

ARCHITECTURE MASTER THESIS

ARCHITECT AS CURATOR: NAVIGATING INTELLIGENCE DATA, AND DESIGN

POLYTECHNIC UNIVERSITY OF TURIN

ARCHITECTURE CONSTRUCTION CITY

STUDENT-RUSTAM MURADOV

ADVISOR-DAVIDE TOMMASO FERRANDO

CO-ADVISOR-GIOVANNI CORBELLINI

JULY 2025

CONTENTS

PAGE 06-14

- 1.0 INTRODUCTION
- 1.1 UNDERSTANDING ARTIFICIAL INTELLIGENCE:: DEFINITIONS AND DEBATES
- 1.2 CLASSIFICATIONS
- 1.3 DIVERSE PATHS TO INTELLIGENCE
- 1.4 GANS AND THE ROLE OF AI IN ARCHITECTURAL DESIGN

PAGE 15-23

- 2.0 AI AND HISTORICAL CONTEXT
- 2.1 ORIGINS AND COMPUTATIONAL THINKING
- 2.2 DEEP BLUE VS. KASPAROV
- 2.3 VISIBILITY, HYPE, AND THE MISUSE OF AI

PAGE 24-43

- 3 HISTORICAL GENEALOGY OF AUTOMATION AND DESIGN: FROM EARLY CYBERNETICS TO THE DIGITAL TURN
- 3.1 THE FIRST “FALSE START” (1960S–1970S): SKETCHPAD, EARLY CAD, CYBERNETIC DISILLUSIONS, AND THE EMERGENCE OF THE AI WINTER
- 3.2. REDISCOVERY THROUGH DRAWING (1980S–1990S)
- 3.3. THE FIRST DIGITAL TURN (EARLY 1990S–2000)
- 3.4. MASS CUSTOMIZATION AND THE RECODING OF FORM (LATE 1990S)
- 3.5. THE SECOND DIGITAL TURN (2000S–2010S)
- 3.6. DISCRETE COMPUTATION AND THE POSTHUMAN AESTHETIC (2010S)

PAGE 44-57

- 4 FROM GLIMMERS TO REALITIES
- 4.1 DEEPHIMMELB(L)AU
- 4.2 MACHINE HALLUCINATION
- 4.2 THOM MAYNE AND MORPHOSIS
- 4.3 EXPANDING CONVERSATIONS ON AI:

INFLUENTIAL PRACTITIONERS
4.4 AI AND FABRICATION

PAGE 58-65

- 5.0 SHAPING ARCHITECTURAL THINKING
- 5.1 LOST IN TRANSLATION: AI, ARCHITECTURE, AND THE PROBLEM OF AUTOMATION
- 5.2 GENERATIVE AI IN ART AND MEDIA
- 5.3 PREDICTIONS FOR THE FUTURE
- 5.4 INTERVIEW 3XN DIGITAL / AI TEAM

PAGE 66-71

- 6.0 DETAILED METHODOLOGY
- 6.1 ESTABLISHING SELECTION CRITERIA FROM PRIOR RESEARCH
- 6.2 COMPARATIVE ANALYSIS
- 6.3 ACKNOWLEDGING ALTERNATIVE WORKFLOWS
- 6.4 CONCLUSION ON TOOL SELECTION

PAGE 72-77

- 7.0 INTRODUCTION TO THE DESIGN WORKFLOW
- 7.1 INITIAL SETUP: DEFINING CONTEXT AND CONSTRAINTS
- 7.2 ITERATIVE GENERATION AND REFINEMENT
- 7.3 ACHIEVING DETAIL AND TRANSITIONING TO PHOTOREALISM
- 7.4 BRIDGING THE GAP: FROM 2D IMAGE TO 3D ARCHITECTURAL MODEL
- 7.5 CONCLUSION: AN AI-AUGMENTED WORKFLOW

PAGE 78-80

8.0 PROMPT ENGINEERING AS DESIGN METHODOLOGY

PAGE 81-84

9.0 CONTEXTUALIZING THE AKITA COMPETITION: TRADITION, CRAFT, AND

PERSONAL ENGAGEMENT
9.1 SITE, PROGRAM, AND TECHNICAL PARAMETERS

PAGE 85-93

- 10.0 CASE STUDIES / UCCA CLAY MUSEUM
- 10.1 CASE STUDIES / YUSUHARA WOODEN BRIDGE MUSEUM
- 10.2 CASE STUDIES / WISDOME STOCKHOLM
- 10.3 CASE STUDIES / SWATCH & OMEGA CAMPUS

PAGE 94-95

11 AKITA IN CONTEXT: LANDSCAPE, HISTORY, AND CONTEMPORARY CONDITION

PAGE 96-105

12.0 MAPS AND URBAN STUDIES

PAGE 106-153

- 13.0 REPRESENTING THE GENERATIVE SEQUENCE
- 13.1 ARCHITECTURAL DRAWINGS AND SPATIAL ORGANIZATION

PAGE 154-155

14.0 CONCLUSION: CRITICAL REFLECTIONS ON GENERATIVE DESIGN AND ARCHITECTURAL CURATORSHIP

PAGE 156-165

THE TABLE OF THE FIGURES
BIBLIOGRAPHY

THESIS OBJECTIVES

1. To contextualize AI's emergence within the broader historical genealogy of automation in architecture. Through a multi-chapter review of key theoretical discourses—from cybernetics and symbolic AI to GANs and diffusion models—the research situates contemporary AI tools within architecture's long-standing engagement with computational thinking.

2. To critically investigate the integration of generative artificial intelligence (AI) at the very outset of architectural ideation. The research considers image-based generative models as speculative tools in the conceptual phase of design. It evaluates the affordances and constraints they provide relative to the existing workflows of architecture.

3. To evaluate the architectural relevance of AI generative images through a design experiment. The thesis produced conceptual outputs with Midjourney for a competition proposal (a crafts museum in Akita, Japan) testing the extent to which AI-generated images can meaningfully inform spatial, structural, and tectonic design development.

4. To consider the role of the architect as a curator in AI-mediated workflows. By engaging with the interpretive work of transforming AI-generated images into architectural drawings and a 3D model, the thesis reframes authorship in computational design, emphasizing the shift from sole creator to editorial and strategic curator.

5. To identify technical, procedural, and epistemological gaps in the current generative AI design workflow. The thesis documents the practical challenges involved in converting AI outputs into coherent architectural representations, highlighting issues of scale ambiguity, image fidelity, iterative instability, and lack of spatial logic.

6. To speculate on future directions for integrating AI into architectural practice.

The thesis proposes that with more structured prompt systems, better 3D integration, and improved dataset curation, generative AI could evolve from a visual provocation tool into a more robust design assistant—potentially enabling simultaneous generation of spatial plans, sections, elevations, and 3D massing.

1.0 INTRODUCTION

Architecture, a profession historically straddling the line of precision and provocation, is now experiencing a new paradigm shift through artificial intelligence. Over fifty years ago, architectural production transitioned from hand drawing to computer-aided design (CAD), introducing speed and iteration into products and rapidly developing more complex formal experimentation through software-based tools. Today's shift is at a deeper conceptual level and harder to define. While text-to-image machine learning (ML) models such as DALL-E and Midjourney may be making much news, the broader impact of AI on architecture is much more than producing images and automating complex geometry. The visual data it can provide is increasingly improved and upscaled. Lev Manovich (2023) contextualizes this generational shift similarly to previous photography and linear perspective paradigm changes. There is little doubt that we are at a crucial juncture in determining how architects define, create, and understand design.

This technological turn is not without precedent. For decades, architects have explored algorithmic strategies—fractals, shape grammars, agent-based systems—that pushed the boundaries of analog design and questioned typological norms. Mario Carpo (2023) has traced the evolving logic of software in shaping architectural language, while Yona Friedman (as cited in Negroponte, 1975) cautioned that no machine could be imagined without the presence of an intelligent observer—an early reminder that computation and human agency are fundamentally intertwined.

By the early 2000s, with the integration of parametric design tools as part of Building

Information Modeling (BIM) systems, established design skills had advanced the ability to control and manipulate architectural data in a way that could be rapidly achieved in real-time through scripting interfaces like Grasshopper or by using parameter driven objects in other software like Revit. Although predefined design skills allow for greater variability of outcomes, these “tools” also impose their logic and limitations, which can further separate designers’ intentions from computational outputs. Current advances in AI promise further efficiencies in even these systems. For instance, AI-assisted scripting will allow a designer with basic Python knowledge to build custom plug-ins for software like Rhino or Revit by automating the algorithmic processes typically undertaken by the designer. Machine learning models can now create efficient responses to briefed design intentions or produce “hallucinated” alternatives that place them in competition with existing workflows and definitions of authorship. Neil Leach (2021) refers to this form of developing design tools as “potentially disruptive,” while others are skeptical of their place in creative autonomy or their ability to interpret across design aspects.

Nicholas Negroponte (1975) proposed speculative visions decades ago, such as “soft architecture machines,” which now seem less hypothetical as computational agents begin to act as collaborators rather than instruments. This thesis does not assume a celebratory stance. Instead, it seeks to investigate how AI tools are currently being integrated into architectural workflows and whether and to what extent they can meaningfully contribute to design processes.

1.1 UNDERSTANDING ARTIFICIAL INTELLIGENCE DEFINITIONS AND DEBATES

Artificial Intelligence (AI) is the domain of computer science that builds systems that can carry out tasks that typically require human-level intelligence. Tasks, depending on the definition, can include processing and synthesizing information, pattern recognition, decision making, and being able to learn from, and adapt to, data (Boden, 2016; Kelleher, 2019).

Currently, most AI applications are what is referred to as “narrow AI”, or “weak AI”, which is to say that narrow or weak AI functions in a specific way and typically has one application, such as voice recognition, image classification, or translation, and cannot think consciously or generalize (Charles, 1996). Narrow AI analyses framed problems using learned and predetermined patterns of behavior. Siri, Google Translate, and medical diagnostic tools are examples of narrow AI. In contrast, “strong AI,” or Artificial General Intelligence (AGI), refers to a theoretical machine that can learn, understand, and adapt to many different intellectual tasks, mimicking and potentially surpassing human-level cognition (Walsh, 2018). AGI may only ever be speculation and exist in science fiction, as with Agent Smith in *The Matrix* or Ava in *Ex Machina*. However, some philosophers theorize that recent emergent capacity in large language models, such as GPT-3 or GPT-4, may signify progress toward AGI more quickly than previously envisaged (Weinberg, 2020).

According to Stuart Russell (as cited in Ford, 2018), despite AI’s record of taking over particular tasks - such as being better than human experts at chess or Go - it might not be long before it is better than humans at more of the intelligence-related tasks we have previously assigned to them. However, being ‘better than’ does not mean ‘understanding

the task.’ As Mitchell (2019) writes, an AI can be faster and better than a human in completing a task, but it can be faster and better without knowledge or context of where it fits.

This raises a much larger question - should AI be judged by how well it performs those tasks, or should consciousness and understanding be the end goal? There is a vast gulf of opinion on that question. Some authors emphasize functional utility, while others consider consciousness the benchmark of real intelligence.

Finally, human cognition is just one intelligence model. It is not certain that our models capture the same aspects as other types of intelligence, which could be computational, biological, or possibly other types that have yet to be imagined. However, to advance the objectives in this thesis, I will concentrate on AI technologies as they exist today, their potential future directions enabling interaction with design, and how to facilitate a symbiotic relationship with human cognition.

1.2 CLASSIFICATIONS

Artificial intelligence (AI) is often approached as a unified technology but contains many different approaches. As Russakovsky (2016) observes, “AI comes in many shapes and sizes,” with each approach representing a different way to imitate or replicate some intelligent behavior. One prominent historical distinction in AI relates to “classical” or “symbolic” AI—sometimes called “Good Old-Fashioned AI” (GOFAI)—and newer data-driven approaches such as machine learning (Nilsson, 2009).

Classical AI systems operate according to rules and logic. Thus, classical AI systems execute tasks only if they are told how to do so; they are designed to imitate the reasoning that takes place symbolically. In contrast, machine learning does not have logic-based determinations. Machine learning allows algorithms to learn from data to improve tasks over time without explicit reprogramming or logic (Mitchell, 1997). Within the machine learning umbrella, deep learning is an advancement. It utilizes artificial neural networks with multiple layers to extract higher and higher degrees of abstraction from raw data (Goodfellow, Bengio, & Courville, 2016). Deep learning methods have accomplished many breakthroughs discussed in the AI realm recently, such as the classification of images, voice recognition, and natural language processing. Thanks to these successful contributions to the scientific field, it is understandable that there is a movement to equate “AI” with deep learning (LeCun, Bengio, & Hinton, 2015).

While these AI systems are highly successful at narrow tasks, they have important limitations. Deep learning models, for example, rely on data and much computing and are task-specific (Dean, 2021). While deep learning

models can generally outperform humans with any stability of their performance on a very narrow task like facial recognition and autonomous driving, they do not represent the general reasoning capacity that is part of human intelligence or is depicted in movies and television shows of intelligent machines (Schmidhuber, 2015).

In short, we can not think of AI as a single field but rather a changing and developing collection of methods. Each method has unique affordances and affordance limitations. All these features are important in determining a definition or pursuing artificial “intelligence.” This understanding is essential for meaningful conversations about the potential for AI or more traditional intelligence in architectural design, and that potential in all fields is unexplored.

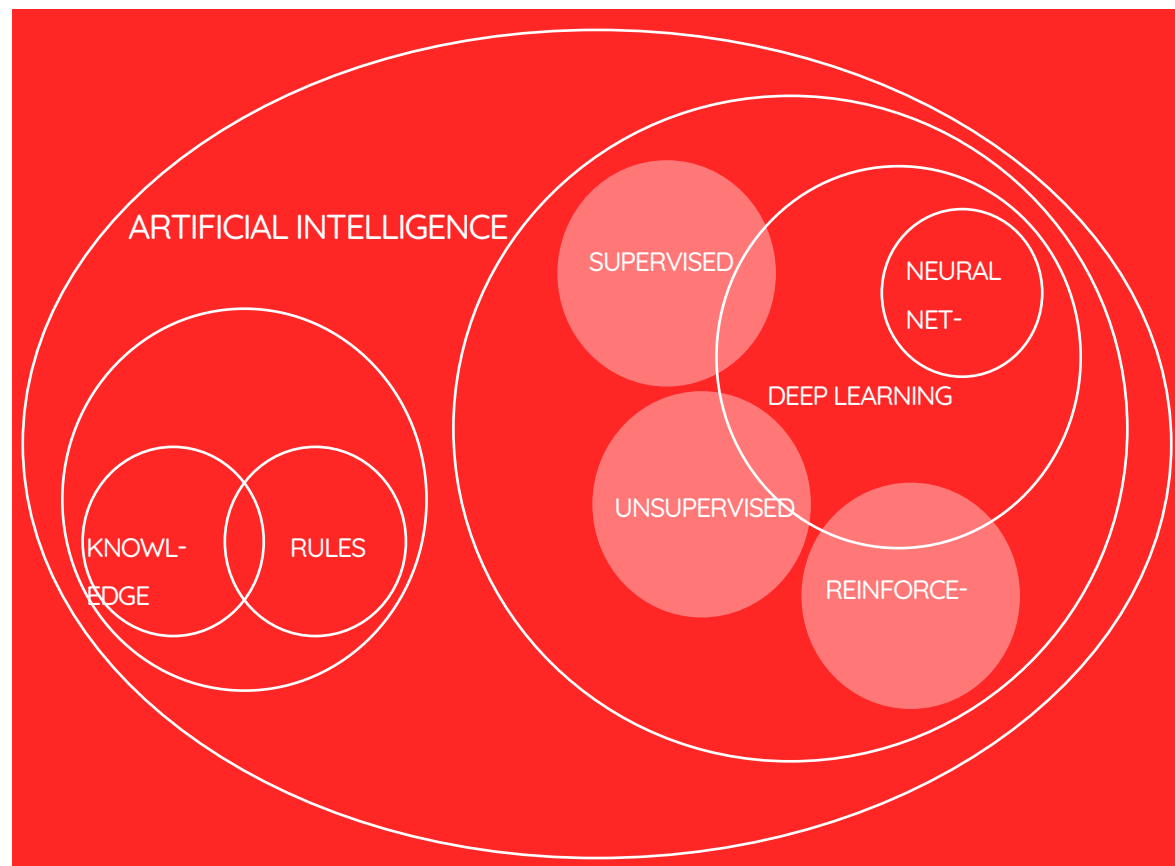


FIGURE 1
Venn diagram showing Artificial Intelligence (AI) encompasses the subsets of Expert Systems, Machine Learning and Deep Learning, in which their methods can be applied to tasks that imitate human decision-making abilities. Reprinted from "Artificial intelli-

1.3 DIVERSE PATHS TO INTELLIGENCE

Machine learning includes various theoretical frameworks and methodologies that explain how machines acquire and use knowledge. Pedro Domingos (2015) identified five core paradigms—or “tribes”—of machine learning: symbolists, connectionists, evolutionaries, Bayesians, and analogizers. Each provides a different explanation of how we might produce intelligent behavior in machines, as each tribe has a distinct contribution to designing AI systems.

Symbolists explain intelligence as the manipulation of symbols and logical rules. This tradition is a remnant of the beginnings of AI research, otherwise known as “A logic-based conception of describing reasoning drives Good Old-Fashioned AI” (GOFAI) and symbolists. Symbolists use mechanisms such as inverse deduction, where systems produce potential causes from known effects, and their goal is to replicate reasoning in a logical, structured, and explainable manner (Nilsson, 2009).

Connectionists focus on artificial neural networks governed by the structure of the human brain. Connectionists learn by manipulating weights between artificial neurons where these neurons are interconnected, enabling machines to learn to recognize patterns of information. This also explains why deep learning techniques rely on successive layers of the neural network representing increasingly higher layers of abstraction in data input (Rumelhart, Hinton & Williams, 1986).

Evolutionaries structure their methods based on scientific traditions found in evolutionary biology. By applying principles such as natural selection and genetic inheritance, the models evolve potential solutions by mimicking reproduction, mutations, and selection, which allows and results in the iterative

improvement of program performance over time (Koza, 1992). Bayesian considers learning to be a type of probabilistic inference. They emphasize the importance of prior knowledge and updating the weighting of possibilities or outcomes using new or added data. They also build AI systems that learn efficiently given incomplete state and uncertainty (Murphy, 2012).

Analogizers rely on similar observations. Their principal tactic is assessing whether new situations are similar to other previous ones, trusting that like problems share like solutions. Analogizers also include a case-based reasoning mechanism, where systems draw from stored system examples rather than say principles or rules (Aamodt & Plaza, 1994). Characterizing that symbolic processes were the predominant models within the early tradition of AI; we see the substantial emergence of deep learning and the availability of high-quality big data and computing resources elevate connectionism as one of the primary paradigms of AI machine learning (Jordan & Mitchell, 2015).

Each one of these approaches or “tribes” is a clear discussion. They introduce different fundamental assumptions about learning that can be extended and contrasted with one another, and each continues to offer some form of theoretical investigation and practical alternatives. The field of AI does not evolve towards a universal model of intelligence; instead, the methods and the appeals against one another drive the conversations and the mapping of the field. Understanding their differences is important for considering the implications for the possibilities and limitations of AI technologies both broadly and intently.




FIGURE 2
Stages of recognition and transformation of a cross. Negroponte, N. (1975). Soft architecture machines. The MIT Press.

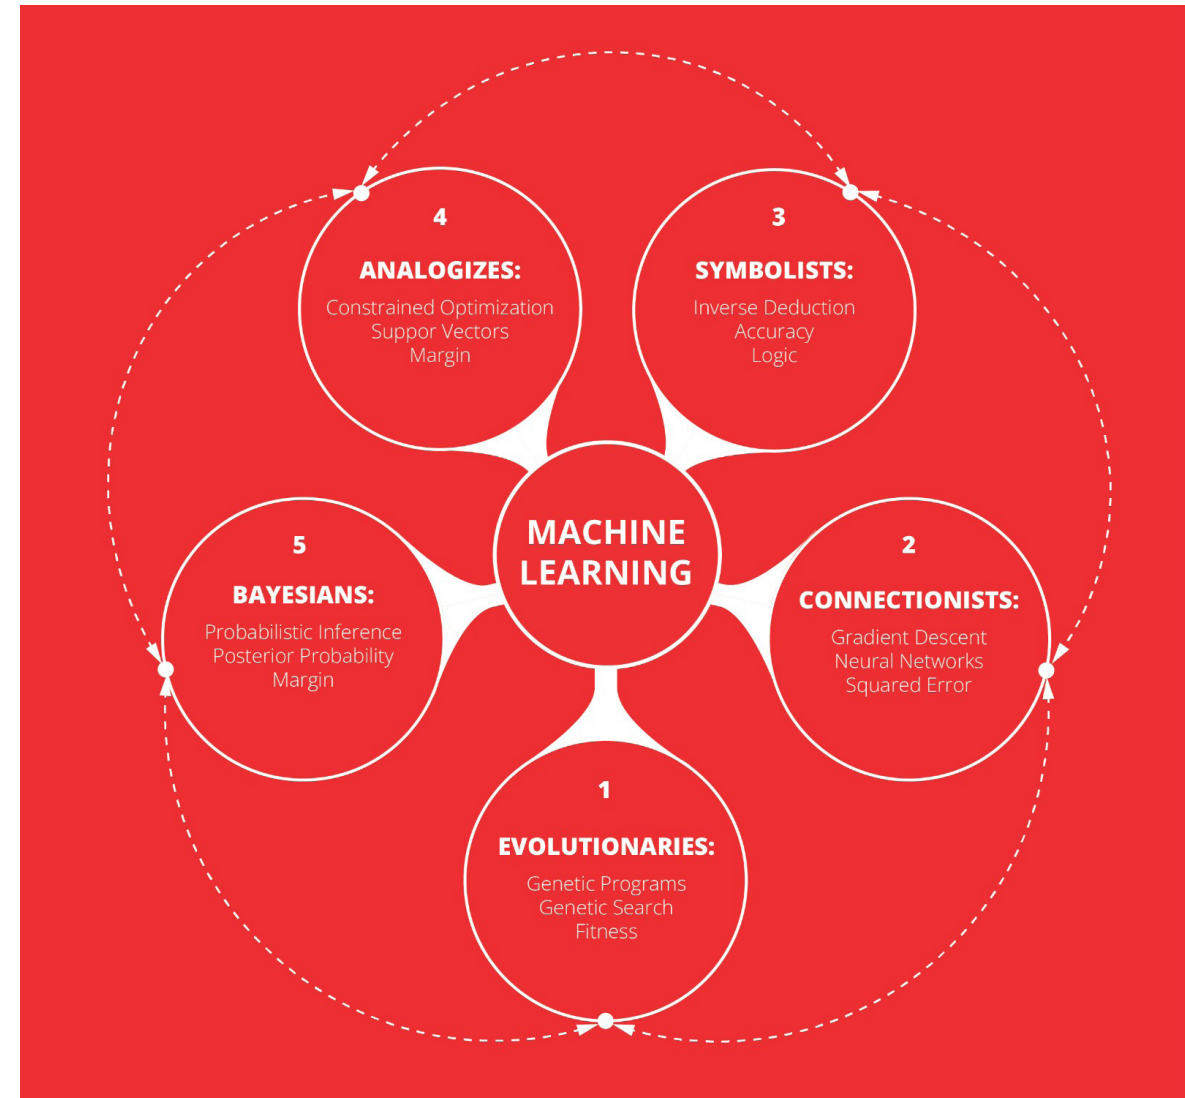


FIGURE 3
AI tribes

1.4 GANS AND THE ROLE OF AI IN ARCHITECTURAL DESIGN

Generative Adversarial Networks (GANs) have offered promising and controversial tools within deep learning, particularly for architecture. These processes generate new images from large datasets of images on the assumption that derivations from those images (synthesizing them or transforming them) can occur. The outputs of GANs blur the distinction between imitation and invention and, with this, challenge settled concerns about authorship and creativity.

Mario Carpo (2023) situates this within a more extended narrative about computational capacity and capability. What was once disregarded as too impractical or too speculative AI models has developed with unprecedented speed because memory and processing become increasingly powerful and available. He states that this has rejuvenated an earlier objective in computer science: to build a “general problem-solving machine” to do significant design work without human intervention. This statement must be held in tension with Carpo’s larger historical framing—aspects contradicted by how the role of computation continues to evolve from dependent to possibly generative.

For most of the late 20th century, computers were “stupendous but always secondary” architectural instruments. Computers contained the architect in calculation, visualization, and production, allowing him or her to exploit new forms, experimenting with “new” means (if only a means few understood), and not in any essential aspect of authorship. Today, AI models like GANs have transcended a position of mere ‘sum.’ New research undertaken by Matias del Campo and Stanislas Chaillou suggests that GANs can generate

new architectural images and execute style transfer. This progression raises questions of originality and agency (Carpo, 2023).

Despite the potential for new visual languages for architects and new workflows, Carpo conveys a sense of cynicism about the real implications of this knowledge or potential. He allows that these systems have developed “automated visual imitation,” which he acknowledges as a significant technical milestone, but cautions against allowing cameras or machines to think; he asks if there are architects who want to “borrow someone else’s intelligence (let alone if it is artificial)” to make derivative architecture.

Carpo also notes a structural failing in the current uses of AI, that the vast majority of machine learning applications are iterative optimization, which operates under the premise of measurable outcomes. This is an issue for architectural work evaluated, qualitatively and contested. Carpo (2023) states, “No one has found a consensual metric to assess values in architectural design to date.”


This mismatch represents a fundamental tension between algorithmic logic and architectural reasoning. In the context of these debates, this thesis takes an exploratory approach to new experimental possibilities. It investigates GANs not as substitutes for design thinking but as experimental collaborators potentially informing, extending, or dilating existing architectural processes. Instead of fixating on the isolated outcomes or speculative claims, this inquiry considers how the processes function in the architectural workflow and whether they dilute authorship, creativity, and/or increase criticality.

2.0 AI AND HISTORICAL CONTEXT


The idea of artificial intelligence is culturally and intellectually deep-rooted. Historians rightfully note its origin in modern computing and ancient myths and stories demonstrating humans’ long-standing and evolving interest in artificial life. The Golem of Jewish folklore and Mary Shelley’s *Frankenstein* (1818) are examples of early imaginative attempts to form non-human actors capable of independent action (Bilski, 2019).

AI additionally owes its development to early developments in mathematics and early computing. In 1854, George Boole began to change the world by using algebraic logic in *An Investigation of the Laws of Thought* to produce Boolean operations; this is significant since much of the advances in digital computations are based on Boolean logic. Charles Babbage began work on the first mechanical computing devices—the Difference Engine and the Analytical Engine—in the early 1830s and 1840s. Ada Lovelace worked with Babbage and, in 1843, is credited with writing the first algorithm for a machine—one of the earliest computer programmers (Fuegi & Francis, 2003).


Alan Turing developed these foundations further in the 1930s. In 1936, he described an idea called a universal computing machine (now known as a Turing Machine), which is capable of simulating any mathematical computation. He also published *Computing Machinery and Intelligence* in 1950, in which he introduced the term “Imitation Game” (now known as the Turing Test) as a criterion for machine intelligence. Without question, Alan Turing’s works were significant in formalizing AI as a scientific field.



1



2



4



3




FIGURE 3
Control room (Architecture Machine Group). Reprinted from *Soft Architecture Machines* (p. 158), by N. Negroponte, 1975, The MIT Press.



FIGURE 4
Institute for Advanced Study. (ca. 1930s). image of Alan Turing walking along the street [Photograph]. Elaine Negroponte collection on John von Neumann and Alan Turing, Shelby White and Leon Levy Archives Center, Institute for Advanced Study

2.1 ORIGINS AND COMPUTATIONAL THINKING

In 1937, British mathematician Alan Turing published the paper *On Computable Numbers, with an Application to the Entscheidungsproblem*, laying the theoretical foundation for modern computing. In it, he introduced the concept of the “a-machine” or “automatic machine,” anticipating a future in which computation would be entirely mechanized (Turing, 1937). After completing his PhD at Princeton in 1938, Turing began working in 1939 at Bletchley Park, a wartime codebreaking center in Buckinghamshire. He had a crucial role in decrypting messages produced by Germany’s military cipher machine, Enigma. In response, Turing developed a counter-machine—the “bombe”—first installed in early 1940 and became instrumental in turning Bletchley Park into an efficient decryption hub (Copeland, 2014).

Turing’s theoretical contributions remained equally important. In 1950, he published *Computing Machinery and Intelligence*. He speculated that by 2000, it should be possible to program a machine so that an “average interrogator” would have less than a 70 percent chance of telling whether, after five minutes, they were conversing with a machine or a human being. This speculation has subsequently become the basis for the Turing Test (Epstein et al., 2008). Turing began with considerations of how machines might execute instructions but eventually morphed into a philosophical discussion: Do machines actually think—and if so, when do they stop being instruments?



FIGURE 5
Nurnberg, W. (Photographer). (1958). Engineers with the early DEUCE computer at English Electric

2.2 DEEP BLUE VS. KASPAROV

Chess has long intrigued scientists and mathematicians, including figures like Charles Babbage and Alan Turing, who saw the game as a way to explore the mechanics of human reasoning. In 1990, Ray Kurzweil predicted that by 2000, a computer would defeat the world chess champion (Diamandis, 2018). This prediction was fulfilled in 1997, when IBM's chess engine Deep Blue defeated Garry Kasparov, the reigning world champion, in a six-game match-winning two games, drawing three and losing one (Newborn, 2011). This event marked a turning point in public perception of artificial intelligence. As Wieder (as cited in McPhee, Baker & Siemaszko, 2015) noted, it was the first time a reigning champion had lost to a machine under standard tournament conditions. It received considerable media coverage and was seen as an example of AI's increasing ability to compete against human abilities in previously merely cognitive areas.

Regardless, the match's outcome raised issues. Kasparov criticized the match's organization, claiming that IBM was a player, organizer, and sponsor, placing Kasparov in a hostile environment (Kasparov, 1997). Researchers Jonathan Schaeffer and Aske Plaat (1997) pointed out that Kasparov was disadvantaged because he could not access any of Deep Blue's prior games, which limited his preparation. Further, his usual anti-machine game (which exploits the assumptions and weaknesses of previous machines) was ineffective against Deep Blue.

Even though some commentators doubted whether Deep Blue could meaningfully play chess or challenge grandmasters consistently, Kasparov still felt this was a watershed

moment. Kasparov said he felt "a new kind of intelligence" and viewed this moment as an important demonstration of machinery's creativity (Kasparov, 1996). He later wrote about the defeat in Deep Thinking as a manifestation of human progress, given that if every un-footed step forward that machines make can be framed as a form of partnership rather than competition, that must represent progress towards a better society (Kasparov, 2017). Today, AI is at the center of chess training and ongoing analysis. At the same time, Magnus Carlsen has recently remarked how AI has democratized elite chess learning, such that now anyone can radically improve their game anywhere in the world (Euronews Next, 2024). AI engines like Stockfish and AlphaZero deploy incredible depth of positional analysis and creativity in their chess moves, thus influencing, altering, and changing the nature of grandmaster-level chess preparation and practice and overcoming privileged biases of elite players. These transitions are examples of the power of AI computing and the opportunities to increase and augment the value and authority of human expertise - insights particularly relevant for domains like architecture; human-machine teams are changing design thinking and authorship, just as they are improving chess.



FIGURE 6
Ewalt, D. M. (2011, May 3). Kasparov vs. Deep Blue [Photograph by George Widman/AP]. Forbes

2.3 INVERSE CAMOUFLAGE AND AI

The history of artificial intelligence is not only technical but also cultural, shaped by cycles of public fascination and misunderstanding. In many cases, AI operates invisibly—powering search engines, social media algorithms, and spam filters—yet remains largely obscure to the average user. Even researchers working with machine learning may not fully grasp the inner workings of complex models. The public's awareness usually increases in response to something that stands out, such as Bill Gates's or Mark Zuckerberg's reports. In one instance, in 2016, Go champion Lee Sedol played a match publicized worldwide against Google's AlphaGo. That match received considerable international exposure (Hassabis & Hui, 2017), but subsequent reports about AlphaZero's self-taught performance were hardly seen. The difference indicates how the public is only interested in AI once it has 'proven' superiority over humans—and loses interest when the claim now lacks the same excitement and validates other narratives.

Other organizations take advantage of the salesmanship of capitalism that surrounds AI without any real ingenuity. This act - which may be considered "inverse camouflage" - is to sell the concept of using AI to generate sales and PR without actually using AI. For example, in 2017, the humanoid robot Sophia, developed by Hanson Robotics, was granted symbolic citizenship in Saudi Arabia and recognized as an "Innovation Champion" by the UN (Weller, 2017; UNDP, 2017). While some critics began to question the extent to which Sophia's capabilities may have been inflated and that her public persona was soft marketing rather than technology (Sinapayen, 2018).

This pattern is not limited to consumer-facing technologies. Several firms have adopted AI-themed branding within architecture without actively incorporating AI into their design processes. In one case, an executive admitted during a conference that the "AI" label boosted visibility more than described technical expertise (Mallard, 2019). This underscores how the label "AI" can function as cultural capital—invoking innovation, futurism, or legitimacy—regardless of usage. Understanding these dynamics is essential for distinguishing between meaningful AI integration and superficial claims. For architectural practice, where design intelligence, authorship, and innovation are deeply valued, critically evaluating the tools and the discourse surrounding them is crucial.

This chapter has traced some of AI's symbolic and strategic uses—real and exaggerated—that shape how the technology enters public consciousness and professional fields. The following sections will shift from narrative and perception to applied practice, examining key experiments and frameworks directly influencing architectural workflows today.



FIGURE 7
Riccio, T. (Photographer). (2016, March). David Hanson and Sophia during a 60 Minutes interview with Charlie Rose, New York City

3 HISTORICAL GENEALOGY OF AUTOMATION AND DESIGN: FROM EARLY CYBERNETICS TO THE DIGITAL TURN

This chapter concludes a three-part historical genealogy of architecture's digital transformation, following the framework defined by Mario Carpo (2017). These chapters are intended to critically map architectural developments across successive technological phases rather than as original periodizations. They offer a structured lens to examine how computational logic has interacted with architectural theory, aesthetics, and practice. Following the path from initial cybernetics through the first and second digital turns, this structure offers a critical view for reading significant architectural thought and practice changes.

In architecture, the genealogy of computational thought is a complicated mess of technological hope, conceptual experimentation, and unfinished critical tensions. The period from the 1940s to the 1960s, when digital technology first emerged through early interactions of artificial intelligence (AI) and cybernetics, established contexts for later practices with automation and design.

The development of computation began between 1943 and 1946 when John Presper Eckert and John Mauchly at the University of Pennsylvania were developing the Electronic Numerical Integrator and Computer (ENIAC). While ENIAC is regarded as a key technological milestone and much fanfare was made regarding its capability, it was Nazi the stored-program architecture since any other programming required manual re-wiring (Martin, 1995). Nevertheless, ENIAC captivated the public imagination as a "giant brain," invoking what appeared to be a new frontier where devices might one day feel, think, and reason like human beings.

This divergence between technical reality and symbolic projection foreshadowed a recurring tension in the discourse surrounding AI and architecture: the tendency to equate computational complexity with cognitive agency. Parallel to these technological developments, Norbert Wiener's articulation of cybernetics in.

Cybernetics Or Control and Communication in the Animal and the Machine (1948) provided a theoretical framework for understanding systems behavior through feedback and regulation (Mindell, 2000).

Wiener argued that mechanical and biological systems could produce goal-directed behavior based on continuous adaptive feedback loops, thus setting the stage for a system-based principle that would later resonate with architectural theory. Wiener also brought a level of caution to the ethical ambiguity of cybernetic systems, in that they could be used as a system of surveillance and social control—this lingering ambivalence plants cybernetics not only as a means of empowerment but as a site for political and ethical struggle. Architectural engagements with cybernetics received vigorous articulation through Cedric Price's Fun Palace project in 1961, developed in collaboration with Joan Littlewood and cyberneticist Gordon Pask. The Fun Palace was imagined as an open-ended infrastructure for user behavior that operationalized cybernetics by proposing a mutable and adaptable field of space rather than an object (Herdt, 2021). Pask's research experiments with adaptive learning complemented Price's aspiration to create spaces adapting to potential user engagement. However, upon reflection much later, Price demonstrated a developing skepticism towards technical systems that used

the more adaptive over social engagement. Price's concerns are attached to an ongoing architectural query: Does technical flexibility alone replace deliberately meaningful design intervention framed by larger cultural and political realities?

As architectural theorists wrestled with the implications of adaptability and control, concurrent developments in AI research sought a computational solution for formalizing human cognition. In "Steps Toward Artificial Intelligence" (1960), Marvin Minsky described a trajectory of symbolic AI predicated on heuristic search, pattern recognition, and structured knowledge representation. Symbolic AI was able to claim that intelligence could be reduced to the manipulation of abstract symbols abiding by formal rules (Minsky, 1960). This reduction produced early algorithmic models of reasoning. However, otherwise, it was structurally inadequate to contend with the tacit knowledge, contextual sensitivity, and intuitive judgment called for in complex cognitive work, including design (AORA and Symbolic AI, 2020). In each case, the limitations were not tangential but fundamental. Architectural creativity—the condition of ambiguity, interpretive richness, and context negotiation—refused to fit into the constraining frames of symbolic reasoning. The early shortcomings of symbolic AI thus anticipated the broader difficulties faced in applying computational models to domains requiring more than procedural rule-following.

Using cybernetic and artificial intelligence theories, Nicholas Negroponte attempted to convert these concepts into a real architectural practice through *The Architecture Machine* (1970). Negroponte proposed a model of

human-machine collaboration in which computational systems could 'learn' from human users to participate in generative activities of design (Negroponte, 1970). Negroponte's vision attempted to entice machines not merely to participate in architecture practice but to let machines build conversations between human users and computational power. However, the Architecture Machine Group's first experiments to argue with the notion of the architecture machine in practice experienced significant epistemological and practical challenges. The learning algorithms used in the experiments could not adapt to the ambiguity, contradiction, and nuance of layered intention that underpins human design activity. Instead of eliciting real collaboration, these machines produced mechanical adaptations without criticality or interpretive richness. Cristian Negroponte's project represented a fundamental tension: while computational systems could improve specific operational capabilities, they did not replicate the interpretive labor of acknowledging the contributions of the human architect.

The period between the 1940s and 1960s produced an ambitious but ambiguous foundation for computational approaches to architecture. Initial technological successes, philosophical positions, and experiments contributed to a horizon of opportunities in which automation and adaptation were recurring themes. However, at each instantiation, critical limitations – whether technical, conceptual, or ethical – emerged to interrupt the seamless absorption of machines into creative and cultural practices.




FIGURE 8
Mario Carpo, Mark Garcia and Steven Hutt, A short but believable history of the digital turn in architecture, 'Prevalence of Computation in Architectural Design', 2023
© Courtesy of the artists and the Jencks Foundation at The Cosmic House

3.1 THE FIRST “FALSE START” (1960S–1970S)

The early confident optimism surrounding the use of computation in architecture experienced a period of reflection and critical evaluation in the 1960s and 1970s. Technological experimentation continued, but the conceptual and practical limitations of early computer-assisted design (CAD) systems, the limitations of cybernetic models of architecture to represent agency, and the fade of early ambitious work of artificial intelligence (AI) led to what would later be called a ‘false start’ in digital practices in architecture. Ivan Edward Sutherland’s work on Sketchpad is still considered an important but problematic marker.

Sketchpad is often referred to as the first graphical communication system between humans and machines, and it represents a remarkably sophisticated entry point for introducing ideas like constraint, object-oriented ideas, and interactive graphical manipulation with a light pen (Sutherland, 1963). While Sketchpad has some significance as a later historical moment in user interface design and human-computer interaction, by definition, it was of little significance to architectural thought. When Sutherland stated that while the system facilitated different methods of geometric operation, it did not alter the cognitive or conceptual framework of design, the notion of direct visual manipulation and that machine operations would follow human creativity was naïve. It reflected the ongoing contradiction, the disjunction, between the abstract nature of computation and the imaginings of architecture.

It is helpful to consider that early usage of CAD in the architectural field echoed these limitations in more practical usage. As Allen and Kouppas (2012) described, the initial CAD

systems primarily automated the workflow of two-dimensional drafting, enhancing technical production at the expense of authentic design ideation of creative productions. CAD improved operational efficiency but risked reinforcing technocratic tendencies in architectural workflows by subordinating creative experimentation to procedural regularities and standardization. Ultimately, the faith in digital tools to deliver freedom for architectural creativity remained largely unrealized, with their use revealing a crux between technological developments and design expectations.

At the same time, AI, more broadly, was in a dramatic decline, with Ballatore and Natale (2023) describing how the origins of AI’s failure in its early stages meaningfully converged to what became known as ‘AI Winter’ in the 1970s. These failures were recognized as primarily structural failures of the symbolic paradigm itself—pluralities resulting from an overreliance on formal rules and an inability to respond to the complexities of the real-world environment. Frustrated by a seeming lack of capacity to fulfill expectations, funders retreated from supporting AI research, and the field suffered shrinkage once again. However, as Ballatore and Natale noted, the myths around artificial intelligence remained prevalent (or, the oscillation between hyperbole and disillusionment maintained some relevance).

In architecture, these larger sets of conditions reinforced critical perspectives that began exploring the limitations of computational models. The mechanistic metaphor that framed early CAD systems and AI systems failed to meaningfully represent the interpretative, situated, and culturally embedded nature of design thinking. Rather than represent or facilitate richer creative practices, digital tools

mostly externalized only those parts of the design context that were abstractly formalizable, leaving the complexities of negotiations around meaning, value, and context, which are meaningful forms of architectural agency, unresolved.

Thus, the 1960s and 1970s did not mark a departure from computational experimentation but a recalibration of expectations. Recognizing the limitations of digital tools began a more tempered and reflexive engagement with technological mediation, allowing for critical inquiry into the relationship between automation and creative authorship. This dynamic would impact architectural discourse in meaningful ways in the following decades.

3.2. REDISCOVERY THROUGH DRAWING (1980S–1990S)

The timeframe from the mid-1980s to the early 1990s is a transition in terms of architecture's relationship with computation, which was marked more by a degree of practicality than a substantial conceptual shift. With the introduction of personal computers and early computer-aided design (CAD) systems becoming more widely available beyond a few research institutions, architecture began to alter its relationship with technology, exploring computational tools for operational-level support rather than cognitive-level support.

The IBM PC—1981 and the Apple Macintosh—1984 allowed a certain degree of computation democratization (Herriman, 2022). However, while it allowed for some freedom, that democratization was primarily for administration or representation rather than a change in approach or method for architectural design. The modularity, standardization, and lower affordability provided access to office management and a drafting system as opposed to a new avenue for speculative testing and design.

A case in point is AutoCAD, which Autodesk published in 1982 and portrayed an entirely pragmatic orientation. AutoCAD was designed to replace drafting and thus did not question the epistemological implications of the design medium. As Herriman noted, the early focus of CAD development was to improve efficiency, accuracy, and reproducibility rather than to innovate conceptual approaches. The transfer of drawing-as-acting into a digital format visually reinforced a representational regime with the machine as a draftsman rather than a shaping assistant.

This period thus witnessed a significant shift: computers ceased to be envisioned as “thinking machines” capable of intelligent

problem-solving, as earlier cybernetic and artificial intelligence paradigms had proposed, and were increasingly deployed as “drawing machines,” instrumentalizing design processes within established paradigms of production and representation. The broader cultural context, as highlighted by Anderson and Bianconi (2017) in their curatorial work at the Museum of Modern Art, suggests a systematic tendency to normalize digital tools within pre-existing operational frameworks rather than engaging them as catalysts for epistemological reconfiguration

Architectural pedagogy responded to these technological changes in uneven and sometimes contradictory ways. The “Paperless Studios” initiative piloted at Columbia University's Graduate School of Architecture, Planning and Preservation (GSAPP) under the leadership of Bernard Tschumi in 1988 is an instructive case study. Tschumi wanted to reformulate architectural education in a way that brought computational media to the center of design pedagogy. Nevertheless, as the narrative of the early Paperless Studios articulates, the project was not so much a coherently theorized program as it was an experimental response to starkly shifting technological circumstances (GSAPP, 1994).

Faculty members like Greg Lynn, Hani Rashid, and Jesse Reiser, later branded as digital avant-gardists, had only a nascent understanding of the technical support for their pedagogical ambitions. Because of this disparity between pedagogical intentions and technological fluency, what came to emerge was an improvised connection between digital and design. The students positioned themselves as technical mediators, finding their way through

new software contexts, competency with a software environment, and other variabilities of particular hardware, as much of the work was improvised with little precedent or guidance. Although the flattening of hierarchies in the form of possible agency that might come from the decentralization of the educational systems offered opportunities for productive play, it also meant that architectural pedagogy was vulnerable to technological fetishism whereby the spectacle of formal novelty offered by digital processes could take precedence over criticality. At least introduced digital modes of representation served as accelerative affordances and limitations.

On the one hand, the capacity for rapid iteration, manipulation of complex geometries, and visualization of speculative forms expanded the operative repertoire of architectural design. On the other hand, the immediacy and fluidity of digital representations tended to privilege surface articulations over programmatic, contextual, and material considerations. As Tschumi later acknowledged, integrating digital tools risked encouraging “image-based” design practices at the expense of critical disciplinary engagement (Tschumi, 2013).

In addition, the portrayal of digital tools as extensions of drawing practices supported a view of architecture as representational rather than generative. The emphasis was still on producing images (drawings, renderings, simulations) as opposed to transforming architectural processes or fundamentally different epistemologies. In this regard, the “rediscovery through drawing” was not a rupture but a repackaging of architecture's historical dependence on visual mediation,

albeit through different technologies.

The broader historical context of this time also complicates narratives of digital innovation. As Herriman (2022) and the Thinking Machines exhibition materials emphasize, the late 20th century saw a rapid and widening instrumentalization of computation across different disciplinary boundaries, where speculative aspirations were subordinated to operational logic. In architecture, this became a technocratic valuing of efficiency, precision, and formal flexibility, often at the expense of serious critical engagement with design's social, cultural, and political dimensions.

Thus, the 1980s and early 1990s constitute a period of ambivalent transition. While digital tools undeniably expanded architectural practice's technical and formal capacities, they also introduced new risks of abstraction, aestheticization, and disciplinary detachment. The pragmatic adoption of computers as drawing machines—rather than as thinking partners—enabled new production modes and constrained the horizon of architectural imagination. This ambivalence would continue to shape architectural discourse into the subsequent decades as the promises and perils of digital technology became ever more deeply embedded within the fabric of architectural thought and practice.




FIGURE 9
Lynn, G. (ca. 1999). Embryological House:
Sketches [Drawing]. Canadian Centre for
Architecture (CCA), Embryological House
fonds.




FIGURE 10
Gehry's 1992 Fish Sculpture in Barcelona
was among the first architectural projects
to use CATIA software, enabling precise
digital modeling from hand-built forms—a
method later applied in the Guggenheim
Bilbao and Disney Concert Hall.
Gehry, F. (1992). Computer and built models
for Gehry's fish sculpture, Barcelona [Pho-
tographs and digital models]. Courtesy of
Gehry Partners, LLP.

3.3. THE FIRST DIGITAL TURN (EARLY 1990S–2000)

The early 1990s witnessed a fundamental shift in the culture of architectural design. The availability of digital tools has opened up new forms of representation and form-making, resulting in shifts in aesthetics and concepts. However, this “First Digital Turn” was not a simple or unthinking acceptance of technology; instead, it emerged with complex and, at times, conflicting positions on digital design’s epistemological, formal, and political implications. Some of the major protagonists in this shift were Greg Lynn, Peter Eisenman, Frank Gehry, and Patrik Schumacher. Each adopted different trajectories regarding how digital techniques would be utilized, resisted, or instrumentalized. Collectively, they create a contested genealogy rather than a cohesive paradigm.

Greg Lynn’s contributions to early digital architecture are inseparable from his theorization of “animate form,” a concept that framed form not as a static entity but as a condition capable of continuous transformation. Lynn’s 1993 essay “Architectural Curvilinearity: The Folded, the Pliant, and the Supple” (Lynn, 1993) positioned him at the forefront of a discourse that wove together formal experimentation with the theoretical apparatus of Gilles Deleuze. In particular, Deleuze’s *The Fold: Leibniz and the Baroque* (1988/1993) served as an intellectual catalyst, providing Lynn with a philosophical rationale for abandoning discrete, typological thinking in favor of topological continuity.

This conceptual shift was facilitated by adopting Non-Uniform Rational B-Splines (NURBS), which allowed architects to digitally model smooth, continuous surfaces beyond the constraints of Euclidean geometry. Lynn’s

approach was not merely technical but methodological: the spline became an instrument for investigating continuity, elasticity, and variation. Nevertheless, this embrace of fluid geometry also opened the door to aestheticization. As Mario Carpo later argued, the early digital avant-garde often prioritized geometrical novelty at the expense of tectonic articulation and socio-spatial critique (Carpo, 2011).

Peter Eisenman’s parallel trajectory during this period offers a sharply different perspective. Rather than celebrating digital tools’ capacity to generate complex geometries, Eisenman interrogated their capacity to disrupt established syntactic and typological systems. His earlier work from the 1970s, such as House I through House VI, already exemplified a process-oriented design logic that prioritized the internal syntax of architectural elements—walls, columns, voids—as autonomous agents within a conceptual system. As he transitioned from functionalism toward structuralism and later to deconstruction, Eisenman progressively abandoned any residual commitment to form as a bearer of programmatic or symbolic meaning. The digital turn did not signal a rupture in Eisenman’s thinking; instead, it was a response and an extension of his long-standing career as a critical architect. His work with the philosophical ideas of Jacques Derrida and Gilles Deleuze informed his critique of representation and the shift to what he would describe as deconstruction. The deconstruction phase of his methodology, seen in projects such as House X, emphasized an analytical mode that invoked the diagram over the image and syntax over the symbol. Eisenman asserted that the design process should occur through a speculative and systematic manipulation of architectural grammar - as he called

it, a “readily reconstructive series of transformations” - rather than intuitive composition or visual coherence (Eisenman, 2006).

Eisenman’s theoretical framework remained skeptical of the instrumentalist claims often associated with digital architecture. He cautioned against reducing digital design to mere parametric control or formal innovation, insisting that digital tools be deployed to interrogate rather than reinforce architectural conventions. As he observed in his critical writings, architecture must preserve its autonomy as a discipline that generates meaning through its internal operations, not through representational excess or technological novelty (Eisenman, n.d.). Frank Gehry’s trajectory provides yet another distinct articulation of the digital turn, one that prioritized technological integration into the fabrication process. Gehry’s office began using CATIA (Computer Aided Three-Dimensional Interactive Application) in the early 1990s—a software originally developed for aerospace engineering—to manage the complex geometries of projects like the Guggenheim Museum Bilbao. It was through this action that Gehry circumvented traditional architectural documentation by establishing a digital continuity from the conceptual model to the complete build. Gehry did not refer to an abstract theory like Lynn or Eisenman, as his use of digital continues to be based in material practice; his workflows were firmly contingent on working materially - former models were made by hand, then turned into a model digitally and manipulated through CATIA. This analog-digital loop worked well to maintain some material materiality while securing incredible formal control.

Nevertheless, Gehry’s work has been celebrated and critiqued for its emphasis on spectacle. Several critics have noted that the seamless translation from digital model to physical form may obscure more profound questions of context, program, and meaning. Stan Allen warned that Gehry’s digitally aided tectonics risked conflating architecture with the sculptural surface, reducing buildings to “complex form devoid of disciplinary depth” (Allen, 2000).

The most explicit attempt to codify a digital design ideology came from Patrik Schumacher, who began articulating the principles of “Parametricism” in the late 1990s and early 2000s. Framing it as a coherent style following Modernism and Postmodernism, Schumacher argued that architecture should embrace computational variability, continuous differentiation, and algorithmic control as its new organizing logic (Schumacher, 2009). Parametricism, he claimed, was not merely a method but a comprehensive approach to urbanism, form generation, and social organization. This assertion, however, has drawn significant criticism.

Schumacher’s texts advocate the conception of digital architecture as efficient, adaptable, and capable of expression through visual means. In contrast, others find in Parametricism a troubling reduction of design to a series of optimization of formal systems (Yuce, 2014). Scholars have identified that the style typically operates in a formal vacuum and generates complex maneuvers that do not consider the real-world circumstances of sites. Furthermore, the claim that Parametricism is intrinsically progressive is challenged by the predominance of its application in speculative

or commercial projects with little to no public engagement. The reduction of architectural agency and the design process to algorithmic editing form crucial questions. Who specifies the parameters? What values undergird the scripts? These questions highlight more troubling ethical and political questions that are too often obscured by the rhetoric of technological development. As Antoine Picon has noted, decreasing the work of architecture to design workflows diminishes the architect's interpretive work and eventually supposes that computation replaces deliberation (Picon, 2010).

The First Digital Turn did not produce a unified movement but rather a constellation of responses to the new affordances of digital technology. Greg Lynn leveraged philosophical theory to propose a new formal paradigm rooted in topological transformation. Peter Eisenman extended his syntactic critique of architectural norms, using digital tools to interrogate the assumptions embedded in the form itself. Frank Gehry redefined the architect's relationship to material production through industrial software, while Patrik Schumacher sought to formalize a design ideology premised on algorithmic logic. These trajectories reveal that digital architecture cannot be understood solely in terms of technique. Instead, it must be examined through its conceptual foundations, disciplinary ambitions, and cultural consequences. The tension between formal innovation and critical reflection remains central. As Picon reminds us, digital tools "do not inherently confer meaning or value"; they merely create new terrains upon which disciplinary arguments must be staged (Picon, 2010).

The First Digital Turn is best understood not as a stylistic epoch but as an epistemological disruption—a moment that redefined how architecture thinks, makes, and represents itself.

3.4. MASS CUSTOMIZATION AND THE RECODING OF FORM (LATE 1990S)

The late 1990s marked a significant transition in architectural thinking and production, shaped by the increasing entwinement of computational technologies, digital fabrication, and shifting theoretical models of authorship and form. Rather than celebrating digital tools for their capacity to generate unprecedented geometries—a tendency that characterized much of the earlier "first digital turn"—this period introduced a more fundamental reconfiguration of design logic. Central to this shift was the rise of what came to be termed "mass customization," a model that disrupted the modernist emphasis on standardization by enabling the cost-effective production of unique, differentiated components within serial systems. At the core of this paradigm was Bernard Cache, whose theory of the "objective" offered a radical redefinition of architectural components. In *Earth Moves: The Furnishing of Territories* (Cache, 1995), Cache advanced the objective as a digitally computed, parametrically variable object—no longer fixed or singular but determined through algorithmic manipulation and rendered producible via numerical control (NC) fabrication. Unlike the industrialized module of modernist architecture, the objective operated as a field of potentiality: each instantiation derived from a standard parametric model capable of expressing continuous variation. This repositioned the architectural object not as a product of formal intention alone but as a computational outcome shaped by variable constraints. In this sense, Cache's work offered a technical solution to form-making in the digital age and a philosophical intervention into the nature of the design object itself.

This notion of variability within seriality resonated strongly with contemporaneous developments in architectural theory and pedagogy. Branko Kolarevic emphasized that the convergence of parametric modeling and digital fabrication technologies allowed designers to overcome the historical trade-off between variety and cost. As Kolarevic (2003) noted, digital fabrication tools such as CNC milling made it "just as easy and cost-effective to produce 1,000 unique objects as to produce 1,000 identical ones." The architectural implications were considerable: variation became operational, not decorative; differentiation was embedded in the logic of production, not appended to it. William J. Mitchell provided a broader epistemological framing of these developments. In *The Logic of Architecture* (1990), Mitchell articulated a vision of computational design grounded in rule-based systems and representational clarity. Nevertheless, by the mid-1990s, in texts such as *City of Bits* (1995) and *e-topia* (1999), his focus shifted toward digital networks' cultural and spatial consequences. For Mitchell (1995), architecture was undergoing a process of despatialization, wherein traditional spatial functions—commerce, education, administration—were increasingly displaced into non-physical, digitally networked environments.

This reorientation required new tools and a new theoretical relationship to the architect's role in distributed interface-based contexts. The complexity of these developments was also refracted through a cultural theory lens. In *The Architecture of the Jumping Universe* (1995), Charles Jencks offered a speculative understanding of architecture's relation to novel scientific paradigms. Jencks used chaos theory, fractal geometry, and systems thinking

as models to suggest that architectural forms accept the discontinuities and recursion of nature and contemporary cosmology. While tools such as CATIA and Maya enabled the formal translation of these metaphors into built form, Jencks's approach raised concerns about epistemological rigor, often straddling the line between allegory and application.

Amid these theoretical developments, digital production's political and ethical dimensions remained contested. The promise of user agency within mass customization frameworks—such as Cache's configurators or participatory interfaces developed by Gramazio & Kohler (2008)—invited scrutiny. While such systems appeared to empower users, the parametric constraints were architect-defined, maintaining the designer's meta-authorial role. This raised enduring questions: What constitutes creative agency in a generative system? Who defines the scope of variation? Moreover, what are the actual limits of user participation? The compression of design and fabrication workflows, praised for its efficiency, also introduced tensions.

As Kolarevic (2003) has pointed out, digital fabrication and its implications often meant high-capital, energy-intensive infrastructures that raised questions of access and sustainability. Digital tools and media provide new forms of making. However, they also require technical knowledge for these forms of making to happen, institutional access, and thus reproduce a set of barriers in another form. We emerged from this engagement not with a style but a conception of architecture as a dynamic system of recalibration, not a static composition. Cache, Mitchell, and Jencks offered different but overlapping conversations

around the ontological, epistemological, and cultural implications of computation and together give rationale for the intellectual heft of architecture in the late 1990s and as of today.

As digital tools permeate all aspects of practice or organization, we must continue to face the challenges of this period: What role does the author play in the generative systems of texts? What ethical boundaries do algorithmic design practices have? Moreover, how do the socio-political implications of the computing turn that architecture is undertaking?

3.5. THE SECOND DIGITAL TURN (2000S–2010S)

The second digital turn in architecture occurred in the early 2000s to 2010s and represented a significant development in the meaning and use of digital elements in design. The first digital turn was about formal innovation with digital and computational modeling and new forms of geometry, whereas the second focused on systems thinking, platforms, and information management. Digital fabrication, Web 2.0, parametric customization, and the professional embrace of Building Information Modeling (BIM) are all the changes that moved architectural discourse from object-making to orchestration with new questions about authorship, agency, and control.

This evolution is seen, theoretically and technically, in the work of Bernard Cache, who, in the book *Earth Moves* (1995), introduced the term "objective." In Cache's later approach, "Tables Projectives," the object responds to parameters in its context through embedded geometry. Emphasizing variability in architectural practice rather than novelty for the sake of form is an issue Cache's systems approach raised, acknowledging variability as a part of designing territory and topology. Cache expressed skepticism about the democratic nature of algorithmic design, noting the issue of the denial of agency, ultimately raising a larger point on the nature of digital design systems that appear to invite users to participate while cordoning off user participation through proprietary parametric frameworks as pre-determined constraints.

This perspective found real-world resonance in collaborative projects like WikiHouse, an open-source building project where users could download elements to customize their assembly of CNC timber frames. Although

WikiHouse aimed to act as a democratizing campaign in housing design, it ultimately illustrated the situated limits of these models. As Thompson (2019) argues, limits to participation were inherently tied to infrastructural and regulatory limitations and even to the understanding of digital fabrication tools. In other words, openness was based more on the system's design capabilities than the participants' potential freedom.

A similar dilemma was demonstrated in the global proliferation of FabLabs. Based on Neil Gershenfeld's work at MIT (2005), FabLabs approached the idea of a navigable personal-making place and positioned themselves as available fabricating spaces. FabLabs represented a shift away from central production to localized, user-driven production. FabLabs was quickly agenda as a mass-customization platform in architecture but was criticized by various authors, including Leach (2014) and Picon (2010). Despite the idea of empowerment and revolutionary access to design and prototyping, they were reliant, to a large extent, on proprietary hardware and software ecosystems additionally on academic institutional funding. Access was limited, and the distribution of resources, networks, and engagements of specific users in different global contexts still occurs.

Additionally, as digital fabrication became more widespread, we also saw a still further significant shift in the professional domain of Building Information Modelling (BIM). Using tools like Autodesk Revit, we aligned all project data and integrated design, engineering, and construction into one model with as little manual handling as possible. BIM is seen as a way of dealing with the complexities of

managing the work of circular geometries, materials, timelines, and costs in a richly layered design model to coordinate amongst disciplines. This lack of constraints and the overall thinking capacity for BIM became problematic! Smith and Tardif (2009) believe BIM supports workflows that facilitate management or institutional purposes, not creative exploration. BIM was primarily complex software with high licensing costs that used monopolistic practices, allowing a small group of professionals to remain as experts.

While optimally tasked to manage attributes, BIM also contributed to quantifying design value. Metrics, simulations, and lifecycle assessments were being embedded into architectural workflows. While increasingly enhancing environmental responsibility and accountability with predictive modeling, it also contracted design-thinking to the possible and calculable certified parameters of representational thought. The aesthetic, cultural, and spatial aspects of architecture were in peril of becoming eclipsed by performance indicators. As Kieran and Timberlake (2004) warned, the digitization of making buildings could mean we exchange architectural judgment for optimization protocols.

In concert, these developments shifted how we think of architectural authorship. Architects were no longer the only authors of an original form but could be thought of as conductors of systems - creating rules, templates, and interfaces that shaped the way anyone else could join in or not. William J. Mitchell (2003) outlined a world in which spatial experience was instead shaped by networked relationships rather than objectified isolation and in which architects would

be curators of relational experiences in both physical and digital spaces. Of course, this change is not without critique, as the highly programmable environments of interactivity and decentralization are often structured through the invisible hand of the author.

In summary, the second digital turn did not introduce entirely new rules than those of the past but represented significant changes in the sets of tools architects draw upon to work and the frames of their making. Results from algorithmically modulated objectives to user-generated building systems, from local fabrication labs to cloud-based BIM models, platforms, protocols, and data-shaped architecture. These systems have promised autonomy and access yet may have, and will likely, perpetuate established hierarchies while accruing other dependencies. Rather than rupture, this period is better understood as a classification of architectural practice that alters the architect and politics of design.

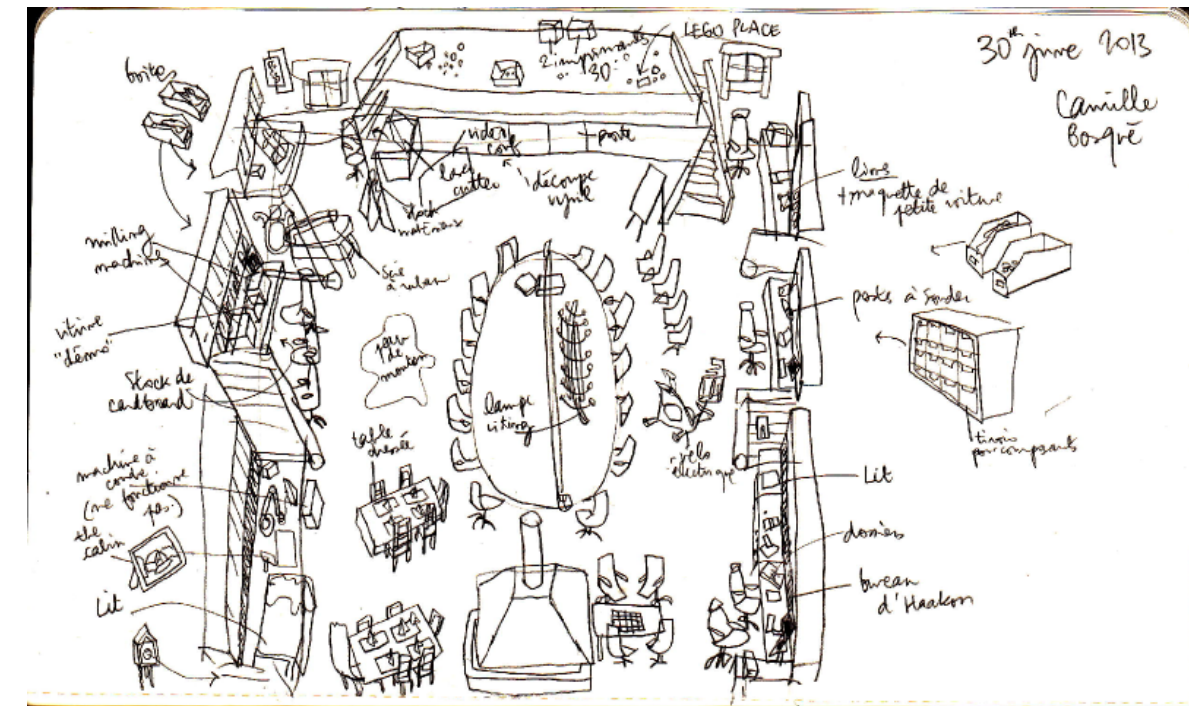


FIGURE 11
Bosque, C. (2013, June 30). MIT-Fablab Norway: Extract from Bosque's sketchbook. In *The Story of MIT-Fablab Norway: Community embedding of peer production*

3.6. DISCRETE COMPUTATION AND THE POSTHUMAN AESTHETIC (2010S)

The 2010s represent the most recent phase in this sequence, distinguished by a move away from seamless digital forms toward discrete, data-intensive processes and materially expressive computation. This transition has reshaped architectural aesthetics and invited renewed scrutiny of authorship, agency, and the boundaries of human cognition in design.

One of the defining projects of this era is *Digital Grotesque* (2013) by Michael Hansmeyer and Benjamin Dillenburger. The major complexities of the present project—260 million surface polygons and 25 billion volumetric voxels—were only possible through high-fidelity additive manufacturing. In contrast to earlier computative projects that obscured their digital logic beneath smooth, biomorphic qualities, *Digital Grotesque* hugely emphasizes its computational beginnings. The aesthetic is the pixelated and voxelized language itself. The work is layered filigree articulated recursively that neither maintains classical proportion nor intuitive legibility.

The advance of discrete geometry was a significant rupture and change from a technical and epistemological perspective, and computer-aided design systems of boundary representations could neither hold this density of information nor produce the resultant new wavelengths of data; volumetric modeling had internal control connected to materials and structures, as significant connections could be opened up via visualization. A designer would establish the behavior of a structure by form, such as self-containment, connecting two opposing dualisms of structure and ornamentation, pace, and performative envelope. With voxel-based algorithms like

Voronoi tessellations and flood fill, designers could engender structural behavior meaningful to their previous operations.

Philosophical shifts also accompanied this technical change. As Carpo (2017) argues, these forms indicate a posthuman design logic whereby the human designer does not dictate every perspective of the design. The form grew out of datasets, scripts, algorithms, and machines. Rather than being an author of forms, an architect is more like a system calibrator, albeit at a distance from the product - a dispersed intention sequencing multiplicities of relationships. This disambiguity or dispersal of intention puts pressure on notions of authorship and critique before a project can be prepared and predicted as a project.

Integrating Big Data and high-performance computing (HPC) into architectural workflows also engaged analysis process possibilities. For example, finite element analysis (FEA) could now simulate real-time performance through a dynamic and probabilistic moment beyond the performance of materials. The assemblages of Masera and Bianchi (2021) represent a convergence of data-rich simulation and machine-led optimization. However, opacity had become a new production landscape: many simulations were of such size and complexity to exceed or bracket fully a human capability of audit or interpretation - aspects of relevance for considerations of trust, validation, and accountability in designerly spaces increasingly dominated otherness defined inferencing and abstractions.

Aesthetically, these shifts aligned with a

return to Brutalist principles—but through a digital lens. This “Neo-Brutalism” adopted monolithic geometries and exaggerated articulation, often prioritizing formal complexity over human scale or social intention. Unlike mid-century Brutalism, which was rooted in material honesty and public housing agendas, the new iteration reflects a fascination with alien materialities and speculative futures. Its rawness is algorithmically derived rather than socially grounded. Persistent claims of technological liberation have accompanied this evolution. Proponents suggest that with algorithms and 3D printing, anyone can design and fabricate unique architecture. In practice, however, such systems remain capital- and expertise-intensive. The computational infrastructure required for projects like *Digital Grotesque* is accessible to only a narrow stratum of global design practice. Far from democratizing design, the aesthetics of voxelization and discretization have often reinforced elite authorship under the guise of technical experimentation.

Therefore, the 2010s do not represent the apex of digital innovation but a reorientation—one that demands scrutiny. The tools used to produce these architectures are increasingly autonomous but not neutral. They are laden with assumptions, constraints, and systemic exclusions. The architectural imagination is no longer bounded only by geometry or software capacity but by broader social, environmental, and epistemological implications.

The three digital turns described here—again following Carpo’s periodization—chart a trajectory from early formalist optimism to systems thinking and, finally, to the

complexities of a posthuman paradigm. As we stand at the edge of another technological transition, marked by artificial intelligence, machine learning, and generative design, a central question looms: Are we entering an actual “age of AI” in architecture? Answering this requires caution. The rhetorical inflation around AI risks obscuring unresolved control, access, and accountability issues. The field must resist the temptation to equate novelty with progress. Rather than forecasting a utopia or dystopia, we should ask what kind of design culture AI will support—and for whom. Whether AI becomes a new architectural epoch or a fleeting phase will depend less on technical possibility than on our capacity to critically engage its tools, frame its impacts, and define its purpose.

4 FROM GLIMMERS TO REALITIES

Generative Adversarial networks (GANs) are typically associated with 2D image tasks, yet they can be surprisingly relevant to architectural design (Goodfellow et al., 2014). Architects rely heavily on two-dimensional representations, such as floor plans, sections, and elevations, to communicate and develop three-dimensional ideas. In this sense, a tool developed to produce and transform images can be adapted to generate new design options.


Architect Stanislas Chaillou's ArchiGAN, developed in 2019, is a prime example of bringing GANs to the architectural field (Chaillou, 2019). Chaillou employs a chain of GAN-based models (inspired by Pix2Pix; Isola et al., 2017) to generate architectural layouts (floor plans) at multiple scales: first, the building outline, then interior partitions, followed by the furniture layout. Each step interacts with the others, so if a user adjusts the building's outline, the system automatically recalculates the partitioning and furnishings. This setup transforms GANs from purely analytical or image-based tools into a design partner that assists with repetitive or time-consuming tasks while leaving creative decisions to the architect. However, practical challenges remain. GANs produce raster (pixel) images, while architectural firms commonly use vector-based software (e.g., CAD or BIM).

It was possible to introduce some error in the process that converts pixel outputs to vectors (Chaillou, 2020). In addition, the datasets and computing power for deep learning are massive. The size of this data could also be restricted to some practices or individual practitioners simply because of cost. Additionally, maintaining the script requires

other skills, such as precise coding. However, GANs show how AI can be an efficiency-generating engine to explore design possibilities that allow design practitioners to explore further afield and do it quickly.

FIGURE 12

Chaillou, S. (2019, July 17). Figure 3. GAN-enabled building layouts [Image]. In ArchiGAN: A generative stack for apartment building design. NVIDIA Developer Blog. <https://developer.nvidia.com/blog/archigan-generative-stack-apartment-building-design/>



4.1 DEEPHIMMELB(L)AU

Wolf Prix, Design Principal and CEO of Coop Himmelb(l)au, has long been known for challenging formal architectural conventions. While initially skeptical of parametricism and its tendency to enclose design within overly deterministic rules (Leach, 2021), Prix has more recently explored how artificial intelligence might offer generative potential without sacrificing creative autonomy.

This shift is evident in the project DeepHimmelb(l)au, an experiment that employs CycleGANs to generate speculative architectural imagery. The system produces hybrid visuals that synthesize geomorphic patterns, architectural references, and material textures by training models on the firm's archive and other datasets. Although these outputs are two-dimensional, they act as provocations—sparking design dialogue rather than dictating outcomes.

The firm's website contextualizes the project by way of a larger exploration into computational creativity, stating that “teaching computers to be creative is very different from how humans create” since human perception is contingent on layers of interpretations and embodied experience (Coop Himmelb(l)au, 2022). This comparison leads to questions about the capacity of architects to employ AI, not just to relieve the designer of repetitive tasks but to access the formal logic or aesthetic strategies beyond human intuition. What types of design intelligence emerge when algorithms learn from and riff off a visual language developed by the studio? How can designers remain critical of their role, and that of AI, between images?

The firm's website frames the project

within a broader investigation of computational creativity, noting that “teaching computers to be creative is very different from how people create” since human perception relies on layered interpretation and embodied experience (Coop Himmelb(l)au, n.d.). This comparison raises questions about how architects might use AI not simply to automate tasks but to discover formal logic or aesthetic strategies that exceed human intuition. What kinds of design intelligence emerge when algorithms learn from and riff on a studio's visual language? How can designers remain critical when navigating such outputs?

However, practical and conceptual limitations remain. Most AI-driven workflows still depend on 2D inputs and outputs, which makes integrating machine-generated visuals into rigorous 3D environments a complex challenge. Furthermore, the sheer novelty of these tools often overshadows their structural limitations or the labor required to adapt them for architectural applications.

The proposal of AI as a breakthrough rather than an area of exploration creates tension in AI projects such as DeepHimmelb(l)au, which serve to evaluate approaches, ask questions, and uncover new tensions. Prix's experiments exemplify how generative models can support architectural exploration while preserving the interpretive design judgment central to the act of design. The larger lesson is not that AI will redefine architecture by default but rather the careful determination of AI's place in the architectural process, which must be done thoughtfully, intentionally, methodically, and under expecting continued authorship.




FIGURE 13-15
Deep Himmelblau — As Coop Himmelb(l)au's semantics and style are not homogeneous... [LinkedIn post]. LinkedIn.

4.2 MACHINE HALLUCINATION

While Refik Anadol is not an architect in the conventional sense—he does not design buildings or spaces—his work has nonetheless influenced architectural discourse by expanding how we might conceptualize space, memory, and data in the built environment. His large-scale installations, often projected onto existing buildings or situated in immersive rooms, demonstrate how machine learning can operate as a design aid and a mode of spatial expression.

Among his most highly publicized work, WDCD Dreams (2018) involved an artificial intelligence (AI) generated animation projected onto Frank Gehry's Walt Disney Concert Hall in Los Angeles. Drawing from the Los Angeles Philharmonic's 77 terabytes of archival material, Anadol trained a neural net to consider the archival material - images, sounds, documents - and produce an animated visual composition across the artwork's changing metal skin (Google Arts & Culture, 2018). Although there was no physical change to the architecture, the hall's facade was transformed into a shifting surface, pushing the boundaries of architecture as media or media as architecture. The installation opened discussions on architecture's relationship with memory and whether the building could act as an interface for culturally processed data from a machine.

Though these interventions remain in visual art, they intersect with architectural concerns: public perception, temporal occupation, and buildings' representational capacity. Matias del Campo has argued that such works, although not spatial in the traditional sense, "reveal unforeseen territories of form and thought" (del Campo, 2022). They open

possibilities for architects to collaborate with data-driven systems in crafting objects and experiences.

Other installations, such as Melting Memories (2018), translate neural activity (via EEG data) into digital sculptures and projections, offering a speculative model for how cognition might be visualized in spatial terms. Similarly, Infinity Room (2015) used mirrors and light to create immersive environments that destabilize spatial boundaries. These projects suggest new vocabularies of atmosphere and perception, which, although not architecture per se, are increasingly relevant to how architects conceive experiential and responsive environments. Projects like Machine Hallucination: NYC (2019) go further in scale, synthesizing more than 100 million images of the city into abstract video sequences. These are not generative design tools in the architectural sense, but they pose conceptual challenges: How might we interpret urban memory through algorithmic recomposition? What possibilities of authorship and interpretation exist when environments are remixed through AI?

Anadol poses AI not as an agent author but as an agent collaborator—an artificial system capable of elucidating hidden patterns in the data. The work exemplifies how machine learning can deliver expressive material and remain within humanity's agency and curatorial function. Thus, his position is similar to some momentums of experimental architectural practice, whereby generative tools act in support of but not replace agency in design.


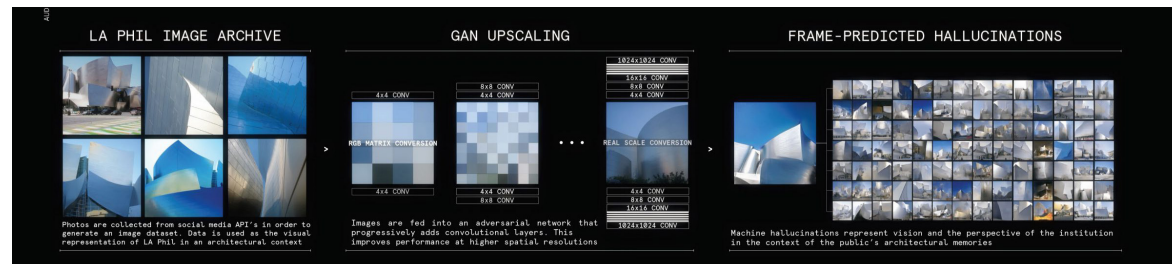
While Anadol's work does not propose architectural solutions, it expands the conceptual space in which architecture operates. His projects suggest a shift from designing objects to curating data and orchestrating media. For

architects, this presents both an opportunity and a caution: to engage AI not only for efficiency or novelty but also to reshape how buildings communicate, respond, and remember.



FIGURE 16
RENAISSANCE DREAMS — PALAZZO STROZZI
The installation by Anadol.R utilizes GAN algorithms trained on Renaissance-era data, creating a "multidimensional shape matching the architecture and infrastructure of MEET"

FIGURE 17
Clemence, S. (2018, September 26). Frank Gehry's Walt Disney Concert Hall to be projected with digital "machine hallucinations". Metropolis Magazine..



4.2 THOM MAYNE AND MORPHOSIS

Thom Mayne, founder of Morphosis and Pritzker Prize winner, has long placed technology at the core of architectural experimentation. Previously, digital technologies allowed the firm to extend its formal vocabulary; Mayne now sees AI as a mechanism to disrupt socialized design habits and create new thinking. He insists, however, that AI does not replace human creativity and is best understood when it induces a strategy of exploration, particularly when intuition may lead to cyclical or self-referential thinking (Mayne, 2019).

Morphosis perceives AI as one strand of a larger shift from intuitive design toward process-based inquiry. Mayne challenges “a priori” solutions and prefers iterative workflows, in which generating a series of options provides insight into the boundaries of form. Morphosis’s global work can include multiple iterations of a concept. In the case of Mayne’s firm, they may use those iterations in exploratory processes, yielding hundreds of formal iterations in the studio. Mayne states, “AI also has become part of that workflow-- as a proxy that will introduce strange geometries, strange patterns, or different associations.” However, Mayne does not fetishize the new technologies. In recent public interviews and lectures, he has reaffirmed that speculative outputs generated by GANs or similar forms must still be made sensible by architectural first principles: material logic, structural feasibility, and programmatic sensibility. For Morphosis, AI is not a mechanized end but an iterative co-author. Internally, the studio culture provides ample opportunity for back-and-forth: firm-generated forms can be chosen, modified, and often discarded according to the practicality threshold.

One example is Morphosis’s experimentation with image-to-form workflows, where AI-generated visual material (inspired by precedents or urban datasets) provokes conceptual starting points. These are then translated into 3D models and tested for spatial and tectonic coherence. Rather than presenting a full design pipeline, the AI serves to “interrupt” the process—introducing dissonance and dislodging default tendencies. Importantly, Mayne acknowledges the opacity of AI systems.

He relates algorithmic black boxes to human intuition; processes resulting from obscure input and attribution are the connection. The difference is liability. Architects must continue to be liable for design decisions despite being guided by what we see or half-remember. AI can be beneficial in allowing us to externalise our preferences and reveal our biases; however, it should not absolve an architect from authorship or critical liability (Leach, 2021; Sugihara, 2019).

In this framework, AI foils complacency—a way to sustain the critical momentum that has defined Morphosis’s work. However, the studio is cautious not to confuse novelty with progress. For Mayne, the integration of AI must serve a larger agenda: to reinvigorate architectural thinking, not to automate it.



FIGURE 18
Bill & Melinda Gates Hall at Cornell University in Ithaca, United States. (Morphosis Architects)

4.3 EXPANDING CONVERSATIONS ON AI: INFLUENTIAL PRACTITIONERS

The increasing integration of Artificial Intelligence (AI) in architecture has established a new generation of specialists and design offices re-evaluating how advanced computation can inform processes ranging from conceptual design to large-scale urban analyses.

Mundane dialogues with various actors characterize the seismic changes AI has introduced into everyday operations and the notion of “professional creativity” in architecture. One prominent actor, Arturo Tedeschi, describes AI as “a fundamentally different style of creativity,” freeing architectural professionals from mouse-driven modeling directed toward an idea-creating process (Parkes, 2023). During our conversation with Tedeschi in Venice at the AI conference around the opening of the 2023 Biennale, he established the notion of “machine hallucinations” being introduced into the early design phase - that stage where architectural designers often feel ambiguous about production pursuits.

Another practitioner, Tim Fu, formerly head of the computational design department at Zaha Hadid Architects, highlighted the emergence of custom generative algorithms as an internal resource within large firms, allowing for deeper exploration beyond conventional parametric tasks. Fredy Fortich is another significant figure in this evolving landscape. An Architect and Engineer specializing in computational design, including BIM coordination, performance-based design, generative design, and machine learning, Fortich is a BIM Coordinator at MVRDV’s French Studio and an AI Researcher on Diffusion Models for MVRDV NEXT (MVRDV, 2023).

In our 2022 conversation with Fredy

Fortich at the early stages of Midjourney, we discovered some valuable ideas about straightforward prompt engineering on the platform. Fortich was experimenting with Generative design with Stable diffusion at our earliest dates, while also bringing AI in as part of the iterative workflow for rapidly producing many design ideas, and managing images along with controlled curation by the right prompt.

At the same time, theorist Neil Leach, whom the author had the opportunity to meet at the Venice Biennale in 2023, reminds architects that meaning is not guaranteed by using AI: “It does not guarantee meaningful architecture, and real-world building constraints and cultural context remain important drivers of design outcomes” (Leach, 2021). This perspective aligns with Henning Larsen’s approach, which primarily sees AI as a tool to manage extensive data, whether labeling, measurements, or environmental analysis. This work often allows architects more freedom to produce conceptual design innovations. However, Henning Larsen highlights ethical and environmental issues, too - data sets can be biased or ignorant, and the carbon footprint of AI can take a significant toll on the environment, so it must be used responsibly (Henning Larsen, 2023).

Amidst this changing landscape is Snøhetta, who reiterates, Input: Output – Curating creative intelligence, describing AI as a part of an overall design philosophy. Instead of thinking of AI as a stand-in for human creativity, Snøhetta facilitates large language models and generative ways, producing “deep, unexpected” results that continue



FIGURE 19-20
In Copenhagen's Refshaleøen district, artificial intelligence was utilized to analyze data and uncover concealed patterns, which were then translated into creative visual interpretations of potential futures. This approach created a meaningful connection between factual data and the emotional experiences of both citizens and stakeholders, showcasing AI's potential to stimulate fresh perspectives and innovative visual storytelling (Henning Larsen). Midjourney

to reflect social stories (Snøhetta, 2024). This aligns with Tedeschi's idea of hybrid human-computer workflows and Henning Larsen's idea to use AI to "prepare" massive amounts of data, later showing it to the user, therefore not giving up their important decisions to an algorithm and letting them inform the human design process.

The same hybrid approach can be seen in Snøhetta's work on projects such as the Bokhus bookshelf, the Collective Oslo brand identity, and a custom-coded interactive media generator for sound based on the dominant colors within an image, transformed into sound sequences composed of site-specific recordings in the area around the officially docked hotel on the Oslo waterfront. Snøhetta uses these projects for integrating brand identity, contextual data, and multisensory thinking through the AI tool. However, whether these experiments offer substantial knowledge on how architecture operates or distinguishes itself within the conceptual branch of creative activities remains. The consistent thread witnessed throughout the discussions with Tedeschi, Fu, and Fortich was that although AI can meaningfully fuel creativity, computing cannot and should not replace the complexity of context, ethics, and aesthetics in creating meaningful architecture. Henning Larsen highlights that AI can save time and invite new design exploration, but responsibly using these opportunities should be a habit of continuously reflecting on issues such as bias, data quality, and environmental impact (Henning Larsen, 2023).

The works produced by practices embracing AI as part of their architectural workflows,

like Snøhetta, Henning Larsen, Zaha Hadid Architects, and MVRDV, are exploring how machine intelligence could aid creative exploration. It is a pertinent question if AI can be helpful through iterative form-making or complex data comprehension, enabling increasing forms of architectural inquiry. Regardless, the implications stem less from AI as an emerging tool than if the computational methods generate situationally and materially significant outcomes in the built environment.

Ultimately, these unique experiences and studio examples demonstrate that AI's most profound emergent promise is through dialogue: a reciprocal relationship between human and computing intelligence. From stage design or branding and parametric forms, the interaction of the architectural mind and machine intelligence redefines creativity for future generations, not limiting its scope but expanding design agency.

4.4 AI AND FABRICATION

Deep learning excels at perception-based tasks—classifying images, recognizing faces, etc. However, its applicability to more complex activities remains limited. As Stuart Russell notes, orchestrating large-scale endeavors like building a factory surpasses what purely deep-learning “black box” systems can manage (Russell, 2018). The stumbling block is straightforward: no vast dataset exists to show a network of every possible way of constructing a genuine factory. Moreover, the process requires knowledge, reasoning about obstacles, structuring a plan, and capacities beyond pattern recognition.

Architecture shares similar challenges. Although AI may help analyze designs and optimize specific tasks, materials remain analog. Regarding “digital” designs, these are only intangible models on a screen. Actual buildings require physical assembly. What we often call “digital fabrication” relies on additive or subtractive processes (e.g., 3D printing, CNC milling)—techniques that have, in principle, existed for millennia, like laying bricks or carving stones (Leach, 2017). While AI can control tool paths or optimize the workflow, fabricating still hinges on physical actions.

Even when robots step in, difficulties abound. Humans can easily pick up and place a brick, but a robot requires elusive fine motor skills and perceptual acuity (Brooks, 2018). Thus far, robots excel at routine tasks, like repetitive assembly lines in a factory, but struggle with unpredictable on-site conditions. Consequently, full-scale “AI fabrication” remains a distant prospect: algorithms themselves do not build anything; they merely control the process (Brown, 2019)

That being said, progress is still being made. For instance, Autodesk’s AI research includes training robotic arms in simulation, “The robotic arms are stacking virtual LEGO bricks millions of times before we transfer the learned in simulation behavior onto the real machines.” (Terdiman, 2018) While the robotic arm is gathering experiences very quickly, they are applying reinforcement learning to the robotic arm as also applied to Deep Blue. There may come a time when the robot stacks bricks if it is not already stacking parts of buildings. The challenge is taking the physical-digital leap from a real active construction site, where the conditions may change or unknowns arise.

In the meantime, AI can help with parts of fabrication, such as controlling print speeds, digitally monitoring assembly lines, or imitative production using simulation, but it will be. The real difficulty is automating a task that requires human skills in real-time or situational awareness to understand the surroundings. Until robotic hands ultimately produce us a beer, as Robotician Rodney Brooks sarcastically put it, we are at least some realistic distance from AI fabrication in architecture being a fully autonomous process. The narrative of autonomous AI fabrication of architecture is still an enthusiasm for an event that is not currently realized.




FIGURE 21-23

Shortlisted in August 2021 for Melbourne’s Merinda Station Integrated Art Project, this public artwork investigates geological processes as inspiration for innovative construction techniques and expressive forms. By combining 3D-printed sandstone with intricately cast metal inlays, the work explores themes of deposition, erosion, and the dynamic interplay between solid and void. Sand layering, metal casting, and material contrast evoke natural tectonic formations while introducing striking visual and textural complexity (Snooks, Harper, & Gibson, 2021).

5.0 SHAPING ARCHITECTURAL THINKING

In today's world, advances in computation and artificial intelligence (AI) have changed architectural practice. Think about the way we write today: instead of wandering around shelves in a library, we typically start by doing a quick Google search and have access to an entire library instantly. Understanding this transition highlights an important aspect of how computation shapes our knowledge: with computation, we scan a large amount of information instead of reading it using a digital medium. Similarly, digital images have redefined our conception of photography: instead of composing one image, we now make many images and select the best while we edit.

AI extends this logic into design processes. Instead of developing a single output, architects may now generate various potential design representations. Tools like Grasshopper allow designers to define constraints and "search" for forms by moving slider values, while AI-driven systems can do this automatically and at far greater speed. Consequently, design becomes a matter of identifying constraints, generating expansive sets of possibilities, and then using human judgment to select and refine an appropriate path forward (Leach, 2014).

This "search-based" approach has multiple ramifications. It changes the architect's role from the unique author of form to the operator of systems and workflows. As a designer, the critical role is framing relevant parameters and curating outcomes—not generating one unique answer (Shea, 2004). Hence, the architect's aesthetic sense is still important. Whether algorithms produce thousands of permutations, the architect's interpretive lens identifies which resonates with

the project's cultural, spatial, or functional objectives. Some studios have embraced computational unpredictability—"directed randomness"—to explore emergent possibilities within constrained systems (Leach, 2017). Others investigate how AI can systematize exploration beyond human bias, producing iterations not limited by familiarity or habit. Still, such processes require human insight to guide decision-making and contextual refinement.

Ultimately, this reframes how architects understand design agency. The process moves from singular authorship to informed negotiation with expansive, computationally defined design spaces. While algorithms contribute to speed and breadth, architects remain the curators of meaning, navigating between machine-generated alternatives and the embodied knowledge required to turn abstract variation into relevant, contextually anchored architecture.

5.1 LOST IN TRANSLATION: AI, ARCHITECTURE, AND THE PROBLEM OF AUTOMATION

The act of translation between conceptual design and material realization has always been important for architecture. Giovanni Corbellini (2020), in his questioning of contemporary architecture approaches, argues that architectural production is an intimate association of negotiation, reinterpretation, and transformation. He offers the metaphor of architecture as a "window," providing a view of the external cultural and technological dynamics and a "mirror" to reflect on the internal change of disciplinary practices within. This dual conceptualization establishes translation not as a secondary technical problem but a constitutive dimension of architectural thinking. Understanding this dynamic becomes particularly urgent in the context of emerging artificial intelligence (AI) technologies, which promise to radically alter the mechanisms mediating between architectural thought and built form.

Historical attempts to achieve precise architectural translation can be traced back to the Renaissance. In the fifteenth century, Leon Battista Alberti created orthographic projection systems to help maintain design intentions across time and space. Alberti's hope, as later interpreted by Robin Evans (1986), was to develop a representational language that could minimize interpretation of the discrepancies between conception and construction. However, Evans also argues that architectural drawings and objects belong to different semiotic systems and that, fundamentally and structurally, it is impossible to have the same fidelity between design and its realization. Every act of construction thus inevitably reinterprets and transforms the original

design, making negotiation an intrinsic condition of architectural production. In contemporary discourse, Mario Carpo (2011) revisits these historical ambitions in the context of digital technologies. Carpo observes that the dream of precise replication, first articulated by Alberti, has been radically reconfigured by computational tools that automate design and construction processes.

Building Information Modeling (BIM), parametric design, and AI generative systems are trying to lock as much information about structural, code, and material behaviors into design files, reducing translation and human contingency. In Corbellini's view, the transition towards computation is more than just a method; it is a reformulation of the project and ambition of architecture. The drive towards industry automation, aimed at countering the uncertainty of translation, may eclipse the heritage of uncertainty on a theoretical 'wrinkle' of architecture. As Umberto Eco (2003) reminds us, we lose something every time we translate meaning; negotiation includes loss. We should also be wary of any idea of perfection and unification - whether through automation or, better yet, centralization.

While automation promises efficiency, it poses a greater threat to authorship and cultural identity within architecture. As Marcos, Fernández-Álvarez, and Pak (2024) suggest, digital architecture more often prioritizes computational elegance than tectonic richness, which effectively separates the form of architecture from the socio-cultural contexts that have, historically, lent architectural form meaning. Similarly, Kwon and Ahn (2024) warn that increasing reliance on AI for building design

poses the risk of deskilling architects, the homogenizing of creative output, the loss of individual authorship, and the diminishment of deep intentional cognition. These critiques underscore the ethical ambiguities within, as well as the ethical entanglements that exist within the computational configuration of architectural production.

Wigley (2015) further reminds us that technological interventions have always transformed architecture's relationship to reality, often to obscure the contingencies and negotiations as part of design. The risk concerning AI is not only in the automation of design but also obscures the interpretive work that has historically defined architectural practice. Outputs are the (seemingly seamless) outputs of AI systems; the systems risk operating under the veils of contingency inherent within meaningful architectural work, which relies on decision-making, trade-offs, and reinterpretations.

Integrating artificial intelligence into architectural practice thus accentuates rather than resolves the historical complexities inherent in the translation from drawing to realized building. While computational tools offer new forms of control and precision, they simultaneously introduce opaque decision-making mechanisms and challenge established notions of authorship and critical agency. Therefore, the automation of architectural processes must be understood not as a neutral technological progression but as a profound transformation in the epistemology and ethics of design. Aligning with Corbellini's framework, these developments reaffirm that architecture remains an intricate negotiation

between external technological pressures and internal disciplinary meanings. Addressing these shifts requires sustained critical inquiry into how technological mediation alters the conditions under which architectural meaning, value, and responsibility are produced.

5.2 GENERATIVE AI IN ART AND MEDIA

The rise of new images generated by GANs like DALL·E, Midjourney, and Stable Diffusion raises new discussions about how images are created, what is considered creativity, and how we understand authorship in art and media. This chapter builds from the work of Lukas R.A. Wilde, Lev Manovich, and Gabriella Manzenreiter to offer a three-pronged synthesis: (1) a historical lineage from photography and digital art (i.e., computer graphics) to contemporary AI art (2) emergent and distinctive aspects of generative imagery; and (3) the cultural and ethical issues generated by new technology.

Historically, photography represented a new potential for image-making via mechanistic means rather than human labor (Mitchell, 1992). Wilde suggests that generative AI operates on similar logic: machines make images today by learning from large datasets, but each era brings new social and ethical implications, such as the tendencies of the data or changing human involvement. Manovich has gone further, retracing a lineage from early computer graphics, where separate algorithms simulated each visual attribute (reflections, cloth textures, for example), to the contemporary AI systems that effectively "reassemble" billions of visual objects at once (Manovich, 1992). Such developments mark new technical potential but also represent something larger in the evolution of how images come to exist and circulate.

Unlike traditional CGI, generative AI produces images using "latent spaces" that have been derived statistically from the vastness of available image archives, thus producing stochastic, dynamic, or even "creative" results rather than indicating a rendering of fixed

objects (Feyersinger, Kohmann, & Pelzer, 2023). Another essential element of these systems is that they rely on users to write textual prompts, which enables visual creation to become a text-based or conceptual task: users type these instructions, and the AI interprets the instructions to produce visual output (DALL-E, 2020). In doing so, the authorship question is also changed. Is authorship the responsibility of the person who assembles the prompt, the engineer of the AI system, or the AI model? Moreover, generative systems can find ways to produce output that is "hyper-realistic", such that the element of algorithmic support is obscured, or even expresses its "computational" nature by indicating artistic styles, changing how we perceive authenticity and mediating our sense of what is "real" (Bolter & Grusin, 2002).

In debating whether generative AI signals a genuine paradigm shift, Wilde emphasizes that new media forms typically emerge when technological change intersects with altered cultural protocols and industry practices. Like photography and cinema, AI-driven imagery could become a recognized medium if stable norms, platforms, and regulations develop around it (Kember & Zylinska, 2012). Early controversies, such as AI-produced art-winning competitions, show that cultural norms on authorship are still in flux.

Meanwhile, the audience's reception also changes. Some viewers admire AI outputs purely for their visual impact but are uninterested in whether a human produced them. Others feel that "human genius" and personal backstory are essential to meaningful art. This tension echoes earlier debates about art's purpose, whether about mimesis, symbolic

representation, or innovation (Aristotle, 1997). As AI advances, Manovich suggests they may highlight hidden visual patterns in cultural datasets, acting like “cultural theorists” that reveal how mainstream aesthetics are formed (Manovich, 2018).

Beyond theory, there are urgent ethical and sociopolitical dimensions. Generative AI systems often rely on large sets of images labeled by underpaid workers, and those datasets may contain biases that get reproduced in outputs (Williams, Miceli, & Gebru, 2022). Legal questions about fair use, style appropriation, and copyright remain unresolved. Critics claim that AI’s tendency to recombine existing works without explicit permission or credit infringes artists’ rights (Bajohr, 2021). Further concerns arise with “deepfake” imagery that convincingly mimics reality, potentially fueling disinformation and eroding trust in visual evidence (Davis, 2023).

Despite these challenges, the widespread adoption of AI in creative fields seems likely to accelerate. Whether we call it a “new paradigm” depends on how deeply AI reshapes our collective norms and whether it forges distinct cultural roles, much as photography and cinema once did. For some, AI is simply the following tool, albeit powerful, in a long lineage of media technologies. For others, the capacity of machine-learning models to autonomously generate images and even emulate distinctive styles points to a significant shift in how we conceive of creativity, agency, and artistic value. As AI continues to evolve, so will our debates about authorship and authenticity, and the balance between human intention and computational method may be permanently adjusted.

5.3 PREDICTIONS FOR THE FUTURE

AI is fundamentally about predicting patterns and using them to anticipate outcomes. Since the days of Alan Turing, experts have offered predictions about AI’s potential. Some projections were overly optimistic, fueling “AI winters” when promised breakthroughs failed to materialize, yet others have proven remarkably accurate (Simon & Newell, 1958). Humans excel at making predictions, too: our brains might be poor at remembering details or crunching calculations, but we are adept at detecting patterns, whether to catch a ball, evade risk, or profit in financial markets (Clark, 2013).

As AI advances, futurists like Toby Walsh foresee significant societal shifts by 2050 (Walsh, 2018). Many predictions are unsurprising: AI-generated news, phone-based health monitoring, or everyday interactions with “smart” environments. Others push boundaries: Walsh envisions “living on” post-death through personalized AI chatbots that mimic our speech or a ban on human driving once self-driving cars become standard. With incremental software updates (such as those in Tesla vehicles), the transition may occur so slowly that we barely notice or mind the disappearance of manual driving.

Just as humans quickly adapt to new technologies, social attitudes evolve. Andreas Huyssen observes that we live in an “epoch of amnesia,” frequently forgetting what life was like before the internet, smartphones, or even basic automation (Huyssen, 1995). This raises intriguing questions: As voice assistants proliferate, might we gradually rely more on spoken than written communication? Could literacy and writing skills erode, much like driving skills might wane with self-driving cars?

Moreover, if so, would we truly miss them? Such queries become even more crucial if the design process shifts away from hand drawing toward voice commands or gesture controls, a trend some researchers predict. Freed from pen and paper, we could lose the ability to sketch. Nevertheless, human societies have repeatedly surrendered older skills without regret, from ornamental calligraphy to manual mapmaking. The question is not whether AI will bring dramatic changes but whether those changes will happen so gradually and be so convenient that we accept them without hesitation.

INTERVIEW WITH 3XN/GXN DIGITAL / AI TEAM

As part of this thesis, I conducted an expert interview with Sam Sweeney, an architect, designer, computational expert, and researcher. Sam currently serves as Head of BIM and Digital Delivery at 3XN Architects in Copenhagen, where he works at the intersection of design, construction, and technology. With over twelve years of international experience spanning concept design, on-site construction delivery, and advanced computational workflows, his role focuses on managing BIM across 3XN's global offices, implementing parametric and data-driven design, and conducting research into carbon, circular economy tools, and AI in architecture.

His portfolio includes contributions to landmark projects such as the Sydney Fish Market, SAP Garden in Munich, and the 2FA Tower in London. As a designer with deep technical expertise, Sam's approach is rooted in the belief that computation should serve the enhancement of architectural quality and human experience.

This interview, conducted via written correspondence on 25 July 2025, explores critical questions related to the integration of AI in architectural design—its promises, limitations, and the shifting role of the architect in an increasingly automated environment. Sam's reflections provide a grounded, practice-based perspective that complements the theoretical and speculative themes developed throughout this thesis.

I would like to extend my sincere thanks to Sam Sweeney for his time, generosity, and insight, which have greatly enriched the arguments and findings presented in this research.

Interview conducted by Rustam Muradov, 3XN Architects in Copenhagen, 25 July 2025

1. IN YOUR EXPERIENCE, HOW IS THE ROLE OF THE ARCHITECT EVOLVING WITH THE INTEGRATION OF AI INTO THE DESIGN PROCESS?

I think architects should be really excited and curious by how AI can help us as both designers and problem solvers. It has the potential to speed up our processes, expand our access to knowledge and ideas and give us more time to focus on the parts of the process that we enjoy and bring value! I have no doubt that AI will present a huge shift in the way we practice. At the same time, I think we need to be very aware about what the value of the architect actually is. Architecture is not a linear field, it doesn't just go - idea (use mid journey), cost (use chatgpt), document (automate drawings). What makes good architecture and good architects, is being able to join all of these together in complex ways. The hardest part of architecture is being able to deliver buildings and amazing ideas while dealing with client demands, ever shrinking deadlines, budget, city planning, complex coordination and a myriad of personalities. This is much harder than many people think, and if you can walk this delicate tight rope, while being able to deliver an amazing building - then this is the true power of an architect. If we can take parts of AI, and use this to assist in each of these steps, and assist in a positive way, it will just help us to be able to walk this tightrope in an even more effective way.

2. WHAT DO YOU SEE AS THE KEY LIMITATIONS OF CURRENT AI TECHNOLOGIES

WHEN APPLIED TO ARCHITECTURAL DESIGN?

At the moment, I believe that the use of AI in architecture is in its infancy and it shows. In image generation we need to understand if we are doing this to create ideas or are we doing this to create a finished looking idea. I found this very hard, I would always be trying to shoe-horn a creative idea into a very prescriptive massing, and it never really was a good design idea, nor was it a good finished product. A big issue here, is the colleagues you are sharing these ideas with also need to recognise this. In my view, if we are talking about creative tools, you need to be totally loose and design architecture like you would if you were painting a painting, just developing ideas. Where I see the strength currently is just helping us with tasks. Image wise, it is very helpful for render teams, to populate scenes, change lighting etc. The biggest advantage for me at the moment is simply asking it questions for computational tasks, how would I do this in grasshopper and especially writing code. So at the moment it is linear, do this kind of tasks, but in the future, it will help us couple these tasks together, I have no doubt about that.

3. WHAT SPECIFIC SKILLS OR MINDSETS DO YOU BELIEVE ARCHITECTS NEED TO DEVELOP IN ORDER TO CRITICALLY AND EFFECTIVELY COLLABORATE WITH AI TECHNOLOGIES?

The only skill people need is the desire to learn and be curious. This should always be exercised against the question, why did I study architecture and what do I believe architecture is.

4. HOW IS 3XN CURRENTLY INTEGRATING AI AND MACHINE LEARNING INTO ITS ARCHITECTURAL WORKFLOWS?

As I mentioned earlier, we are using it for linear workflows at the moment. I have chat-gpt open constantly, using it as a personal assistant. How can I write that formula in excel, help me with this revit family formula, how can I do this in rhino, help me write this python script to use in revit. I mentioned earlier, but I think the image generation is fantastic, and I have some friends who are doing great work with this, but for a firm like 3XN, right now, I don't think it's totally there yet to be imbedded in each project.

5. IN YOUR VIEW, IS AI PRIMARILY TRANSFORMING THE METHODOLOGY OF ARCHITECTURAL PRACTICE, OR IS IT ALSO RESHAPING THE MEANING AND LANGUAGE OF ARCHITECTURE ITSELF?

In my opinion, right now, AI is changing the way we work, helping us be more efficient, more informed and be able to do linear tasks in a much more efficient way. I believe that the greatest architecture, is architecture that makes people feel something that they can not explain nor do they need a design background to be able to feel it. At this point in time, AI is not a substitute for creating this feel, but in the future, who knows. It's an exciting time to be practicing and to be practicing with an open mind.

6.0 DETAILED METHODOLOGY

The preceding chapters have traced the historical trajectory and theoretical debates surrounding the use of Artificial Intelligence (AI) in architecture, from early symbolic systems to contemporary developments in deep learning. This exploration included perspectives from scholars such as Mario Carpo, Neil Leach, and Lev Manovich and insights from practitioners including Wolf Prix, Thom Mayne, and Refik Anadol. Together, these voices illuminated both the transformative potential of AI in design workflows and the persistent challenges accompanying its integration—among them, the tension between computational optimization and architectural judgment, the opacity of algorithmic processes, and the unresolved question of how digital outputs translate into buildable form.

Against this backdrop, the methodological approach of this thesis is grounded in critical experimentation. Rather than conceptualizing AI as a neutral artifact, the research treats AI as an active participant in the design process characterized by operational biases, aesthetic tendencies, and limitations. This research chose the platform Midjourney, not because it was the most technically exhaustive in an established sense but because it operated effectively to support speculative conceptual exploration. The platform also satisfied the following aspects: It was accessible and responsive to natural language prompts and generated visually coherent visuals with no programming experience. It was an effective tool to question AI's potential at the ideation stage.

The methodology involves iterative prompting, selective image evaluation, comparative assessment with an eye to spatial

consistency and material representation, and how clearly the architect's intent is recorded. The work does not seek to conclude on finished designs, it seeks to understand how generative tools can provoke the early stages of architectural thought. Therefore, the research takes an exploratory stance that sees AI as not just replacing architectural reasoning but as a provocation for rethinking how we create, test, and visualize design ideas.

Notably, the methodological approach also cautions us about the limitations of such systems. While tools like Midjourney represent new potential avenues of visual ideation, they do not constitute architecture in itself; the product of these tools requires critical interpretation, translation, and often a re-assertion of the authorship of architecture. Through foregrounding this labor of interpretation, the methodology seeks to position AI as not a substitute direction in design innovation but a collaborative co-agent in something larger, epistemologically and architecturally.

6.1 ESTABLISHING SELECTION CRITERIA FROM PRIOR RE-SEARCH

The selection of an appropriate AI platform for this thesis was not a technical decision alone but a methodological response to the conceptual and practical challenges identified in prior chapters. Theoretical perspectives and case studies—ranging from symbolic to generative AI—highlighted recurring tensions between automation and authorship, integration and experimentation, and formal novelty versus architectural rigor.

With these contextual articulations in mind, the selection of toolkits entered a delineation of primary importance to align with the research question: How can AI meaningfully enhance architectural ideation?

The criteria below were developed from this process, representing both the selection criteria and a critical lens through which to evaluate the tool in the design experiments:

1. Generativity and Architectural Relevance The platform needed to offer generative abilities, especially generating visually suggestive and spatially engaging images. This criterion was an attempt to go beyond an optimization or procedural scripting approach, with the aim to support early-stage ideation, referring to what Refik Anadol describes as “machine hallucinations” and what Wolf Prix describes in his search for unknown forms. The objective was not resolved design but rather invoking new formal trajectories and speculative thinking.

2. Workflow Compatibility and Integration to Existing Model Here, as discussed around raster-to-vector translation and CAD/BIM interoperability in tools like ArchiGAN, the tool needed to reside within the pragmatic limitations of a thesis. While it was not expected to integrate completely into a professional

workflow, the tool needed to enable a structured and repeatable process without onerous technical demands and in a way that did not obfuscate the research.

3. Balance Between Designer Agency and Algorithmic Autonomy. The literature continually focused on how designers retain agency within generative systems—from Mayne's critique of intuition in the design process to Prix's focus on critical authorship and knowing the designer is tied to the cognitive activity of generation. Similarly, the platform was required to provide implies for meaningful control—such as prompt engineering—that could guide the generative process rather than opaque automation. This way, the architect would be active and engaged, not simply stagnant.

4. Output Quality and Stylistic Coherence Reflecting Mario Carpo's arguments about architectural legibility in the age of mass customization, the tool needed to produce high-resolution, stylistically coherent imagery suitable for aesthetic evaluation. The goal was not to achieve photorealism but to ensure that visual outputs could support iterative comparison and critical discussion—prerequisites for any generative design process to function meaningfully.

5. Flexibility and Iterative Responsiveness Drawing on the exploratory ethos emphasized by Tedeschi, Leach, and Mayne, the selected platform had to accommodate rapid prototyping and a wide range of design intentions. Its responsiveness to varied prompts and capacity for producing divergent outcomes was central to modeling a “search-based” design process—an approach that values iteration over finality.

6. Accessibility and Feasibility of Research. Lastly, the platform also had to be financially and technically accessible. It was important that the tool does not rely on advanced computing hardware or overly complicated installation processes and that there was an active user community for troubleshooting and sharing insights during the research period—for cohort continuity. Combining these criteria legitimized our decision to use Midjourney as the platform sufficiently aligned with the thesis objectives. More than a checklist, they form a framework for reflection from which we can establish our operational requirements to our conceptual objectives for critically assessing the role of AI in architectural ideation.

6.2 COMPARATIVE ANALYSIS

A comparative assessment of three prominent text-to-image generation platforms—Midjourney, Stable Diffusion, and DALL·E—was conducted to evaluate their suitability for architectural conceptualization based on the criteria defined in Section 6.1. While each system presents distinct affordances, Midjourney emerged as the most appropriate for this thesis, particularly due to its usability, output quality, and alignment with generative experimentation in architectural contexts.

Stable Diffusion is an open-source model that demonstrates excellent flexibility and potential for fine-tuning, which aligns with the objectives of early symbolic AI models, which aspired to accountability and user control. However, to effectively engage with Stable Diffusion, a great deal of technical confidence and local computational resources are needed, precluding implementation within this study's parameters. DALL·E provided excellent image-generation capabilities and evident semantic control. This testing instance did not feature the stylistic consistency and user-centric control appropriate for application in iterative design approaches.

Midjourney was adopted based on its balanced accessibility, responsiveness, and visual quality. It operates through a Discord-based chat interface, is based on diffusion modeling technology, and provides an intuitive yet powerful *modus operandi* to create architectural imagery. The system is designed around textual prompting, providing a working mechanism for designers to more or less adjust the outcome in ways similar to prompt-based design frameworks discussed earlier in terms of speculative workflows, as referenced within the work established by

DeepHimmelblau or Anadol's data performances. Thus, prompt engineering emerged as the principal way to govern Midjourney's generative behavior.

This interactive mode of control enables a productive negotiation between algorithmic autonomy and human design intent—a theme echoed in prior reflections by Prix and Mayne. While Midjourney's outputs remain non-editable raster images, their relative stylistic coherence and high resolution provided a sufficient foundation for visual analysis and iterative development. Though not explicitly transparent, patterns in the relationship between textual input and the visual result could be identified through repeated testing, offering a degree of empirical predictability that partially offsets the platform's black-box opacity. In addition, Midjourney's cloud-based format removed the necessity for GPU processing locally, and an active community made the trial-and-error learning process faster. This allowed for the frequency of prompt iterations needed to satisfy the experimental objectives of the thesis. While Midjourney is not directly integrated with vector CAD or BIM environments (similar to ArchiGAN), the ability to generate and iterate conceptual images in such an expedited manner justified its use.

Midjourney's combination of visual fidelity, user engagement, and practical accessibility made it a fitting tool for exploring how generative AI can augment the early stages of architectural ideation.

6.3 ACKNOWLEDGING ALTERNATIVE WORKFLOWS

In assessing AI tools relevant to architectural design, we must consider alternative platforms that have emerged within the practice, in particular, ComfyUI, a node-based graphical user interface for working with Stable Diffusion models.

With ComfyUI, we have the potential to minutely control the image generation pipeline as we continually alter numerous variables in the process of model selection, prompt weighting, upscaling, inpainting, and modules like ControlNet (ComfyU). This pipeline manipulation has practical work applications that require precise expressions, such as stylized consistency and refinements for architectural visualizations. This increases ComfyUI's attractiveness for professional workflows, especially late in the work process and for presenting images. For example, Henning Larsen has publicly indicated they are using state-of-the-art AI tools — in this case, known variants of Stable Diffusion — for rendering quality and visual consistency presentation across projects (Henning Larsen, 2023). Additionally, conversations on professional platforms — like in the ArchiTech network interviewing UNStudio — have indicated that ComfyUI and other similar systems are used to refine design communication outputs (ArchiTech, 2025).

While it has many strong points, ComfyUI was eventually cut from the core method of this thesis due to a difference in focus. While ComfyUI performs strongly to refine the design previously made and gives control to a level suitable to complex visualization workflows, the research in this thesis sits at the beginning of the design workflow. The aim was to look at how generative AI could enhance the brief ideation of architecture through rapid iterations and exhaustive exploration

of form. In this case, the technical setup and steep learning curve accompanying ComfyUI would have limited spontaneous explorative experimentation rather than increasing it. In comparison, Midjourney affords a more accessible, prompt-driven interface better suited to the aims of early-stage conceptual generation. There is less control over output parameters, and outputs are more difficult to control; however, as it can generate a large number of coherent images quickly, Midjourney supports a search-based design method that values breadth over refinement.

The choice to emphasize Midjourney reveals a methodological preference for conceptual provocation and rapid visual exploration—areas where ComfyUI could be leveraged as a shorter-term framing device. This point contrasts another important distinction: AI tools are not interchangeable and correlate more accurately with different design process phases. ComfyUI may be outstanding for visualization and communication that has depth; however, the research that this paper was concerned with, as it pertained to epistemological and creative inquiries on generative AI in early design, fit with Midjourney more naturally.

6.4 CONCLUSION ON TOOL SELECTION

Midjourney, as the primary tool for this thesis, will be considered a methodological choice that originates from the critical discussion about the evolving role of AI in architectural design. After having discussed the context historically and theoretically earlier, especially regarding generative tools, authorship, and the use of exploratory workflows, Midjourney's accessibility, quick visualization of ideas, and reactivity to abstract text prompts made it an ideal choice to research how AI might contribute to architect's early stages of developing ideas.

However, the decision also recognized several significant constraints of the Midjourney platform. First, Midjourney operates entirely in a 2D, raster environment, meaning that any form of output from the tool could not be directly translated into the vector-based systems of CAD or BIM, something noted in earlier critiques of systems like ArchiGAN. Second, Midjourney is still a proprietary, black-box system, which means users have little or no knowledge of the training dataset, model architecture, or how it functions internally, raising ethical awareness of transparency and authorship. In addition, the selected visual outcomes are aesthetically analogous but have a highly curated charged output, which can create an added level of homogenization across outputs and ostensibly orient toward specific stylistic norms that structure design outcomes rather than enabling open-ended exploration.

Despite these constraints, Midjourney's generative capacity, low technical barrier, and iterative flexibility made it a productive tool for exploring the thesis's central question: How might AI serve as a co-creative partner in

architectural design, not by producing finished solutions but by provoking new directions in form-making? The following sections document how this tool was integrated into a defined workflow, assessing its contributions and limitations in shaping the speculative design outcomes of this research.

7.0 INTRODUCTION TO THE DESIGN WORKFLOW

This chapter details the design workflow established for the architectural experimentation within this thesis. Building on the methodological rationale outlined in Chapter 6, the workflow focuses on the intentional use of Midjourney as a generative and speculative engaging tool early in the conceptual design process. The process does not consider AI-generated images to be final solutions but somewhat speculative provocations—visual prompts that create pathways of alternative design narratives. The use of Midjourney is established during the ideation phase of the process, where its ability to produce high-resolution, compelling images from abstract text prompts is seen as a way to engage in architectural thinking. While these images are often hyperrealistic and stylistically curated, they are uncritically not integrated into the design process. The rationale of the workflow is related to broadening the designer's assumptions, indicating unthought-of spatial types, and continuing to broaden the designer's list of possibilities.

The second phase in the process is converting these chosen AI-generated images into three-dimensional architectural models. This process—transforming a 2D raster output to a 3D spatial it's a reinterpretation that is thoughtfully dealing with the in-between of speculative visualization and architectural representation. Thus, the workflow does have fidelity in mind, but does not expect to be the same, instead, using AI outputs as conceptual scaffolding to be extrapolated and developed further by conventional modelling tools. This approach allows the thesis to explore AI not simply as a tool for formal generation but as a co-author in the iterative design process—one whose suggestions must be critically mediated

through architectural knowledge, spatial reasoning, and material logic.

7.1 INITIAL SETUP: DEFINING CONTEXT AND CONSTRAINTS

The design workflow sets the spatial and contextual parameters for the AI's generative process. One of the most substantial initial decisions will be choosing the viewpoint of the image generation. After experimenting with conventional architectural representations, plans, sections, and axonometric projections, a birds-eye view was the most effective viewpoint to work with Midjourney. This viewpoint provided the most comprehensive visibility of the site and context. It also provided the AI with adequate visual language to convey spaces and urban morphology.

To initiate the generation process, two reference images are prepared: an aerial photograph of the project site, typically sourced from publicly available satellite imagery platforms such as Google Earth, and a rendered bird's-eye view from a basic 3D digital model of the site. This model is created using common architectural software such as Rhino, Revit, or SketchUp. The modeling platform itself is secondary; what matters is that the representation includes the massing and immediate context of the proposed architectural intervention.

These placeholder volumes are not intended to be more than schematic. They aim to define the project site and construction envelope and provide the intended spatial boundaries in the AI frame of reference. It is important to ensure that the model renders and the aerial photograph align—specifically in viewpoint and scale, as it will allow the AI to produce more coherent visual serves inputs. Both images are then uploaded into Midjourney as reference images. At this stage, no text prompt will be entered; this lack of input will require the AI output solely

by visual inference from the reference images and training data. It also affords some visibility into the model's latent bias and generative dispositions while enabling the designer to see how it performs when free from any textual prompt. Adjustments are made to parameters such as aspect ratio and image weight (that is, the relative contribution of each image to the output) from Midjourney's prompt settings (Midjourney, 2024). These variables offer the designer a limited but meaningful degree of control—setting the stage for more directed experimentation in subsequent stages.

7.2 ITERATIVE GENERATION AND REFINEMENT

Following the initial reference-based generation, Midjourney produces four image variations. These typically depict abstract massings—often rendered as white volumetric forms—situated within a synthesized context extrapolated from the input bird's-eye images. It is common for slight distortions to exist in the model's interpretation of the site's geometry or urban context because the generative model has its interpretative possibilities. It may take several iterations to capture the expected level of contextual fidelity. At this early developmental stage, we are more concerned with the accuracy of the larger site environment than the specific form of massing, which we want to see as a shapeable condition for future development. We begin a generative design workflow by selecting one image from the first generation that we feel is promising. It serves as a generative seed (Midjourney, n.d.-a) and will be the first of a series of progressive iterations. The workflow is done in three phases: Selection of a seed image from the first generation. Subsequent generation cycles are built on the selected image—incremental textual prompt adjustments, either immediately or in later cycles.

Refinement is conducted using Midjourney's "Remix Mode" (Midjourney, 2024-b), which allows the prompt to be modified between iterations. Each adjusted prompt influences the following output, enabling the designer to guide the development process through focused interventions. If the prompt remains static, the AI tends to produce minor visual refinements based on its initial interpretation, often gravitating toward generic or aesthetically average outcomes. Prompt modulation becomes essential for steering the evolution of the image, allowing the designer

to articulate complex intentions across multiple cycles in manageable, interpretable steps.

This operational format showcases the previous "search-based" design method: the designer is not simply another form giver; they act like an explorer in a broad generative space and influence the course of the exploration through comments and controlled prompts. The designer has to carefully control two related strategies to make this iterative process work: Incremental Prompt Change: Large changes with prompts often create strange or confusing outputs. On the other hand, more minor and gradual changes to a prompt introduced every 10-15 iterations allow the main visual ideas to remain intact while encouraging formal development.

Ongoing Expert Evaluation: The designer must interpret and curate the outputs at each stage, selecting the most promising trajectory and crafting the following prompt accordingly. This evaluative role demands architectural expertise, aesthetic sensibility, and a clear understanding of project goals. While often visually engaging, AI outputs are not inherently meaningful; they represent statistically plausible image constructs rather than resolved architectural proposals. Thus, the designer's role remains central: not to produce images by hand, but to critically interrogate, filter, and guide machine-generated variations within a conceptually structured workflow.

7.3 ACHIEVING DETAIL AND TRANSITIONING TO PHOTO-REALISM

As the iterative process progresses, the design focus shifts from general massing and formal articulation toward finer details, including architectural elements, materiality, lighting, and contextual integration. At this point, textual prompts begin to specify more aspects of qualities themselves—including descriptions of, for example, facade articulation, surface finishes, or atmospheric condition—to help refine and visually enrich the output.

That said, distinguishing between visual realism and architectural photorealism is important. Midjourney is very effective at providing hyper-realistic imagery with high visual impact and photo-like ambiguity, but the outputs do not usually include specificity suitable for the level of construction detail. Architectural imperatives—such as material joints, detailed window framing, or structural logic—are often abstracted or stylized references.

This aspect is inherent in an image-modeled platform with a statistical image-synthesis capability to operate in a 2D image space that cannot interrogate volumetric references such as sectional views or maintain orthographic projections. Moreover, while effective for contextual legibility, the bird's-eye viewpoint that anchors the workflow restricts detailed engagement with human-scale elements or spatial sequencing. Consequently, Midjourney's outputs at this stage are best understood as provocative visualizations rather than resolved architectural renderings. To mitigate these constraints and refine localized aspects of the design, Midjourney provides a targeted editing feature known as "Vary (Region)" (Midjourney, n.d.). This enables the user to select specific portions of an image, such as

a part of a facade or landscape area, and then update the prompts selectively. The tool is intended for more firmly focused visual development, eliminates the entire regeneration of the image, and means tighter control over the composition and density of detail in specific areas of the overall visual. (See FIGURES 59–61 for representative examples.) While this type of editorial engagement is interesting and valuable, it is limited to the generative model of the platform. It is not a replacement for more thoroughly documented architectural modeling. The examples support a degree of elaboration on narratives and atmospheric quality of the concept; a more vigorous interpretation as an architectural representation could be completed with more traditional CAD or BIM tools outside of the Midjourney space.

7.4 BRIDGING THE GAP: FROM 2D IMAGE TO 3D ARCHITECTURAL MODEL

After developing an interesting 2D output through the iterative nature of the AI process, the next crucial transition must happen: transforming the AI-created image into a conventional 3D architectural model. This is important for moving from the speculative visualization stage of the design into an intelligible and buildable spatial form. It allows for architectural evaluative measures of spatial organization, structural logic, and material qualities, areas where 2D generative tools like Midjourney remain inherently restricted.

SPATIAL REASONING AND INTERPRETIVE TRANSLATION.

This is a deliberate interpretative translation. The AI output does not have explicit depth, sectional logic, or buildable detail. It is a suggestive artifact—a visual prompt rather than a design intent. The human designer must remove spatial clues within the image to imaginatively derive a three-dimensional arrangement, inferred dimensions, and architectural logic. The designer is resuming elements such as massing hierarchies, site relationships, and material qualities by judgment, experience, and model making.

MODELING PROCESS.

Usually, the intent for modeling begins with volumetric studies in a 3D software program. The primary volumes are identified based on the overall proportions evident in the AI image, updated iteratively to have a structural capacity, a reasonable form of materials and operation appropriately, and make sense in the context of a space. Although previous prompts may provide some deliberate direction in early iterations (for

example, suggesting materials or aspects of facades), the explicit architecture as an articulation of a lettered form is rarely translated. Instead, the model allows for a form of reconciliation—between the image suggestive of the AI and the restraints as an architect.

INTEGRATING PROGRAMMATIC LOGIC.

It is critical to note that functional layout and circulation are not merely appended to a finalized form. Instead, these aspects must be envisioned throughout the iterative image generation process. While Midjourney operates at the level of visual exteriority, the architect must continuously project potential internal configurations onto emerging forms. This entails a simultaneous mental rehearsal of functional zoning, access strategies, and volumetric subdivision. As such, selecting or refining an AI-generated image is already informed by a latent spatial rationale—even before explicit 3D modeling begins.

SOFTWARE ENVIRONMENT AND WORKFLOW CONTINUITY

For this thesis, Autodesk Revit was chosen due to its capacity for detailed parametric modeling and Building Information Modeling (BIM) integration. Beyond facilitating the reconstruction of form, Revit enables the production of conventional architectural representations—floor plans, sections, and elevations—bridging the gap between speculative images and professional design deliverables. This ability to produce intelligible documents highlights the importance of human involvement in making visual provocation into architectural particularity.

7.5 CONCLUSION: AN AI-AUGMENTED WORKFLOW

In this chapter, we have discussed a methodical approach to applying Midjourney in the initial conceptual stages of architectural design. While AI is not to be considered a fully autonomous generator of resolved products, this workflow depicts an image-generating position as a collaborative device—enabling new visual thinking but explicitly subject to continual human interpretation and critical management.

The entire process—from identifying spatial limitations and seeding a set of first inputs through a series of iterative steps to translating the outcome into a three-dimensional architectural mode—has shown how designers can implement AI tools into a designer-led process that balances generative capability with architectural reasoning and makes use of the benefits of computational image-making, all without abdicating authorship or disciplinary integrity. That said, the way we utilized the AI was not without limitations. Whilst Midjourney opens up potential wider access into generative design, we still need to quickly become accustomed to the platform while recognizing that proper use of the platform requires a high level of precision in prompt engineering, an acute sense of aesthetics, and a good understanding of managing iterations. Achieving consistent and contextually meaningful outputs often demands dozens—if not hundreds—of trial prompts and an ability to anticipate and interpret the tool's often unpredictable behavior. These factors underscore that AI tools, despite their apparent accessibility, necessitate a specialized form of digital literacy that blends architectural expertise with emerging computational fluency.

The following chapter builds upon this workflow to explore its practical implications through a series of experimental case studies, testing the extent to which AI-augmented processes might meaningfully contribute to contemporary architectural form-finding

8.0 PROMPT ENGINEERING AS DESIGN METHODOLOGY

This chapter illustrates prompt engineering in detail as a fundamental design logic of the thesis. In design processes, specifically generative AI, prompt engineering refers to purposeful textual inputs that trigger how image-generating models (e.g., Midjourney) interpret and visualize architectural concepts. Midjourney, in their manual, states that the difference between parametric control (i.e., design logic that is coded explicitly) and prompt engineering is that prompt engineering uses language as a form of influence—indirectly, probabilistically, but structured—the prompt acts as a bridge between conceptual speculation and algorithmic rendering.

An extended experiment with Midjourney, conducted between July 13 and October 13, 2023, yielded 1,915 distinct prompt generations, each yielding a 4-image grid with approximately 7,660 distinct visual iterations. Rather than completing a design workflow, this research was inductive, testing, evaluating, and modifying prompts and processes to backtrack toward a credible method. Although arbitrary, this empirical iteration provided insight into the values and significance embedded in Midjourney's interpretation of architectural vernacular and compositional cues.

MIDJOURNEY'S PROMPT MECHANICS

Midjourney runs through a straightforward command line—primarily through the /imagine command in Discord. It instructs to create a grid of four AI-generated images based on a text-based prompt. Generally, Midjourney does not “understand” language in a traditional grammatical sense; it can statistically parse words and phrases to interpolate them against its internal weights and

styles it learned during training.

Successful prompts on Midjourney are those that balance descriptive richness with structural clarity. The software does not interpret full-sentence grammar, so commands like “a minimalist museum in Akita with a charred wood facade” are not read as prose but as a weighted list of keywords. Therefore, specific nouns (e.g., “museum,” “charred wood,” “Akita”) and modifiers (e.g., “minimalist,” “hyper-realistic”) are used to signal Midjourney's algorithm towards certain visual conventions. Key prompt mechanics include prompt length and simplicity, which are short prompts that default to Midjourney's aesthetic biases. Descriptive, moderately long prompts yield more directed results, but overly complex inputs can dilute the influence of each keyword.

Weighting Syntax (::): Users can specify the weight of specific words by entering a:: followed by a number (e.g., AerialView::3), thus increasing the weight of aerial view.

Negative Prompts (--no): Negative prompts will remove unwanted things (e.g., --no people).

Aspect Ratio (--ar): This specifies the image's aspect ratio (e.g., --ar 16:9). It helps compose a particularly designed perspective that prioritizes architecture.

Styling and stylization options (--style, --stylize): The styling and stylization options define how strongly Midjourney gets applied through its own aesthetic engine.

Version Control (--v): The version control is defined what version of the AI model you are using.

These prompt design rules are not merely technical settings but compositional strategies. Unlike software like Grasshopper, where

logic is codified through nodes and scripts, Midjourney's design process is more speculative. The prompt becomes an architectural sketch in text, balancing control and ambiguity. In the course of this research, early prompts that relied on generic architectural terminology (e.g., “modern Japanese crafts museum, flat site, minimalist form”) produced erratic and unconvincing results. As the iterations progressed, prompts were restructured to introduce architectural references (e.g., “inspired by Terunobu Fujimori,” “traditional Japanese silhouettes,” “Shou Sugi Ban”) and compositional directives (e.g., “AerialView::3,” “bird's-eye perspective”).

As described, these modifications pushed Midjourney to yield more architecturally consistent results. Prompt design developed alongside the design coordination. In the early stages (around prompts 1-35), the focus was on designing the spatial organization and massing rules. The prompts had more prosaic descriptions of the site, plan, and volume:

“Akita Crafts Museum::2, Aerial View::3, birdseye view, Kengo Kuma-inspired, zen-like minimalism, modern Japanese architecture, flat landscape”.

This syntax used Midjourney's tagging hierarchy (::2) to weigh each reference, while aerial perspective ensured a constant framing technique.

From prompts 36–70, the emphasis shifted toward typological articulation and architectural detailing. Influences from architects like Terunobu Fujimori, Studio KO, Peter Zumthor, and SANAA were explicitly referenced to modulate materiality, fenestration, and facade rhythm. Phrases like “Shou Sugi Ban cladding,” “charred timber,” and “minimalist

silhouette” began yielding consistently legible tectonic languages. At this stage, prompt engineering began operating less as thematic input and more as an embedded design grammar.

Future prompts continued over the horizon for the richness of the interior atmosphere, sequence of spaces, and environmental narratives in the text. Some prompts referenced elements like “Olson Kundig-inspired interiors,” “Issey Miyake textile layering,” and “kinetic facade elements” that served both as compositional prompts and conceptual frames. Of note, and due to Midjourney's responsiveness, agent names or studios needed to be adjusted again: entities such as “Fujimori::3” or “SANAA::1” needed to be adjusted to those levels to retain influence but not overly constrain results.


In the final phase (iterations 70–130+), focus turned to landscape integration and environmental framing. Terms such as “Dan Pearson-inspired wild gardens,” “central courtyard,” “reflecting pond,” and “landscape:2” enabled a continuity between architectural objects and ecological settings. Contextual realism was reinforced with directives like “daylight,” “hyper-realism,” and selective material tags (e.g., “slate and basalt,” “vintage timber”). The final operative prompt—refined across iterative testing—was: “ModernMuseum:: of crafts, AerialView::1 of the Akita Crafts Museum inspired by TerunobuFujimori::3 in a flat landscape, featuring a sculptural architectural form with an intriguing Facade. The architecture fuses transformative designs from StudioKo::2, playful geometries by TezukaArchitects::2, and biophilic concepts of Vo Trong Nghia Architects. Internally, spaces echo OlsonKundig's::2 warmth and mechanics, using reclaimed dark

timber. The museum's facade showcases innovative elements like Shou Sugi Ban, a deeply charred wood, creating a textured, captivating surface. Traditional Japanese roof silhouettes blend intricate wooden joinery with modern detailing. Inspired by visual artist James Turrell's², sound chambers enhance the visitor's experience. Vintage timber contrasts with innovative materials such as ferrofluid sculptures and kinetic Walls—monumental¹ wooden beams reminiscent of ancient temples, playing with light and shadow. The design draws inspiration from Issey Miyake's textile craftwork, weaving intricate patterns and layers, highlighting the essence of hand-made crafts. Limited fenestration, accentuated by slate and basalt² elements, refines the interplay of artificial and natural light. The landscape² marries Dan Pearson's² wild aesthetic."

Although richly visual, such prompts also emphasized Midjourney's interpretive opacity. Outcomes typically remained unpredictable, often requiring repeated graphics adjustment and regeneration. The number of iterations resulted from exploring design and a necessary consequence of Midjourney's inherent unpredictability and heavy stylistic influence.

Prompt engineering in this thesis ultimately functioned as both tool and terrain—an evolving space where architectural intentions were negotiated through language. The process demanded design literacy, linguistic sensitivity, and a deep understanding of how syntax and reference hierarchy guide generative AI's visual logic.

FIGURE 24
ITERATIVE EVALUATION FROM THE SKETCH
created with Midjourney by Tim Fu



9.0 CONTEXTUALIZING THE AKITA COMPETITION: TRADITION, CRAFT, AND PERSONAL ENGAGEMENT

This chapter situates the design investigation of this thesis within the framework of an open international architecture competition held in Akita, Japan. Organized around preserving and revitalizing Japan's intangible cultural heritage, the competition called for designing a crafts museum—a space to sustain the country's endangered artisanal traditions in a time increasingly dominated by speed, automation, and cultural flattening.

In addition to its functional program, the brief raised fundamental inquiries about the relationship between built form and the transmission of memory: How can architecture be a tool for cultural continuation, not just representation? What might episodic spatial experience suggest about practices of slowness, repetition, and embodied knowing?

The rationale for participating in this particular competition was motivated by personal interests and not by strategic merit. It proposed an investigation into a personal long-standing interest in Japanese architecture—not for the eye's sake, but that it stages values such as focus, discipline, materiality, and restraint. These values resonated with the craft's ethos, and the brief invited an architectural exploration of that relationship. While the project could have been based on a similar brief in another context, Japan's rich, deeply layered character of craft past and present, and the last 150 years of the very institutionalization and evolution of craft, offered a rich context for critical and creative engagement.

What made the Akita competition so significant was that it framed tradition not as nostalgia but as an active, adaptive, participatory practice. The brief clarifies that traditional

craft has many current-day threats—decreasing apprenticeship, rural depopulation, and changing economic focus—but that it expects an architectural answer that can intervene actively. The intention is not to create a museum in the literal sense but a blended cultural infrastructure through which knowledge can be made, passed on, and transformed.

This intention aligned closely with the methodological concerns of the thesis, which sought to examine how computational and AI-assisted workflows could address not only form-

making, but also the cultural narratives embedded within architectural production.

Akita, as a site, brings specificity to these themes. Known for four traditional crafts officially recognized by Japan's Ministry of Economy, Trade and Industry—Kabazaiku (cherry bark work), Kawatsura lacquerware, Odate Magewappa (bentwood vessels), and Akita cedar cooperage—the region embodies an intensely local expression of Japan's broader craft heritage. By grounding the project in Akita, the competition rooted its ambitions in place-based authenticity while participating in the broader discourse on cultural sustainability and "soft power" diplomacy—where tradition becomes a medium through which a nation negotiates its global identity.

9.1 SITE, PROGRAM, AND TECHNICAL PARAMETERS

The competition brief of the Akita Crafts Museum included a very particular set of parameters in terms of the spatial, regulatory, and functional envelopes. These limitations also played a role in the iterative logic of design that was employed in this thesis because they constrained not only the architectural massing, but also the proposal's infrastructural and experiential conditions.

The proposed project site in Akita, Japan, is approximately 11,398 square meters. Its position falls between residential housing and farmland, presenting a semi-rural condition that must be considered regarding scale, integration, and visual permeability. The zoning only allows for a maximum height of 25 meters, a maximum ground coverage of 50%, and a maximum FAR of 1.0. These constraints defined the volume that could be built reasonably around an area of 11,398 square meters, and only half of that could be built in footprint on the ground. The low-rise restrictions particularly influenced the vertical ways massing studies could be articulated and the architectural language of large horizontal spans that arose from the Midjourney iterations.

Program Overview The brief required a multi-programmed cultural institution that would not merely exhibit crafts but act as a hub for learning, production, and communal interaction.

The total functional program includes spaces for making, displaying, performing, studying, eating, dwelling, and resting. These spaces must coexist without hierarchy—each program

plays a role in supporting the continuity and transmission of traditional Japanese craft.

The anticipated space requirements include:

Artist Accommodation: 10–20 guest rooms 25 sqm = 250–500 sqm. These accommodation units facilitate an extended commitment from practitioners that promotes intense apprenticeship and cross-cultural interactions.

Auditorium: approximately 1,000 sqm, to include stage, seating, mechanical space, green room, storage, loading dock, lounge, and areas for front of the house. The auditorium is intended for performance, lecture, demonstration, and symposium-related craft culture.

Exhibition Space / Art Gallery: ~300 sqm. It is intended to feature historical materials and contemporary re-interpretations that straddle static display and active product(s) formalities.

Workshops / Craft Bazaar : 400-500 sqm, a main functional component of the museum that facilitates visitor engagement and professional quality products. The layout should have ventilation, movable furniture, and storage for tools.

Library: 200 sqm, conceived as a quieter place for mediation and reading, to archive or develop research, for the general public and faculty generating craft-based scholarship.

Food Court: 600 sqm, to include eatery, kitchen, and back-of-house service space. This area provides sustenance, enables social interaction, and may also be attractive for public engagement in the facility.

Amphitheatre: approximately 800 sqm, designed for outdoor or semi-covered uses. It supports seasonal events, open-air

demonstrations, and show exhibitions for community audiences.

Souvenir Shop: ~60 sqm, focused on high-quality, locally made craft items. This space also serves as a secondary point of education and dissemination.

Visitor Centre: 60–100 sqm, including a ticket booth, reception area, and information kiosk. As the threshold space between the city and the institution, this area must express the museum’s identity in microcosm.

Administration Offices: ~200 sqm, including a director’s office, curator’s room, conference space, storage, and pantry. These spaces must be functionally efficient yet spatially discreet.

Circulation and Services: Estimated at 30% of the total built area, including vertical and horizontal movement systems, restrooms, service corridors, and maintenance access.

Parking and Landscape Areas: The project calls for landscape integration, including a terrace garden, green roof systems, and curated ground-level gardens informed by Japanese landscape principles. These components are not decorative, but integral—designed to extend the sensorial and educational dimensions of the museum.

10 CASE STUDIES / UCCA CLAY MUSEUM

Architects: Kengo Kuma & Associates
Total Area: 3,437 m²
Completion: 2024
Photographer: Fangfang Tian, Eiichi Kano

The UCCA Clay Museum, developed by Kengo Kuma & Associates and completed in 2024, is located in the city of Jingdezhen, historically revered as the center of porcelain production in China. The 3,437 m² museum encapsulates Kuma’s long-held “anti-object” philosophy, wherein the architecture is concerned with dissolving into its environment and not asserting dominance over it. As Kengo Kuma (2008) argues, “Architecture should no longer stand as a monumental object but should instead dissolve into the environment. The use of materials and transparency can achieve this goal.”

MATERIAL ENGAGEMENT AND CONCEPTUAL FRAMEWORK

Jingdezhen’s historical legacy with ceramics is not merely represented in the museum—it is materially and procedurally embedded within the architecture itself. Rather than presenting porcelain solely as an artifact, the museum emphasizes process and transformation. Amanda Game (2016) observes that “Kuma’s porcelain museum design integrates the crafting process as part of the visitor experience, allowing them to witness and understand the transformation from raw clay to delicate porcelain” (p. X). This design decision foregrounds a pedagogical narrative wherein making and exhibition co-exist.

The use of clay, both as a thematic reference and literal building material, affirms the

architectural ambition to engage in what Barry W. Brownell (2023) refers to as “material empathy”—a condition in which the act of making becomes central to architectural expression. Brownell explains, “In Kuma’s approach, the act of making is as important as the finished object, transforming museum spaces into working studios where visitors witness the clay’s metamorphosis”.

SPATIAL STRATEGY AND VISITOR EXPERIENCE

Internally, the spatial organization avoids conventional gallery compartmentalization. Instead, the museum is structured around open and semi-open sequences, allowing fluid visual

access to crafting. These spatial strategies encourage a reconsideration of the museum typology—not merely as a container of objects, but as a dynamic interface between the material, maker, and viewer. The architecture inhabits the artifacts but also acts as a didactic instrument that amplifies the craft’s cultural significance through its form and materiality.

o



FIGURE 25
Kengo Kuma & Associates. (2024, October 30). UCCA Clay Museum. ArchDaily. <https://www.archdaily.com/1022949/ucca-clay-museum-kengo-kuma-and-associates>



FIGURE 26
Kengo Kuma & Associates. (2024, October 30). UCCA Clay Museum. ArchDaily. <https://www.archdaily.com/1022949/ucca-clay-museum-kengo-kuma-and-associates>

10.1 CASE STUDIES / YUSUHARA WOODEN BRIDGE MUSEUM

Architects: Kengo Kuma & Associates
Total Area: 14,736 m²
Completion Year: 2011
Photography: Takumi Ota Photography

The Yusuvara Wooden Bridge Museum, completed in 2011 by Kengo Kuma & Associates, is located in the mountainous region of Kochi Prefecture, Japan. Spanning 14,736 m², the project creates connections between several functional, spatial, and symbolic roles in one architectural gesture. It is both a connecting infrastructure—connecting public buildings that have been separated in different zones before—and a hybrid cultural facility containing artist residency spaces, workshops, and exhibition functions all in one project.

TYPOLOGY AND PROGRAM

Contrary to a typical project where the museum typology is enacted spatially and schematically, the Yusuvara project is a linear bridge structure that crosses over a roadway, which is simultaneous circulation and inhabitation. According to the project documentation (Divisare, 2024.), “This is a plan to link two public buildings with a bridge-typed facility... It functions not only as a passage between the two facilities

but also as an accommodation and workshop” (para. 1). The architecture thereby embodies a programmatic interweaving of civic infrastructure and cultural program, challenging conventional separations between public service and cultural production.

MATERIAL STRATEGY AND CONSTRUCTION TECHNIQUES

The design emphasizes locally sourced timber and references traditional Japanese joinery methods, particularly through the repetition of the *hangi* joint—a variation of the traditional Tokyo bracket system used in temple architecture. As the architects explain, “By repeating this method persistently to the scale of a bridge, we were able to present a completely new expression of architecture, while keeping its bona fide Asian appearance” (Divisare, 2024, para. 2). This approach reinforces material continuity with vernacular traditions while exploring new structural configurations.

SPATIAL EXPERIENCE AND CULTURAL POSITIONING

Rather than acting as a passive vessel for static displays, the structure accommodates participatory cultural activities, including artist residencies and craft production. The spatial sequencing—alternating between enclosed and semi-open conditions—affords varied experiential encounters with the material logic of timber. Mass timber is not aestheticized as surface treatment but contextualized as a tectonic and expressive system in which craft plays an active spatial and structural quality.



FIGURE 27

Kengo Kuma & Associates. (2012, February 16). Yusuvara Wooden Bridge Museum. ArchDaily. <https://www.archdaily.com/199906/yusuvara-wooden-bridge-museum-kengo-kuma-associates>

10.2 CASE STUDIES / WIS- DOME STOCKHOLM

Architects: Elding Oscarson
Area: 1,325 m²
Year: 2023
Photographs: Mikael Olsson

The Wisdome project, housed at the Tekniska Museet, marks an opportunity for the museum in Stockholm to carry staple immersive science communication experiences. Designed by Swedish architecture firm Elding Oscarson, the Wisdome project represents a landmark digitally fabricated timber architecture project. As Block, Boller, DeWolf, Pauli, and Kaufmann (2024) note, the project also employs extensive amounts of parametric modeling, which allows precise control of complex curvature and structural requirements.

Architecturally, the Wisdome project is distinct from typical museum typologies because the dome is incorporated into the design as part of the overall building, routing the dome into a free-form laminated timber roof of 26 by 48 meters. As seen with ArchDaily (2024), this not only provides visual coherence to the building but incorporates the dome as the focal spatial and symbolic element: “Conventionally, the program would generate a low volume with a protruding dome, but the dome is given a focal position under a free-form timber structure.”

This typological inversion enables a stronger architectural identity aligned with programmatic ambitions and structural clarity.

Experientially, the dome offers a 360-degree projection environment that immerses visitors in educational content related to science and nature. The timber

structure around it softens the technological aesthetic and anchors the project in sustainable construction logic. In its fusion of computational design, engineered wood, and educational narrative, Wisdome Stockholm exemplifies a forward-thinking synergy between content, form, and technique (ArchDaily, 2024).

FIGURE 28-29

Elding Oscarson. (2024, March 22). Wisdome Stockholm. ArchDaily. <https://www.archdaily.com/1014815/wisdome-stockholm-elding-oscarson>



10.3 CASE STUDIES / SWATCH & OMEGA CAMPUS

Architects: Shigeru Ban Architects + Jean de Gastines
Location: Biel/Bienne, Switzerland
Area: 25,000 m²
Year: 2019
Photographers: Didier Boy de la Tour, Philippe Zinniker.

The Swatch & Omega Campus, designed by Shigeru Ban Architects in partnership with Jean de Gastines, is located in Biel/Bienne, Switzerland, and accommodates private and public functions. The project includes the Swatch Headquarters, the Omega Factory, and the Cité du Temps, a museum and conference center.

A key architectural component is the timber gridshell roof system, defining the space of the three buildings and making the three separate buildings feel entirely distinct in formal identity. Place and Wharf (2016) noted that Shigeru Ban and de Gastines have been instrumental in furthering the conversation around large-scale architecture and gridshell types. In this instance, 7,700 individually programmed timber elements create the curves of the shell, each one programmed with particular objectives through bespoke computational tools to maximize accuracy and limit waste (ArchDaily, 2019). The architectural work strikes a balance of structural innovation and experiential clarity: the form of Swatch is playful and organic, while the Omega Factory is more orthogonal, and the Cité du Temps acts as a mediator element that goes beyond the requirements of separate buildings, while still providing space for exhibitions. The architects use these spatial and formal strategies to weave together parametric experimentation with straightforward brand narratives.

The project exemplifies how computational timber structure can advance corporate and cultural agendas. It meets sustainability goals, offers an expressive materiality, and fosters immersive visitor experiences, all through a unifying design ethos centered on innovation (ArchDaily, 2019).

FIGURE 30-31

Shigeru Ban Architects. (2019, October 10). Swatch and Omega Campus. ArchDaily. <https://www.archdaily.com/926166/swatch-and-omega-campus-shigeru-ban-architects>



11 AKITA IN CONTEXT: LANDSCAPE, HISTORY, AND CONTEMPORARY CONDITION

Understanding Akita beyond the frame of the competition brief requires an examination of its geographical setting, demographic dynamics, cultural geography, and regional socioeconomic transformations. This chapter analyzes Akita not as a passive backdrop but as an active agent in shaping architectural strategies. The site's environmental forces, population trends, and historical context directly inform the design framework established in this thesis.

11.1 GEOGRAPHICAL SETTING AND LANDSCAPE CHARACTER

Akita Prefecture is located in the northwestern part of Japan's Tōhoku region, facing the Sea of Japan. The region is characterized by a mountainous inland topography and fertile coastal plains, creating a landscape defined by firm seasonal shifts, significant snowfall, and a deep relationship with forest and water systems (Yoshino, 1992).

The site of the project, framed by a residential area and agriculture fields still actively farmed, is representative of Akita's semi-rural morphology, where modern infrastructure meets agricultural land.

Not only the natural environment and landforms can do interesting work in organizing how we build and use space - while a site may only be modeled after the familiar geography of Akita, the Akita cedar is not simply a botanical artifact but an economic artifact and cultural marker that is central to the craft traditions of cooperage (barrel making) and bentwood vessels. Climatic constraints such as heavy winter snow loads and humid summers influence building orientation, envelope

strategy, and landscape treatment (Japan Meteorological Agency, 2023).

11.2 DEMOGRAPHIC AND SOCIOECONOMIC CONDITIONS

Akita has one of the fastest shrinking and aging populations in Japan. According to the Statistics Bureau of Japan (2023), the prefecture's population decreased from over 1.2 million in 1995, to less than 950,000 in 2023, with projections illustrating few conditions to reverse this downward trajectory. With over 37% of Akita's population aged 65 and over, this demographic imbalance is acute. In this context, considering public infrastructure means examining the changing urban fabric rather than a growing urban fabric, as this demographic trend has spatial implications through under-occupied buildings, a shrinking urban area, and declining public infrastructure. In these settings, architecture is less concerned with developments and invention and more with operating with less and a different social order. Though small, the competition site becomes a site to test a design approach that can accommodate resistance, adaptation, and cultural longevity within these more general pressures.

11.3 HISTORICAL ROLE IN CULTURAL PRODUCTION

Though often overshadowed by Japan's urban centers, Akita has long contributed to the national cultural fabric. It has been a site of artisanal production, aesthetic refinement, and rural innovation. The region's craftsmanship developed under necessity and cultural patronage, notably during the Edo

period when local domains promoted regional goods to support self-sufficiency and prestige (Moeran, 1997).

Crafts like Kabazaiku and Kawatsura lacquerware represent deep engagements with local materials—cherry bark, cedar, urushi resin—and processes refined over centuries. These practices were artistic and infrastructural: they enabled sustainable economies based on long-term resource management, apprenticeship, and community labor (Kamogawa, 2015). Understanding this craft ecology is essential to situating contemporary design not as an imposed system but as a continuation of local modes of making.

11.4 CULTURAL IDENTITY AND PERIPHERAL NARRATIVES

Akita also occupies a peripheral status in the Japanese imaginary. Often stereotyped as rural, snowy, and remote, it has paradoxically become a national discourse symbol of nostalgia and neglect. This marginality, however, grants it a certain autonomy—allowing room for experimental redefinitions of what constitutes cultural value and spatial innovation. As scholars like Morris-Suzuki (1998) and Befu (2001) have argued, national identities are often negotiated at the margins. In this sense, designing in Akita engages with Japan's "invisible center" of cultural identity—where craft, nature, and time operate at different speeds than in the metropole. The design project engages this condition not as deficit but as a resource, drawing upon the spatial, temporal, and material logic that Akita's peripheral status affords.

This chapter presents a series of spatial investigations developed using GIS-based mapping and data analysis. Each cartographic layer contributes to understanding the site's physical, demographic, and infrastructural conditions. These studies informed early design positioning, site relationships, and environmental performance.

1. NATIONAL CONTEXT

Japan and the Akita Prefecture The first map illustrates Japan's regional structure, highlighting Akita Prefecture in the northern Tōhoku region of Honshu. Japan has 47 prefectures, and Akita is one of the most sparsely populated and rural provinces (Statistics Bureau of Japan, 2020). The map is intended to be a geopolitical locator regarding Akita's placement relative to the national cultural and economic activity centers.

2. POPULATION DENSITY

y A population density map of Japan gives further context to Akita's demographic state. In 2020, Akita Prefecture's population density was approximately 77 persons per square kilometer, well below the national average of 340 persons per square kilometer (Statistics Bureau of Japan, 2020). This demographic prematurity has consequences for local infrastructure demands and urban expansion conditions.

3. EARTHQUAKE AND TSUNAMI RISK ASSESSMENT

Japan is in a seismically active zone due to its position along the Pacific Ring of Fire. The third map visualizes seismic hazard zones

and tsunami inundation risks. According to the Japan Meteorological Agency (2021), Akita lies in a moderately active seismic zone with lower tsunami risk than coastal Pacific regions. However, historical records include earthquakes such as the 1896 Rikuu earthquake (M 7.2), which affected inland Akita (Utsu, 2002). These parameters were considered in terms of foundational planning and structural stability.

4. AKITA PREFECTURE

Regional View This map zooms into Akita Prefecture, distinguishing municipal boundaries and highlighting Akita City as the administrative and population center. The city houses key transportation links, including the Akita Shinkansen and the Port of Akita. This regional-scale View delineates Akita City's administrative boundary within the larger prefecture, indicating a concentration of services in an otherwise low-density prefecture.

5. AKITA CITY AND PROJECT VICINITY

The following map presents a plan of Akita City, locating the competition site within the city's southern fringe. The topography of the area is mostly flat, which reduces complications associated with terrain but introduces issues around stormwater management and exposure. The surrounding context is a mixture of developed residential and agricultural land typical of the peri-urban transitioning areas.

6. SITE-SCALE LOCATION

A site location map provides a focused

site overlay map that pinpoints the actual site of the proposed Crafts Museum within the area. It identifies the road network, parcel boundaries, and proximity to surrounding uses. The immediate surroundings to the west are residential and agricultural to the east, which complements the description of the semi-rural edge given in the brief.

7. FUNCTIONAL ANALYSIS MAP

This diagram illustrates the land use classification of the site's surrounding area: The retail zones are represented in blue and are predominantly located along the arterial road. The medical facilities are represented in dark brown and distributed throughout the area.

The educational sites are represented in light beige. Public parks appear in green. A diagonal hatch pattern denotes agricultural and rural zones. This functional breakdown provides an overview of urban rhythms and programmatic adjacencies that may influence visitor accessibility and community engagement strategies.

8. MICRO-CONTEXTUAL STUDY

Sun Path and Photographic Documentation The final situational analysis includes solar path diagrams illustrating seasonal sun angle and duration variations. This data is critical for daylighting and passive heating considerations. Supplementary site photographs will be inserted to provide visual documentation of terrain, vegetation, and built context.

FIGURE 31
National Context map

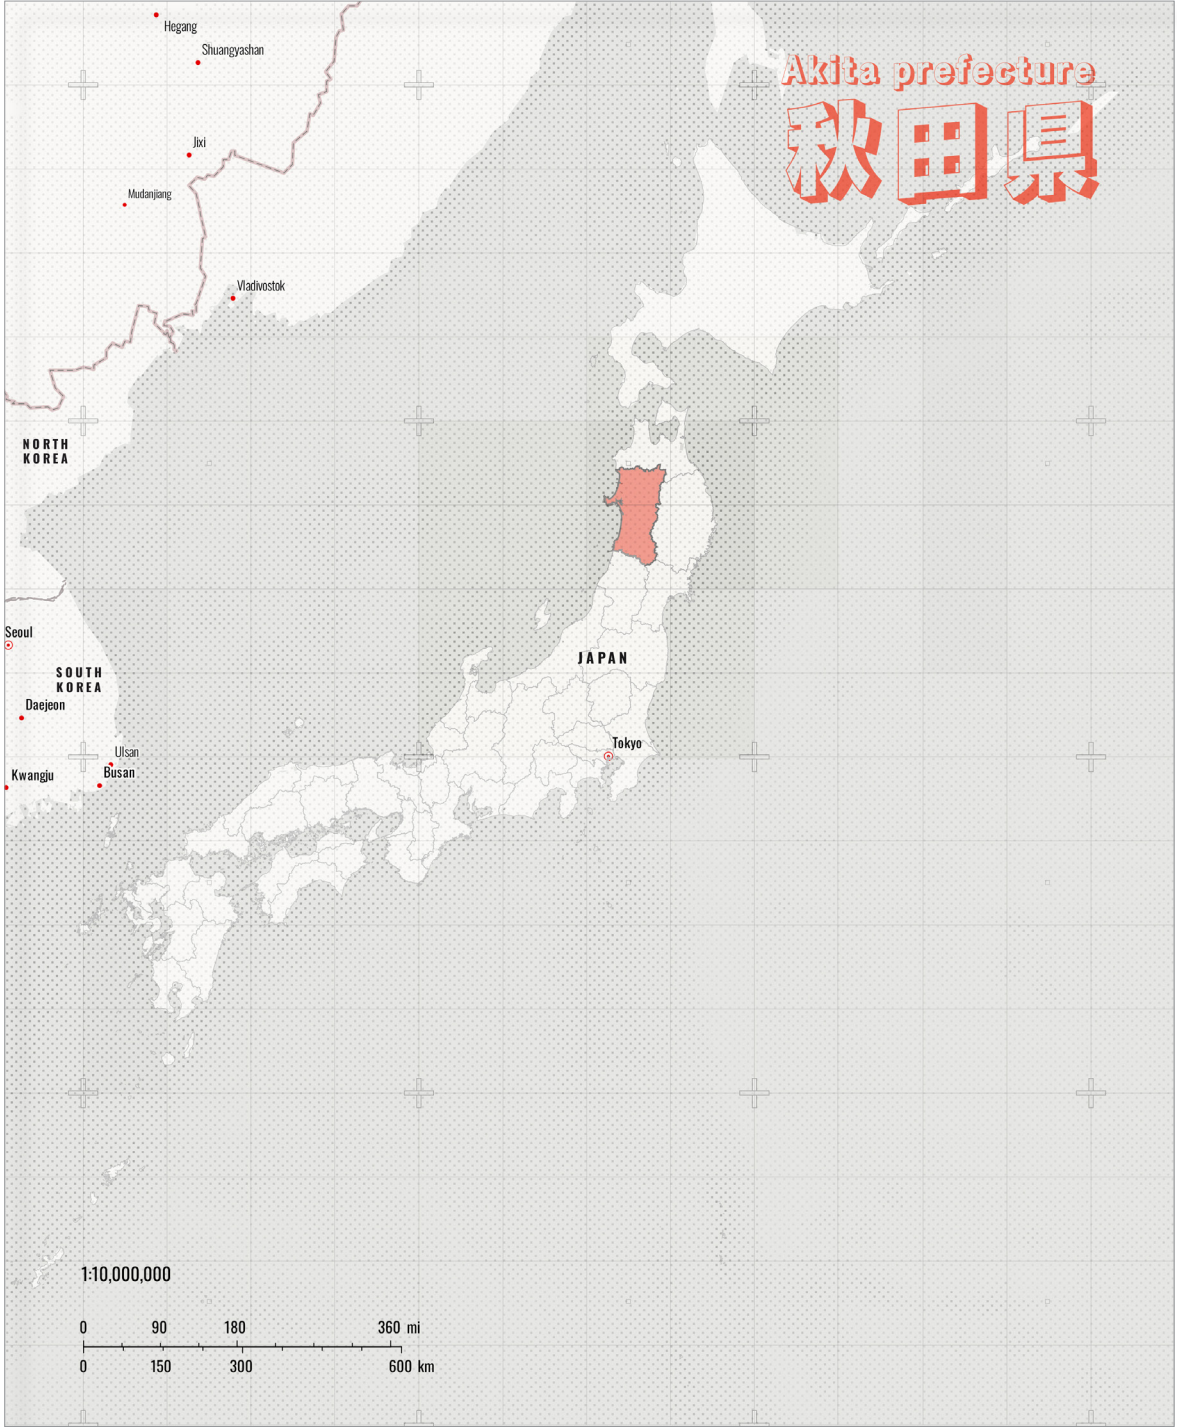


FIGURE 32
Population Density map

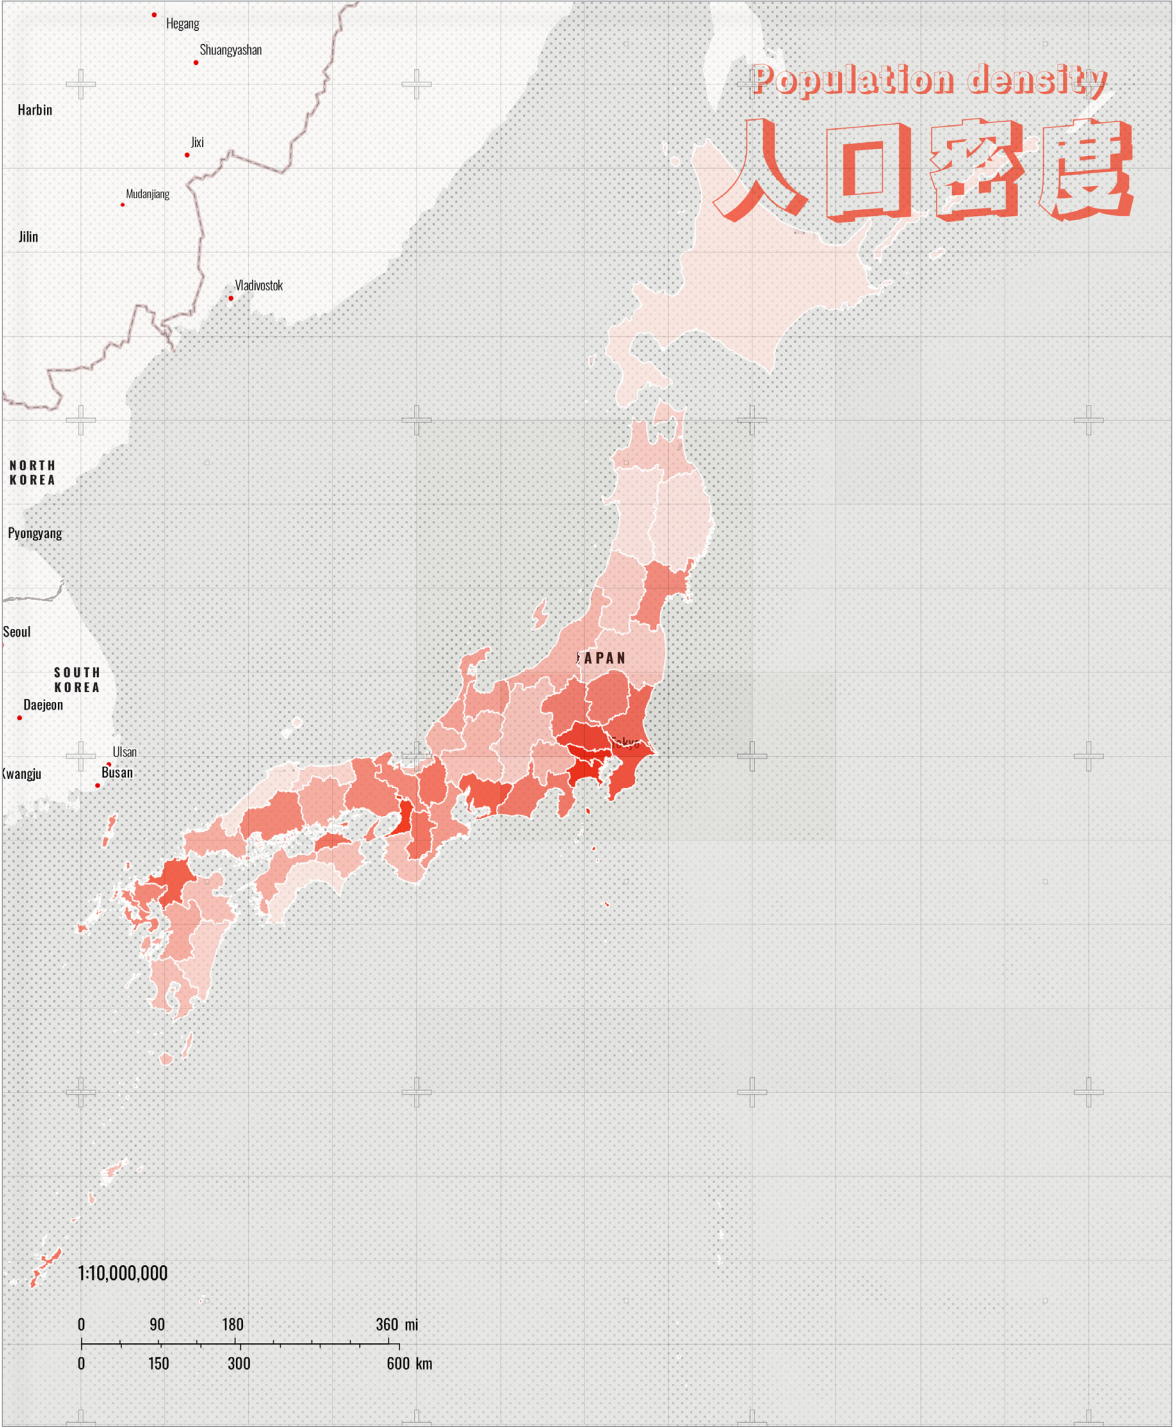


FIGURE 33
Earthquake and Tsunami Risk Assessment
map

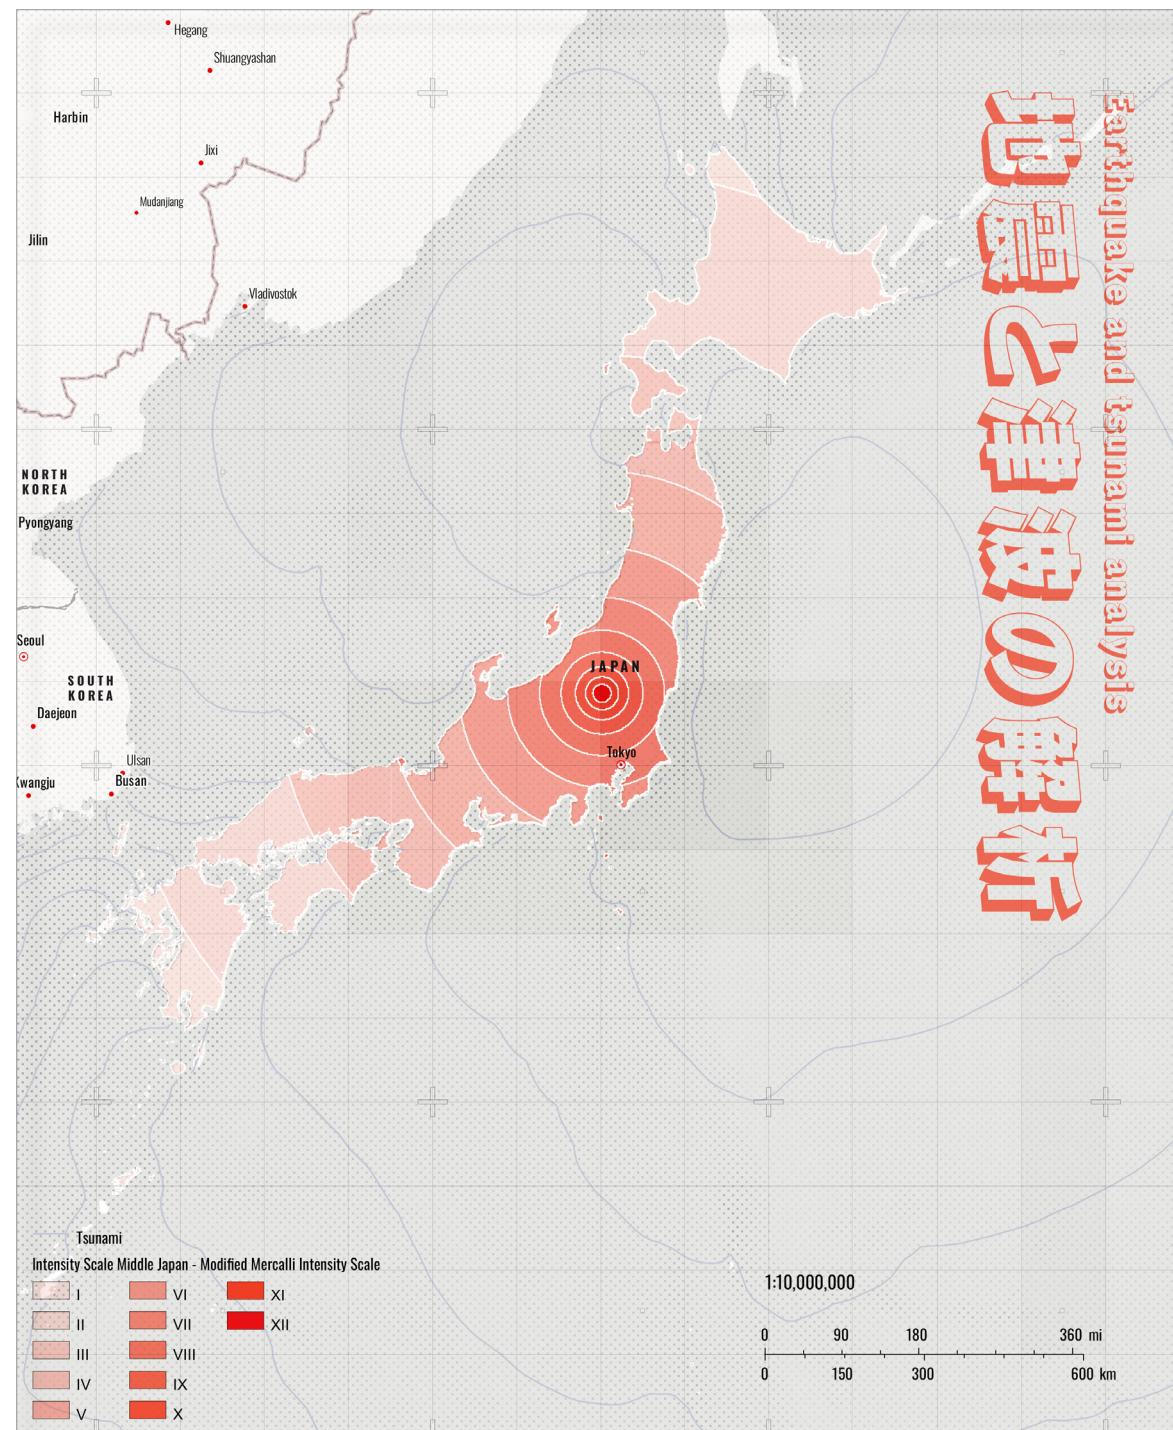


FIGURE 34
Akita Prefecture map

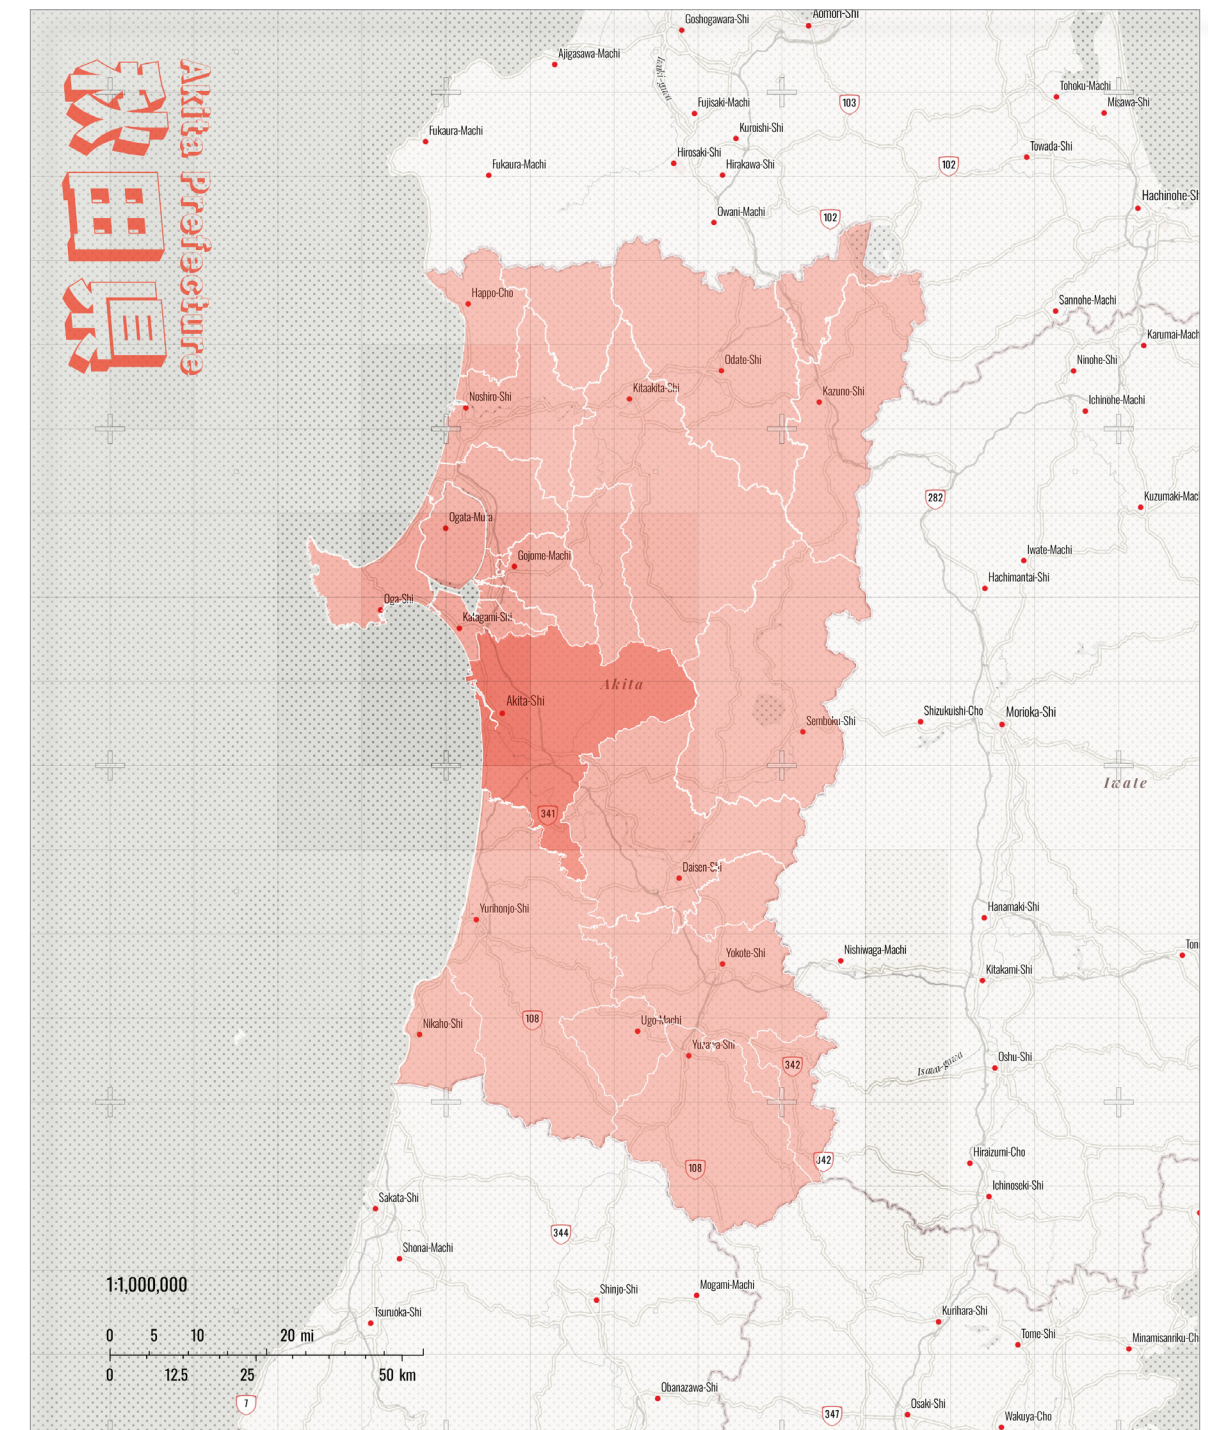


FIGURE 35
Akita City and Project Vicinity map

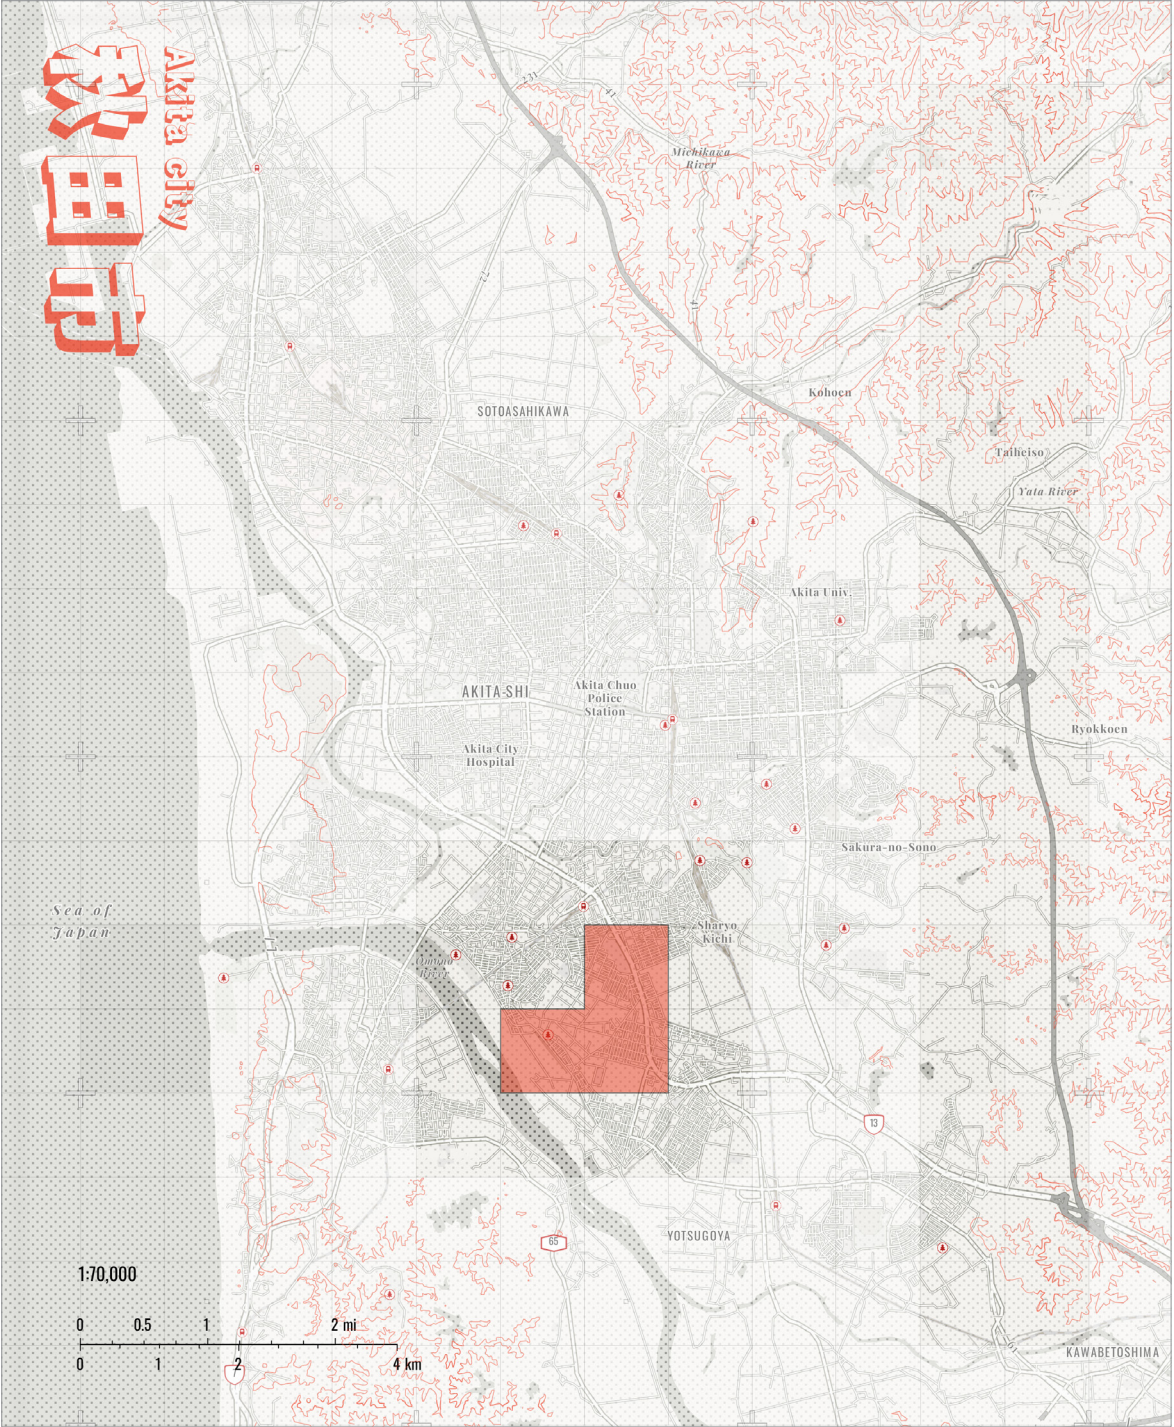


FIGURE 36
Site-Scale Location map

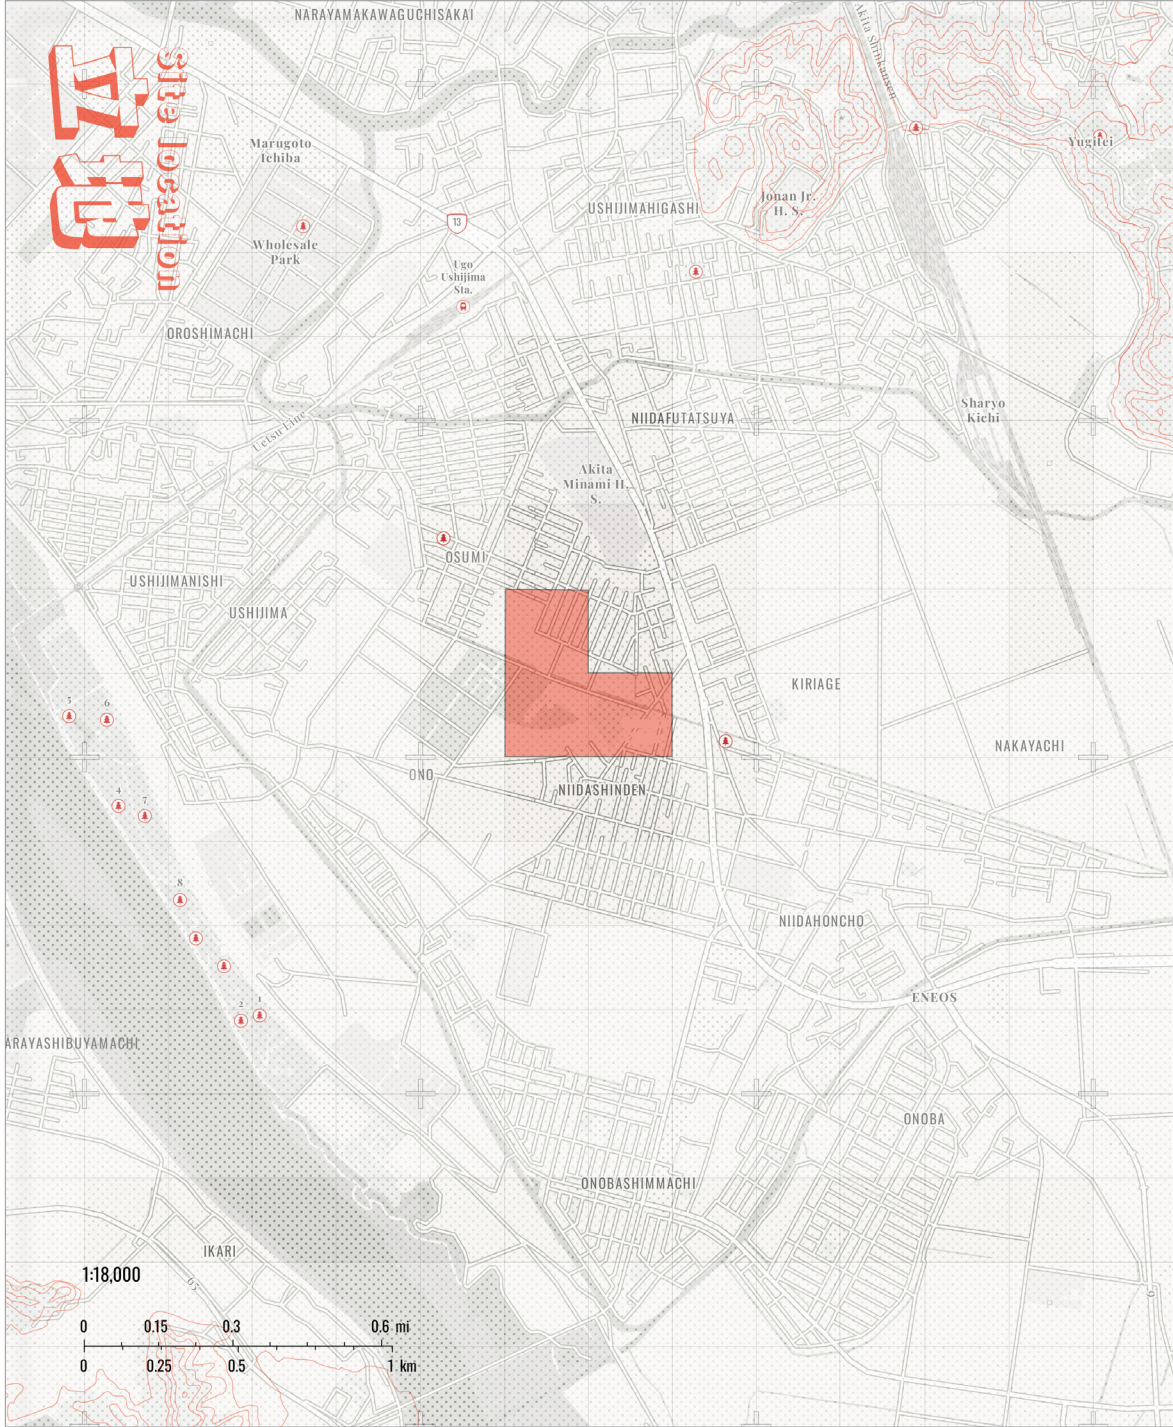


FIGURE 37
Functional Analysis Map

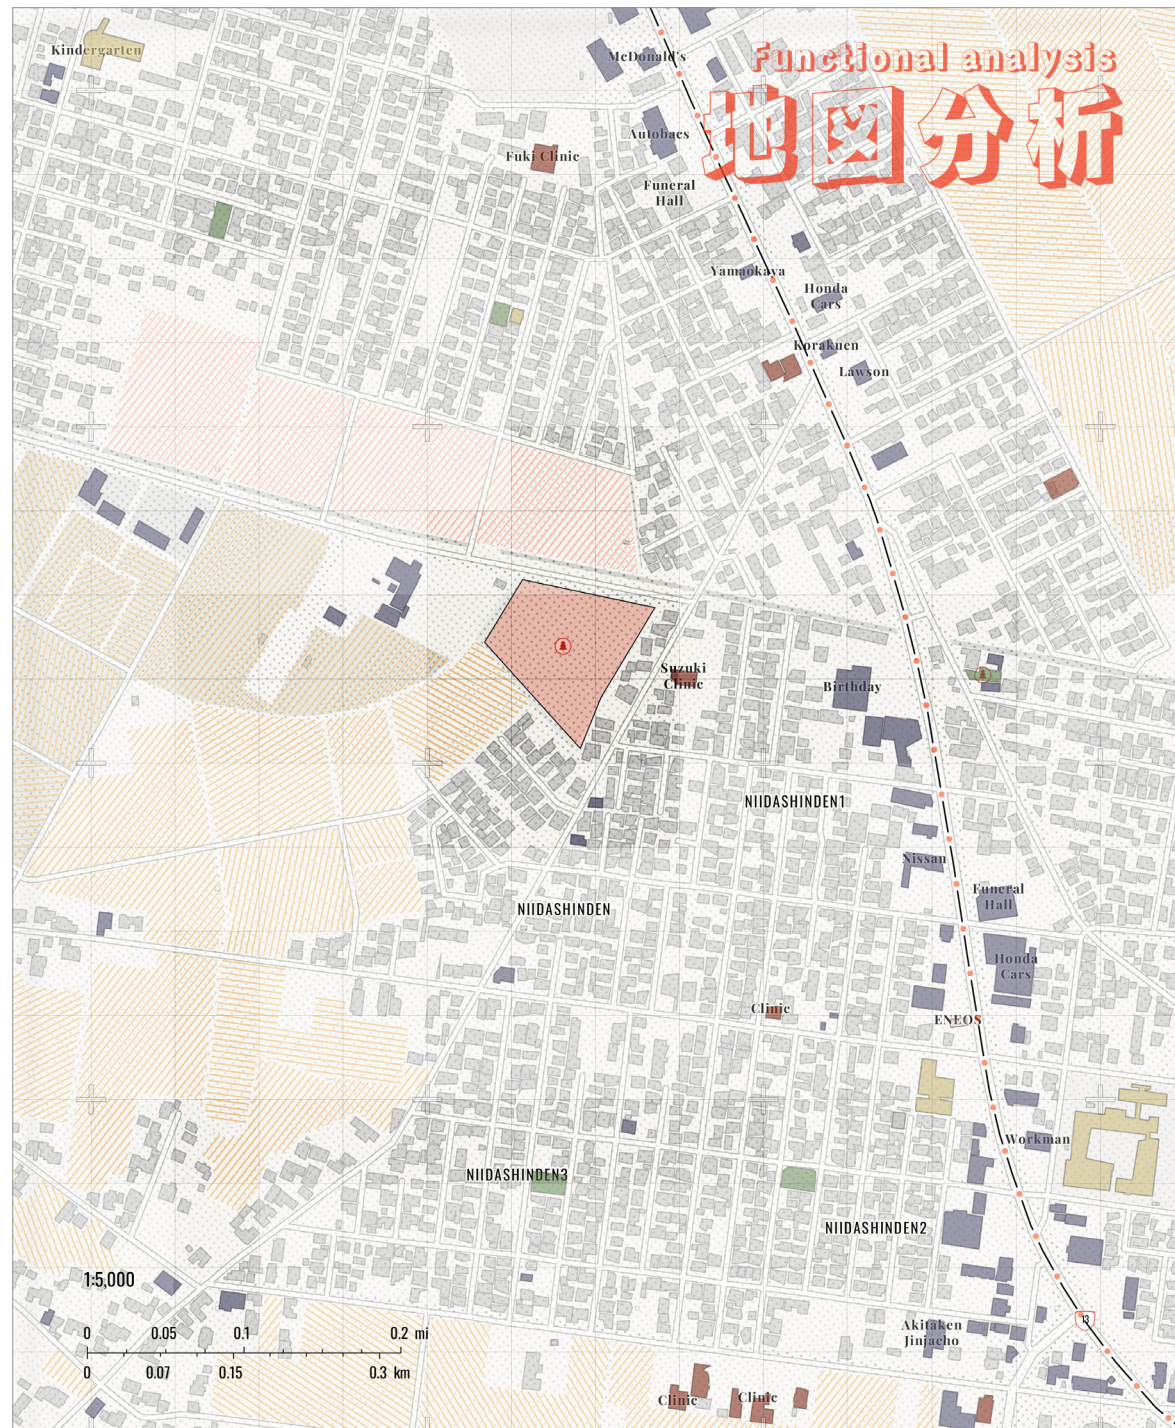
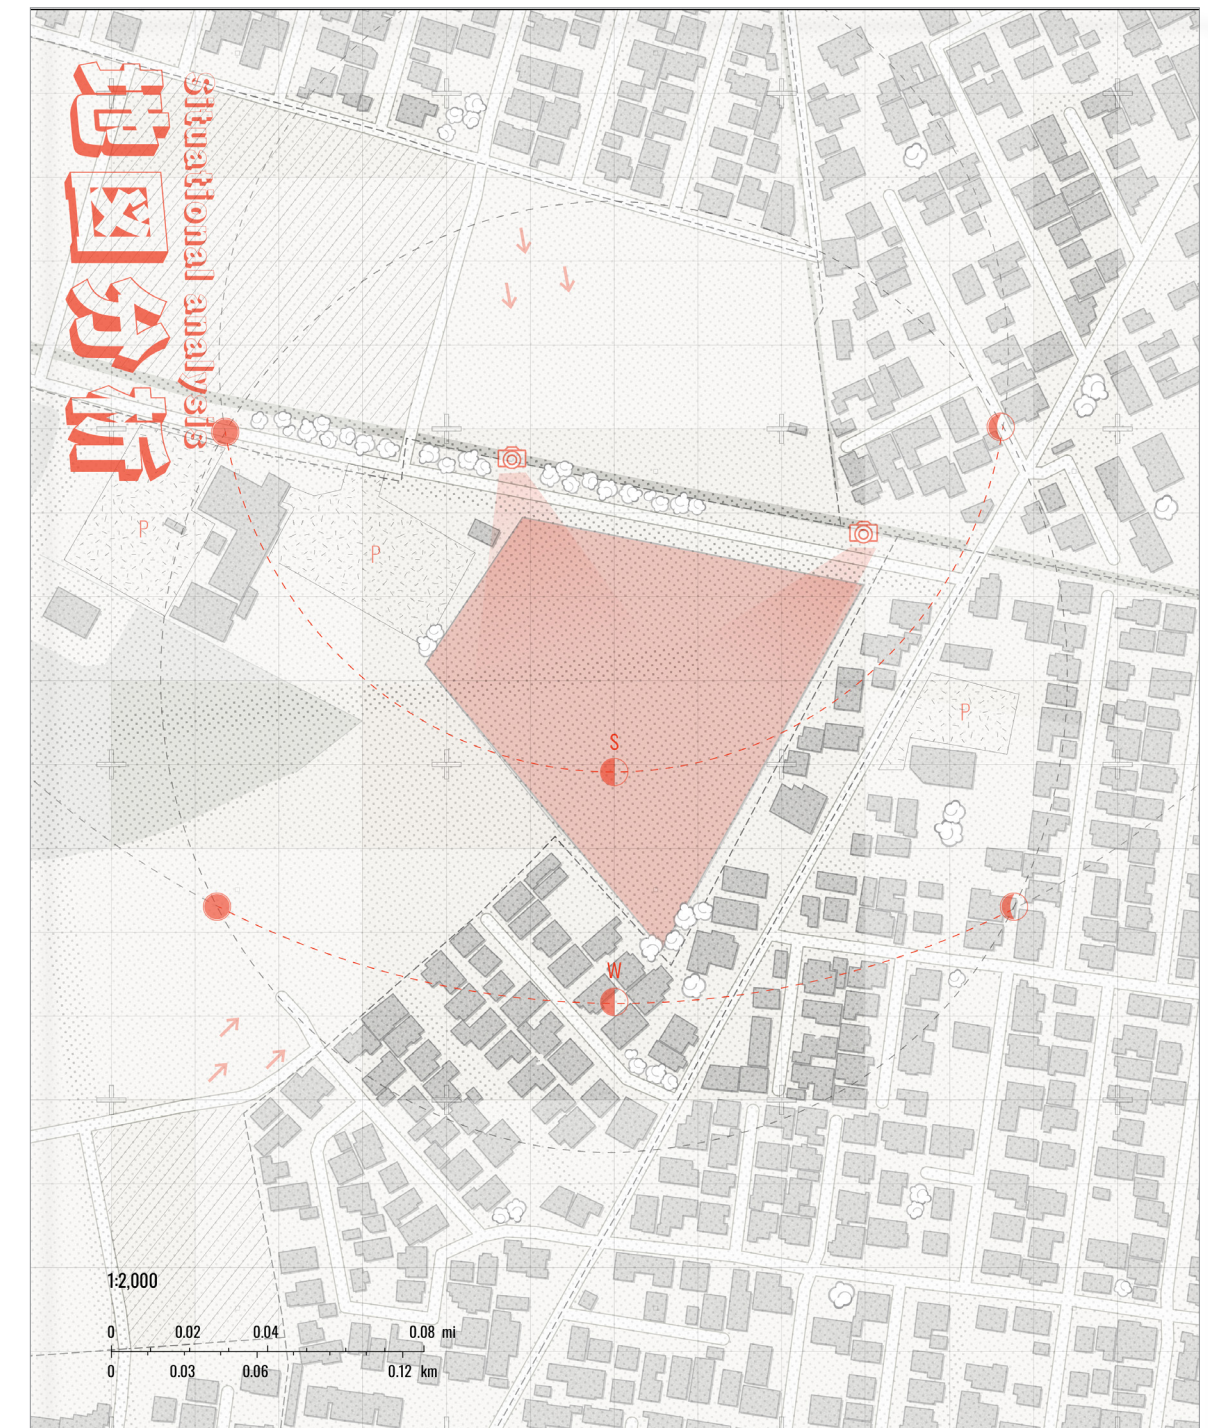


FIGURE 38
Micro-Contextual Study map



13.0 REPRESENTING THE GENERATIVE SEQUENCE

This chapter chronicles the visual development of the generative design sequence I arrived at through iterations of experimentation with Midjourney. While previous chapters have focused on the conceptual underpinnings and methodological framework of prompt engineering, including the crucial element of how this mediates cultural, aesthetic, and compositional intent, this chapter will draw attention to the representational outcome of that process.

The sequence began with more generalized morphological studies, including spatial massing, volumetric articulation, and site responsiveness. These first iterations were not presented as definitive proposals but rather a series of visual hypotheses, propositions to be tested and developed through repeated prompting. I purposely positioned the language to be very open, focusing primarily on terms like craft museum, flat site, traditional shape, and minimalist form. As the previous chapters suggested, this generic language led to sporadic outputs and a methodological pivot toward referential specificity.

The core of the generative sequence spans approximately 7,660 image outputs across 1,915 prompt iterations (see Chapter 7.0), with each stage reflecting increasing resolution in spatial, material, and programmatic terms. Rather than rely on a static idea of the project, each prompt was constructed in direct response to the visual feedback of the previous iteration. References to figures such as Kengo Kuma, Terunobu Fujimori, and Vo Trong Nghia were gradually embedded into the syntax, alongside descriptors relating to landscape integration, material tactility, and atmospheric control.

The image sequence is neither a linear design development nor a stylistic evolution. It functions as a visual archive of computational interpretation—a parallel to sketching or model-making in traditional practice. Midjourney's outputs were evaluated not on their realism or polish but on their ability to register architectural themes relevant to the project: enclosure vs. openness, surface vs. depth, and artifact vs. infrastructure.

The selected images presented in this chapter represent key turning points in the design logic. Some mark the first emergence of a cohesive massing language; others show successful integrations of landscape typologies or lighting strategies. In several instances, unexpected outcomes—such as hybrid roof geometries or spatial overlaps—were retroactively integrated into subsequent prompts, reinforcing the non-linear nature of the sequence.

The visual sequence is organized thematically and not chronologically. It begins with distant aerial views meant to establish territorial context, then moves to studies at the scale of buildings that foreground buildings' material and tectonic qualities, and finally narrows to interior atmospheres and spatial junctions where the relationships between program and architectural language are legible. Each image is annotated with its respective prompt string or a redacted version when required and indexed against the broader dataset for ease of reference.

In this sense, the generative sequence is not simply representational, but it is also an architectural inquiry in spatial thinking since it was produced by generating images. The

visual archive resulting from the generative sequence is supportive but also contesting and sometimes counterproductive to the design intentions set out beforehand. Thus, the visual archive serves as a record of negotiation between authorial control and algorithmic mediation.

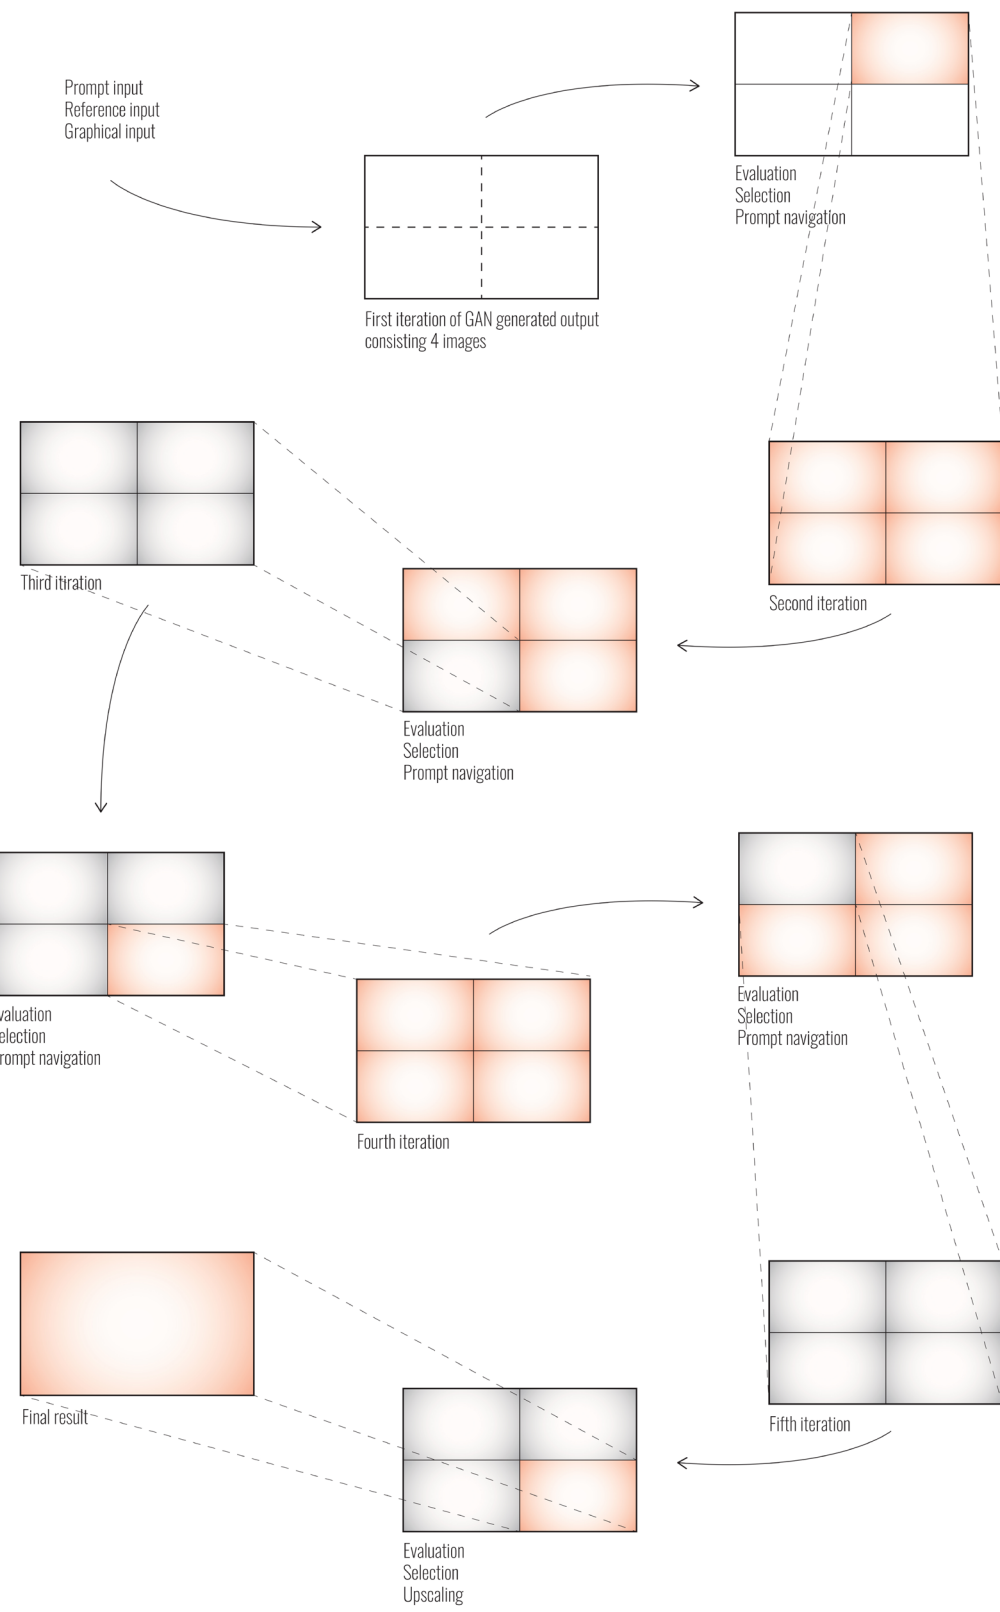


FIGURE 39-41
Starting point for generative sequence








































FIGURE 42-120
Generative sequence









END OF EXERCISE.

The generative sequence resulted in the image above - an aerial composition that synthesized the material, spatial, and organizational directives developed in earlier iterations. The resulting image signals architectural clarity at the scale of massing and form but with a conscious resolution limitation. Indeed, like much of Midjourney's principal output, the imagery produced through the generative sequence favors atmospheric coherence and a formal language rather than overt evidential architectural detail. Edges dissolve into shadow, fenestration dissipates into texture, and landscaping transitions into obscurity rather than providing an evident resolution. Consequently, the image produced does not offer a definitive architectural proposal but represents a conceptual scaffold.

However, this ambiguity is not simply a technical failing. It exposes a structural condition of generative workflows: they excel at articulating intent at the scale of material logic and compositional rhythms, while poor at articulating operable joints, thresholds, or distinct functional demarcations. In this regard, the output must not be construed as a vision fat singular, but as an interpreted surface; something to be translated, judged and adapted. In practice, this leads to a condition where architects are compelled to improvise at finer scales. The image's low resolution and aestheticized bias necessitate further authorial intervention to define apertures, resolve circulation, or integrate service cores. These unresolved zones become the terrain for critical design thinking rather than automated reproduction. In this way, the limitations of the tool redirect attention toward the role of the architect as an editor—curating, refining, and sometimes contradicting the AI's

suggestion to suit the demands of structure, climate, and user experience.

Thus, the final image is best read not as a rendered building but as a generative field of constraints and opportunities. It visualizes the architectural thesis as a constellation of values—material restraint, formal sensitivity, and cultural citation—without prescribing their execution. The work follows is one of translation: from suggestion to drawing, from atmospheric form to spatial precision.



FIGURE 121
Final iteration of generative sequence

13.0 ARCHITECTURAL DRAWINGS AND SPATIAL ORGANIZATION

This chapter presents the architectural drawings developed from the generative process and contextual analyses previously discussed. The focus is on how spatial strategies and material articulation respond to the cultural, technical, and functional requirements outlined in the brief—translated here into grounded architectural forms through axonometric projections, plans, and sectional studies.

The ground floor is a semi-public interface, engaging circulation, leisure, and cultural introduction. Upon entering the space, users are met with a reception that acts as the fulcrum for the entire moveable sequence of the museum, and from this point, the moveable disperses into two linear paths. One goes toward the temporary exhibition space, and the other is available to the left wing, where workshops and the public library is housed.

The layout of the museum places hospitality functions as a key focal area. The café, restaurant, and Japanese teahouse are clustered together activating the edge of the main courtyard, while also preserving a level of permeability into the remainder of the museum. Internally, these hospitality areas are connected to the artist residency wing thus encouraging a porous boundary between public and semi-private spaces.

It is important to note that the café and the artist accommodations share a direct internal vertical connection to the -1 level through stair connections. Therefore, this adjacency contributes to operational efficiency in the café while allowing artist residencies to access the workshop areas below. At the same time, users can roam freely from

lower galleries to the café above.

The -1 level contains the core curatorial areas that generate an experiential continuum. The -1 level is mostly arranged around a long linear pathway that organizes the permanent exhibition space. This pathway rests on low ambient light levels along with trimmed skylight and skylight walls creating a sequence of patterned moments of orientations and pauses.

A key feature to this level is an underwater viewing gallery. Located in the middle of the exhibition pathway, the viewing gallery provides entitled viewing inside the adjacent courtyard pond. The glazed boundary of the pond provides a moment of movement, ecology, and temporal experience for the interior.

The -1 level also contains the amphitheater, which connects to the ground floor. Natural light penetrates the space from a skylight opening above the amphitheater stage, providing a dramatic experience into the otherwise darkly lit lower volume. The amphitheater level has access from both levels, with scene loading occurring from the -1 level.

The gallery extends into a more flexible volume at the end of the permanent exhibition path. This plaza-like condition allows for informal movement, social gatherings or specific installations while allowing for a tempo change from the linearity of the curated path. Also located next to this space are technical rooms and craft restoration labs that are publicly viewable through large glass walls. This transparency supports the educational agenda of the museum by foregrounding processes normally kept out of sight. Beneath the library,

the basement houses archival storage in a stable, light-controlled environment. Similarly, portions of the exhibition's object storage and preparation spaces are positioned here, reinforcing the division between visitor-accessible and support zones.

Level +2 The sloping constraints and formal attributes of the roofscape—the sloped roof styling draws from traditional Japanese morphology—results in a small second level. This elevates the library to a new vertical location with a series of museum administrative offices. The museum and library are intended to have operational separation, with circulation through an interior stair at the entry and through the ground-floor service corridor adjacent to reception.

STRUCTURAL SECTION AND ROOF SYSTEM

The sections demonstrate the architectural strategy to roof as parametric timber shell. The primary structure comprises laminated timber pieces placed horizontally relative to a vertical primary column. The central column gradually grows towards the top of the space, with the horizontal components stacked in a radial pattern for a spiraling assembly logic. The central column acts structurally (as a column) and behaves as a daylight distributor. Daylight enters through concealed openings at numerous junctures in the timber structure, eventually delivered into the workshop and temporary exhibition spaces. Nature daylight contributes, is controllable, establishes wayfinding, and enhances the impactful attributes of material in the realms we are focusing on, which is craft (tactile and visual). The use of glulam allows structural performance and aesthetic continuity that

complements the material language of craft. The structure engages the visitor as part of the exhibition experience, as an architectural artifact representing the artisanal logic of levels of assembly and joinery.




FIGURE 121.1
Functional diagram

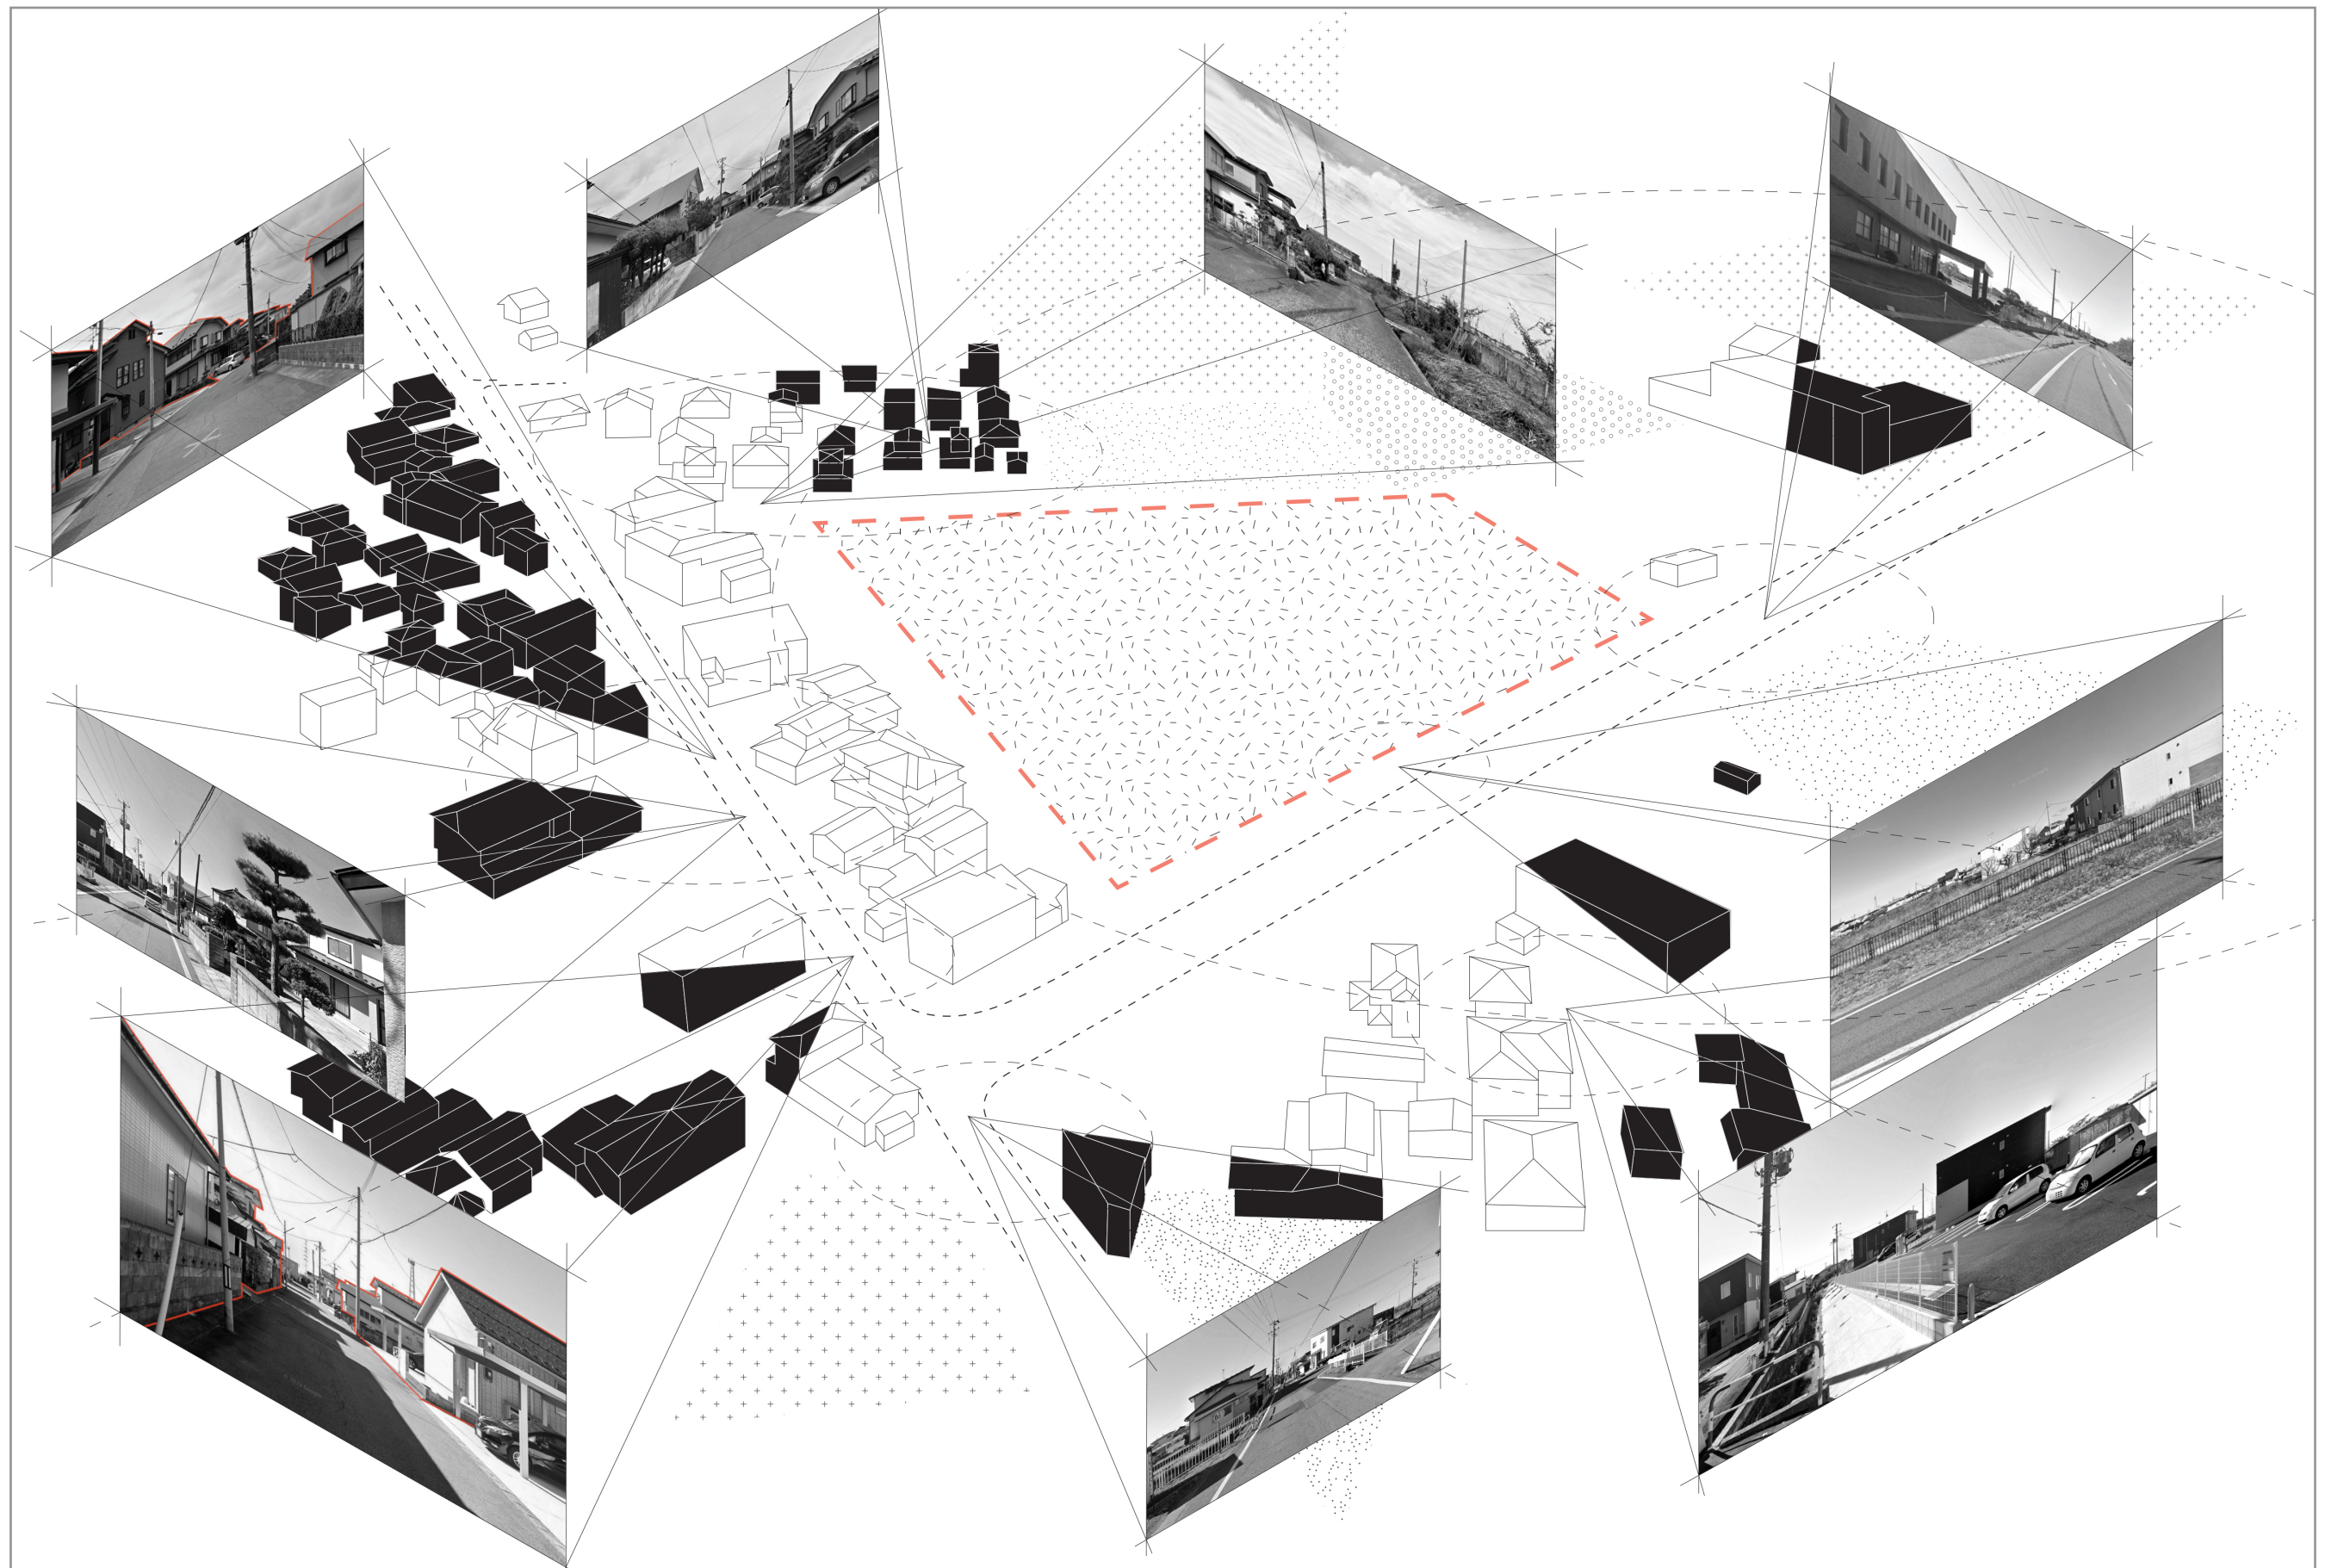


FIGURE 121.2
Context diagram




FIGURE 122
Axonometric view



FIGURE 123
Ground floor plan

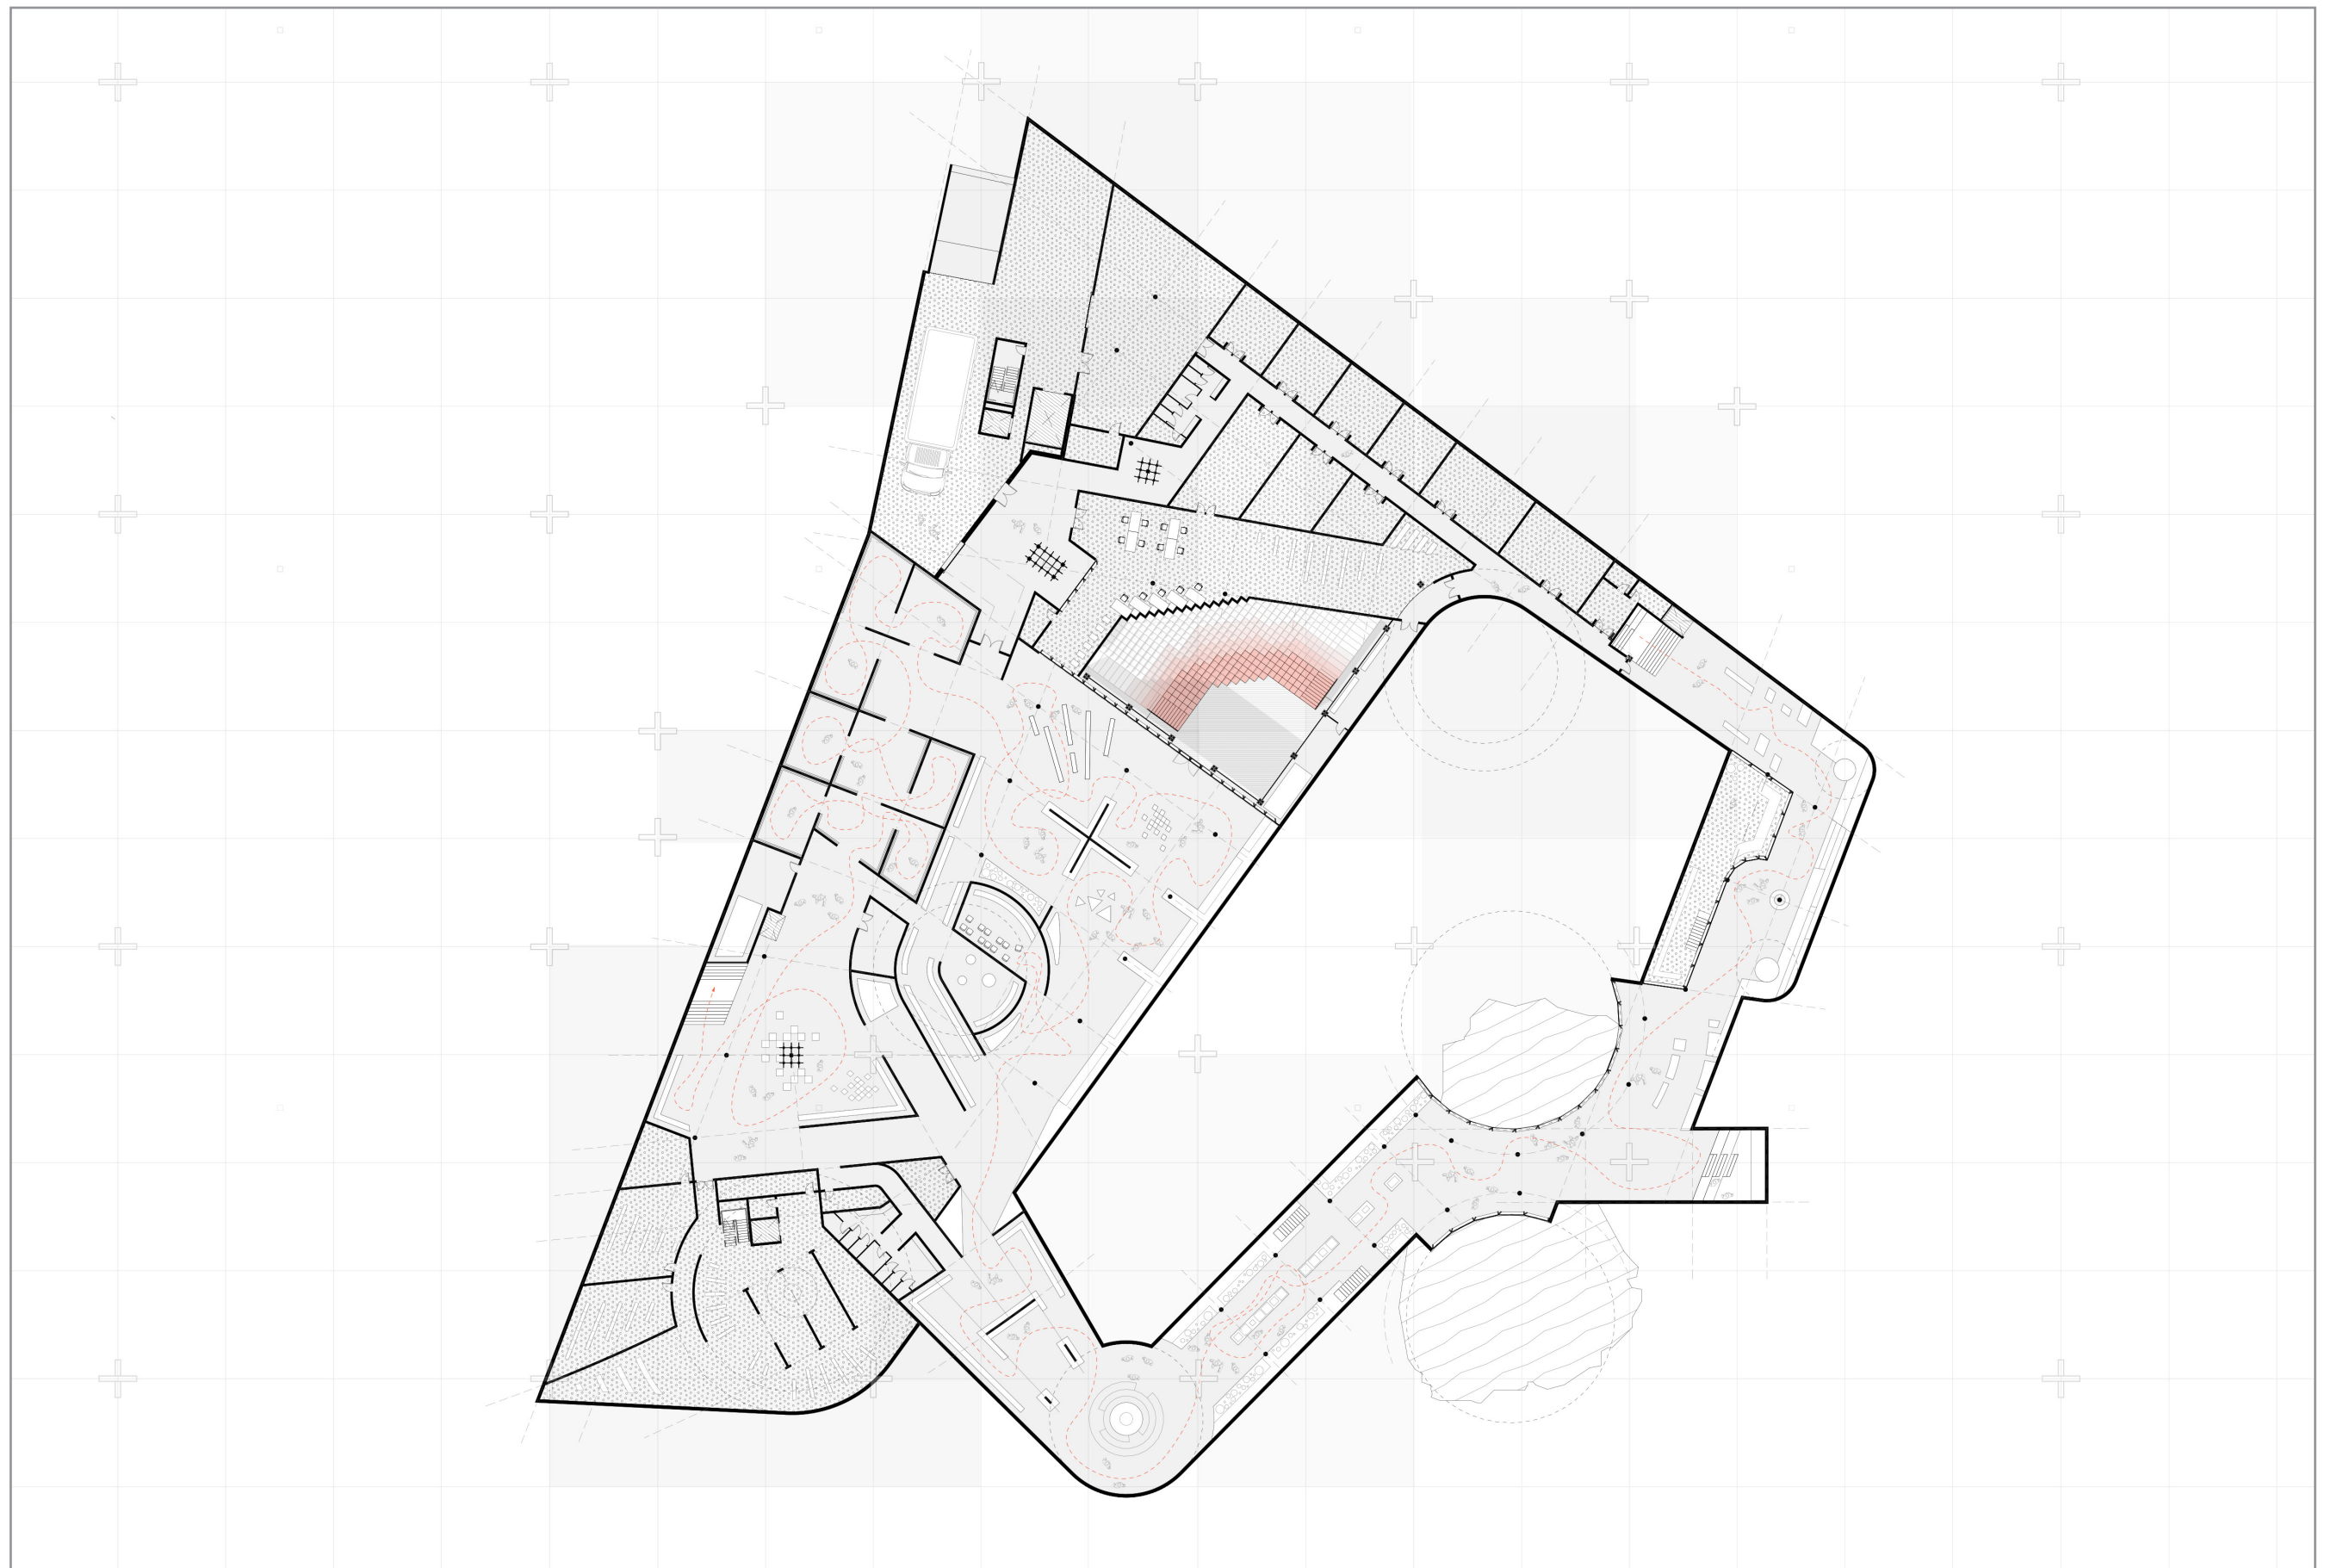


FIGURE 124
-1vl plan




FIGURE 125
2nd lvl plan




FIGURE 126-127
Section views



FIGURE 128
Rendered view



FIGURE 129
Rendered view




FIGURE 130
Rendered view




FIGURE 131
Rendered view



FIGURE 132
Rendered view



FIGURE 133
Rendered view




FIGURE 134
Rendered view



FIGURE 135-136
Rendered view



CONCLUSION: CRITICAL REFLECTIONS ON GENERATIVE DESIGN AND ARCHITECTURAL CURATORSHIP

This thesis has presented a comprehensive study of the possibilities and limitations of incorporating generative artificial intelligence, in this case, Midjourney, into architectural ideation processes. This process involved considerable prompt engineering, GIS-based contextual considerations, and thorough case study research. The computational workflows were deliberately conceived as both speculative and practical tools. The investigation was a critical consideration of the issues and contradictions of using an algorithmic process of AI, to architectural questions and problems, instead of positioning AI as a solution.

An important outcome of this research is discovering how troublesome it is to introduce generative AI into architectural practice. It is not a simple or a purely facilitative experience. While Midjourney creates all kinds of visualizations with ease from the prompt and enables rapid iteration of an idea, it is the ambiguity of scale (the images, of course, vary in size), not to mention the errors in materials, details, and consideration of building and architectural contexts, (e.g., structural forms, building code regulations and material realities), that characterizes this process. Thus, the visuals from AI, in this case, midjourney, was a provocation to think differently about space or concept (if we privilege the outcome) rather than being drawn with precision in architectural representation - it required considerable reframing, rewriting, and transforming.

Significant portions of the design development, where architectural practice as a profession is concerned with (space, spaces, spatial order, structure, and circulation) were generated from AI - though precariously suggestive. Therefore, the AI, in this case, worked

as a generator/producer of an atmosphere or idea. In contrast, architects must take on the more active role of a curator, who will consider, critique, and work with the ambiguous outcomes produced by AI.

An important finding was the surprising commitment of time required to produce legitimate and contextual visualizations. Although the rendering process has been more automated, achieving acceptable results still meant iterating over processes, carefully making prompts, and additionally curating the result. The amount of human effort mainly involved canceled any efficiency encouraging AI-driven design, often making it a more complex experience than in previous early stages of design, where AI has enabled quicker entry into design development with [little] complexity. The thesis also made clear distinctions between the surface aesthetics created by AI and the architectural logic beneath.

Although Midjourney produced visually interesting textures, potential dramatic lighting, and expressive forms, essential architectural aspects such as structural cohesion, logical circulation, and code compliance were frequently misrepresented or missing altogether. This reiterated the need for architects to mediate actively between AI generative imagery and real-life architectural design.

Considering the architect role, the research called for architects to act as curators by facilitating and taking a curatorial role in the generative dialogue of AI. As a result of this research, these workflows are more likely to become a reality when potential prompt outcomes could become more regimented, reliable, and systematic. By structuring prompts in particular ways, referencing architectural

precedents specifically, and defining the fundamental elements of a design prompt, generative AI could offer a more well-polished and definitive regimen of integrated work. If we were to imagine the notion of simultaneously articulating the 3D model and rendered image — in effect, allowing plans, sections, elevations, and three-dimensional representations to develop in parallel — this would be a phenomenal leap in practice; however, we are not there yet, and creating coherent architecture from translating 2D AI imagery takes extreme imagination and iterative validation, resulting in considerable time spent developing prototypes, and ultimately not being a viable workflow.

Moving forward, this thesis acknowledges that generative AI workflows hold promise. This promise lies with generative technical improvements, where the generative imagery can be more directly translated into accurate, actionable 3D architectural models. With these objectives in place, the potential for an integrated workflow could significantly improve the overall efficiency and accuracy of architectural design; however, it should be highlighted that in this phase of the learning curve, many iterative attempts, prompt trials and re-translations will need to be made. This thesis does not end with naive optimism or outright rejection of generative AI in architecture; instead, it provides a critical review of the complex relationship between computational generativity, architectural authorship and real-world limitations. Thoughtfully and critically employed, generative AI has the potential to enrich architectural discourse and practice. Without careful consideration, however, it risks superficial outcomes and inefficient workflows. The challenge and opportunity lie

in refining methodologies to harness AI's speculative strengths within disciplined, spatially coherent, and technically sound architectural frameworks.

THE TABLE OF THE FIGURES

FIGURE 1 - PAGE 10

Venn diagram showing Artificial Intelligence (AI) encompasses the subsets of Expert Systems, Machine Learning and Deep Learning, in which their methods can be applied to tasks that imitate human decision-making abilities. Reprinted from “Artificial intelli

FIGURE 2 - PAGE 12

Stages of recognition and transformation of a cross. Negroponte, N. (1975). Soft architecture machines. The MIT Press.

FIGURE 3.1 - PAGE 13

AI tribes

FIGURE 3 - PAGE 16

Control room (Architecture Machine Group). Reprinted from Soft Architecture Machines (p. 158), by N. Negroponte, 1975, The MIT Press.

FIGURE 4 - PAGE 17

Institute for Advanced Study. (ca. 1930s). image of Alan Turing walking along the street [Photograph]. Elaine Negroponte collection on John von Neumann and Alan Turing, Shelby White and Leon Levy Archives Center, Institute for Advanced Study

FIGURE 5 - PAGE 19

Nurnberg, W. (Photographer). (1958). Engineers with the early DEUCE computer at English Electric

FIGURE 6 - PAGE 21

Ewalt, D. M. (2011, May 3). Kasparov vs. Deep Blue [Photograph by George Widman/AP]. Forbes

FIGURE 7 - PAGE 23

Riccio, T. (Photographer). (2016, March). David Hanson and Sophia during a 60

Minutes interview with Charlie Rose, New York City

FIGURE 8 - PAGE 26

Mario Carpo, Mark Garcia and Steven Hutt, A short but believable history of the digital turn in architecture, ‘Prevalence of Computation in Architectural Design’, 2023

© Courtesy of the artists and the Jencks Foundation at The Cosmic House

FIGURE 9 - PAGE 32

Lynn, G. (ca. 1999). Embryological House: Sketches [Drawing]. Canadian Centre for Architecture (CCA), Embryological House fonds.

FIGURE 10 - PAGE 33

Gehry’s 1992 Fish Sculpture in Barcelona was among the first architectural projects to use CATIA software, enabling precise digital modeling from hand-built forms—a method later applied in the Guggenheim Bilbao and Disney Concert Hall. Gehry, F. (1992). Computer and built models for Gehry’s fish sculpture, Barcelona [Photographs and digital models]. Courtesy of Gehry Partners, LLP.

FIGURE 11 - PAGE 41

Bosque, C. (2013, June 30). MIT-Fablab Norway: Extract from Bosque’s sketchbook. In The Story of MIT-Fablab Norway: Community embedding of peer production. ResearchGate.

FIGURE 12 - PAGE 45

Chaillou, S. (2019, July 17). Figure 3. GAN-enabled building layouts [Image]. In ArchiGAN: A generative stack for apartment building design. NVIDIA Developer

Blog. <https://developer.nvidia.com/blog/archigan-generative-stack-apartment-building-design/>

FIGURE 13-15 - PAGE 47

Deep Himmelblau — As Coop Himmelb(l)au’s semantics and style are not homogeneous... [LinkedIn post]. LinkedIn.

FIGURE 16 - PAGE 49

RENAISSANCE DREAMS — PALAZZO STROZZI

The installation by Anadol.R utilizes GAN algorithms trained on Renaissance-era data, creating a “multidimensional shape matching the architecture and infrastructure of MEET”

FIGURE 17 - PAGE 50

Clemence, S. (2018, September 26). Frank Gehry’s Walt Disney Concert Hall to be projected with digital “machine hallucinations”. Metropolis Magazine..

FIGURE 18 - PAGE 52

Bill & Melinda Gates Hall at Cornell University in Ithaca, United States. (Morphosis Architects)

FIGURE 19-20 - PAGE 54

In Copenhagen’s Refshaleøen district, artificial intelligence was utilized to analyze data and uncover concealed patterns, which were then translated into creative visual interpretations of potential futures. This approach created a meaningful connection between factual data and the emotional experiences of both citizens and stakeholders, showcasing AI’s potential to stimulate fresh perspectives and innovative visual storytelling (Henning Larsen). Midjourney

FIGURE 21-23 - PAGE 57

Shortlisted in August 2021 for Melbourne’s Merinda Station Integrated Art Project, this public artwork investigates geological processes as inspiration for innovative construction techniques and expressive forms. By combining 3D-printed sandstone with intricately cast metal inlays, the work explores themes of deposition, erosion, and the dynamic interplay between solid and void. Sand layering, metal casting, and material contrast evoke natural tectonic formations while introducing striking visual and textural complexity (Snooks, Harper, & Gibson, 2021).

FIGURE 24 - PAGE 81

ITERATIVE EVALUATION FROM THE SKETCH created with MidJourney by Tim Fu

FIGURE 25 - PAGE 86

Kengo Kuma & Associates. (2024, October 30). UCCA Clay Museum. ArchDaily. <https://www.archdaily.com/1022949/ucca-clay-museum-kengo-kuma-and-associates>

FIGURE 26 - PAGE 87

Kengo Kuma & Associates. (2024, October 30). UCCA Clay Museum. ArchDaily. <https://www.archdaily.com/1022949/ucca-clay-museum-kengo-kuma-and-associates>

FIGURE 27 - PAGE 89

Kengo Kuma & Associates. (2012, February 16). Yusuhara Wooden Bridge Museum. ArchDaily. <https://www.archdaily.com/199906/yusuhara-wooden-bridge-museum-kengo-kuma-associates>

FIGURE 28-29 - PAGE 91

Elding Oscarson. (2024, March 22). Wisdome Stockholm. ArchDaily. <https://www.archdaily.com/1014815/wisdome-sto->

FIGURE 30-31 - PAGE 93

Shigeru Ban Architects. (2019, October 10). Swatch and Omega Campus. ArchDaily. <https://www.archdaily.com/926166/swatch-and-omega-campus-shigeru-ban-architects>

FIGURE 31 - PAGE 98

National Context map

FIGURE 32 - PAGE 99

Population Density map

FIGURE 33 - PAGE 100

Earthquake and Tsunami Risk Assessment map

FIGURE 34 - PAGE 101

Akita Prefecture map

FIGURE 35 - PAGE 102

Akita City and Project Vicinity map

FIGURE 36 - PAGE 103

Site-Scale Location map

FIGURE 37 - PAGE 104

Functional Analysis Map

FIGURE 38 - PAGE 105

Micro-Contextual Study map

FIGURE 39-41 - PAGE 109

Starting point for generative sequence

FIGURE 42-120 - PAGE 110-117

Generative sequence

FIGURE 121 - PAGE 119

Final iteration of generative sequence

FIGURE 121.1 - PAGE 122

Functional diagram

FIGURE 121.2 - PAGE 125

Context diagram

FIGURE 122 - PAGE 126

Axonometric view

FIGURE 123 - PAGE 128

Ground floor plan

FIGURE 124 - PAGE 130

-1lvl plan

FIGURE 125 - PAGE 132

2nd lvl plan

FIGURE 126-127 - PAGE 134

Section views

FIGURE 128 - PAGE 136

Rendered view

FIGURE 129 - PAGE 138

Rendered view

FIGURE 130 - PAGE 140

Rendered view

FIGURE 131 - PAGE 142

Rendered view

FIGURE 132 - PAGE 144

Rendered view

FIGURE 133 - PAGE 146

Rendered view

FIGURE 134 - PAGE 148

Rendered view

FIGURE 135 - PAGE 150

Rendered view

FIGURE 136 - PAGE 152

Rendered view

BIBLIOGRAPHY

Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1), 39–59.

Anadol, R. (2020, July). *Art in the age of machine intelligence* [Video]. TED. https://www.ted.com/talks/refik_anadol_art_in_the_age_of_machine_intelligence

Anadol, R. (2020, July). *DigitalFUTURES: AI and neuroscience* [Video]. YouTube. <https://www.youtube.com/watch?v=L1H7eL8pk5k>

ArchDaily. (2019, October 10). *Swatch and Omega Campus / Shigeru Ban Architects*. <https://www.archdaily.com/926166/swatch-and-omega-campus-shigeru-ban-architects>

ArchDaily. (2024, March 22). *Wisdom Stockholm / Elding Oscarson*. <https://www.archdaily.com/1014815/wisdome-stockholm-elding-oscarson>

ArchiTech. (n.d.). *EP 16 | ArchiTech office tours | UNStudio* [Video]. YouTube. https://www.youtube.com/watch?v=Hbt1bG_nb3Y

Aristotle. (1997). *Poetics, Categories, Hermeneutics* (J. Sarkady, Trans.). Kossuth. (Original works published ca. 335 BCE)

Bajohr, H. (2021). The gestalt of AI: Beyond the holism-atomism divide. *Interface Critique*, 3, 13–35.

Befu, H. (2001). *Hegemony of homogeneity: An anthropological analysis of Nihonjinron*. Trans Pacific Press.

Bilski, E. (2019). Artificial intelligence avant la lettre: The golem of Jewish mysticism, legend and art. In C. Woods, S. Livingston, &

M. Uchida (Eds.), *AI: More than human* (p. 210). Barbican International Enterprises.

Block, P., Boller, G., DeWolf, C., Pauli, J., & Kaufmann, W. (Eds.). (2024). *Proceedings of the International Association for Shell and Spatial Structures (IASS 2024)*.

Boden, M. A. (2016). *AI: Its nature and future* (p. 1). Oxford University Press.

Bolter, J. D., & Grusin, R. (2002). *Remediation: Understanding new media*. MIT Press.

Brown, G. (2019, August 23). 3 fundamentals of high-speed, high-quality 3D printing. *Medium*. <https://medium.com/ai-build-techblog/3-fundamentals-of-high-speed-high-quality-3d-printing-1f9b384cfe70>

Brownell, B. (Ed.). (2023). *The Routledge companion to contemporary materials*. Routledge.

Brooks, R. (2018). In M. Ford (Ed.), *Architects of intelligence: The truth about AI from the people building it* (p. 435). Packt.

Carpo, M. (2023, March). A short but believable history of the digital turn in architecture. *E-flux Architecture*. <https://www.e-flux.com/architecture/chronograms/528659/a-short-but-believable-history-of-the-digital-turn-in-architecture/>

Chaddad, A., Wu, Y., Kateb, R., & Bouridane, A. (2023). Electroencephalography signal processing: A comprehensive review and analysis of methods and techniques. *Sensors*, 23(14), 6434. <https://doi.org/10.3390/s23146434>

Chaillou, S. (2019, September 17). The advent

of architectural AI. *Towards Data Science*. <https://towardsdatascience.com/the-advent-of-architectural-ai-706046960140>

Chaillou, S. (2020). *ArchiGAN: Machine learning for architecture* (Master's thesis, Harvard Graduate School of Design).

Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford University Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.

Corbellini, G. (2020). *Window and mirror: Translating architecture*. *Villardjournal*, 002, 5–14. <https://doi.org/10.2307/j.ctv160btcm.5>

Copeland, B. J. (2014). *Turing: Pioneer of the information age*. Oxford University Press.

ComfyUI. (n.d.). *Comfy – Home*. <https://www.comfy.org/>

Daily Nous. (2020, July 30). Philosophers on GPT-3 (with replies by GPT-3). <https://dailynous.com/2020/07/30/philosophers-gpt-3/>

Dean, J. (2021). Building large neural networks. *Communications of the ACM*, 64(4), 36–38.

Diamandis, P. H. (2018, August 22). Ray Kurzweil's mind-boggling predictions for the next 25 years. *Medium*. https://medium.com/@singularity_41680/ray-kurzweils-mind-boggling-predictions-for-the-next-25years-ce3c9163588b

Domingos, P. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books.

DALL-E. (2020). *OpenAI*. <https://openai.com>

Davis, B. (2023, December 26). 10 predictions about the unexpected impact of AI on art (Part 1 of 2). *Artnet News*.

Divisare. (n.d.). *Yusuhara Wooden Bridge Museum*. <https://divisare.com/projects/310486-kengo-kuma-associates-takumi-ota-yusuhara-wooden-bridge-museum>

Eco, U. (2003). *Mouse or rat? Translation as negotiation*. Phoenix.

Evans, R. (1986). *Translations from drawing to building and other essays*. MIT Press.

Epstein, R., Roberts, G., & Beber, G. (Eds.). (2008). *Parsing the Turing test: Philosophical and methodological issues in the quest for the thinking computer*. Springer.

Fourtané, S. (2019, June 23). Refik Anadol: The Leonardo da Vinci of the 21st century. *Interesting Engineering*. <https://interestingengineering.com/refik-anadol-the-leonardo-da-vinci-of-the-21st-century>

Game, A. (2016). *Intimate ecologies: An exploration of the languages of contemporary exhibitions and making in museums and related cultural spaces in the UK* [Doctoral dissertation, University of Leicester].

Google Arts & Culture. (2018). *WDCH Dreams*. <https://artsandculture.google.com/exhibit/wdch-dreams-laphil/yQlyh25RSGAtLg?hl=en>

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672–2680).

Henning Larsen. (2023). Rethinking our relationship with AI. <https://henninglarsen.com/news/rethinking-our-relationship-with-ai>

Hassabis, D., & Hui, F. (2017, April 10). Exploring the mysteries of Go with AlphaGo and top Chinese Go players. *DeepMind Blog*. <https://deepmind.com/blog/article/exploring-mysteries-alphago>

Huwe, K. (2018, November 20). Refik Anadol: “WDCH Dreams”. *Flaunt*. <http://www.flaunt.com/content/refik-anadol>

Huyssen, A. (1995). *Twilight memories: Marking time in a culture of amnesia*. Routledge.

Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1125–1134).

Japan Meteorological Agency. (2021). *Seismic Intensity Database*. <https://www.data.jma.go.jp/svd/eqev/data/kyoshin/jishin/index.html>

Japan Meteorological Agency. (2023). *Climate statistics for Akita*. <https://www.data.jma.go.jp>

Jencks, C. (2002). *The new paradigm in architecture: The language of postmodernism*. Yale University Press.

Johnson, P., & Wigley, M. (1988). *Deconstructivist architecture*. Museum of

Modern Art.

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260.

Kamogawa, S. (2015). *A new guide to modern Japanese craft arts: The beauty of Japan through selected masterpieces*. The National Museum of Modern Art.

Kasparov, G. (1996, March 25). The day that I sensed a new kind of intelligence. *Time*, 13.

Kasparov, G. (1997, May 26). IBM owes me a rematch. *Time*.

Kasparov, G. (2017). *Deep thinking*. PublicAffairs.

Kelleher, J. (2019). *Deep learning* (p. 251). MIT Press.

Kember, S., & Zylinska, J. (2012). *Life after new media: Mediation as a vital process*. MIT Press.

Koza, J. R. (1992). *Genetic programming: On the programming of computers by means of natural selection*. MIT Press.

Koolhaas, R. (1995). *S, M, L, XL*. Monacelli Press.

Kuma, K. (2008). *Anti-object: The dissolution and disintegration of architecture*. Architectural Association.

Kwon, Y., & Ahn, M. (2024). Artificial intelligence and the ethics of architectural authorship. *Journal of Architectural Education*, 78(1), 58–69. <https://doi.org/10.1080/10464883.2023.2266112>

Leach, N. (2014). There is no such thing as digital design. In D. Gerber & M. Ibanez (Eds.), *Paradigms in computing: Making, machines, and models for design agency in architecture* (pp. 148–158). eVolo Press.

Leach, N. (2017). Introduction. In P. Yuan, A. Menges, & N. Leach (Eds.), *Digital fabrication* (pp. 10–13). Tongji University Press.

Leach, N. (2021). *Architecture in the age of artificial intelligence: An introduction to AI for architects*. Bloomsbury Visual Arts.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

Marcos, A., Fernández-Álvarez, J., & Pak, B. (2024). Rethinking authorship and materiality in post-digital architecture. *Architectural Design*, 94(2), 36–43. <https://doi.org/10.1002/ad.2876>

Manovich, L. (1992). Assembling reality: Myths of computer graphics.

Manovich, L. (2018). *AI aesthetics*.

Manovich, L. (2023). AI image and generative media: Notes on ongoing revolution. In L. Manovich & E. Arielli (Eds.), *Artificial aesthetics* (Chap. 5, p. 2). <https://manovich.net/content/04-projects/175-artificial-aesthetics/manovich-ai-image-and-generative-media.pdf>

Manzenreiter, G. (2023). Memory, Draw. *The future of art or the art of the future?*

Mayne, T. (2019). Interview. In C. Ahrens & A. Sprecher (Eds.), *Instabilities and potentialities: Notes on the nature of knowledge in digital architecture* (pp. 115–117). Routledge.

McPhee, M., Baker, K. C., & Siemaszko, C. (2015, May 10). Deep Blue, IBM's supercomputer, defeats chess champion Garry Kasparov in 1997. *New York Daily News*. <https://www.nydailynews.com/news/world/kasparov-deep-blues-losingchess-champ-rooke-article-1.762264>

Midjourney. (n.d.). *Home*. <https://www.midjourney.com/home>

Midjourney. (n.d.). Image weight – prompt parameter reference. *Midjourney Documentation*. <https://docs.midjourney.com/docs/image-prompts#image-weight>

Midjourney. (n.d.). Remix mode – prompting with variations. *Midjourney Documentation*. <https://docs.midjourney.com/docs/remix>

Midjourney. (n.d.). Seed – prompt parameter reference. *Midjourney Documentation*. <https://docs.midjourney.com/docs/seed>

Mitchell, M. (2019). *Artificial intelligence: A guide for thinking humans* (p. 20). Farrar, Straus and Giroux.

Mitchell, T. M. (1997). *Machine learning*. McGraw-Hill.

Mitchell, W. J. T. (1992). *The reconfigured eye: Visual truth in the post-photographic era*. MIT Press.

Moeran, B. (1997). *Folk art potters of Japan: Beyond an anthropology of aesthetics*. Curzon Press.

Morris-Suzuki, T. (1998). *Re-inventing Japan: Time, space, nation*. M. E. Sharpe.

Murphy, K. P. (2012). *Machine learning: A*

probabilistic perspective. MIT Press.

MVRDV. (2025, March 27). Fredy Fortich at PA Academy's Generative Architecture with AI SDflux workshop. <https://www.mvrdv.com/events/4670/fredy-fortich-at-pa-academys-generative-architecture-with-ai-sdflux-workshop>

Newborn, M. (2011). *Beyond Deep Blue: Chess in the stratosphere*. Springer-Verlag.

Nilsson, N. J. (2009). *The quest for artificial intelligence: A history of ideas and achievements*. Cambridge University Press.

Parkes, J. (2023, August 18). Rise of AI marks the “first time in history where we stop being monkeys,” says Arturo Tedeschi. *Dezeen*. <https://www.dezeen.com/2023/08/18/arturo-teseschi-ai-architecture-interview/>

Place, C., & Wharf, C. (2016). *Timber gridshells: Architecture, structure, and craft*. Routledge.

Popper, K. (1995). *The open society and its enemies*. Routledge.

Russakovsky, O. (2016, August 23). AI's research rut. *MIT Technology Review*. <https://www.technologyreview.com/2016/08/23/157971/ais-research-rut/>

Russell, S. (2018). In M. Ford (Ed.), *Architects of intelligence: The truth about AI from the people building it* (pp. 40–45). Packt.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536.

Schaeffer, J., & Plaat, A. (1997). Kasparov

versus Deep Blue: The rematch. *ICGA Journal*, 20(2), 95–101. <https://doi.org/10.3233/ICG-1997-20209>

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85–117.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–457.

Shakespeare, W. (1599). *Julius Caesar* (Act 1, Scene 2). http://www.online-literature.com/shakespeare/julius_caesar/3/

Shea, K. (2004). Directed randomness. In N. Leach, D. Turnbull, & C. Williams (Eds.), *Digital tectonics* (pp. 82–88). Wiley.

Simon, H., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, 6(1), 7–8. https://iiif.library.cmu.edu/file/Simon_box00064_fld04874_bdl0001_doc0001/Simon_box00064_fld04874_bdl0001_doc0001.pdf

Sinapayen, L. (2018, November 20). Sophia the robot, more marketing machine than AI marvel. *Skynet Today*. <https://www.skynettoday.com/briefs/sophia>

Snøhetta. (2024, January 1). Input: Output – Curating creative intelligence. <https://www.snohetta.com/perspectives/harnessing-generative-design-and-creative-technology>

Standage, T. (2002). *The Turk: The life and times of the famous eighteenth-century chess-playing machine*. Walker.

Statistics Bureau of Japan. (2020). *Statistical handbook of Japan*. <https://www.stat.go.jp/english/data/handbook/index.html>

[english/data/handbook/index.html](https://www.stat.go.jp/english/data/handbook/index.html)
Statistics Bureau of Japan. (2023). *Population statistics by prefecture*. <https://www.stat.go.jp/english/>

Sugihara, S. (2019). Design outside of the frame: A role of architects in the era of artificial intelligence. In C. Ahrens & A. Sprecher (Eds.), *Instabilities and potentialities: Notes on the nature of knowledge in digital architecture* (pp. 177–178). Routledge.

Terdiman, D. (2018, July 20). Autodesk's Lego model-building robot is the future of manufacturing. *Fast Company*. <https://www.fastcompany.com/90204615/autodesks-lego-model-building-robot-is-the-future-of-manufacturing>

Turing, A. M. (1937). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2–42(1), 230–265. <https://doi.org/10.1112/plms/s2-42.1.230>

UNDP Asia and the Pacific. (2017, November 22). UNDP appoints world's first non-human innovation champion.

Utsu, T. (2002). *Catalog of damaging earthquakes in the world (through 2000)*. University of Tokyo Press.

Walsh, T. (2018). *Machines that think: The future of artificial intelligence*. Prometheus Books.

Watanabe, M. (2019). AI Tect: Can AI make designs? In C. Ahrens & A. Sprecher (Eds.), *Instabilities and potentialities: Notes on the nature of knowledge in digital architecture* (pp. 68–71). Routledge.

Weinberg, J. (2020, July 30). Philosophers on GPT-3 (with replies by GPT-3). *Daily Nous*. <https://dailynous.com/2020/07/30/philosophers-gpt-3/>

Weller, C. (2017, October 27). Meet the first-ever robot citizen—A humanoid robot named Sophia that once said it would “destroy humans.” *Business Insider*. <https://www.businessinsider.com/meet-the-first-robot-citizen-sophia-animatronic-humanoid-2017-10>

Williams, A., Miceli, M., & Gebru, T. (2022). The exploited labor behind artificial intelligence. *Noema*. <https://www.noemamag.com/the-exploited-labor-behind-artificial-intelligence>

Williams, D. E. (1988). Ideology as dystopia: An interpretation of *Blade Runner*. *International Political Science Review*, 9(4), 381–394. <http://www.jstor.org/stable/1600763>

Wigley, M. (2015). The architecture of atmosphere. *Log*, 33, 96–111.

Yoshino, M. M. (Ed.). (1992). *Climatic variability and ecosystem response in Japan*. In M. M. Yoshino (Ed.), *Global environmental change in the Asia-Pacific region* (pp. 63–78). United Nations University Press.