

Deep Learning as heuristic approach for architectural concept generation

Daniel Bolojan; Emmanouil Vermisso

Florida Atlantic University
School of Architecture
111 East Las Olas Boulevard
Fort Lauderdale, FL 33308
dbolojan@fau.edu; evermiss@fau.edu

Abstract

This work discusses the use of machine learning for extending the boundaries of architectural concepts. The use of deep learning algorithms is exploited for transferring stylistic attributes from one visual (architectural) domain to another (architectural or otherwise). The combination of various semantic features into hybrid results helps locate possibilities for further exploration of one domain. The process is regarded as open-ended, taking advantage of a generative adversarial network's (cycleGAN) capacity to "imagine" new results which may set the base for new morphological or structural architectural paradigms. We are interested to situate the training results within the context of human and computational creativity, embracing the capacity of automated heuristic methods to produce novel outcomes which can augment human decision making in design and architecture.

Creativity and Artificial Intelligence

The project discussed herewith aims to examine the potential of AI methods for architectural design, acknowledging that the latter warrants a substantial degree of creative thinking, that could be enhanced by machine learning processes. In spite of early interest in machine intelligence and automation in the 19th century, the scientific community only became seriously involved in Artificial Intelligence during the late 1940s and '50s, prompted by Alan Turing's work during the Second World War. It seems that the turning point for strong interest from the general public came quite later; most likely, Garry Kasparov's defeat to IBM's Deep Blue algorithm in chess in 1997 signified a rupture in the collective perception of the capabilities of human mind and the limitations of machines. (Kasparov, 2017) Although wide-spread claim of a "Computational Creativity" might have been premature, that match set a tone for what was possible in the following years, considering the loss of a human grandmaster had not yet been anticipated. Machine cognition and chess had preoccupied scientists for a long time: in the last chapter of the original 1948 "Cybernetics", MIT Professor Norbert Wiener contemplated "...whether it is possible to construct a chess-playing machine and whether this sort of ability represents an essen-

tial difference between the potentialities of the machine and the mind. Note that we need not raise the question as to whether it is possible to construct a machine which will play an optimum game in the sense of von Neumann. Not even the best human brain approximates to this." (Wiener, 1965)

Wiener foreshadowed an aspect of AI that would come to the foreground of academic interest much later; a machine that could merely play chess was possible, but to challenge a human player of the highest level, a machine would require some degree of ingenuity - in other words, (computational) creativity. According to experts, Computational Creativity is defined as "...an emerging field of research within AI that focuses on the capacity of machines to both generate and evaluate novel outputs that would, if produced by a human, be considered creative." (Veale, Cardoso, & Pérez, 2019) Our engagement in defining and assessing creativity within the context of machines can help critique our existing preconceptions about creativity: "...the artifacts that are produced also serve as empirical tests of the adequacy of scientific theories of creativity".

By extension, understanding the process by which machines learn to produce novel outcomes may help us analyze our own cognitive processes more fundamentally. The notion of "Creativity" is a considerable topic in 20th century academic inquiry, and a number of insights about the behavioral structure that nurtures it have been offered. Mathematician Henri Poincaré proposed a four-stage model for achieving high level thinking: Preparation (conscious thought); Incubation (unconscious thought); Illumination and Verification. (Boden, 2004) These stages are sequential and necessary to stimulate the mind towards creative problem-solving. Margaret Boden has explained that creativity manifests in primarily three different forms, progressing from the most common to the least achievable: Combinatorial, Exploratory and Transformational Creativity. (Boden, 2004) Boden's work focuses on understanding the limitations of machine "creativity", described as the context "...in which the computer at least appears to be creative to some degree." Through this task, Boden believes that "Computational ideas can help us to understand how human creativity is possible." (Boden, 2004)

Technological Paradigms in Architectural Design Method

The progressive development of architectural design processes which are informed by digital tools has led to the inevitability of incorporating - to some degree - artificial intelligence as the next step in exploring design ideas; starting in the late 1990s, the exponential growth of digital technologies like automated drafting, modeling and fabrication tools (CAD/CAM) changed the architectural arena and rapidly permeated both architectural education and research, as well as practice. It is noteworthy that although the recent themes of state-of-the-art research venues in architectural design computation reflect architects' intent to be critical and integrative in light of new technologies, tools and processes, the results still largely focus on automated fabrication protocols and material research. While this is important and necessary to advance our practice and probe the industry to innovate, it has inadvertently suppressed a more introspective need to reflect on the design process itself, not only the technical underpinnings of design innovation, but also the cognitive ones, and how these can inform each other.

Symptomatically, technology-driven architectural discourse of the past 10 years describes a tendency to increasingly focus on "synthetic" workflows which use a variety of tools, modes and often, expertise. "Integration" seems a common theme and indeed reflects the desire, and often, necessity on our end to collaborate with other disciplines to fulfill research objectives. The emergence of "parametric modeling" software and its adoption in design (post-2010) has not sufficed to address the numerous and complex problems resulting from the possibilities it offered in the first place. Nevertheless, architects were able to start regarding the design steps as a sequence of parameterization, and the building as a system, with several possible solutions within a large design space, similar to evolutionary models in nature, where only the fittest design survives.

Recently, the nexus of simulation-fabrication has been enhanced by simulated material performance and virtual and augmented reality tools (VR/AR) which increase our participation in various stages of the design workflow ("persistent modeling"). Still, the most promising and disruptive factor so far - albeit not yet exploited - remains the introduction of "Data" as an important contribution for design. This has remained latent, explicitly addressed within limited domains of scholarly inquiry.

The introduction of Artificial Intelligence as a theme for recent specialized architectural research venues reinforces our latent intention to examine the possibility for AI-driven processes to inform design thinking. Among the few examples of AI as a generative tool for imagining architecture is the work of media artist Refik Anadol. Anadol curates "hybrid" environments in existing architectural space using media arts and AI to direct "machine hallucinations". His work is enabled by large data sets which have been compiled and sorted using machine learning algorithms. (Anadol, 2017)

The emergence of Machine Learning: Connectionism in AI as precursor to Deep Learning

The process we discuss herewith involves training a deep neural network to recognize image features and "imagine" new outcomes. Deep Learning is an area of AI that evolved from Cybernetics and early Connectionism. The transition from the early paradigm of classical Artificial Intelligence (GOFAI), which was inspired by logic and required hard-coded instructions for any task, to the biologically inspired paradigm of neural networks (Connectionism), took longer than originally expected. The earliest work on neural networks began in the 1940s by Warren McCulloch and Walter Pitts, while the first artificial neural network ("Perceptron") was created by psychologist Frank Rosenblatt in the late 1950s. The use of neural networks started the "connectionist" (biological) paradigm in artificial intelligence. Despite their insight, early A.I. researchers had difficulty predicting the evolution of automated procedures, because these relied on factors like data availability, data accessibility and representational protocols. Furthermore, connectionism was criticized by MIT Professors Marvin Minsky and Seymour Papert in "Perceptrons: An introduction to computational geometry" (1969). Their critique, focusing on "...the lack of adequate basic theories" to support the connectionist approach to machine learning persisted during twenty years; in a revised edition of the book, machine learning expert Léon Bottou pointed the persisting limitations of mathematically describing meaning, even in certain machine learning areas where progress had been made (i.e. computer vision): "*Although we can conceivably prove that a program fulfills the heuristic specification, this does not guarantee that the program can perform the task of interest. In the case of computer vision, scientists have devised many ingenious ways to leverage physical and geometrical insights about the nature of images. However, absent a mathematical definition of what makes a cat look like a cat, the logical gap remains. Almost a proof is no proof.*" (Minsky & Papert, 2017 (1st ed. 1969))

While heuristics are imperfect for problem-solving, there are certain advantages in this approach. The importance of Machine Learning is understood when we consider the difference between the two AI paradigms; traditional hard-coded tools are problematic when the numbers of possibilities increase exponentially (i.e. Chess vs GO). Machine learning enables an algorithm to not do exhaustive search by brute force, but "learn" to recognize probability through pattern identification.

Margaret Boden discusses the advantages of heuristics when using neural networks. Although the networks' architectures are known to their designers, their detailed performance-driven adjustments sometimes remain a "black box". Still, a heuristic pursuit of a problem is useful for pruning the search-tree. (Boden, 2004) According to Professor Michael Dertouzos, late director of the MIT Computer Science Lab, this research strategy does not guarantee an optimal solution but can be useful for beginning to address difficult problems. In his foreword to the book

“What Will Be”, Dertouzos identified a particular tendency in certain cultures, to approach a given problem with a mix of systematic vs. loose thinking. (Dertouzos, 1997) This involves a convoluted way of seeking answers rapidly, then revising errors and refining until it is adequate. This kind of “undisciplined” structure of thinking (in comparison to very systematic work process among i.e. Japanese or Swiss cultures, where most parameters are sought beforehand), can be advantageous for working with Computer Science problems. This is an important observation with regards to building working contexts that can nurture creative thinking and can likely prove beneficial to Architectural Design, because it is an inherently non-linear process.

Deep Networks Architecture

The affinity between unsupervised learning in neural networks and human learning process underlines the importance of contextualizing unsupervised learning that could help simulate creative thinking in domains like Design and Architecture. Consultation across disciplines and dissemination of such works among the fields of Computer Science, Neuroscience, and Architecture points to integrative schemas of architectural research. In response to recent demand from the architectural community for generative design inquiry, we tested the use of cycleGAN during a 2-day workshop on machine learning for architectural design. The workshop - titled Gaudi’s Hallucinations - examined a synthetic design process which reviewed aspects of architect Antoni Gaudí’s work in light of other mathematical filters (i.e. mathematical inversion) and subsequently extracted large sets of data (2,000-3,000 images) to train a Generative Adversarial Network (GAN) for generating new visuals referencing the Sagrada Familia. This process capitalizes on one of several types of deep neural networks of the GAN family, first introduced by Ian Goodfellow in 2014, featuring competing networks for training generative models by unsupervised learning. (Goodfellow, et al., 2014) GANs, as described by Yann LeCunn, are “the most interesting idea in the last 10 years in Machine Learning”.

A typical GAN is defined by two independent deep neural networks that compete against each other: a discriminator network (D) and a generator network (G). The two networks compete against each other by engaging in an adversarial learning process of generating fake samples that are incrementally more realistic (generator) and a process of classifying fake samples from real samples (discriminator). While the generator network tries to predict features given a certain category, the discriminator network tries to predict a category given the feature of an instance of data.

The deep structure of such generator and discriminator networks allows more efficient learning process by adding more layers, where semantic learning is distributed. A deep learning network uses its hidden layer architecture, to learn incrementally categories that characterize low-level features like a mullion, then it learns gradually higher-level features like a frame and afterwards even higher features like a window. This is a major advantage of using deep

learning networks, compared with other machine learning techniques that require domain expertise and intense feature extraction. Every node and neuron in a deep learning network describes one aspect of the image being learned and all nodes and neurons together provide the overall semantic representation of the image. The network’s nodes are adjusted using backpropagation, where after each forward-pass through the network, the backpropagation performs a backward-pass while adjusting the model’s weights. Backpropagation fine-tunes the weights of the network based on a loss function resulting from the previous training epoch. Lower loss function levels are ensured by properly tuning the weights of the network, which in turn improves model generalization.

In the case of GANs the optimization process is not aimed at finding a minimum (i.e. minimizing the loss function), but an equilibrium between the competing networks. The process of training two competing networks, (G) and (D) simultaneously, is inherently unstable, as the improvement of one network comes at the expense of the other network: “Training GANs consists in finding a Nash equilibrium to a two-player non-cooperative game. Each player wishes to minimize its own cost function [...] unfortunately, finding Nash equilibria is a very difficult problem. Algorithms exist for specialized cases, but we are not aware of any that are feasible to apply to the GAN game, where the cost functions are non-convex, the parameters are continuous, and the parameter space is extremely high-dimensional” (Goodfellow, et al., 2016)

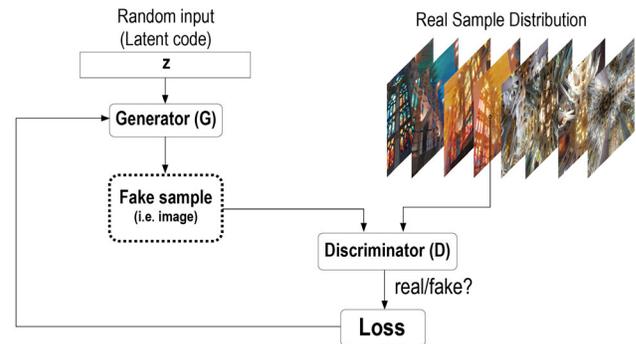


Figure 1. Typical architecture of a GAN

A statistical approach for architectural concept generation

The process tried to look into people’s own perception of what the building symbolizes in the collective unconscious or individually (a forest of columns or the “Gothic” style - in spite of Gaudi’s own regard of the building as Classical), and map these specific impressions on the built space visuals to generate versions of a mutated space based on these alternative realities. The Sagrada Familia was selected as case-study because the building demonstrates substantial spatial complexity, which manifests in a blurring of clear tectonic distinction of parts (Figure 2). For example, the typical architectonic paradigm of post and lintel, columns supporting beams, which in turn support slabs, gives

way to branching columns which decrease in size, reaching out to support multiple combined vaults which compose the ceiling. As a result of the surface connection complexity, the semantic recognition of features by a neural network, is harder, especially in an unsupervised learning environment, where labeled data is not provided. Due to the complexity of the first domain we want to take advantage of robust process like machine learning which can examine thousands of samples. At the same time, the use of statistical method for visual results is no typical to designing, and makes outcome evaluation harder, as we discuss later.



Figure 2. Exterior-Interior views of the Sagrada Familia, showing the complexity of surface topology forming the window-ceiling connections (image credit: E.Vermisso).

Various interfaces were used for the respective stages of data preparation and training. Two original data-sets of Sagrada Familia interior perspectives were extracted and augmented from video files, to serve as primary data. These were trained separately with four data collections, each one combined with one of two Sagrada Familia training data-sets (SF1,SF2), to identify the response of the algorithm to image groups of different semantic expression (i.e. lines, specific geometric patterns, etc.), and therefore, their respective suitability for future training. Two data-sets contained about 3,000 unfiltered samples each from real life (Gothic cathedral images; Forest images), while the other two were generated using 3D modeling software.

1. Sagrada Familia image collection “SF1” (A) is trained with “Gothic cathedral” (B) images.
2. Sagrada Familia image collection “SF1” (A) is trained with “Forest” (B) images.
3. Sagrada Familia image collection “SF2” (A) is trained with mathematical Inversion Line (B) drawings.
4. Sagrada Familia image collection “SF2” (A) is trained with 3D surfaces (B) from Boolean subtraction.

Pytorch deep learning libraries were used within the Anaconda environment to augment the datasets and perform the training. For a better performing model, augmentation methods were applied by creating new, synthetic but reasonable examples from the initial input domains A and B. The size of image samples in the dataset was 512x512 pixels. We divided the dataset into 85% training and 15% testing sets for each A and B domains, using 2550 image samples for the training and 450 image samples for the testing.

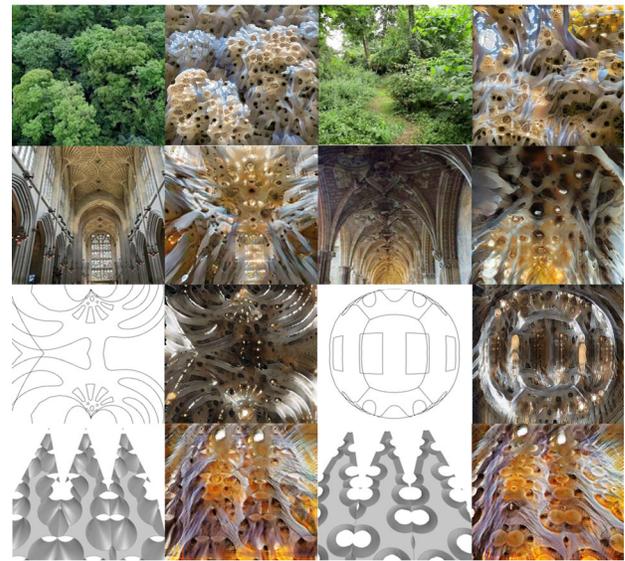
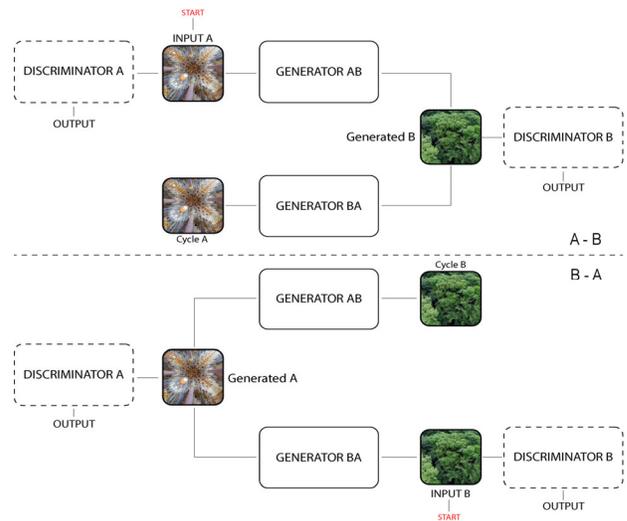


Figure 3. Typical cycleGAN architecture showing the cycle consistency approach: the network has to prove it is not only capable of learning the domain translation from A to B but also from B to A; examples from four data-sets paired with respective examples from the Sagrada Familia data-sets.

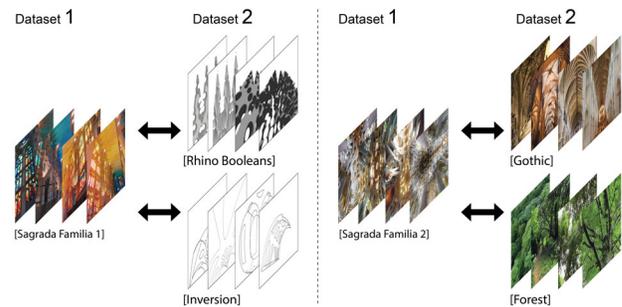


Figure 4. Design workflow showing the software environments used to generate, process and feed data to the neural networks for training; two original Sagrada Familia data-sets were used for the separate training experiments (4 experiments total).

Results: Objective vs. Subjective Evaluation

What benchmark should we use for selection of “successful” results? Two criteria we can consider are how well the process of domain transfer is applied, and the visual quality of results relative to our specific problem search. The former is an objective factor while the latter is subjective. Can these factors be cumulatively considered? We should keep in mind that the best domain translation does not also guarantee the optimum solution for a design problem. Those outcomes which demonstrate a “good” workflow, like adequate transfer of desired features from domain A to B may not be the ones we favor in terms of other, more subjective requirements (i.e. aesthetics, novel configurations etc).

We isolated a few samples which exemplify the results’ diversity well (Fig.5). The “real A” and “fake A” samples are shown next to each other, for three domains: Gothic to SagradaFamilia, SagradaFamilia to Forest and Inversion to SagradaFamilia. Using metrics like “Structural Similarity Index” (SSIM) we can obtain some understanding of the results’ accuracy. However, these results seem to deviate substantially from human assessment: “While it is nearly effortless for humans to quickly assess the perceptual similarity between two images, the underlying processes are thought to be quite complex. Despite this, the most widely used perceptual metrics today, such as PSNR and SSIM, are simple, shallow functions, and fail to account for many nuances of human perception.” (Zhang, Isola, Efros, Shechtman, & Wang, 2018) Another metric used to evaluate the results uses VVG networks. We have annotated the “perceptual similarity” factors of our chosen samples, following those measured using SSIM (Fig.5). Can we reconcile objective and subjective means of assessment? At this point -at least for architecture- the use of automated evaluation should be complemented by human assessment: “One intuitive metric of performance can be obtained by having human annotators judge the visual quality of samples” (Goodfellow, et al., 2016). At the end of figure 5 we have manually annotated three samples we find interesting with respect to possible spatial transformations, indicating the directionality of the domain application based on the kind of features we feel were prioritized by the network.

The first example (epoch 15 of Gothic to SF) indicates some learning is already taking place and (G) is trying to apply semantic features from the real distribution of domain B (Gothic), like the perspectival effect of a Gothic church aisle to output a fake example of a more fluid space. The second example (epoch 33) uses a (real) input with similar perspectival point of view as the previous example, but the result is very different, indicating a bulging, or expansion of the vaulted ceiling, warping the columns around it. The reason may be the difference in the domain A (SF) example. Example 3 (epoch 16 of SF to Forest) is especially intriguing given the early training stage and the sample from the real “B” domain; it is only a patch of green, but the network managed to apply the intricate grass texture well on the vaulted ceiling to create a clear directionality of what looks like structural ribs. Overall, automated assessment clearly differs from human assessment.

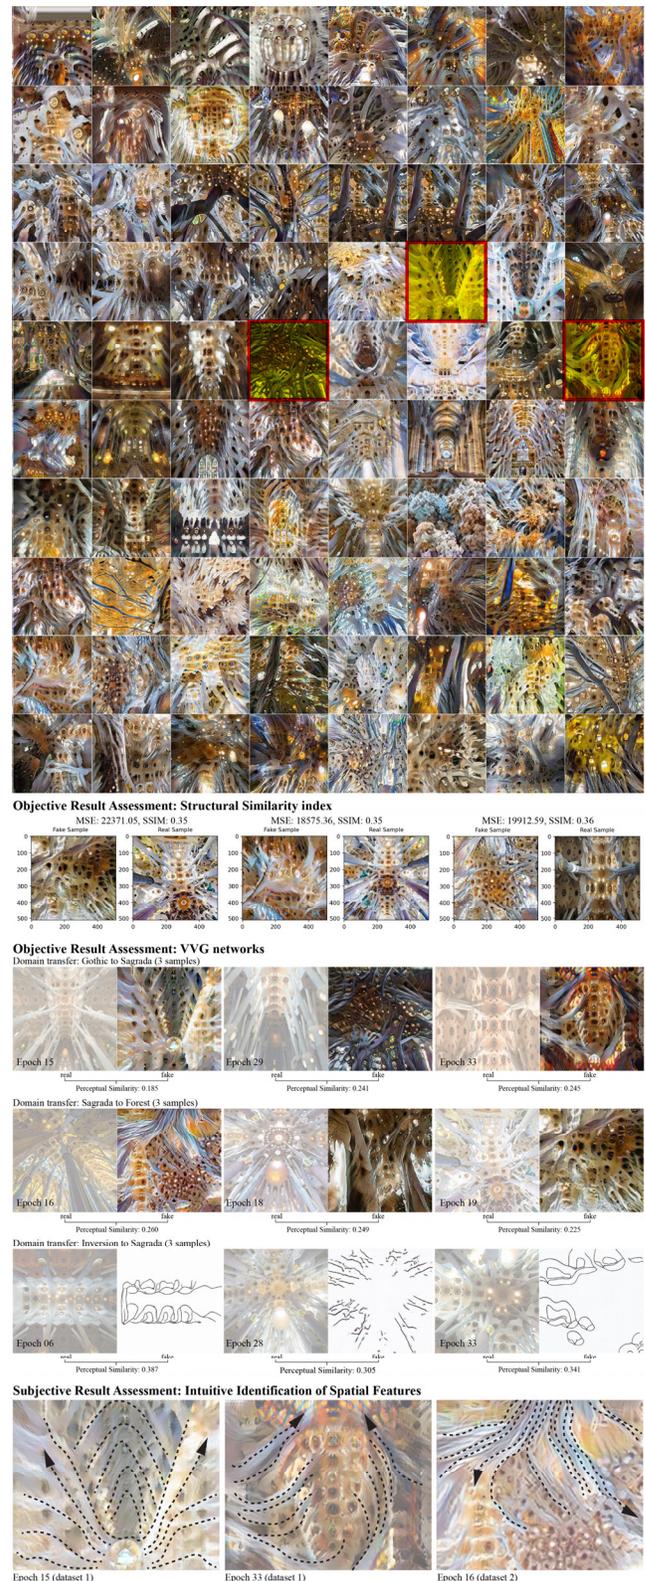


Figure 5. Interior space deformations using domain-transfer with datasets of Gothic, Forest and Inversion images; automated evaluation methods using SSIM & VVG networks (Perceptual Similarity) for assessment in 9 selected pairs; Subjective evaluation.

While a network primarily looks at domain transfer without knowing the significance of specific semantic features, it may assign a low (good) value to a fake output if the transfer of a high-res texture or pattern is for example, successful, even if the image content seems it cannot stand up (may be rotated, or tilted). A designer would look at more features at the same time, like orientation or proportion to classify a result as successful. As a result, we feel that automated metrics at this point can be a complementary feedback to the more intuitive assessment necessary to filter results and revise the network architecture. Such metrics can include for example “Species Explorer” developed by Andy Lomas and Jon McCormack (2020).

Despite lacking a particular theory to provide a clearer trajectory of the experiment in advance, an intuitive result evaluation indicates that -given adequate training- the network can sample “Forest” examples and apply the kind of “branching” structure found in trees to the existing structure of the SF, rendering an even more organic architectural vocabulary. These imagined design alternatives can help architects expand their morphological/structural design repertoire by pushing the perception of what may be possible, while providing insights on the importance of specific architectonic elements based on the network’s reading.

Improved results using high resolution data

We should note that these experiments are an open-ended attempt to explore unsupervised learning as a heuristic workflow for a new conceptual understanding of already known semantic features of architectural space. The learned features which seem to be preferred and carried over by the neural network in the “hallucinated” output matrix may lead us to narrow training data-set towards “curating” specific results. After the workshop, the networks were further developed and trained to investigate outcomes of higher resolution (Fig.6) including larger data-sets that better represented both domains (3000 samples of 1024x1024 each domain). The improved networks were trained for 240 epochs, using Adam optimizer with learning rate of 0.0002 for the first half of the training, and then reduced linearly to 0.0000 over the remaining iterations. We used a 512x512 input resolution, a batch size of 1 and a pool size of 50 to reduce model oscillation. (Shrivastava, et al., 2016) The model used a LSGAN mode and features a 70x70 PatchGAN discriminator architecture, a fully convolutional neural network looking at a patch of the input sample for which it outputs the probability of the sample being “real” or “fake”. It uses a resNet 9 blocks generator architecture; each layer is followed by an instance normalization and a ReLU layer. Resolutions of 256-512-740 were trained with batch sizes of 2-1-1 and pool size of 50-60-70. Figure 6 indicates improvement in the network’s ability to learn detailed feature representation from the input data and translate the semantic representation of a “Forest” domain to the “SF” domain. A remaining challenge, when working with GANs is the lack of objective metrics for evaluating network performance.



Figure 6. Results from training the cycleGAN network with different resolutions and different sets of hyperparameters. The experiments used the same input datasets (Sagrada-Forest) but have different input resolutions, leading to larger and deeper networks which allow for a higher amount of semantic information to be learned. The three testing results use the same image sample input highlighting the network performance requirements.

Creative design potential in AI training

It is interesting that the new samples created by the “generator” (G) network are labelled as “fake”; to us they are valuable, because they did not exist prior to the training procedure, and some of these manifest quite original, not typically considered configurations. Given the short timeframe of the initial experiment (2 days) we feel that an iterative process of training and revising the data-sets to achieve high quality training examples may increase both the fidelity of the results (resolution, appropriate semantic selection) as well as their potential for novelty, according to Boden’s criteria for assessing creative achievement: Novelty, Surprise, Value and Acceptability. (Boden 2004) Developing this work further, we can identify those examples that seem more promising with regard to a particular aspect of space, like the conception of a new structural system, or merely regarding morphological possibilities for surface manipulation. In this respect, these first outcomes, as well as their successors constitute a kind of creative outcome of the network. Still, without a particular goal in mind at the start of the process, it is hard to predict the outcomes, especially during later epochs, when results become more refined. Consequently, the value of the new images lies in their nature as outcomes of a “heuristic” process which has to be tested before assuming a meaning and enabling explanation. This type of “exploratory” process can potentially lead to innovation: “Exploration is the start of non-combinatorial creativity. Indeed, if the style of thought is an interesting one...then even just exploring it will lead to many novelties and may reasonably be regarded as ‘creative’“. (Boden, 2004) In the absence of a clear research question, a process of simulating “play” (open-ended investigation, Fig.7) can often engender high probability for creative output. Boden’s notion of P and H Creativity (psychological, historical) is not clearly applied anymore, as new

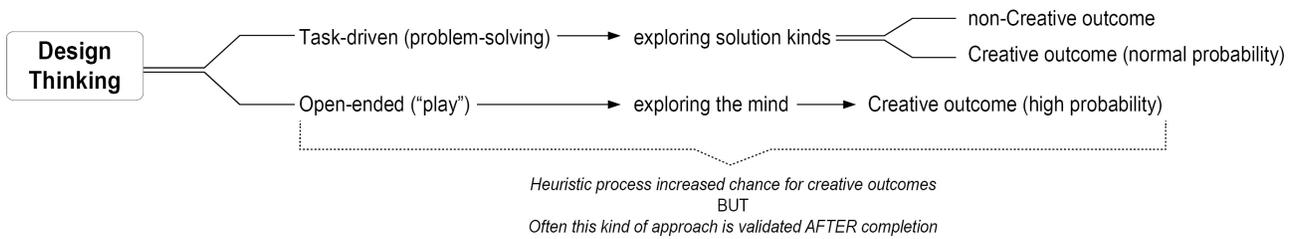


Figure 7. Heuristics usefulness for novelty, according to Boden

outcomes depend on an artificial neural network; while a designer may have encountered prior similar results, the network has not. With regards to any discussion on human-machine combined creativity, these relationships should probably be considered.

Within the Design domain, “...Newell and Simon (1972) describe a general theory of human problem solving”, where “...humans address problems by searching in problem spaces, where a problem space is defined by the goal and the domain knowledge, in the form of operators that enable the search...applied to design, the goals are part of the design requirements that express design variables and the ranges of values they can take. Some design theorists have characterized design creativity in terms of problem-space search: if the design variables and the ranges of values they can take remain fixed during design problem solving, the design is routine; if the design variables remain fixed but the ranges of values change, the design is innovative; and if the design variables and the range of values both change, the design is creative.” (Goel, 2019) As previously mentioned, in a heuristic search for design solutions (like this) questions are somewhat open. By correlation to Newell and Simon, our search space is very large because the task is loosely defined (task = domain-transfer); we might say that the variables here are the semantic features which the network may recognize, and the values they assume are the weights of the particular layer in the network which range from -1 to 1. As we are not yet aware of the particular internal interactions of the network we cannot make an educated guess on the weight adjustment, however we know that the variables/features are not fixed, but depend on factors like the “quality” of the training examples (resolution, diversity, lack of bias). This structure of some degree of uncertainty, as defined in the Newell-Simon theory aligns with Boden’s claim of promise in heuristic processes, which we have here undertaken.

Application Limitations: Extrapolation from Two to Three dimensions

This experiment has assessed a domain translation process using GANs and identified relevant constraints. Training data is a critical factor which largely affects the network’s performance and the anticipated outcomes. The data used currently relies on recognition of 3-dimensional features from 2-dimensional images by an algorithm which does not understand it is looking at two-dimensional data. The network performance is also affected by how well the two

domains are represented throughout the dataset. This can be observed in cases where a network learns a fair share of semantic characteristics and is able to perform successful domain translations from domain A to domain B but fails to perform translations from B to A. In this case, one of the domains is not represented well through the image samples, leading to unsuccessful translation in one direction. To ensure the trained model can generalize well, various augmentation methods were applied. (Hastie, Tibshirani, & Friedman, 2016) We observed that the network’s performance improves greatly if the dataset samples have some features consistent (i.e. background). Trained networks generalize well when the new input samples are reasonably different examples of the initial domains.

Can machines imagine space yet? The possibility for manipulating the third dimension to successfully generate appropriate training data remains a critical component in using AI for architecture. Deep Mind recently used a “GQN” algorithm (Generative Query Network, Eslami, et al., 2018), to solve the so called “Inverse Graphics Problem”. The network was fed a limited number of two-dimensional images of a three-dimensional space and was able to generate a new set of images of the space from a different viewpoint. Potentially, this can be used to generate 3D space from a handful of 2-dimensional inputs. While this seems promising for Architectural Design, the current output resolution is limited to simple spatial layouts (i.e. 3D shapes inside a room) because AI systems don’t yet fully understand the semantics of a scene (separation, background, and foreground) to model complex configurations. In order to use this as a generative tool for “imagining” new architectural space (i.e. drawing one elevation of a building and asking the algorithm to generate another), we need to refine the network’s architecture.

Conclusion

Notwithstanding the evident practical applications of statistical learning in medicine, economics etc., consideration of corresponding implications for Design Thinking is critical. While machine learning re-emerged as a growing area of Artificial Intelligence during the last 15 years, it remains nascent within the architecture community. It is necessary to clarify the contribution of AI in the design disciplines, and assess the requirement for human input within collaborative human-AI workflows. The execution of this project in an architectural context poses some challenges in terms of technical skills; however, this kind of workshops can operate well, given more time. This can improve the investigation by providing the opportunity to discuss the particu-

lar design intent of the training, so the data is explicitly curated by the participants.

Furthermore, evaluation of AI-generated results is a challenging task and can happen in several ways; from a subjective perspective, any results of legible resolution which encompass an adequate amount of semantic recognition from the network can serve as guidelines to steer conceptual development of new surface morphologies or structural iterations based on the imagined alternatives of Gaudí's work offered by the network. Technically, a challenge remains, of transferring diverse complex two-dimensional visuals into three-dimensional models. In certain cases, this may be hard; a secondary assessment is necessary to understand how to categorize and control learning of particular features, to output legible design possibilities.

From an automated process standpoint, adequate objective metrics also need to be incorporated, as current methods like SSIM and Perceptual Similarity factor are not very effective in sorting out what designers might consider "optimal" solutions. As we were interested in augmenting design ability at early conceptual stage, the objective metrics seem secondary -but not entirely unnecessary- at this point. For future development, we can optimize both objective and subjective methods of evaluation. In relation to our own design intuition, we could ask a number of designers to select the "best" or most interesting results, cross-reference their preference and assess the diversity of the subjective evaluation. Then we can use additional tools for automating personal aesthetic judgment, like "Species Explorer". This provides an interface to categorize computationally generated outcomes and identify evolutionary parameters which tend to give results similar to the designers' programmed preference. (McCormack & Lomas, 2020)

From an architectural point of view, considering the current limitations of machine learning algorithms and the human bias embedded in their programming, it is fundamental to regard these outcomes as complementary, rather than antagonistic, to human thinking. By extension, we may regard the role of machine intelligence as positively disruptive, as early cyberneticists imagined it could be; Norbert Wiener's chess-machine discussion mentions characteristically that "*...it is unquestionably possible to construct a machine that will play chess in the sense of following the rules of the game, irrespective of the merit of the play. This is essentially no more difficult than the construction of a system of interlocking signals for a railway signal tower. The real problem is intermediate: to construct a machine which shall offer interesting opposition to a player at some one of the many levels at which human chess players find themselves.*" (Wiener, 1965) Garry Kasparov identified the same limitation in early chess software, which demonstrated substantial fluctuation in skill, incorporating brilliant moves and serious errors in the same game. (Kasparov, 2017) Both Wiener and Kasparov express a humanistic perspective on human-machine interaction, identifying the assistive role of AI within the broader scope of human endeavor. People and Learning machines can reciprocally improve through interaction. Computa-

tional creativity - even if this is not identical to human creativity - can help us become better designers and thinkers, augmenting our cognition through automated feedback.

Acknowledgements

The authors thank the "Gaudí's Hallucinations" workshop participants (eCAADe/Sigradi 2019) for contributing to the data-set generation: Alexander Grasser, Begum Hamzaoglu, Daiki Kido, Samim Mehdizadeh, Hongmei Wu.

References

- Anadol, R., 2017. Archive Dreaming. [Online] Available at: <http://refikanadol.com/works/archive-dreaming/>
- Boden, M., 2004. The Creative Mind: Myths and Mechanisms. Oxon: Routledge.
- Dertouzos, M., 1997. What Will Be: How the World of Information Will Change (How the New World of Information Will Change Our Lives). New York: HarperOne.
- Eslami, S. M. A. et al., 2018. Neural scene representation and rendering. Science, Vol. 360, Issue 6394, pp. 1204-1210.
- Goel, A., 2019. Computational Design, Analogy and Creativity. In: The Philosophy and Engineering of Autonomously Creative Systems. Cham: Springer Nature Switzerland, pp. 141-158.
- Goodfellow, I. J. et al., 2014. Generative Adversarial Nets. [Online] Available at: <https://arxiv.org/pdf/1406.2661.pdf>
- Goodfellow, I. J. et al., 2016. Improved Techniques for Training GANs. [Online] Available at: <https://arxiv.org/pdf/1606.03498.pdf>
- Hastie, T., Tibshirani, R. & Friedman, J., 2016. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Second Edition ed. New York: Springer Science & Business Media.
- Kasparov, G., 2017. Human Plus Machine. In: Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins. New York: Public Affairs, pp. 221-248.
- McCormack, J. & Lomas, A., 2020. Understanding Aesthetic Evaluation using DeepLearning. [Online] Available at: <https://arxiv.org/pdf/2004.06874.pdf>
- Minsky, M. & Papert, S., 2017 (1st ed. 1969). Perceptrons: An Introduction to Computational Geometry. Cambridge, MA: MIT Press.
- Ritchie, G., 2019. The Evaluation of Creative Systems. In: The Philosophy and Engineering of Autonomously Creative Systems. Cham: Springer Nature Switzerland, pp. 159-194.
- Shrivastava, A. et al., 2016. Learning from Simulated and Unsupervised Images through Adversarial Training. [Online] Available at: <https://arxiv.org/pdf/1612.07828.pdf>
- Veale, T., Cardoso, F. A. & Pérez, R. P. y., 2019. Systematizing Creativity: A Computational View. In: The Philosophy and Engineering of Autonomously Creative Systems. Cham: Springer Nature Switzerland, pp. 1-19.
- Wiener, N., 1965. Cybernetics or Control and Communication in the Animal and the Machine. Cambridge, MA: MIT Press.
- Zhang, R. et al., 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. [Online] Available at: <https://arxiv.org/pdf/1801.03924.pdf>