This City Does not Exist
An attempt at a theory of Neural Urban Design

MATIAS DEL CAMPO
Taubman College for Architecture and Urban Planning

2015 was a turning point in the application of techniques derived from AI research to the arts. This year saw the introduction of the Generative Adversarial Network (GAN) by Ian Goodfellow1 and the paper *A Neural Algorithm of Artistic Style* by Leon Gatys et al[2] In recent years, these novel methods have taken hold in the arts[3] and music4. The newly emerging artform is fittingly named Neural Art[5]. The source of this term can be found in the title of a paper by Leon Gatys[6], which forms the base for the work of several of the most prolific Neural Artists, such as Mario Klingemann (who describes himself as a *Neurographer*) and Sofia Crespo, who's series of works named *Neural* Zoo reflects a keen interest in the estrangement and defamiliarization of deep-sea creatures. There are many more artists in this area[7]. How is the result of these artists related to architecture and urban design? Maybe an example will help to clarify how provocative this novel method, in a good way, is for architecture. Mario Klingemann uses databases of Western art, mainly portraits, as a base for his StyleGAN[8] applications. Thousands of images from the Renaissance to the 19th century are fed through a StyleGAN algorithm. In a conventional case, StyleGAN would be used to create images that convincingly represent a known object to the observer. Most famously demonstrated with examples like *This Person does not exist* (Fig.1), a website that generates compelling images of persons based on an extensive database of portraits. It might be necessary here to understand that Neural Networks are Function Approximation Algorithms[9]; they will always strive to approximate the number 1. Or in other words, in building the training dataset, it seeks to adapt a curve to a given condition based on input weights. A curve can also describe function approximation; thus, the famous quote "Machine Learning is just glorified 'Curve Fitting'"[10]. As much as this approach can produce convincing images of objects, it is the area outside the perfect fit of the curve that produces the more interesting results in that they maintain a certain familiarity, despite their alien appearance. To come back to the work of Mario Klingemann as an example of what I mean by that: His images and animations maintain elements of the database informing the StyleGAN. This results in images that show contorted bodies and distorted faces that have a surreal quality to them – bizarre Janus heads with multiple faces, strange Cyclops, monstrous Chimeras between humans and animals. (Fig.2) The trained eye will still recognize features of historical paintings and drawings. A bit of Goya here, some Whistler there, glimpses of Jeanne-Etienne Liotard, Eduard Magnus, Lenbach, Winterhalter. But none of the images is designed to approximate the particular artists, as 'This Person does not Exist' would do. Instead, it renders the features recognized by the Neural Network and re-combines the pixels into a new image outside the conventions as to how we as humans understand the depicted object. In doing so, the emergent pieces of art provoke questions of authorship and agency. In addition, the question of the value of a sensibility that was created somewhere between human input and machinic output is raised. Is this the art of the posthuman age? Does this development take into consideration the possibility to reconsider the role of humans in a world where the boundaries between human and nonhuman creativity are blurred? In architecture, we can observe a similar tendency with architects increasingly

Figure 1. Philip Wang, A grid of faces of people that do not exist - https://thispersondoesnotexist.com/

Figure 2. Mario Klingemann, Memories of Passersby I, 2018. (Credit: Quasimodo)

picking up novel techniques in machine learning and machine vision. I would describe this new tendency in Architecture as ***Neural Architecture,*** borrowing from Computer Science and Neuroscience from the language used in the arts and music (Neural Art, Neural Music). Increasing the scale allows, of course, to speculate about the value of a ***Neural Urban Design.***

Familiar but Strange: Bits, Pieces, Features & Neurons

I have relied on Neural Networks to train databases, exploring them in a reverse flow of information in my work, which means that algorithms that are typically trained to recognize objects, such as in automated cars, etc., are used to perform *generatively* instead of *analytically.* A similar technique to the ones used by the artists mentioned above. Unbeknown to myself, for example, Sofia Crespo uses a very similar 2D to 3D Style transfer technique to the one I have been using in recent years[11].To give you a more concrete example, let me rely on an image-based approach: Using a database with several thousand images of Baroque plans (Fig.3), we experimented how a Neural Style transfer would imprint the learned features from the baroque plans to a modern plan. Echoing the conversation about the work of Neural Artists, we achieve results that are familiar but strange at the same time. The nature of Baroque architecture reflected in the dynamic opposition and intersection of elliptical volumes of space, the thick poches, undulating walls, and rich surface treatment are features that are learned by the Neural Network and applied to modern plans (Fig.4). Even though this is a simple experiment, it was a proof of concept as to how this method does not simply imitate a style or create a crude collage of elements. Still, instead, it produces unexpected, highly provocative, and yes – inspiring images.

With these new toolsets come a whole set of questions regarding their ontological and epistemological qualities. How do they change the way that architects conceive of their projects?

What does that mean? To unpack this problem, we need to divide the problem into its specific components. The assemblage of elements needs to be laid out. Knolling the problem, so to speak. In neatly laying out the particular parts of the conversation, we hope to achieve a clear overview of the challenge that urban design is currently facing in the gestalt of a novel agent in design. Unpacking the pieces, unwrapping them, and cleanly positioning them on a flat surface allows for a clear overview. We will unpack the problem by closely looking into the meaning of Neural Networks and slowly and methodically working ourselves towards the aspect of the Epistemology of the object we are observing and the elements of paradigm and theory that they produce.

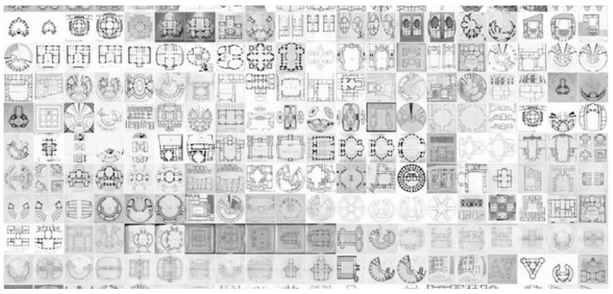

Figure 3. Matias del Campo, selection of images from the Baroque Data base, 2020 (Credit: Matias del Campo)

---

[1]Goodfellow, Ian J., Jean Pouget-Abadie, M. Mirza, B. Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. "Generative Adversarial Networks." ArXiv abs/1406.2661 (2014): n. pag.

[2]Gatys, Leon A., Alexander S. Ecker and M. Bethge. "A Neural Algorithm of Artistic Style." ArXiv abs/1508.06576 (2015): n. pag.

[3]Elgammal, A., Bingchen Liu, Mohamed Elhoseiny, and M. Mazzone. "CAN: Creative Adversarial Networks, Generating "Art" by Learning About Styles and Deviating from Style Norms." ICCC (2017).

[4]Schwartz, H.. "Evolving Chord Progressions as Neural Networks." (2006).

[5]Artificial Neural Networks and Paintings: What is Neural Art. https://www.artdex.com/blog/art-world/ann-artificial-neural-networks-paintings-neural-art/ (visited 01/15/2021)

[6]Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "A neural algorithm of artistic style." arXiv preprint arXiv:1508.06576 (2015).

[7]See for an extensive list of artists here: https://aiartists.org/ (visited 01/15/2021)

[8]Karras, Tero, S. Laine and Timo Aila. "A Style-Based Generator Architecture for Generative Adversarial Networks." 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019): 4396-4405.

[9]Brownlee, J., "Neural Networks are Function Approximation Algorithms" in https://machinelearningmastery.com/neural-networks-are-function-approximators/ (visited 01/15/2021)

[10]J. Pearl. 2015a. Causes of Effects and Effects of Causes. Journal of Sociological Methods and Research 44 (2015), 149–164. Issue 1.

[11]Campo, Matias del, Alexandra Carlson, and Sandra Manninger. "Towards Hallucinating Machines - Designing with Computational Vision." International Journal of Architectural Computing, (October 2020). https://doi.org/10.1177/1478077120963366.

Ontology

Neural Networks are abstract Objects.

As described above, NN's are prolific in recognizing relationships within a set of data. This is one of the reasons why this method is such a provocation for urban design. The traditional role of the designer is not very different. Architects, for example, are trained to recognize, understand and process a large set of correlated information to assemble them in a project. Just think of an urban design competition situation. There is a call for a competition; a group of rules in the form of programmatic, technical, social, and economic aspects is established. The specific program calls for the urban planner's knowledge, experience, and expertise to synthesize the rules in the form of a design idea. How does this *knowledge* emerge? Let us be concrete about the problem by assuming it is a specific competition; how about picking up a historical example. How about the "Park de la Villette" competition? The result of the competition can be described as a very concrete object[12]. To achieve this concrete object, a set of abstract objects have to be established first, such as plans and sections, which are based on calculations, numbers, and collections of diagrams demonstrating the spatial relationships. This produces a set of ontological dependencies, not only between concrete (material in the form of a building) and abstract (design, plans, sections) object but also between the concrete objects (such as the Follies, in the case of the Park de La Villette) and the set of data present in the architect's mind. How many parks did he/she design before, how many did he/she see, how many images of parks were absorbed? Admittedly, this is a weak ontological dependency, but a dependency regardless. With this in mind, we can interrogate these entities, these examples of parks in the architect's mind, by their properties.

Properties

What is meant by properties? Properties allow us to compare entities with each other. To stay with the example of the park: The Park de La Villette[13] (Fig.5) in Paris has a playground and a cafe. So does the Central Park in New York[14], the Fuxing Park in Shanghai[15,] and the Stadtpark in Vienna[16]. So we can assume that playgrounds and cafés are properties of a park in that they are universally present. Any number of parks can share them. To that extent, we can consider the park as an object with a bundle of properties. (It would be easy to name more properties present in parks in general, such as benches, trees, fountains, monuments, walkways, bandstands, pergolas, etc.) This might be true for every urban object that operates along with a specific program which ultimately represents a compresence-relation responsible for the bundling. Of course, it is necessary to distinguish between the categorical and dispositional properties of an urban object. When discussing definite concerns, we discuss what something *is like*, meaning its qualities. Dispositional qualities discuss the potentialities of an object (which can be tied to Manuel DeLandas' conversation on p*otentialities*[17]). The Park de la Villette has a specific shape, a definite property, while its tendency to provoke joy is a dispositional property. Concerning the discussion on AI and architecture, we are primarily interested in the categorical properties as these can be examined by Neural Networks. Shape, dimension, plan, section, umriss, pixels, etc. For this reason, we would also include color in the area of the categorical properties as Neural Networks divide color into their RGB values, into numbers, that can be interrogated for their properties, thus positioning themselves within the conversation of abstract objects.

Relations

When talking about the properties of a piece of architecture, it seems evident also to discuss the relationships of these objects to each other. Traditionally in architecture, the discussion of the relationships of buildings to their environment, whether it is urban or rural, is described as the contextualization of a building. I do not know how many hours I spent in Hans Hollein's office discussing the contextualization of a specific structure. How does this axis meet that axis? How does the sun rotate around the building? How are the proportions and scales of the buildings around the project? etc. etc. I had to build dozens of models, physical models in blue foam, with variations of how a building is contextualized in its surroundings (Fig.5). Later I applied the same method in a digital context, allowing me to increase the possible variations into the hundreds today it is databases with thousands of examples (Fig.6). This short story demonstrates the transition from the age of data scarcity to the era of data abundance and the struggle to harness the inherent information. Or as Mario Carpo put it:

"The collection, transmission, and processing of data have been laborious and expensive operations since the beginning of civilization. Writing, print, and other media technologies have made information more easily available, more reliable, and cheaper over time. Yet, until a few years ago, the culture and economics of data were strangled by a permanent, timeless, and apparently inevitable scarcity of supply: we always needed more data than we had. Today, for the first time, we seem to have more data than we need. So much so that often we do not know what to do with them, and we struggle to come to terms with our unprecedented, unexpected, and almost miraculous data opulence." [18]

Circling back to the statement about architects operating as data miners, the response to Carpo's concern about how to harness the qualitative information of big data to inform the architectural project at hand is the use of Neural Networks. Not only do NN's process enormous amounts of data, but they also need enormous amounts of data to learn anything of value, whether it is relationships (contextualization), features (properties), or behavioral patterns. To this extent, the collision of the wealth of data produced by the city and its history of urban design with processing methods borrowed from Artificial Intelligence results in Neural Urban Design.

About Features of things and how to capture them

What does this mean for the relations between these urban objects? We can find a close connection here to the aspects of properties, in that both -relationships and properties- describe

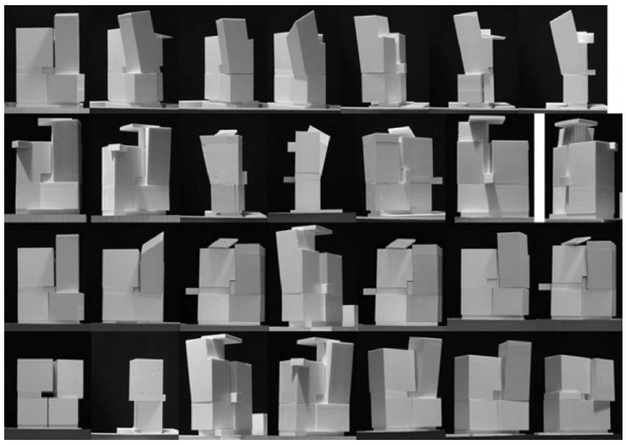

Figure 5. Hans Hollein, Series of Models for the Saturn Tower, Vienna, Austria, 2004 (credit: Sina Baniahmad)

the character of the things they are applied to. Occasionally properties are described as a special case of relations[19]. The example of the park as an urban object described in this essay allows us to differentiate between internal and external relationships. The Park de la Villette and the Stadtpark are in an internal relationship of similarity to each other as both are *parks*.

The relationship is determined by the relata of similar *Features*. As evident and almost simplistic as this insight seems to be, it is crucial for the operation and inner workings of Neural Networks. Features and the relationships to each other (for example, in an image) make it possible to train a Neural Network to recognize a park – or anything else it is trained to recognize, for that matter. Currently, this method is utilized in machine vision tasks that allow self-driving cars to recognize other vehicles, the street, signage, people, dogs, cats…– by understanding a family of internal relationships within a large number of images. In architecture, we can use this ability in an analytical sense and in a creative sense in that the flow of information can be reversed to generate a park instead of rec*ognizing* one. I would like to emphasize the impact of such a motion, as it opens up questions about agency, authorship, aesthetics, and sensibility. Apart from the consideration of internal relationships, there are also the aspects of external relationships. As briefly discussed before in architecture, this discussion would gravitate around the *contextualization* of a built object. In contrast to internal relationships, external relationships in architecture are not defined by similarities.

Things, Facts and the Ontology of Neural Networks

After defining that Neural Networks are abstract objects defined by their properties and their relationships and dabbling in abbreviated definitions of some of the terms used in this essay, one more question remains: Is Neural Urban Design a thing or a fact? Why is this relevant within the conversation laid out in this essay? It is very alluring to declare that every architecture is a thing, and thus every aspect and part of it is a thing too. But is this really so? What about plans, calculations, proportions, simulations? Are those things or facts? Well, maybe it would help first to briefly define what is meant by things and facts in this conversation. Thing ontologies discuss the possibility that the universe is made of a plurality of discrete objects[20]. To this extent, the paradigm example of a thing ontology would be atomism21 in contrast to non-thing ontologies, which consider that an object is not necessarily an assembly of objects but could also be one continuous object. For our conversation, let us stick to the idea of Neural Networks being things in that they represent an assembly of objects. These discrete objects can be the dataset of individual images or individual obj[22] models (as in the research on GraphCNN design methods[23]) or the numerical data of images (in the form of their RGB values) or the layers in the networks that fulfill different tasks such as edge recognition et al. How are these things and not facts? Or are they both? Let's take a Materialist perspective (and Materialist is interchangeable with Realist in this conversation). We can assume that Neural Networks' calculation is based on material processes: Electrically charged matter. moves, creating a set of physical phenomena resulting in an electric current which in turn is used to run computers and their respective GPU's (Graphic Processor Units), made of Silicon, Tantalum, and Palladium, that calculate the numbers laid out by a set of algorithms known as Neural Networks. All of which are things operating in discrete processes to create a result. The result can be the successful recognition of a car, a person, a dog, a sign, or a voice – or the creation of an image that we describe as art (Sofia Crespo, Mario Klingemann), music (Dadabots, YACHT, Holly Herndon) or architecture (SPAN (Fig 9), Daniel Bolojan (Fig.10), Immanuel Koh (Fig.11), etc.). What if Neural Networks are not things but facts? What is the difference? Things are, in general, put in contrast with the properties and relations they instantiate[24].

---

[12]Lowe, E. J. "The Metaphysics of Abstract Objects." The Journal of Philosophy 92, no. 10 (1995): 509-24. Accessed January 15, 2021. doi:10.2307/2940785.

[13]Tschumi, B., Park de la Villette, Paris, 1987

[14]Olmsted, F.L. and Vaux, C., Central Park, New York, 1876

[15]Selleny, J., Stadtpark, Vienna, 1862

[16]Built by the French in the French Concession, Shanghai, 1905

[17]DeLanda, M. Assemblage Theory, Edinburg University Press, Edinburg UK, 2016, pp. 73

[18]Carpo, M., The Second Digital Turn, The MIT Press, Cambridge, Massachusetts, USA, pp. 9

[19]Orilia, F., and Michele Paolini Paoletti, "Properties". The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University 2020. https://plato.stanford.edu/entries/properties/ (visited 01/15/2020)

[20]Michael E., University of Lausanne, Department of Philosophy Michael-Andreas.Esfeld@unil.ch, www.michaelesfeld.com (draft 21 December 2017 visited 01/15/2021)

[21]Cartledge, P., 1999. Democritus. New York: Routledge.

[22]Green, E. J. & Quilty-Dunn, Jake (forthcoming). What is an object file? _British Journal for the Philosophy of Science_:axx055.

[23]del Campo M., Manninger S., Carlson A., 3D Graph Convolutional Neural Networks in Architecture Design in the proceedings of the ACADIA Conference 2020, Distributed Proximities.

[24]Rettler, Bradley; Bailey, Andrew M. (2017). "Object". The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University https://plato.stanford.edu/entries/object/ (visited 01/15/21)

[25]DeLanda, M. Assemblage Theory, Edinburg University Press, Edinburg UK, 2016, pp. 79

Figure 6. SPAN, Matias del Campo, Sandra Manninger, Urban Fiction - Vienna 2020 (credit: Matias del Campo)

In contrast to facts described as adhering to the things in combination with their properties and relations as their *constituents.* In short, facts swallow things and their properties/relations versus things in themselves, maintaining independency in a flat ontology.

What about Aesthetics, Agency, and Authorship?

There are three takeaways here for the conversation on Neural Networks and their meaning for the architectural discourse. NN's are things; they are abstract objects with properties. They maintain internal and external relationships within a flat ontology, and they constitute discrete assemblies[25.] After interrogating the ontological frame in Neural Architecture, it might be time to ask if we actually gathered any knowledge from the frame we laid out. Do we know more now that we have explored questions such as objects, properties, and relations between elements at play in Neural Architecture? Most importantly, the previous interrogation already hinted towards some of the epistemological questions that arise out of the work with Neural Networks. Such as the Aesthetics, Agency, and Authorship of the resulting project. Relevant for the frame of conversation in this essay are the positions of Roland Barthes and Michel Foucault concerning the nature of Authorship and Author at large. These critics interrogated the role and relevance of authorship about the interpretation or meaning of text - this can be expanded to all areas of artistic production, such as architecture, in the present text of this essay. Barthes, for instance, attributed meaning to the language and not the author of the text. Instead of relying on the legal authority to exude authorship[26], Barthes assigns authority to the words and language itself. Foucault's critical position vs. the author can be found in the argument he presents in his essay *What is an Author*[27]. Foucault argues that all authors are writers, but not all writers are authors - echoing a broadly discussed sentiment in architecture: ***not every architecture is a building, and not every building is architecture.***

Sensibility of Neural Architecture:

When I talk about the term sensibility, it is specifically geared towards aspects of artistic sensibility. This assertion discusses the arts of the past as much as the arts of the present, recognizing the aesthetic value as a specific feature of all experience. Or, as Arnold Berleant put it: *Such a generalized aesthetic enables us to acknowledge the presence of a pervasive aesthetic aspect in every occasion, whether uplifting or demeaning, exalting or brutal. It makes the constant expansion of the range of architectural and aesthetic experience both plausible and comprehensible*[28]. How is it possible to scrutinize the term Sensibility when discussing aspects of urban design and architecture? What I mean by sensibility is the perceptual awareness developed and guided through training and exercise. To this extent, it is certainly more than simple sensual perception and closer to something like a trained or *educated* sensation. Education has to be continuously fostered, polished, and extended through encounters and activities to maintain the ability to execute tasks with an *aesthetic sensibility.* This ability is attributed in the Western traditions primarily to the arts -to painting, sculpture, music, literature, and so on - with architecture being this strange animal living somewhere between engineering and the arts. We can interrogate these territories of human production concerning their fashions, styles, etiquettes, and changing behavioral patterns resulting from transforming sensibilities. Giving rise to novel movements and entire epochs in the arts in general and architecture in particular.

Neural Architecture is a New Paradigm

New paradigms have the habit of emerging when the existing paradigm has run its course – but what is the current paradigm? That question alone can fill tomes, so we will leave that question open for now instead of focusing on what Neural Architecture brings to the table regarding a substantial Innovation in the Architecture discipline. For one, it critically interrogates the role of the architect in the creative process of architecture design. Neural Architecture embraces the possibility of a design method that is deeply informed by exiting information in the form of databases and understands that the artificial modeling of neural processes can aid in harnessing the information in Big Data. Interestingly the results do not resemble historical examples, and thus the methodology is not a repetition of Postmodern tropes, such as collage, quote, and ironic assemblies. The results instead construct a frame around aspects of ***defamiliarization and estrangement,*** in that we can recognize certain features without it being a copy. We are still at the beginning of this new paradigm of architectural production; the first built examples are currently emerging[28].

---

[26]Barthes, R., The Death of the Author, Essay 1967, in Sontag S., ed. A Barthes Reader. New York: Hill and Wang, 1982

[27]Foucault, M., What is an Author? in Faubion, J.D ed, Aesthetics, Method and Epistemology, The New Press, New York, USA 1998, pp.205-222

[28]Berleant, A., Aesthetic Sensibility. Ambiances - International Journal of Sensory Environment, Architecture and Urban Space, 2015