AESTHETICS AS A CRITERION: NAVIGATING SOLUTION SPACES UTILIZING COMPUTER VISION, THE AESTHETIC MEASURE, AND ARTIFICIAL NEURAL NETWORKS

Victor Sardenberg
Mirco Becker

Leibniz Universität Hannover
Welfengarten 1, 30167 Hannover, GERMANY
{sardenberg, becker}@iat.uni-hannover.de

ABSTRACT

A computational framework is proposed to quantify the aesthetic experience of perceiving architecture. This framework is used as the criterion for navigating a design solution space. A parametric model's solution space is represented as a set of computer-generated images. We utilize computer vision to recognize parts and the relations between parts and the whole. These values are inserted into an adaptation of Birkhoff's aesthetic measure formula, calibrated via crowdsourced hedonic response. An artificial neural network (ANN) model is trained to predict the hedonic response of images using (A) parts relations as inputs and (B) the average hedonic responses as target outputs. The prediction of the ANN is used as a fitness function for optimization. The same public evaluates the parametric model's output to compare the ANN model's effectiveness in predicting their response. The method is useful in translating formal qualities into quantities and navigating solution spaces, especially with surrogate models.

Keywords: Computational Aesthetics, Artificial Neural Networks, Computer Vision.

1 INTRODUCTION

This research is put in a context where generating thousands of design options became fast and cheap, primarily through parametric modeling where an algorithm automatically produces multiple design options varying its parameters, resulting in a collection of all possible designs named solution space. A solution space, in general, can be *"understood to be a set of all feasible solution elements, which can be used for the development of the requested solutions"* (Vajna, Kittel, and Bercsey 2010).

The current set of methods for navigating a solution space efficiently always requires numeric values. This important characteristic of optimization algorithms has privileged criteria closer to other disciplines, like structural, financial, and environmental engineering, than architectural design. For example, finite element analysis has quantified structural behavior; radiation analysis has quantified environmental performance, and isovists and floor meterage have quantified the market value of properties.

The premise of this research is that aesthetics - the sensorial perception of artifacts - and its evaluation could enrich solution space navigation. Aesthetics has been incorporated as a criterion through the involvement of the users as part of the optimization process via their manual input responding to multiple design solutions. These previous methods have as drawbacks the slowness and lack of availability of human evaluation during extended periods, besides user response fatigue (Takagi 2001). Our proposed method is similar to Newton's (2018) proposal for developing a qualitative objective function that mathematically represents qualities using a neural network to evaluate qualitative aspects of 3D voxel compositions. However, we focus on automating the quantification of the aesthetics of images as 2D bitmaps, which is closer to how a subject experience visually an object or environment.

There is a tradition of quantifying aesthetics, beginning with G. D. Birkhoff (1933), who proposed the following mathematical formula for the aesthetic measure (M) that correlated the order (O) and complexity (C) of artifacts: Aesthetic Measure = Order / Complexity.

Other precedents are Bense (1965), that proposed the concept of Informational Aesthetics, and Maser (1971), who developed the concepts of Order and Complexity more objectively. At the beginning of our Century, Computational Aesthetics emerged in computer science as the *"research of computational methods that can make applicable aesthetic decisions in a similar fashion as humans can"* (Hoenig 2005).

Our methods apply these architectural precedents, focusing on how parts relate to the whole. There is a tradition of considering this problem essential to aesthetically evaluating buildings. Alberti defined beauty as *"a definite proportional relationship among all parts of a thing, so that nothing can be added, reduced, or changed, without making that thing less deserving of approval"* (Alberti 1991).

We utilize Computer Vision (CV), more specifically Maximally Stable Extremal Regions (MSER), to recognize parts in perspective views (Fig. 2), both photographs and computer-generated imagery, and analyze these parts in isolation, in a diagram of scaled parts (Fig. 3), or their relations, in a diagram of connectivity graph (Fig. 4). We focus on perspective views because it approximates how a subject experience the built environment better than other kinds of drawings like elevations and floor plans. These diagrams are quantified, and the number of parts and their relations are inserted in an adapted version of Birkhoff's formula to calculate an aesthetic measure of each design.

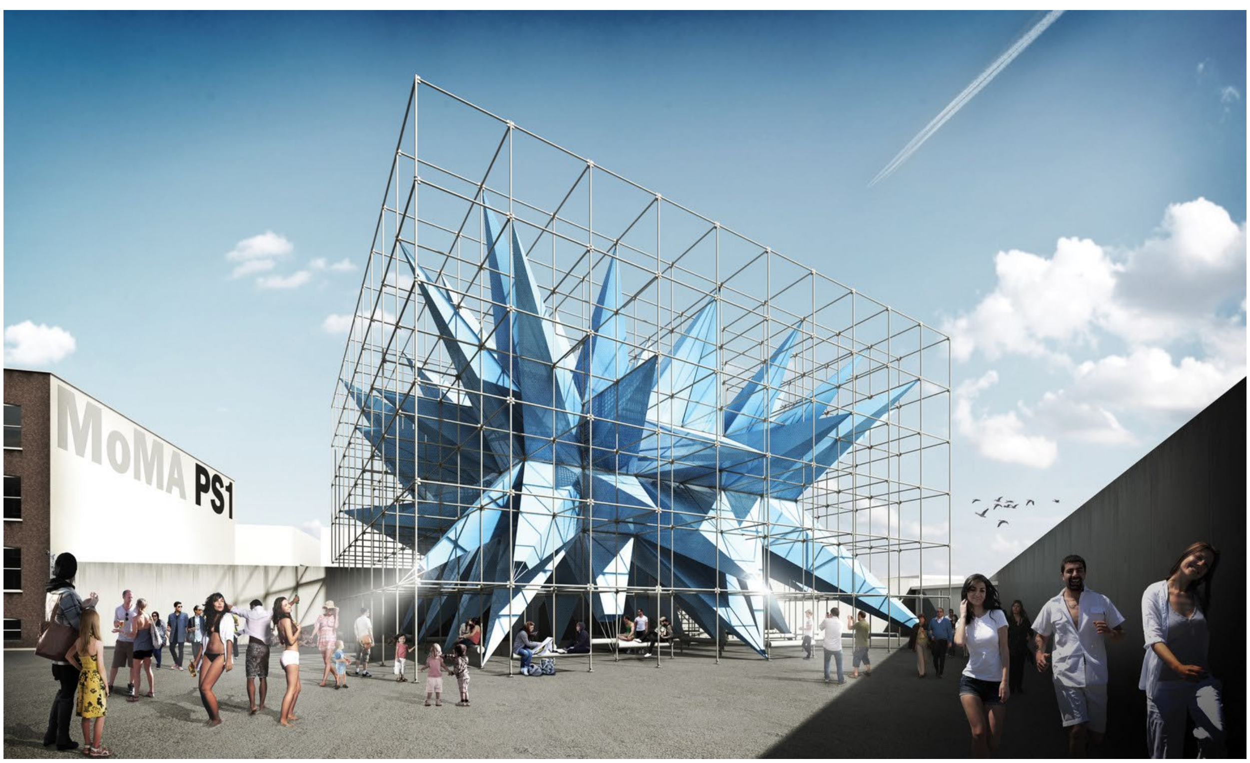

Figure 1: Computer-generated image of Wendy. Author: HWKN Architecture.

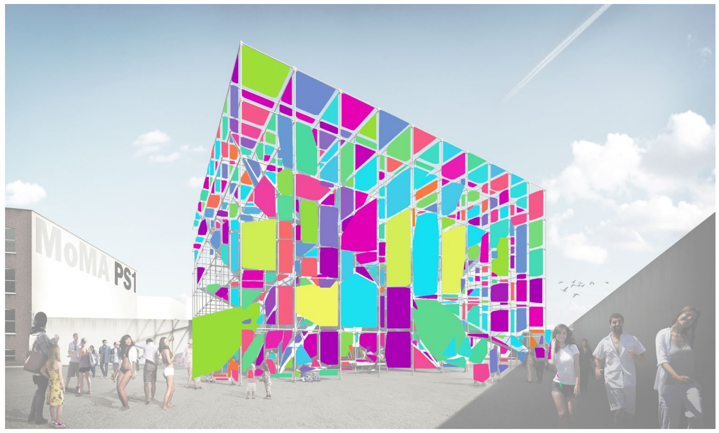

Figure 2: Parts recognized by MSER CV algorithm.

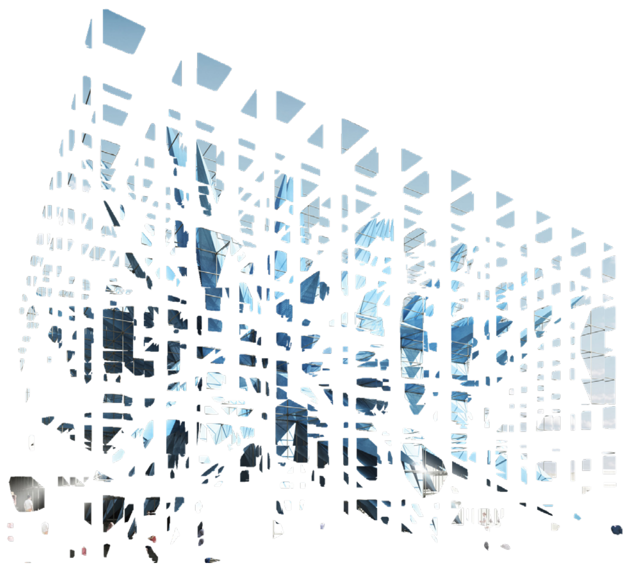

Figure 3: Diagram of scaled parts.

Figure 4: Diagram of connectivity graph.

To avoid (A) a universal aesthetic judgment and/or (B) understanding perception purely in formal relations in objects disconnected from the observer, each audience's aesthetic preferences and biases are incorporated into the framework via the calibration of the adapted Birkhoff's formula and training of the ANN. To predict the preferences of a specific audience, a set of images are scored on social media according to liking and disliking. It is not the focus of this experiment to understand or describe why and why not a particular audience likes or dislikes design solutions but to predict the same average response in a computational framework. This hedonic score is used as target output, and the aesthetic measure, the number of parts, and their relations as inputs to train an ANN to predict this audience's hedonic response to new images.

As a case study, an audience scored all perspective views of all proposals for the MoMA PS1 Young Architects Program, and this dataset was used to train an ANN. A parametric model of one of the built proposals, Wendy, by HWKN in 2012 (Fig. 1), was created to generate a solution space. The genetic algorithm and simulated annealing in Galapagos and the multi-objective optimization algorithm RBFMOpt in Opossum are used to navigate the solution space using as criteria the aesthetic measure and the predicted response from the ANN model. The design solutions are again scored on social media to compare their average hedonic response with the predicted hedonic response from the ANN. Finally, we discuss possible future developments.

1.1 Aesthetic Measure

The mathematician G. D. Birkhoff proposed the first aesthetic formula in his 1933 book *Aesthetic Measure*. In this book, his project is *"to bring the basic formal side of art within the purview of the simple mathematical formula defining aesthetic measure"* (Birkhoff 1933). Each kind of aesthetic object gives rise to aesthetic feelings, which are suis generis. Birkhoff asserts that *"the fundamental problem of aesthetics [is] to determine, within each class of aesthetic objects, those specific attributes upon which the aesthetic value depends"* (Birkhoff 1933).

Birkhoff divides the aesthetic experience into three phases: {1} an initial effort of attention proportional to the complexity (C) of the artifact, {2} an aesthetic feeling of value, or aesthetic measure (M), that is a reward of the initial effort, and {3} the perception that the artifact is characterized by particular order (O) that allows the emergence of the aesthetic feeling. The density of order relations in the aesthetic object determines the basic formula for the aesthetic measure: Aesthetic Measure = Order / Complexity.

For Birkhoff, the precise rules to determine O, C, and M are necessarily empirical because the symbols O and C represent social values. A higher aesthetic measure does not necessarily mean that the effect is more substantial but that the artifact produces more associations concerning the effort to perceive it.

1.2 Computational Aesthetics

The Informational Aesthetics project, in the '60s, attempted to define aesthetics with mathematical rigor in Germany, around Bense and the Stuttgart School, and in France, around Abraham Moles (1968). According to Nake, its *"concepts turned out to be reductionist and schematic, which we argue led to its eventual disappearance, if not failure"* (Nake 2012). At the beginning of the 21^{st} Century, the term Computational Aesthetics surfaced. The first known printed use was in the 1993 paper Computational Esthetics by Remko Scha and Rens Bod (Scha and Bod 1993; Greenfield 2005).

It is possible to map its initial popularization as the product of the "Eurographics workshop on computational aesthetics in graphics, visualization, and imaging" in 2005 in Girona, Spain. The workshop defined computational aesthetics *"as an experimentally based scientific field and not a philosophically based modern version of art"* (Neumann et al. 2006), with divergent goals and methods. The research method proposed was an *"iterative two-part process [which] consists of using artistic computer graphics techniques to enhance the presentation of important data features, then conducting perceptual studies to evaluate the effectiveness of the resulting imagery"* (Neumann et al. 2006).

Florian Hoenig used the opportunity to publish a paper to define the discipline of Computational Aesthetics in the context of computer science, where he reflected and synthesized the contributions and discussions of the previously mentioned event. One similar discipline is empirical aesthetics, with Gustav Theodor Fechner's Vorschule deer Aesthetik as its foundation and whose followers, like David Berlyne (Greenfield 2005), founded the International Association of Empirical Aesthetics in 1965. However, their work is a subbranch of empirical psychology, and *"their main aim is to apply these methods to collect data upon which aesthetic theories can be tested"* (Hoenig 2005). This goal contrasts with computational aesthetics, which Hoenig defines as *"the research of computational methods that can make applicable aesthetic decisions in a similar fashion as humans can."* On the same occasion, Frieder Nake paraphrased Max Bense to define the objective of computational aesthetics as *"to obtain a scalar or vector measurement of the aesthetics of a work of art"* (Greenfield 2005).

1.3 MoMA PS1 YAP and the Wendy Pavilion

This paper demonstrates the application of our computational aesthetics framework to navigate a solution space. MoMA PS1 Young Architects Program proposals were selected as a case study to produce a solution space, focusing more specifically on the Wendy Pavilion by HWKN from 2012. The choice for the YAP was made based on the availability of 87 projects for one specific location. This set of images was previously used to train an Artificial Neural Network model, described in detail by Sardenberg and Becker (2022a). Since the ANN model was trained only in a dataset of YAP projects, it is only reliable to evaluate new projects in the same context. Therefore, a parametric model was created to mimic the Wendy pavilion, varying the number and dimensions of spans of the 3D grid that structures the pavilion, resulting in 6 parameters that generate a solution space: The number of cells and their dimensions in X, Y, and Z.

1.4 Parts to Whole Relations

How parts relate to and produce a whole is a vital architectural problem. Vitruvius, in the first treatise in architecture history from c. 25 BCE, defined: *"[...] Beauty [is] when the appearance of the work is pleasing and in good taste, and when its members are in due proportion according to correct principles of symmetry"* (Vitruvius Pollio 1914).

This aspect received particular attention in Alberti's treatise on Architecture, published in 1452. Alberti defines beauty as *"a definite proportional relationship among all parts of a thing, so that nothing can be added, reduced, or changed, without making that thing less deserving of approval"* (Alberti 1991).

Palladio also wrote regarding it in n similar terms in his treatise from 1570: *"Beauty will derive from a graceful shape and the relationship of the whole to the parts, and of the parts among themselves and to the whole, because buildings must appear to be like complete and well-defined bodies, of which one member matches another and all the members are necessary for what is required"* (Palladio 2002).

Our current understanding of how humans perceive objects is based on the idea of the analysis of their parts. Two well-established divergent theories exist: Feature Integration Theory and Recognition-by-Components theory (Goldstein 2011). Feature Integration Theory argues that we recognize objects in two stages: the preattentive stage, where we realize features of its parts like lines, curves, and colors. These features are combined in the second stage, where we perceive an object.

In the Recognition-by-Components theory, features are three-dimensional volumes known as Geons. Our cognitive system can understand these volumetric parts and combine them to perceive objects and guess what they could be based on their context, previous knowledge, and features of their parts.

2 METHODS

2.1 Computer Vision – MSER

Our method employs Maximally stable extremal regions (MSER) to recognize regions in images that coincide with parts of buildings. It converts multi-tone grayscale images into black or white binary pixels using multiple thresholds from entirely black to completely white (Matas et al. 2004). Therefore, it recognizes regions, and when they are consistent among various thresholds, they are outputted. We calibrated such a system to coincide with architectural elements in photos. We utilize the recognized parts to insert them into two diagrams: a Diagram of scaled parts and a diagram of connectivity graph. Sardenberg and Becker (2022b) describe these diagrams in more detail.

2.2 Diagram of Scaled Parts and Connectivity Graph

The diagram of scaled parts is intended to analyze parts in isolation (Fig. 3). It is the result of scaling all parts by half using each centroid as the scaling reference point. It is possible here to perceive each part's size and proportion and to quantify the number of parts. The diagram of the connectivity graph is used to analyze how parts relate to each other and the whole (Fig 4). If any part pixel is shared with others, resulting from tangency to full containment, an edge is drawn from each part's centroid. Applying this diagram enables quantifying the number and length of connections.

2.3 Birkhoff's Adapted Formula and its Calibration

We propose to adopt Birkhoff's formula to analyze perspective images of architecture, considering the number of connections between parts and their length as indexes of order and the number of parts as an index of complexity. To normalize across different image resolutions, we divide it by the square root of the number of pixels.

Using MSER with the diagraming and applying this adapted formula can score an aesthetic measure for each image. Considering that each term of the formula may have a different weight on the perception of images, we inquired via social media an audience to score how they liked or disliked each image that was scored by our method. This process is detailed by Sardenberg and Becker (2022a). Comparing the average hedonic response by this specific audience with the aesthetic measure, the accuracy was 33%. We utilized Galapagos genetic algorithm to find the following weight values and reach a 66% accuracy:

$$\text{Aesthetic Measure} = \frac{0.82\ \text{Connection Length Average} \times 0.96\ \text{Number of Connections}}{0.29\ \text{Number of Parts} \times \sqrt{\text{Number of Pixels}}}$$

2.4 Artificial Neural Network Model

We utilize supervised machine learning to train an ANN. The data contains the connection length average, number of connections, number of parts, aesthetic measure, calibrated aesthetic measure as inputs, and the average hedonic response as the target output. The data was split as 75% for the training set and 25% for the testing set, and the model was trained for 3000 steps. This model has an 85% accuracy rate. For further information regarding the model, refer to Sardenberg and Becker (Sardenberg and Becker 2022).

2.5 Solution Space Navigation: Genetic Algorithm, Annealing Simulation, and Surrogate Model Methods

With parametric modeling, it became fast to generate thousands of alternatives. However, evaluating them may be time-consuming, depending on the criteria chosen. Therefore, it is crucial to navigate the solution space efficiently. This paper will focus on using Grasshopper's plug-ins, Galapagos and Opossum.

Galapagos is McNeel's native plug-in and enables designers to utilize metaheuristics solvers genetic algorithm (GA) - that mimics biological evolution, crossover, and survival - and simulated annealing (SA) - that mimics the slow cooling of metal to decrease the probability of accepting worse solutions (Rutten n.d.). Galapagos is only able to use a single objective.

Opossum uses surrogate models that are a computationally faster approximation of a simulation. A sample of design solutions is simulated and used to train a surrogate model that replaces the original time-consuming simulation (Wortmann and Nannicini 2017). The approximated model is then used to navigate the solution space and point toward optimum solutions that are further simulated. We used RBFOpt for sing-objective optimization and RBFMOpt for multi-objective optimization.

3 EXPERIMENTS AND RESULTS

We propose three aesthetic criteria: Predicted hedonic response (PHR) from the trained ANN, aesthetic measure (AM), and calibrated aesthetic measure (cAM). They are used in isolation or combination as evaluation criteria to navigate a design solution space.

The first set of experiments (Table 1, next page) compared how GA, SA, and SM navigate a solution space of design options considering the PHR as the single criterion. The time consumed to generate and score a design solution is approximately 11 seconds, so it is crucial to develop solutions efficiently.

On the top of the next page, we have a map that displays each solution as a point, the centroid of a polygon defined by vertices in each axis representing a parameter. The coldest colors are the lowest values of the criterion, and the warmest colors are the highest. It is possible to notice that this is a multimodal solution space where multiple regions can produce good solutions.

Below it, we have a graph representing how the normalized input values (vertical axis) change for each iteration (horizontal axis), and further below, a graph depicting the normalized output of our aesthetic measure and calibrated aesthetic measure formulas and the predicted hedonic response (vertical axis) in each iteration (horizontal axis). On the bottom, there is the highest-scoring solution.

The GA distributed the solution space evenly for 1000 iterations until it found a region to explore further and stabilize. The PHR increases sharply at the beginning of the optimization and remains within an extensive range while the solutions evolve. After running for 4552 iterations, it stopped, finding a peak value of 7.67 (iteration 2027) and averaging 2.98.

The SA briefly navigates globally to choose an area to explore for a thousand iterations to optimize locally. Then, it guides globally again to locate another area and repeats the process repeatedly. It is clear to see that this run explored three areas visible in all graphs. In this specific optimization run, the second area explored contained outstanding solutions. Finding several regions may take many areas with thousands of iterations, making it prohibitively time-consuming. In this specific run, the peak solution scored 7.79 in iteration 2157, and the average solution achieved 3.50 after 4169 iterations.

The SM navigated a more extensive solution space and mapped a more prominent area more detailedly, even if it ran for a shorter time, producing 1600 iterations with an optimum solution scoring 7.61 (iteration 1002) and an average of 4.09. Compared to the other methods, it is interesting that it never stabilizes in a specific range of values for the inputs.

Table 1: Comparison between GA, AS, and SM for design optimization maximizing the PHR.

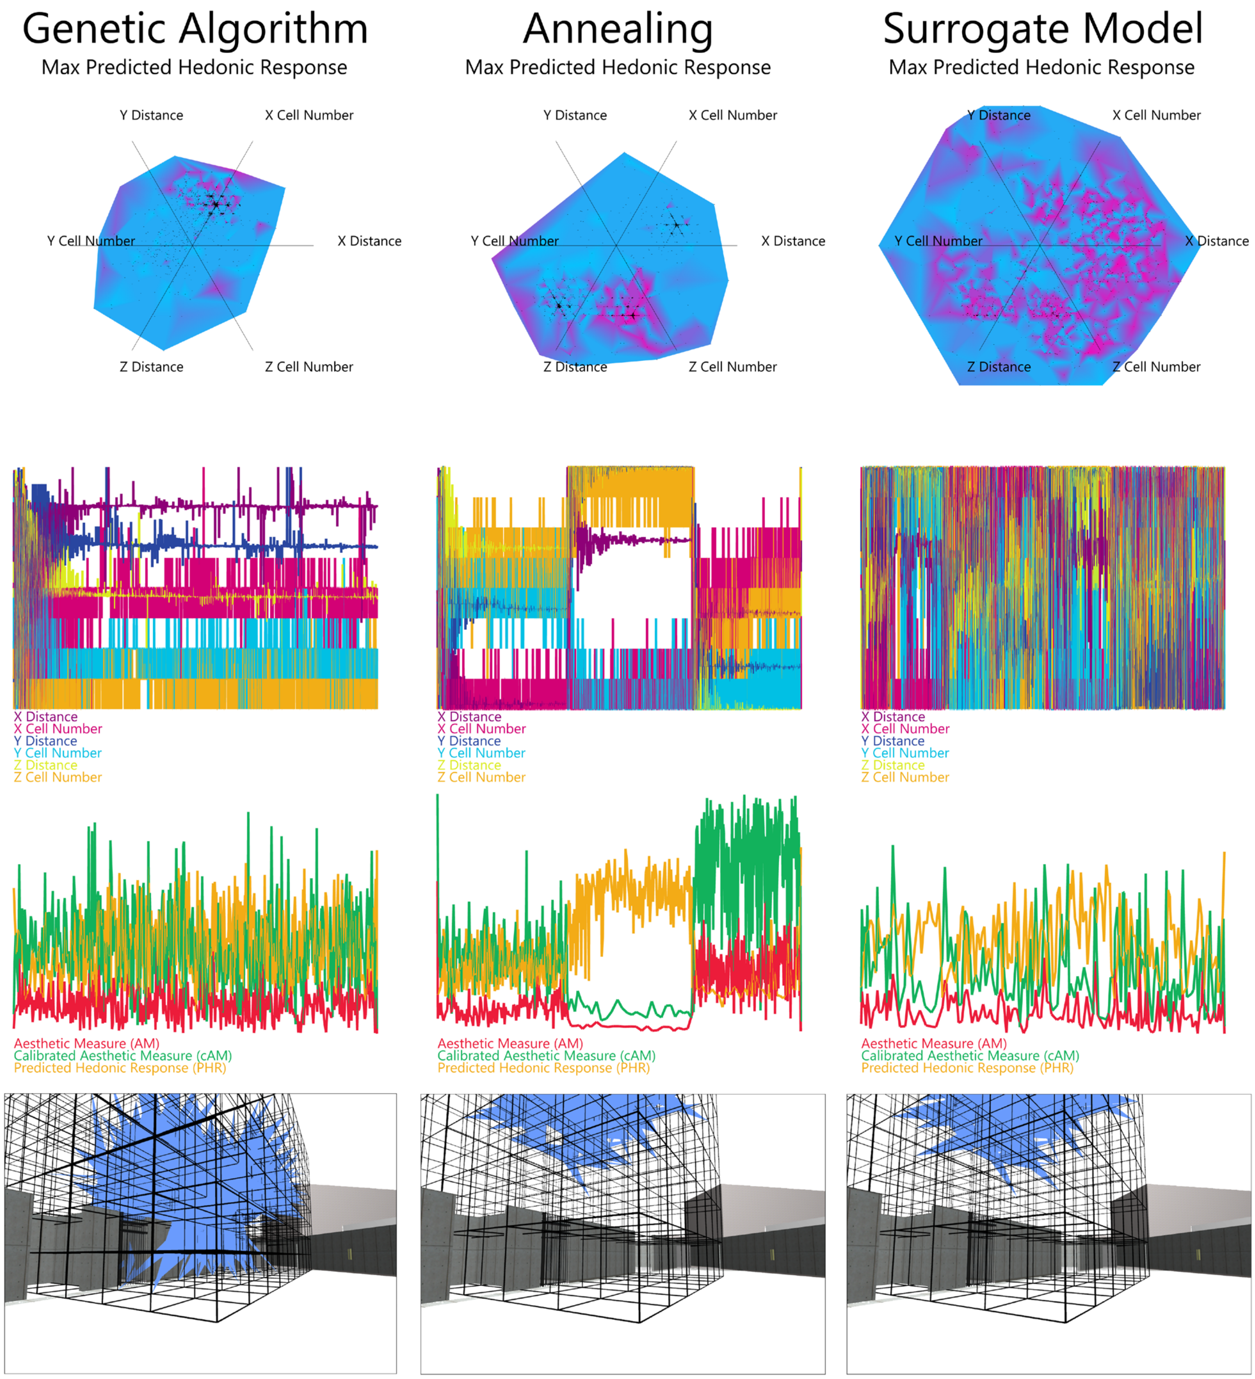

The second batch of design optimization experiments is shown on the next page (Table 2, next page). It includes the following: (A) utilizing the GA to maximize an average of the AM, cAM, and PHR, (B) using the multi-objective SM capabilities of Opossum (RBFMOpt) to maximize AM, cAM, and PHR, and (C) to use GA to optimize only the AM.

Table 2: Comparison between GA and SM maximizing different criteria.

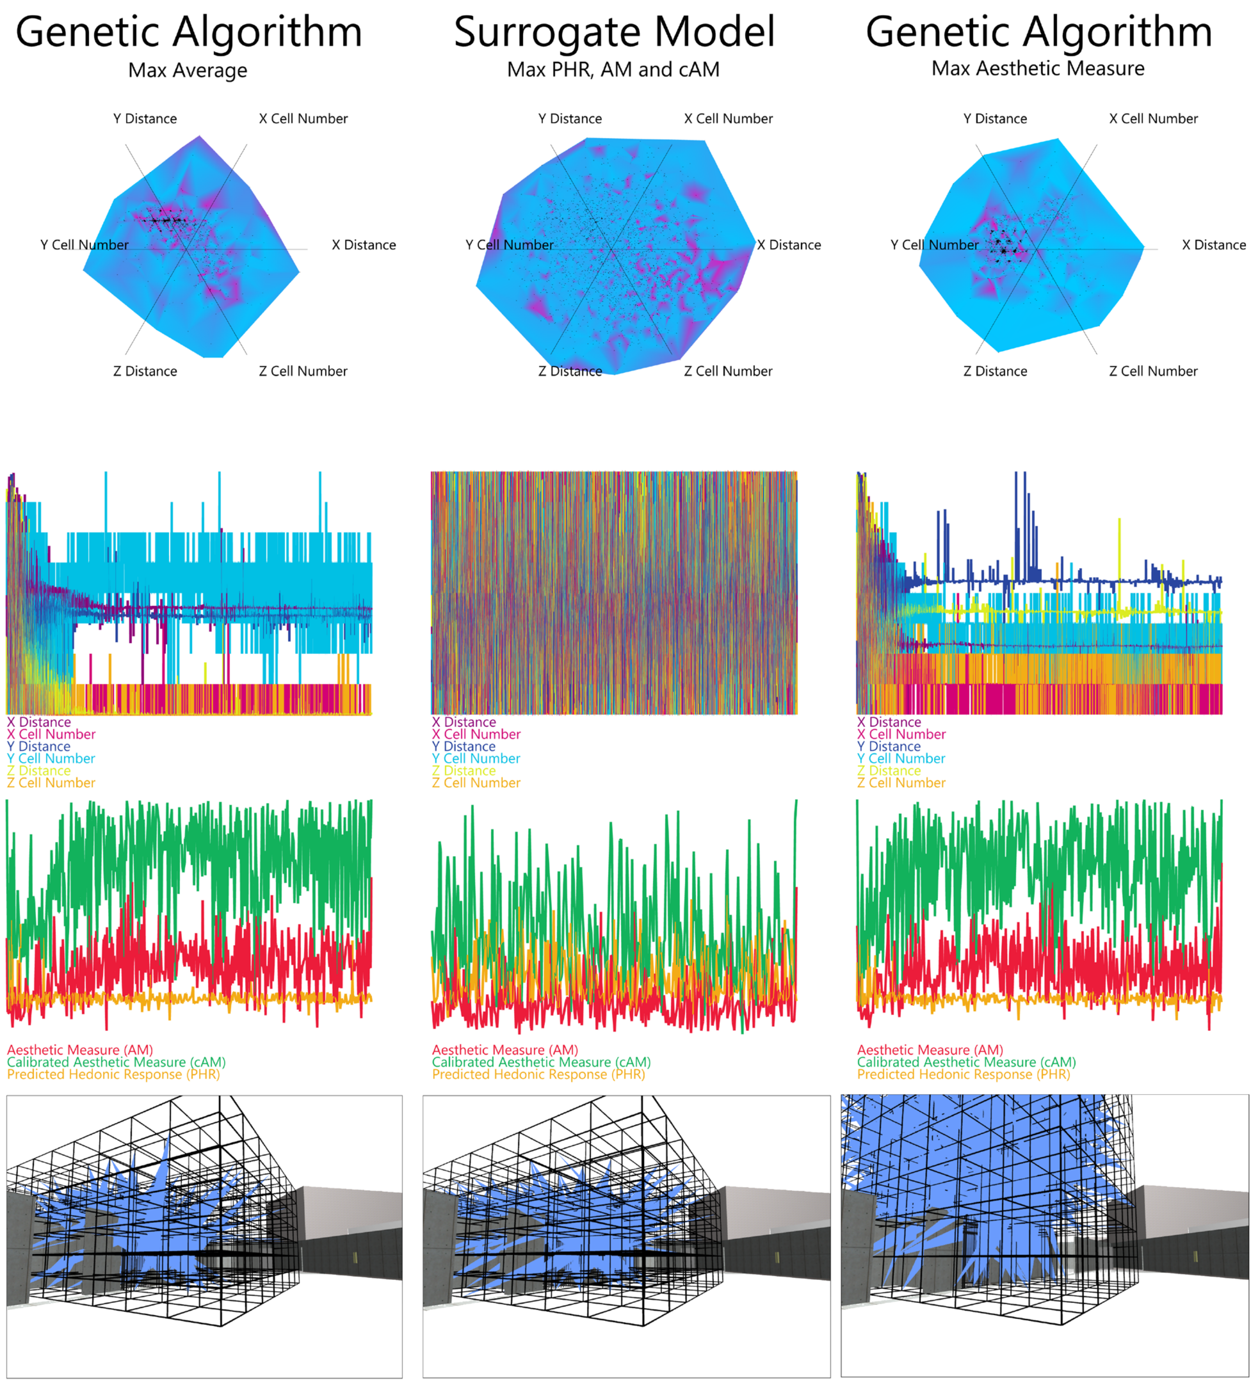

In (A), the GA globally explored the solution space for 1000 iterations when it stabilized in a region where the PHR was low and the AM and the cAM were maximized. Considering all 4852 solutions generated before Galapagos stopped, the maximum PHR was 5.25 (average 1.77), the maximum AM was 0.68 (average 0.30), and the maximum cAM was 1.87 (average 0.82). All values were normalized and included in this evaluation function: $F = (\log AM + \log cAM + \log PHR)/3$

The log of each value was applied to avoid a single criterion being very high to the detriment of other criteria. The peak solution was generated in iteration 2305, scoring 0.68 for the AM, 1.87 for the cAM, and

1.66 for the PHR. It is possible to argue that because the AM and the cAM are correlated, the function privileged its maximization in prejudice of the PHR.

Regarding (B), we took advantage of the Opossum's ability to maximize multiple criteria for AM, cAM, and PHR. In 2807 iterations, the maximum PHR was 7.08 (average 2.09), the top cAM was 1.75 (average 0.38), and the maximum AM was 0.64 (average 0.13). The peak solution ranked according to approximation to the true Pareto front, scored 0.64 for AM, 1.75 for cAM, and 1.7 for the PHS, in a very similar solution to the experiment (A).

In experiment (C), we focused on maximizing the AM to produce a design solution that creates more associations about the effort to perceive it and understand its link to the PHR. After 4553 iterations, a peak solution of 0.74 was generated in iteration 2021. The PHR of this solution was 1.65, and its cAM was 2.01.

3.1 Crowdsourced Evaluation

To achieve the two-part process of using computation to enhance aesthetic qualities and conducting perceptual studies to study its effectiveness, we generated six design variations to create six images that were scored by our ANN and by the same audience as previously. Regarding the design solutions (Table 3), the ANN was close to estimating the average hedonic response from the audience, with an absolute average inaccuracy of 1.30.

It is worth noticing that the highest scoring solution in PHR (7.61) is, in fact, very distant from the average hedonic response from the audience (3.00). However, 4 of 6 predictions were close to the average response, deviating less than 0.50 units.

Table 3: Comparison between the PHR and the average response of different designs.

	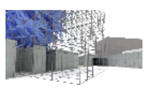	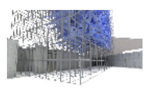	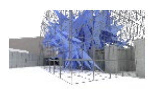	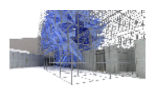	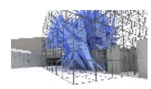	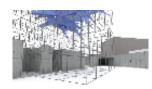	Average
PHR	3.01	4.60	5.11	5.98	7.21	7.61	5.58
Average response	3.50	4.50	7.50	6.00	7.00	3.00	5.25
Absolute difference	0.49	0.10	2.39	0.02	0.21	4.61	1.30
Differ-ence %	15.05%	2.19%	37.90%	0.33%	2.95%	86.89%	27.50%
PHR rank-ing	1	2	3	4	5	6	
Average response ranking	2	3	6	4	5	1	
Ranking difference	1	1	3	0	0	5	1.66

4 CONCLUSION AND DISCUSSION

When considering the deviations between the PHR and the average hedonic response, one could ask about the truthfulness of this ANN model. According to the statistician George Box, *"it would be very remarkable if any system existing in the real world could be exactly represented by any simple model. However, cunningly chosen parsimonious models often do provide remarkably useful approximations"* (Box 1976). In this spirit, the question of truthiness is replaced by a question of usefulness. As shown in our experiments, our methods of PHR, AM, and cAM successfully translated aesthetic qualities of perspective views of an architectural proposal into a numerical score that was used to navigate solution spaces with different algorithms. In this sense, it is useful to apply it to complement other criteria, like structural and environmental performance. The framework can translate aesthetic qualities of perspective views of an architectural proposal into a numerical score.

Although not the focus of this paper, it is possible to argue that if the presented accuracy of the ANN model requires increasing in its accuracy, we could incorporate more models to judge images, like Semantic Segmentation, Object Recognition and counting, color and brightness histograms, and spatial semantics, all currently being done with CV. Recent developments in ML have been pointing to combining various models to increase accuracy (Jumper et al. 2021). All these criteria can be used parallelly in multi-objective optimization (Wortmann 2017).

When comparing the use of GA, SA, and SM, there are a few reasons why we would favor the latter. First, considering that each optimization took around 6 hours to go through a couple of dozen iterations, we would favor a solution space navigation algorithm that could initially rapidly explore many different designs. Opossum's RBFOpt algorithm consistently searched the broadest space and, at the early stages, resulting in high-scoring solutions. Besides, the performance explorer implemented recently is a massive differential in our case. It presents a GUI with an interactive map of the solution space that allows the user to explore high-scoring alternatives compared to Galapagos, which only displays the top-ranked solutions that are mostly identical. Visualizing multiple divergent design variations is valuable in solution space navigation using aesthetics as a criterion.

Finally, we propose as future work to devise an experiment where instead of pointing towards a single optimum solution, we explore the solution space mapping multiple solutions produced by a wide range of combinations of parameters that score high in the PHR and present it to the designer. The goal of the research is to develop a method for gaining a comprehensive understanding of the solution space allowing the designer to make efficient choices and acknowledge the tradeoffs inherent in each design decision.

REFERENCES

Alberti, L. B. 1991. *On the Art of Building in Ten Books*. Translated by Joseph Rykwert. New ed edition. Cambridge, Mass London: The MIT Press.

Bense, M. 1965. *Aesthetica*. Agis-Verlag.

Birkhoff, G. D. 1933. *Aesthetic Measure* Cambridge, MA: Harvard University Press.

Box, G. E. P. 1976. “Science and Statistics.” *Journal of the American Statistical Association* 71 (356): 791–99.

Goldstein, E. B. 2011. *Cognitive Psychology*. 3. ed., International student ed., Special ed. [Belmont, Calif.]: Wadsworth, Cengage Learning.

Greenfield, G. 2005. "On the Origins of the Term Computational Aesthetics". The Eurographics Association.

Hoenig, F. 2005. "Defining Computational Aesthetics". The Eurographics Association.

Jumper, J., R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, et al. 2021. “Highly Accurate Protein Structure Prediction with AlphaFold.” *Nature* 596 (7873): 583–89.

Maser, S. 1971. *Numerische Ästhetik: neue mathematische Verfahren zur quantitativen Beschreibung und Bewertung ästhetischer Zustände*. Krämer.

Matas, J, O. Chum, M. Urban, and T. Pajdla. 2004. “Robust Wide-Baseline Stereo from Maximally Stable Extremal Regions.” *Image and Vision Computing*, British Machine Vision Computing 2002, 22 (10): 761–67.

Moles, A. A. 1968. *Information Theory and Esthetic Perception*. University of Illinois Press.

Nake, F. 2012. “Information Aesthetics: An Heroic Experiment.” *Journal of Mathematics and the Arts* 6 (2–3): 65–75.

Neumann, L., M. Sbert, B. Gooch, and W. Purgathofer. 2006. “Computational Aesthetics 2005 Eurographics Workshop on Computational Aesthetics in Graphics, Visualization and Imaging Girona, Spain, 18–20 May 2005.” *Computer Graphics Forum* 25 (1): 145–46.

Newton, D. 2018. “Multi-Objective Qualitative Optimization (MOQO) in Architectural Design.” In *Kepczynska-Walczak, A, Bialkowski, S (Eds.), Computing for a Better Tomorrow - Proceedings of the 36th ECAADe Conference - Volume 1, Lodz University of Technology, Lodz, Poland, 19-21 September 2018, Pp. 187-196*. CUMINCAD. http://papers.cumincad.org/cgi-bin/works/paper/ecaade2018_323.

Palladio, A. 2002. *The Four Books on Architecture*. Translated by Richard Schofield and Robert Tavernor. Reprint edition. Cambridge, Mass: The MIT Press.

Rutten, D. n.d. “Evolutionary Principles Applied to Problem Solving.” Accessed January 17, 2023. https://www.grasshopper3d.com/profiles/blogs/evolutionary-principles.

Sardenberg, V., and M. Becker. 2022. “Computational Quantitative Aesthetics Evaluation - Evaluating Architectural Images Using Computer Vision, Machine Learning and Social Media.” In *Pak, B, Wurzer, G and Stouffs, R (Eds.), Co-Creating the Future: Inclusion in and through Design - Proceedings of the 40th Conference on Education and Research in Computer Aided Architectural Design in Europe (ECAADe 2022) - Volume 2, Ghent, 13-16 September 2022, Pp. 567–574*. CUMINCAD. http://papers.cumincad.org/cgi-bin/works/paper/ecaade2022_75.

Scha, R., and R. Bod. 1993. “Computationele Esthetica.” *Informatie En Informatiebeleid* 11 (1): 54–63.

Takagi, H. 2001. “Interactive Evolutionary Computation: Fusion of the Capabilities of EC Optimization and Human Evaluation.” *Proceedings of the IEEE* 89 (9): 1275–96.

Vajna, S., K. Kittel, and T. Bercsey. 2010. “Designing the Solution Space for the Autogenetic Design Theory (ADT).” In . https://www.semanticscholar.org/paper/DESIGNING-THE-SOLUTION-SPACE-FOR-THE-AUTOGENETIC-Vajna-Kittel/77b1d79d0f66edfcf308ed7c1dc3fd6ddb046a7c.

Vitruvius Pollio. 1914. *Vitruvius, the Ten Books on Architecture*. Cambridge : Harvard University Press. http://archive.org/details/vitruviustenbook00vitr_0.

Wortmann, T. 2017. “Opossum - Introducing and Evaluating a Model-Based Optimization Tool for Grasshopper.” In *P. Janssen, P. Loh, A. Raonic, M. A. Schnabel (eds.), Protocols, Flows, and Glitches - Proceedings of the 22nd CAADRIA Conference*, Xi'an Jiaotong-Liverpool University, Suzhou, China, 5-8 April 2017, pp. 283-292

Wortmann, T., and G. Nannicini. 2017. “Introduction to Architectural Design Optimization.” In , edited by *Athanasia Karakitsiou, Athanasios Migdalas, Stamatina Th. Rassia, and Panos M. Pardalos, 128:259–78. Springer Optimization and Its Applications.* Cham: Springer International Publishing.

AUTHOR BIOGRAPHIES

VICTOR SARDENBERG is a Ph.D. candidate at the Leibniz Universität Hannover, where he is an associated researcher in the Chair of Digital Methods in Architecture (dMA), and the Universidade Presbiteriana Mackenzie. His thesis is titled "Computational Framework for Quantifying the Architectural Aesthetic Experience." His e-mail address is sardenberg@iat.uni-hannover.de

MIRCO BECKER is a full professor at Leibniz University Hannover, where he holds the Chair for Digital Methods in Architecture (dMA). He has led the computational design on two airport projects in practice and has a keen research interest in miniaturized architectural robots. His e-mail address is becker@iat.uni-hannover.de