

Projections: On the Design Process in AI Architecture

Roberto Bottazzi

AI and architecture: Innovation and invariance

Challenges posed by technical innovations to artistic disciplines are never just “technical.” Their significance is proportional to their capacity to exceed the purely technical domain and reconfigure the tenets of a field to furnish new representations, reconsider practices, and develop new agendas. The integration of AI in architecture and urbanism is no exception. As we question what architectural intelligence could become in the light of the introduction of AI, the task is not so much technical—the penetration and improvement of AI models in architecture will continue regardless—as conceptual. Mathematics offers fruitful analogies for setting the terms of this inquiry. We can think of AI and architecture as two series, each with their own characteristics, and our role as crafting the instruments for their convergence and resonance. To extend the mathematical analogy, we could think of the elements of connections between series as invariants: properties that remain unchanged when a system of objects such as a series undergoes transformation. To frame the relation between AI and architecture through series and invariants will not subsume either field under the other; rather, it will allow us to understand the specificity of each, and redefine their operations on each other.

This essay explores how to think of invariants in the burgeoning relation between AI and architecture. The technical logic of AI must guide this search but not limit it, as the challenge is how to think of AI models within the design process. As Alejandro Zaera-Polos notes, “nothing gets built that isn’t transposable onto AutoCAD,”¹ indicating not only that the history of architecture is also

1 Quoted in Bernard Cache, “Towards a Non-Standard Mode of Production,” in *Projectiles: Architecture Words 6* (Architectural Association, 2011), 61.

the history of their design technologies, but also that design agendas emerges from the conceptualization of the technological instruments at hand. In so doing, we aim to stimulate a more proactive approach to the integration of AI in architecture, which at present seems to be concerned with measuring architecture's ability to comply with the principles of AI.

Beyond visual mimicry: The technical logic of AI

Technically, Deep Learning (DL) models—a subset of AI that we will refer to for the arguments put forward in this essay—produce an output from a vast collection of input instances. Broadly speaking, DL models are programmed to devise a number of steps to turn input data into an output. Inputs and outputs are separated by several layers of parameters—representing the “depth” of the model—which adjust after each input in order to tune inputs and outputs.² In other words, during the tuning of parameters—known as training—DL models generate processes that can eventually be applied to a vast range of issues (training and application are, in fact, two distinct phases). This technical characteristic marks a clear departure from previous computational generative methods in which users were tasked with designing processes by formalizing knowledge into code. Now inputs and outputs are the main points of intervention for designers, whereas process (what sits in between) can only be inspected or altered indirectly. Central to this reconfiguration of the creative process is the mechanism known as “backpropagation”: the differential feedback function that adjusts the parameters of a DL model to align inputs and outputs during the training stage.³ Seen from the point of view of design processes, by moving upstream from outputs back to inputs, backpropagation shifts the user's agency from process to output. It is through the evaluation of outputs that users can adjust input data, or, as we will see, charge the output with new, additional design ambitions. Another key technical aspect that DL models share with other AI models is vectorization: That is, to be processed by the model, input data must be abstracted into vectors, which become the actual objects manipulated by DL models. As all input

2 Ian Goodfellow, Joshua Bengio, and Aaron Courville, *Deep Learning* (MIT Press, 2016).

3 Arthur E. Bryson, “A gradient method for optimizing multi-stage allocation processes,” in *Proceedings of the Harvard University Symposium on digital computers and their applications*, 3–6 April 1961 (Harvard University Press, 1962).

data is eventually rewritten in a vector, within the DL model translations and projections between different media are technically possible. In perhaps the most familiar example, text-to-image AI tools generate images from text, that is, the model projects a vectorized string of text onto an image database to return a new image.

From the point of view of design, architects are endowed with the technical means to project data onto each other and “backpropagate” their intentions onto the design process; the cumulative effect of these two conditions brings design closer to curatorial practices. If in traditional algorithmic approaches designers reified knowledge into code to solve problems, with DL models the designers’ task is to critique, question—in short, curate—the outputs of a DL model. The actions require designers to critically investigate the space between two technical and the architectural series to foreground what Aby Warburg called the “iconology of the interval,”⁴ the speculative space between existing things.

The exploration of the aesthetic of intervals is further complexified and enhanced by the logic of vectorization. Though technically impressive, the current overreliance of AI tools on image production could also represent a potential involution for the aesthetic of architecture. Images are in fact only representational devices that rewrite numerical probabilistic distributions computed by DL models; in short, they are visualizations of vectors. It is therefore limiting to think of the capacities of AI platforms in terms of visual outputs or mimicry, as their technical logic is much broader and more complex, involving mathematical operations of abstraction, recombination, and projection. It is along these lines that an aesthetic for AI architecture should be articulated. Vector-based projections shift creativity away from pure image-making towards a more curatorial and strategic approach. Again, it is the relation between datasets, the questions posed to tease out relations that come to the fore to provide an alternative approach to visual mimicry. How might we start redefining creative processes as projections?

Projections: How to explore the interval

The following examples are part of a larger list of precedents to suggest a strategic use of AI in architecture. As the issue is conceptual rather than technical,

4 Matthewa Rampley, “Iconology of the Interval,” *Word & Image* 17, no. 4 (2001): 303–24.

these precedents are understood as speculative instruments through which to attend to issues of curation and projection in architectural design. Given that what is being suggested is a way of thinking rather than a specific method, the list of precedents that follows deliberately avoids architecture. By referring to art or music, operations of projection and curation, literal applications will unavoidably give way to critical projections.

In 2016 former *New York Times* music critic Ben Ratliff published *Every Song Ever: Twenty Ways to Listen to Music Now*.⁵ The “now” Ratliff concentrates on is the age of digital platforms in which a very large database of songs (more than any individual would previously have had access to, yet only a portion of all the recordings there are, as Ratliff himself acknowledges) is available through our smartphones. Each chapter in the book focuses on a theme (virtuosity, density, closeness, etc.) and strings together a series of examples that furnish an open map of how musicians (from different periods and traditions) have interpreted or developed such themes. The result is not only a rich landscape of musical experiments (a true learning journey, regardless of one’s musical taste and knowledge), but also an empowerment of the reader/listener whose agency over digital platforms is emboldened. Readers of this essay who pick up *Every Song Ever* with the expectation of finding a compelling analysis of digital technologies for music consumption will probably be disappointed. The book in fact entirely bypasses issues such as data, algorithms, etc.; that is, the themes featuring in most of the literature on AI. Architecture does not feature at all, whereas cities only appear as backdrops to the development of a particular piece of music, genre, etc. Why is Ratliff’s book so relevant in the context of AI and architecture, then? Rather than engaging in endless principled battles on the morality of AI, Ratliff moves the conversation one step forward by “naturalizing” it and examining what space for creativity and learning AI might usher. The premise of the book is to recognize the limitations of previous models for listening to music (such as categorizing the work by author or album) and explore the opportunities provided to listen better or differently. This is not to say that the book does not offer elements of critique; on the contrary, the twenty strategies animating each chapter can be seen as forms of resistance to algorithm-guided choices.⁶

5 Ben Ratliff, *Every Song Ever: Twenty Ways to Listen to Music Now* (Penguin, 2016).

6 Ben Ratliff, “Listening, inefficiency, and value,” lecture delivered as part of the B-Pro Prospective lecture series, The Bartlett School of Architecture, UCL, November 16, 2023.

To tease out relations implies a procedural change that Ratliff articulates through a series of ideas that indicate practical and conceptual ways to move transversally and exercise intelligence as we navigate through vast amounts of data. On the one hand, the structure of the book closely mirrors the technical logic of the processes criticized when we listen to music on digital platforms: given a starting point, we quickly fall into a “rabbit hole” of free associations, emotional responses, rational progressions, philological inquiries, etc., that take us through the massive archive available in a transversal fashion. On the other, while the logic of digital platforms is embraced, a whole series of other considerations are overlaid, or, shall we say, projected and curated, onto the act of listening: these regard the theme of each chapter and unfold irrespective of traditional classifications. It can be a detail, a hidden connection between two musicians who are part of a larger ensemble, a sustained engagement with a piece of music to be the source of discovery for new readings. If we drew an analogy between Ratliff’s approach and DL models, we could say that his is a call for small data, for precision and definition. This analogy would entail a non-mimetic and disjunctive relationship between the technology (which can thrive in managing colossal datasets) and our interaction with it that would be based on understanding and exploiting the technological affordances provided, but would complement, rather than mimic, them.

Finally, Ratliff is addressing us, the listeners, not the technological apparatus we use to access music: he proposes a form of digital literacy in the shape of an intellectual gymnastics afforded by digital platforms and their ability to let us travel across the widest spectrum of musical production ever accessible. It exploits the possibility to break boundaries, redraw them, or even follow existing ones, suggesting that the key skill required is not technological prowess but intellectual curiosity. The implicit message seems to be that digital literacy will benefit from the contribution of any other discipline *but* the strictly technical ones. Ratliff’s book is thus an ideal for our list of conceptual approaches to AI in architecture. Each chapter in the book offers a series of instruments that can charge an apparently passive activity such as listening to music or inspecting the outputs of a DL model in architecture into a creative moment in its own right to define an aesthetic of curation.

Roni Horn’s artistic production has often concentrated on paired objects, on the art of “distant doubles.”⁷ This theme is present in several pieces such

7 Christy Lange, “Clowd and Cloun (Blue),” in *Roni Horn*, ed. Ingvild Goetz, Larissa Michelberger, and Rainald Schumacher (Hatje Cantz, 2013), 179.

as *You are the Weather*⁸ (1994–96) or the *Pigment Drawings*⁹ series (1984–2012) in which objects appear doubled up or in pairs. The work *Cloud and Clown (Blue)* (2000–01) consists of thirty-two photographs alternating images of a blurred portrait of a clown and that of a shifting white cloud in an otherwise clear blue sky. Arrayed either in a line or a grid, the subject in each of the two series varies as we move from photograph to photograph, albeit not in a continuous fashion.

The work plays in complex ways with themes we have already encountered. The notion of the series is used to represent each of the subjects—i.e., clouds and the portrait of the clown—which are however both portrayed through the medium of photography. The viewer is invited to focus on the “interval,” the “nameless” space connecting the two disparate series. The notion of non-mimetic representation is also at play. The two series do not converge on a descriptive, narrative, or, most importantly, visual plane, but rather on a linguistic one (the least obvious); the title of the piece in fact plays with the similarities between the obsolete spelling of the words “clown” and “cloud.” It is, again, the interval between the images that gives rise to an open, ambiguous condition opening up a conceptual space for repositioning.

Similar conditions had already been explored by Horn in other pieces such as *You Are the Weather*, in which interaction between the format of the portrait and the weather leaves the viewer wondering who is influencing whom. The piece consists of a series of 100 photographs of the same model as she takes a daily dip in different geothermal pools in Iceland. The effect is that of a projected portrait in which we seem to be invited to see one element (the face of the subject portrayed) through the eyes of another (the weather), and vice versa. In both pieces, it is the consistent use of the medium—the photographic series—that is tasked to be the technical support to allow the viewer to see projectively. Is the cloud being portrayed on the face of the clown or vice versa? Analogies with earlier considerations on the projective logic of DL models can be drawn as well as considerations on the technical process of vectorization. Through Horn's work, we can appreciate how projections allow for the object/subject relation to be mobilized, swapped, and expanded in order to charge the act of reading the piece into a creative moment in its own right.

8 Roni Horn, *You are the Weather* (Scale, 1997).

9 Amy S. Wilkins, ed., *Roni Horn: 153 Drawings* (Hauser & Wirth, 2013).

Intelligence in AI architecture could thus be one in which multiple points of view can be instrumentalized to dislodge preconceived positions and conjure up a more inclusive, uncertain, and even fragile architecture. Some of the experiments we have been developing with students that play with the notion of projection and curation through DL models have allowed them to assume unconventional points of view to instrumentalize non-architectural themes such as language, sounds, visual clues, etc. as subjects of their design. These entry points allow designers to reread the city and architecture to reconceive it from non-human points of view.¹⁰ Again, Christy Lange's words on Horn's piece are illuminating:

Both clowns and clouds are spaces of projection, more fleeting and intangible than definitive objects. By pairing the two within the same representational system—the photographic series—Horn suggests how that system of representation might be as arbitrary or mutable as the linguistic system used to name them. *Clowd and Cloun (Blue)* forms a circulatory, contingent web of signs, sounds, words, images, and meanings—all dependent on each other, and tenuously linked by resemblance. In it, we see the gaps between word and image, and, by extension, between the image and their subject, between the image of ourselves, and our selves.¹¹

In his famous *Simulations and Simulacra*,¹² Jean Baudrillard had already warned that a technical apparatus able to indefinitely produce and reproduce images would devoid the notion of reality and its rationality to replace it with purely operational procedures. Though in a less apocalyptic tone, some of Baudrillard's warnings still resonate with us as we witness the relentless production of AI-generated images of architecture. To counter this trend, this essay foregrounded the importance of a conceptual approach to AI in architecture in the form of an open list of precedents. The proposition put forward is that what the discipline of architecture needs now is not visual inflation through endless production of images or a technical subsumption of architecture to AI. Rather, what the debate on AI architecture is missing is the ability to conceptualize the technical logic of AI within the disciplinary

10 Roberto Bottazzi, Mollie Claypool, and Tyson Hosmer, "Disruptive Ecologies: Design with Nonhuman Intelligences," *Architectural Design* 94, no. 1 (2024): 30–37.

11 Lange, "Clowd and Cloun (Blue)," 180–81.

12 Jean Baudrillard, *Simulations and Simulacra*, trans. Sheila Faria Glaser (University of Michigan Press, 1981).

repertoire of architecture: to indicate references, directions, practices; in short, a conceptual agenda for the discipline in the light of the penetration of AI in architecture. Beyond visual mimicry, the potential of AI in architecture is far greater and deeper: to redefine architecture's conceptual and technical operations.

Fig. 32: Sensory Balance uses DL models to design the sensory (visual, olfactory, and auditory) experiences of urban environments. DL models and data spatial analytics are deployed to survey and generate different ephemeral spatial qualities concentrating on sensorial experiences rather formal innovation. For instance, DL models are trained on film images to provide a palette of colors, spatial arrangements, and materials to induce particular emotions in the user.




Fig. 33: Accent Diffusion utilizes data analysis and DL models to generate an urbanism based on accents, projecting the immaterial qualities of the 270 languages making up London's cultural landscape onto a database of physical artefacts. The massive catalogue of morphologies generated was used to represent the complex and hybrid cultural landscape of London.



Fig. 34: Ebb and Evolve explores the use of DL models to develop an urban strategy for flood-prone areas in East London. Urban, social, and environmental data are analyzed to develop a time-based strategy in which programs and structures can change to adapt to raising water levels. A series of inflatable structures are used to provide emerging spaces, protect important buildings, and provide a communication infrastructure for the local population.




Fig. 35: Attenomy City deploys DL models to study and design spaces around the notions of attention and spatial intelligibility to redesign Euston Station in London. The most dramatic element of the project is the large roof spanning over different parts of the site. The roof performs different roles such as guiding pedestrians as they cross the public areas and controlling environmental conditions. Its overall effect is, however, much greater than any of its functions: the dynamic features of the roof animate and guide the users' experience.

