approach to built environmental authorship and design, such as that exhibited by AlphaZero in chess, is also nonexistent. In many respects, the state of AI in architecture rests in the Deep Blue era, with many of the current tools functioning as slight, certainly not radical, advancements of the prior parametric period. Consider, for instance, current platforms which accept basic inputs and conditions from the user regarding site, program, location, orientation, certain basics of urban planning regulations, etc., and then output sets of fully detailed BIM models of proposals, oftentimes with carbon and energy analyses attached—e.g., “Architechtures,” “Maket,” “Arkdesign,” to name a few. When placed within a broader context, these platforms appear to be slight refinements to an already well-trodden lineage. In the 2010s in particular, the parametric era produced quite a range of platforms focusing on precisely this task of generating quantitative performance data in step with early design development choices—e.g., site plan massing and its impacts on energy performance metrics,[44] the multivariate layers surrounding urban sustainability that include energy and/or carbon metrics but also social and infrastructural considerations,[45] and even more niche, but quite complex areas such as urban ventilation potentials.[46] Even two decades prior to these efforts, California Polytechnic University, Lawrence Berkeley Laboratory, and the University of Oregon were developing a “Building Massing Intelligent Design Tool,” under the broader auspices of the “Advanced Energy Design and Operation Technologies Research Project” to address a comparable task—namely, the development of a rule-based pattern-matching software to assess the potential energy performance of early-design-phase floor- or site-plan configurations.[47]

The “AI-powered” platforms currently taking on these tasks (and others) certainly boast impressive output numbers and can be mobilized with significant constructive utility within the design process. However, two issues emerge. First, output capacity has its practical limits. For instance, a weak in-house AI platform can generate 100,000 potential floorplans every twenty-seven hours.[48] Certainly an eye-catching quantity to consider, but in terms of workflow, given that the idea is that these plans are then secondarily filtered by human discretion—if a singular person were to try to merely skim through these 100,000 plans at a generously fast pace of three seconds per plan, this would require 83.3 work hours to complete the task in full.

This process could of course be automated further, but filtering is already a significant dynamic of generating AI outputs. Within Generative Adversarial Networks, for instance, one network is assigned the task of generating output, while the other is allocated the inspection or filtering part of the process.[49] The former tries to produce coherent and meaningful output that passes the secondary network’s filtering discretion, while the latter maintains high standards to keep the former on its toes. The filtering process can certainly be strengthened to reduce the quantity and increase the meaningful quality produced.

However, the current architectural discursive landscape may have difficulty in navigating the subtleties of quality, considering that its assignation of the adjective “revolutionary” to various AI platforms appears to be dominantly reliant on the presence of vast levels of output quantities. This further indicates that even in terms of the vocabulary in common architectural use, the discursive domain hasn’t quite moved beyond the Deep Blue era’s pursuit of pure output metrics (e.g., the capacity to evaluate 200 million moves/positions per second, or to produce approximately 100,000 plans per day). The AlphaZero (and MuZero era) still appears to be on the horizon. This points to a rather uncomfortable possibility. The purported architectural AI revolution may in effect be the extension of the prior parametric era, simply caught in the current adrenaline storm surrounding artificial intelligence as a whole.

> The most obvious [issue with cyclically-occurring seasons of high interest for A.I.] is that it creates this ripple every few years of what have sometimes been called A.I. winters, where there's all this overpromising that A.I.'s will be about to do this or that. It might be to become fully autonomous driving vehicles instead of only partially autonomous, or it might be being able to fully have a conversation as opposed to only having a useful part of a conversation to help you interface with the device. This kind of overpromise then leads to disappointment because it was premature, and then that leads to reduced funding and startups crashing and careers destroyed, and this happens periodically, and it's a shame. It hurt a lot of careers. [...] It's just immature and ridiculous, and I wish that cycle could be shut down. And that's a widely shared criticism. I'm not saying anything at all unusual.[50]

Will an AI system eventually be able to produce architecture? A readily available defense of the built environmental domain is that the shaping of the built world requires an intuitive, nuanced, layered, and qualitative approach, in order to engage at a high level. Incidentally, these are variations of the very terms once ascribed to grandmaster-level chess, used to explain why an amalgam of processors and semiconductors would not be able to match the human capacity for the game. That barrier was overcome in the late 1990s. It was in this same timeframe that an artificial intelligence system named EMI (Experiments in Musical Intelligence), developed by Dr. Steve Larson out of the University of Oregon, began composing increasingly complex music, another area of human endeavor that was partnered with the above adjectives as an innate barrier for non-human entry.[51] These same words were then shifted to the game of Go as the next bastion of defense against the rise of artificial intelligence. This barrier too was overcome in 2016 with the emergent dominance of AlphaGo, also developed by DeepMind Technologies.[52]

Will AI produce a higher caliber architectural product than what homo sapiens is capable of? This question, which is one of the cornerstones of panic within contemporary discussions concerning artificial intelligence, seems to lean on a rather fundamental misconception. Architecture does not have an Elo rating. There is a "fairly" objective framework through which one can distinguish, for lack of a better word, great works of architecture from poor works of architecture. However, the weight of objectivity clearly begins to dwindle when a micro-scale of granularity is pursued with such categorization. Is poor-work-A poorer than poor-work-B? Or is great-work-A greater than great-work-B? Is Le Corbusier's Ronchamp a higher example of architectural merit than Gaudi's Sagrada Familia? While this discussion requires a full-length philosophical treatise to delve into in full, this level of granularity in gauging architectural caliber clearly bears too much of the influence of the observer's own alignment (or misalignment) with the work's, or the author's, underpinning worldview, ideology, aesthetic and spatial sensibilities, and so forth, to be able to qualify as an objective measure.

Will AI produce the greatest works of architecture? The landscape of the built world does not unfold in this manner. Unlike chess, there is no potential ELO rating of 3,546 to be achieved within the domain of built environmental authorship. AI may produce an architecture that it values at a higher level, but just as a human agent cannot author, objectively and undisputedly, the "greatest" work of architecture, AI authorship of the built world will invariably rest within the same restrictions.

Another layer underpinning this disciplinary fear is anchored within the HAL-9000-oriented speculation of an AI architectural office simply outcompeting their human counterparts in the free market. Namely, that AI will be able to generate alternative, more cost effective (both in terms of design and construction costs), and more efficient pathways to the shaping of the built environment. In reality, however, for the clients, builders, developers, architects, municipalities, etc., interested in accelerating the pace through which projects are run through the developmental pipeline while sacrificing certain standards, there are already numerous available pathways to do so. The built world, globally, has no paucity of copy+paste or shoddy developments, whether of a suburban or urban nature, that are produced at a bewildering pace and a bewildering substandard of quality. The production of the built environment can already be accelerated when the correct politico-economic levers are leaned upon. There is no need or even perhaps capacity, for artificial intelligence to hasten this runaway process any further.

Within this context, there is also the wrinkle of accountability to take into consideration. An architect is a licensed professional, bearing a professional stamp, and as such open to litigation for a wide range of conditions, from breach of contract to professional malpractice or negligence. A district judge in Washington D.C. recently concluded that a work of art generated by artificial intelligence, without any human input, cannot be copyrighted.[53] It is difficult to extrapolate how an AI system, anchored in a sitting computer network, which cannot legally claim ownership over an artwork, would eventually be granted the capacity to not only achieve professional standing but also be able to claim legal accountability over a work of architecture or a built environmental project of even larger scale.

There is another layer though to the question of speed. If the integrity of the architectural office is maintained (and not lost to synthetic competition), will artificial intelligence, operating under the JARVIS model, accelerate the pace of the production of architecture (e.g., modeling, rendering, structural/environmental systems design and analyses, carbon footprint quantifications, etc.) to higher limits? While some acceleration is likely at hand, an exponential trend is difficult to imagine, for the simple reason that the JARVIS model appears to be a rather straightforward extension of, not a radical leap from, the already extant lineage of parametric and quantitative analytical platforms. Many of these activities—modeling, rendering, analyzing, quantifying, etc.—can already be mobilized with a near-zero wait time. For instance, there is a vast range of readily accessible platforms (e.g., SimScale, Ansys, Flowsquare, etc.) that can run some of the most complicated calculations inherent to the built environment, namely three-dimensional computational fluid dynamics, at nearly instantaneous velocities.

The main time investment for such a calculation is the development of a thoroughly detailed BIM model to be used within such simulations. In constructing a digital model in turn, the actions of extruding or lofting walls, floors, surfaces, roofs, etc., while finicky, do not consume a great deal of time. Most of the hours spent in front of the screen involve specifying materials, methods of construction, thermal performance values, and so forth. Even if an AI system were relegated to these tasks, a significant back-and-forth interaction with a human counterpart or team would be required to make these specifications. This isn't much different from the on-screen input-based interaction that already takes place within digital modeling platforms. Above all, the fact remains that within current-day architectural practices, even the most complex BIM model can be produced in a rather short period of time by even an average team at work.

The acceleration of the processes surrounding built environmental authorship, setting aside the lumbering pace of construction, is not so permeable. In the garment industry, a similar outcome was encountered while trying to ramp up sewing machine efficiency, and by extension garment production, in the twentieth century. Increases in mechanical speed (stitches per minute) initially led to increased production in quite a straightforward manner. However, after some point, this trajectory began to flatline. It turns out that most of the time spent operating a sewing machine is actually spent handling, moving, and preparing the fabrics and textiles about to be sewn (estimated at 63% of total time).[54] Only 37% of total time is spent sewing. Increasing the stitching speed thus quickly hits a limit, much the same as accelerated processing speeds, or vastly expanded output quantities, hit a hurdle within the production of architecture. Is the frenzied discussion within which the discipline finds itself then all for nought? Will AI have no radical impact on the architectural landscape?

Much of human play within chess, at least within the early- and mid-game, relies on strategic storylines. The English Opening, the Sicilian Defense, King's Indian, Queen's Gambit, Dutch Defense, the Catalan Opening, etc.—these are defensive and offensive patterns of moves, which start to lock the game into certain narratives, and in turn, limit the emergence of other alternative narratives. Within the documented games between AlphaZero and Stockfish, for instance, AlphaZero is often observed playing with the English Opening. At least in the start, this narrative structure is fairly recognizable, however, once AlphaZero's deviant approach to the game begins to unveil itself, significant aberrations to this structural storyline begin to take place and a fundamentally unrecognizable gameplay then emerges, sometimes abruptly, sometimes gradually.

Within the practice and teaching of contemporary architecture as well, there is a predominant (but certainly not ubiquitous) stress placed on the importance of narratives. These structures are leaned upon in order to legibly explain the conceptual framework underpinning a completed work, and establish a recognizable story arc as to how the design unfolded. Within the design process itself, architects often rigidly anchor theoretical frameworks, storylines, ideologies, worldviews, etc., that constrain how a project takes shape and is refined. Deviations from the set narrative are quickly pruned. Storylines are maintained.

If the story arc of artificial intelligence in chess is taken as precedent, one nearly certain upheaval that will take place if there is the necessary investment to develop an AI architectural system trained via a self-play deep reinforcement learning model is the uprooting of narrative. Namely, it will likely emerge that the stress placed upon the importance of narratives and structural storylines within the architectural design process is not due to their innate value to the design process in itself, but rather because they behave as a crutch to the limits of human foresight. Just as AlphaZero plays chess as a differently abled cognitive species would, with fundamentally different understandings of foresight, positional advantages, and long-term strategy, an artificial intelligence system geared toward the shaping of the built environment will likely approach the design process with a significant loosening of the notion of narrative, or an abandonment of such concepts altogether. It may not be able to articulate the different process that it will navigate in terms that homo sapiens can comprehend. Or even if such clear articulation is achieved, the cognitive and cerebral limits of our own species will likely keep any potential cross-species deployment of such novel discursive territory at bay. In any case, a radical questioning of the notion of narrative, as well as other presuppositions of the design process currently taken for granted, will likely be at hand.

By extrapolation, significant upheavals in terms of understanding fundamental notions of cognition, intelligence, and creativity will likely emerge on the horizon. These are significant disarrangements of the discourse and discipline of architecture, but they are not quite the upheavals with which the overarching discussions seem to be concerned. Underpinning much, not all, of this speculation, however, is a significant "if." The emergence of a specialized artificial intelligence system geared toward built environmental authorship will likely arise if the necessary investments for its development are put in motion. Given the potential lack of significant impact to the profession and market at large, and the vast amounts of funding that such a specialized AI system would require, a primitive query arises—how will a return on investment be achieved? What appears more likely, in fact, is the emergence of a general AI system, or artificial general intelligence system, which is then asked to author architecture. This general AI may indeed engage with architecture as an intellectual intrigue, but on the other hand, just as many members of our own species, it may find that its interests lie elsewhere.

Notes

1. Martin Giles, "The GANfather: The Man Who's Given Machines the Gift of Imagination," *MIT Technology Review*, February 21, 2018 – https://www.technologyreview.com/2018/02/21/145289/the-ganfather-the-man-whos-given-machines-the-gift-of-imagination/, accessed October 24, 2023.
2. Krystal Hu, "ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note," *Reuters*, February 2, 2023 – https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/, accessed October 15, 2023.
3. *Iron Man,* directed by Jon Favreau (Paramount Pictures, 2008).
4. *2001: A Space Odyssey*, directed by Stanley Kubrick (Metro-Goldwyn-Mayer, 1968).
5. Søren Kierkegaard, "Søren Kierkegaards Skrifter," *Journalen JJ* (1843): 167. Translated from the original Danish, "Det er ganske sandt, hvad Philosophien siger, at Livet maa forstaaes baglaends. Men derover glemmer man den anden Saetning, at det maa leves forlaends."
6. Choulwoong Kwon, "AI and the Future of Architecture: A Smart Secretary, Revolutionary Tool, or a Cause for Concern?," *International Journal of Sustainable Building Technology and Urban Development* 14, no.1 (2023): 130.
7. Stanislas Chaillou, "Archigan: Artificial Intelligence X Architecture," *Architectural Intelligence: Selected Papers from the 1st International Conference on Computational Design and Robotic Fabrication CDRF 2019* (Singapore: Springer Nature Singapore, 2020).
8. Ibid.
9. Ibid.
10. Semen Andreevich Budennyy et al., "eco2AI: Carbon Emissions Tracking of Machine Learning Models as the First Step Towards Sustainable AI," *Doklady Mathematics* 106, no. 1 supplement (December 2022); Carole-Jean Wu et al., "Sustainable AI: Environmental Implications, Challenges and Opportunities," *Proceedings of Machine Learning and Systems* 4 (2022): 795–813.
11. Sanguk Park et al., "Reinforcement Learning-Based BEMS Architecture for Energy Usage Optimization," *Sensors* 20, no. 17 (2020): 4918; Yassine Himeur et al., "AI-Big Data Analytics for Building Automation and Management Systems: A Survey, Actual Challenges and Future Perspectives," *Artificial Intelligence Review* 56, no. 6 (2023): 4929–5021.
12. Hadi Salehi and Rigoberto Burgueño, "Emerging Artificial Intelligence Methods in Structural Engineering," *Engineering Structures* 171 (2018): 184; Kaan Bingöl et al., "Artificial Intelligence Applications in Earthquake Resistant Architectural Design: Determination of Irregular Structural Systems with Deep Learning and ImageAI Method," *Journal of the Faculty of Engineering and Architecture of Gazi University* 35, no. 4 (2020): 2197–2209.
13. Sibel Macit Ilal and H. Murat Günaydın, "Computer Representation of Building Codes for

Automated Compliance Checking," *Automation in Construction* 82 (2017): 43–58; Hamed Nabizadeh Rafsanjani and Amir Hossein Nabizadeh, "Towards Human-Centered Artificial Intelligence (AI) In Architecture, Engineering, and Construction (AEC) Industry," *Computers in Human Behavior Reports* 11 (2023): 100319.
14. Alexandros Nikitas et al., "Artificial Intelligence, Transport and the Smart City: Definitions and Dimensions of a New Mobility Era," *Sustainability* 12, no. 7 (2020): 2789; Yang Zhang et al., "StreamCollab: A Streaming Crowd-AI Collaborative System to Smart Urban Infrastructure Monitoring in Social Sensing," in *Proceedings of the Ninth AAAI Conference on Human Computation and Crowdsourcing*, eds. Ece Kamar and Kurt Luther (Palo Alto CA, USA: The AAAI Press, 2021).
15. Farah Nayeri, "How A.I. Is Helping Architects Change Workplace Design," *The New York Times*, June 20, 2023 – https://www.nytimes.com/2023/06/15/business/workplace-design-zhai-ai.html, accessed October 24, 2023.
16. Daniel Cardoso Llach, *Builders of the Vision: Software and the Imagination of Design* (Abingdon, UK: Routledge, 2015); Linda Neuhauser and Gary L. Kreps, "Participatory Design and Artificial Intelligence: Strategies to Improve Health Communication for Diverse Audiences" (paper presented at the 2011 AAAI Spring Symposium, Stanford CA, USA, March 2011).
17. H. James Wilson and Paul Daugherty, "How Humans and AI Are Working Together in 1,500 Companies," *Harvard Business Review* – https://hbr.org/2018/07/collaborative-intelligence-humans-and-ai-are-joining-forces, accessed November 19, 2023.
18. Nat Barker, "ZHA Developing 'Most' Projects Using AI-Generated Images Says Patrik Schumacher," *Dezeen*, April 26, 2023 – https://www.dezeen.com/2023/04/26/zaha-hadid-architects-patrik-schumacher-ai-dalle-midjourney/, accessed October 24, 2023.
19. UNStudio, "Technology Report: The Many (Inter)Faces of Technology in the Built Environment," UNStudio website – https://www.unstudio.com/en/page/16764/technology-report-the-many-inter-faces-of-technology-in-the-built, accessed September 5, 2023.
20. "Using AI as a Design Tool in My Architecture Practice," YouTube video, posted by "30x40 Design Works," June 3, 2023 – https://www.youtube.com/watch?v=zhKb9l6LBSo.
21. Gina Kolata, "With Major Math Proof, Computers Show Flashes of Reasoning Power," *The New York Times*, December 10, 1996, C1.
22. Neil Leach, "AI Is Putting Our Jobs as Architects Unquestionably at Risk," *Dezeen*, February 13, 2023 – https://www.dezeen.com/2023/02/13/ai-architecture-jobs-risk-neil-leach-opinion/, accessed April 5, 2023.
23. Ibid.
24. Ibid.
25. Beatriz Colomina and Mark Wigley, *Are We Human?: The Archaeology of Design* (Zurich: Lars Müller Publishers, 2016).
26. Jonathan Crary, *24/7: Late Capitalism and the Ends of Sleep* (London, New York: Verso, 2014).
27. Eduardo Porter, "Jobs Threatened by Machines: A Once 'Stupid' Concern Gains Respect," *The New York Times*, June 8, 2016 – https://www.nytimes.com/2016/06/08/business/economy/threatened-by-machines-a-once-stupid-concern-gains-respect.html, accessed August 4, 2023.
28. Joseph Williams, "Chess Computers Move Toward Mastery Over Humans," *The New York Times*, September 23, 1980, C1.
29. Paul Hoffman, "Retooling Machine and Man for Next Big Chess Faceoff," *The New York Times*, January 21, 2003, F1.
30. Ibid.
31. Steve Strogatz, "One Giant Step for a Chess Playing Machine," *The New York Times*, December 26, 2018.
32. Ibid.
33. Johnnie Rios, "What Did Magnus Carlsen Learn from the Most Advanced Chess Engine?" *Medium*, May 24, 2023 – https://medium.com/@jr4216232/what-did-magnus-carlsen-learn-from-the-most-advanced-chess-engine-d8907c47e091.
34. Ibid.
35. "Magnus Carlsen: Greatest Chess Player of All Time," YouTube video, posted by "Lex Fridman," August 27, 2022 – https://youtu.be/0ZO28NtkwwQ.
36. "AlphaZero vs Stockfish Analysis: Neural Nakamura Part 1," YouTube video, posted by

"GMHikaru," September 4, 2019 – https://youtu.be/W1LehcujmDY.
37. Rios, "What Did Magnus Carlsen Learn."
38. Nakamura, "AlphaZero vs Stockfish Analysis."
39. Ludwig Wittgenstein, *Philosophical Investigations* (Hoboken NJ, USA: John Wiley & Sons, 2010).
40. Fridman, "Magnus Carlsen."
41. Ibid.
42. Ibid.
43. Leonardo Espinosa Leal, Anthony Chapman, and Magnus Westerlund, "Reinforcement Learning for Extended Reality: Designing Self-Play Scenarios" (proceedings of the 52nd Hawaii International Conference on System Sciences, Grand Wailea HA, USA, January 2019) – http://hdl.handle.net/10125/59456.
44. Christoph Reinhart et al., "Umi – An Urban Simulation Environment for Building Energy Use, Daylighting and Walkability," *Proceedings of BS 2013: 13th Conference of the International Building Performance Simulation Association* (2013) – doi: 10.26868/25222708.2013.1404.
45. Rick Titulaer, Alexandros Christodoulou, and Matthew Vola, "Multi-Objective Parametric Massing Optimisation for a Sustainable and Healthy Urban Living Environment," in *Impact: Design With All Senses: Proceedings of the Design Modelling Symposium, Berlin 2019*, ed. Christoph Gengnagel et al. (New York: Springer International Publishing, 2020), 737–751; Dogan, Timur, "Procedures for Automated Building Energy Model Production for Urban and Early Design" (PhD diss., Massachusetts Institute of Technology, 2015).
46. Jaekun Lim and Riyozo Ooka, "Building Arrangement Optimization for Urban Ventilation Potential Using Genetic Algorithm and CFD Simulation," *ISARC. Proceedings of the International Symposium on Automation and Robotics in Construction* 31 (2014): 1.
47. G. Z. Brown et al. "A Building Massing Intelligent Design Tool" (report no. CONF-9208196-1, InterSymp-92, Baden-Baden, Ger., August 1992).
48. Nayeri, "How A.I. Is Helping Architects."
49. Giles, "The GANfather."
50. Jaron Lanier, "The Myth of AI," *The Edge*, November 14, 2014 – https://www.edge.org/conversation/jaron_lanier-the-myth-of-ai, accessed October 25, 2023.
51. George Johnson, "Undiscovered Bach? No, a Computer Wrote It," *The New York Times*, November 11, 1997, F1.
52. "Artificial Intelligence: Go Master Lee Se-Dol Wins against Alphago Program," *BBC News Online*, March 13, 2016 – https://www.bbc.com/news/technology-35797102, retrieved August 13, 2023.
53. Blake Brittain, "AI-Generated Art Cannot Receive Copyrights, US Court Says," *Reuters*, August 21, 2023 – https://www.reuters.com/legal/ai-generated-art-cannot-receive-copyrights-us-court-says-2023-08-21/, retrieved September 8, 2023.
54. Nancy Green, *Ready-To-Wear and Ready-To-Work: A Century of Industry and Immigrants in Paris and New York* (Durham NC, USA: Duke University Press, 1997).

Credits

Figure 1: Midjourney generated imagery, prompted by Cem S. Kayatekin.

Cem S. Kayatekin is an Assistant Professor at IE School of Architecture and Design (Segovia, Spain). He holds a BArch from Auburn University, an MArch from Cornell University, and a PhD in Architecture from the University of Oregon, completed under the tutelage of Howard Davis. His teaching and research interests are focused on how the physical, socioeconomic, and ecological dynamics of a city operate and overlap.
E-mail: Cem.Kayatekin@ie.edu

June Aoun is a graduate student at IE School of Architecture and Design. She is interested in the pairing of advanced technologies combined with low-tech processes to develop solutions that are creative, sustainable, and regenerative. Her undergraduate thesis was focused on the topic of artificial intelligence and the architectural profession.
E-mail: june.aoun@student.ie.edu

Yusuf Sühan Bozkurt is a fifth-year architecture student at IE School of Architecture and Design. He is currently the creative director at *Prologue* and *Twist* magazines. He is interested in the subject of process, whether it is architecture, graphic design, or magazine-making. E-mail: ybozkurt.ieu2019@student.ie.edu

Daphné Fournel is a fifth-year architecture student at IE School of Architecture and Design. Her current interests include ecological/sustainable architecture, archaeology, and the role of sociocultural dynamics in architecture. E-mail: dfournel.ieu2019@student.ie.edu