Harnessing Game-Inspired Content Creation for Intuitive Generative Design and Optimization

Lorenzo Villaggi[1(✉)], James Stoddart[1], and Adam Gaier[2]

[1] Autodesk Research – AEC Industry Futures, New York, US
lorenzo.villaggi@autodesk.com
[2] Autodesk Research – AI Lab, New York, US

Abstract. A generalizable and example-based model for multi-scale generative design is presented. The model adapts the Wave Function Collapse (WFC) algorithm, a procedural approach popularized in game development, to a quality-diversity (QD) framework, a state-of-the-art multi-solution optimization approach. QD enables the search of high-performing solutions not only against objectives, but along a set of qualitative features -- explicitly ensuring diversity within the solutions. We demonstrate the challenges and opportunities in applying these novel methodologies to AEC-focused problems through a real-world residential complex case study.

Keywords: Procedural modeling · Quality diversity · Generative design

1 Background

1.1 Quality-Diversity (QD)

In AEC, optimization is often used to understand the constraints and possibilities of a design space and discover novel solutions by exploring a diverse set of high-performing alternatives [1]. The most common methodology is multi-objective optimization (MOO) [2, 3]. Though MOO is often used for exploratory analysis, in many ways it is poorly suited for it. MOO requires the definition of objectives to be minimized—design exploration calls for features to be explored.

Quality-Diversity [4] algorithms move beyond MOO to produce sets of high performing designs organized by high-level features better suited to the judgement of domain experts. QD searches explicitly for high performing solutions with varied qualities, such as the perimeter of a building or the number of bedrooms in a unit. In contrast to a pareto curve of non-dominated solutions, the most widely used QD algorithm MAP-Elites [5] produces a grid or 'map' of the solutions – with each axis corresponding to a feature. This map provides an intuitive overview of the performance potential for each region of the feature space. Though originally designed for applications in robotics [6] and artificial life [7], QD techniques have begun to be applied in design applications such as engineering optimization [8, 9] and procedural content generation [10–12].

C. Gengnagel et al. (Eds.): DMS 2022, *Towards Radical Regeneration*, pp. 149–160, 2023.
https://doi.org/10.1007/978-3-031-13249-0_13

1.2 Wave Function Collapse (WFC)

Though generative design is gaining broader adoption in the AEC industry, its impact is limited by the level of technical skill required to operate computational design tools and the challenge of building generalizable applications[13]. We address both issues with a versatile design space model for semi-constrained designed systems, like modular or prefab, compatible with traditional design methods.

Our design space model adapts WFC[14], a texture synthesis approach popular in the game development community. WFC is a constraint-based procedural content generation method which extracts local patterns from a sparse set of examples and transforms them into a set of local constraints. These constraints drive generation and ensure that every local patch of the output also exists in the set of input examples. The inner workings of the algorithm have been extensively described [14, 15].

We extend the WFC algorithm to architectural applications (Fig. 1) where discrete architectural tiles are manually composed into larger assemblies and supplied to the algorithm as design examples (Fig. 2). This example-based approach makes this methodology compatible with traditional architectural design workflows where experienced designers can *show and teach* what good designs look like and have the computer replicate virtually infinite variations of the provided examples. As discussed by Karth and Smith, the WFC "is particularly suited to non-programmers" [15], an uncommon feature among many advanced computational and generative design frameworks.

While in the original implementation the probability of certain patterns to appear in the final output is determined by pixel frequencies in the design samples [14], in our work that probability is guided through exposed normalized *weights* assigned to each individual tile.

Related to texture synthesis, model synthesis is one of the earliest applications of procedural constraint solving for 3D environments [16]. Recently, a renewed interest in such methods has attracted designers beyond game design applications including urban and building scale applications in conjunction with machine learning methodologies [17–19]. Our work further extends these by allowing a search algorithm to manipulate the probability-weights and placement of fixed tiles to control diversity of output and optimization along a set of objectives and features. Despite the observed growing interest in procedural constraint solving methods, viable applications for architectural design are still unexplored.

Observed Limitations. WFC is remarkable for its simplicity but, despite some work on extending its functionality [20] has several limitations:

- WFC does not offer control over global constraints.
- Constraints are purely spatial (adjacency).
- Lacks input controls for a search algorithm.
- Lacks domain specific constraints.

Our approach addresses these limitations via:

- Control of formal massing via global performance metrics (e.g., natural ventilation and noise) and global geometric features (e.g., building façade area) via integration with a QD optimization framework.
- Dynamic weighting for tile unit selection as optimization controls.
- Dynamic pre-constraining of tiles for improved searchability.
- Fixed pre-constraining with boundary solution tiles for design-domain ease of use.

The manual nature of crafting design examples and building a catalog of units makes this design space model highly versatile, accessible, and compatible with traditional modeling techniques and design approaches.

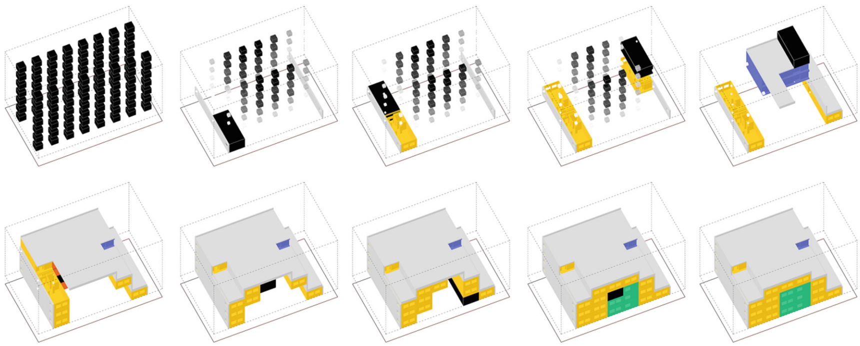

Fig. 1. Collapsing process of the WFC. Refer to this video featuring this process in action: https://vimeo.com/668784164

2 Methods and Data

2.1 Geometry System

The tiles catalog supplied to our model includes basic building components: façade, apartment, and stair core tiles. We also use empty tiles to govern building boundaries and open spaces (Fig. 2). This tile set is then used to manually create a defined set of example designs that can represent the kinds of desired variations and provide the WFC algorithm with tile-to-tile adjacency rules. Using this set of design examples in conjunction with the WFC algorithm we automatically generate a wide variety of site building layouts (Fig. 3).

To improve control over the WFC output we extend the algorithm's basic functionality with tile probability weights and variable tile pre-constraints. The set of weights, one for each tile type, can be varied to control the probability of the associated tile to appear in the WFC output (Fig. 4 top). Variable pre-constraints – tiles which are fixed at the start of WFC – 'lock-in' parts of the design while leaving the rest to the WFC generation process. These fixed tiles are added and removed from solutions as part of the search process (Fig. 4 bottom).

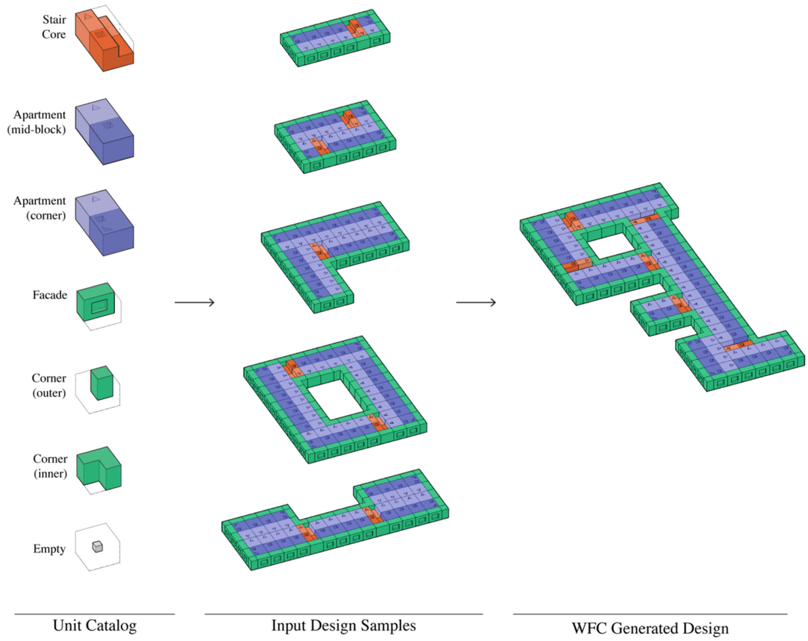

Fig. 2. WFC steps: definition of a catalog of units (left), design of examples (middle), and example WFC solution based on provided examples (right).

2.2 Features and Objectives

The QD algorithm MAP-Elites produces designs that are high performing along a set of objectives and diverse along a set of features (Fig. 5). While objectives are functions to be minimized or maximized (such as ventilation or site noise), features are quantifiable design characteristics to be fully explored (such as perimeter length or number of buildings). In our case study project, we included sustainability, livability, and penalty types of objectives and features that define geometric attributes.

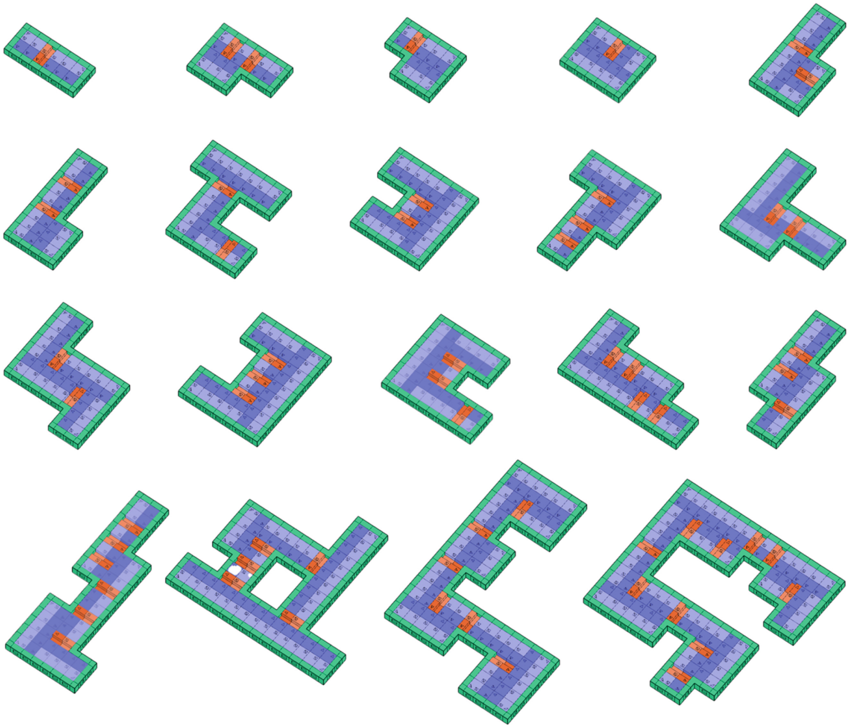

Fig. 3. Sample WFC outputs. Varied and diverse results that comply with the rules encoded by the example designs. This represents a new way of generating industry-specific design solutions beyond typical parametric approaches.

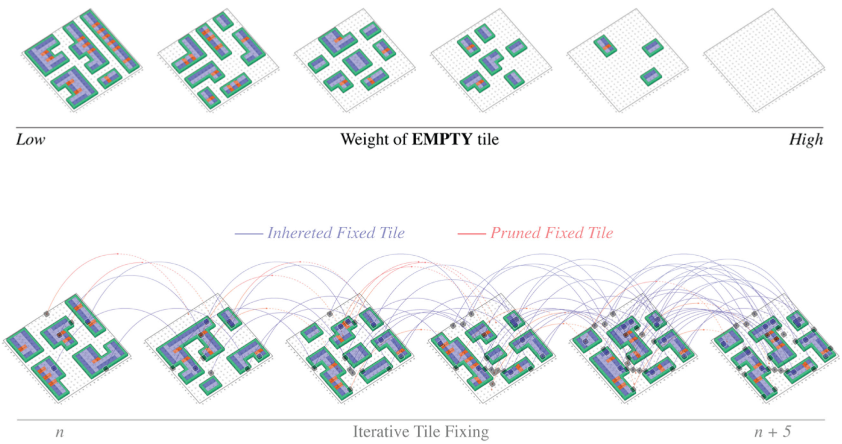

Fig. 4. Model parameters: (above) weights are assigned per tile type and drive the probability of it appearing in the final output. (below) Variable pre-constraints: location, addition and removal of pre-determined tiles are additional variables to influence collapsing process.

Sustainability Objectives

- *Indoor Ventilation.* This metric defines the natural indoor ventilation potential for each apartment. A simplified version of the air flow network (AFN) methodology is used—the connectivity distance of each room to the apartment's windows.
- *Landscape Capacity for Carbon Sequestration.* This metric measures the potential capacity for outdoor green areas to store and avoid carbon [21]. We approximate this capacity from a 'clearance' metric, the amount of clear space green areas have form adjacent buildings. This metric values larger areas of clear space, which can support greater levels of vegetation and trees, more highly – differentiating it from the total area of open space.

Livability Objectives

- *Site noise.* This metric is defined as the percentage of tiles on the site with a noise level of less than 50db. To accelerate optimization, we estimate this site-specific measurement using a surrogate model trained on a large set of noise analysis simulations performed on apartment complex designs designed manually by customers. Noise sources are highways and surface roads near the actual site.

Penalty Objectives. The penalty metrics are introduced to help steer the optimization towards viable and acceptable design solutions.

- *Number of Apartment Units.* Count of apartment units.
- *Proximity of Units to Building Cores.* Tile distance of each apartment unit to the building's cores. The distance to a core must be less than 5 tiles.

Features. This set was deliberately chosen to promote as much diversity as possible based on the model's geometry system. Three features are chosen as the number of feature pairs (three pairs, versus six for four features) can be succinctly communicated.

- *Façade Length.* Ratio of number of façade tiles to number of units
- *Number of Buildings.* Count of complete buildings.
- *Total Size of Open Spaces.* Surface area of open spaces. Open spaces are computed as the number of empty tiles not occupied by buildings excluding 1 tile corridors between buildings.

2.3 Encoding and Optimization

MAP-Elites. The QD algorithm MAP-Elites [5] is used to optimize an encoding composed of two parts: a vector of tile weights and a set of fixed tiles. MAP-Elites first divides the feature space into a set of discrete bins, or map. The map houses the population, with each bin in the map holding a single individual. When a new solution is evaluated, it is assigned a bin based on its features and, if that bin is empty, it is added

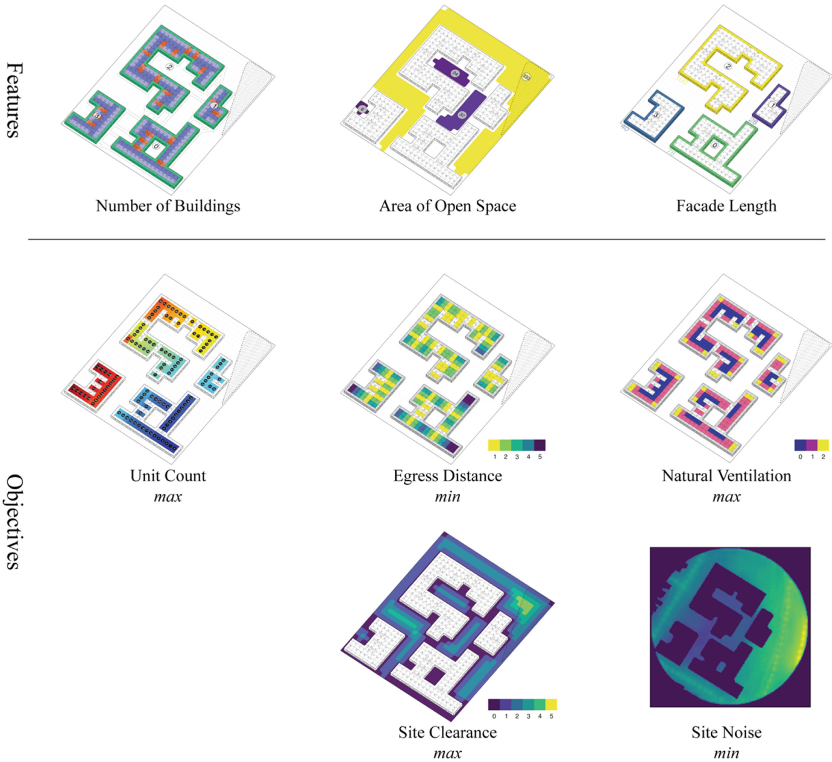

Fig. 5. Features (above) and objectives (below). Greyed out region indicates area unavailable for development.

to the map. If the bin is occupied by another solution the solution with a higher fitness is kept in the bin and the other discarded. In this way each bin contains the best solution ever found for that combination of features. These best solutions are known as elites.

To produce new solutions parents are chosen randomly from the elites, mutated, evaluated, and assigned a bin based on their features. Child solutions have two ways of joining the map: discovering an unoccupied bin, or out-competing an existing solution for its bin. Repetition of this process produces an increasingly explored feature space and an increasingly optimal collection of solutions. The optimization process is illustrated in Fig. 6.

Multiple objectives are optimized using the T-Domino [22] variant of MAP-elites. T-DominO ranks solutions according to the number of other solutions in the map that are dominated on each objective – rewarding solutions with balanced performance over those which excel at only a single objective. Solutions which follow constraints are always preferred over those which do not.

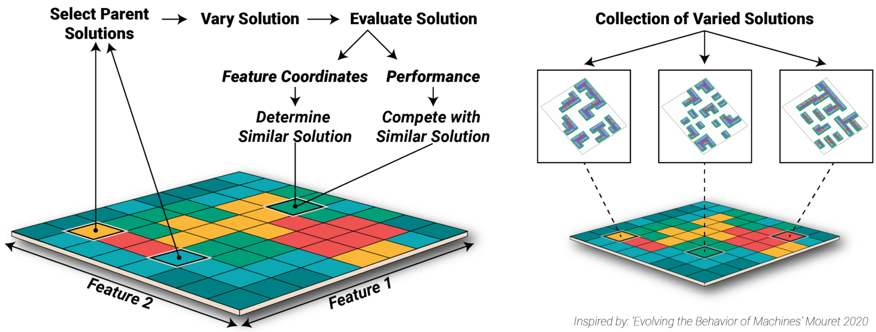

Fig. 6. Optimization of diverse solutions with Quality-Diversity.

Tile Weights. In WFC generated tiles are chosen from a set of valid tiles probabilistically – if the weighting of two valid tiles is 3:1, the first will be chosen 75% of the time and the second 25% (Fig. 4 top). Differences in this weighting has broad effects, but alone has limited effectiveness for optimization (Fig. 7).

Fixed Tiles. To achieve the fine-grained control necessary for optimization, solutions are encoded with a set of fixed tiles. MAP-Elites is an evolutionary algorithm, which produces new child solutions by altering existing parent solutions. Children inherit these fixed tiles from parents, in addition to fixing an additional tile from the design produced by the parent or removing one of the tiles that were fixed by the parent. Fixing tiles freezes key portions of the parent design and saves progress toward interesting designs – while still allowing substantial deviation from the parent, as the rest of the tiles are generated stochastically with WFC (Fig. 4 bottom).

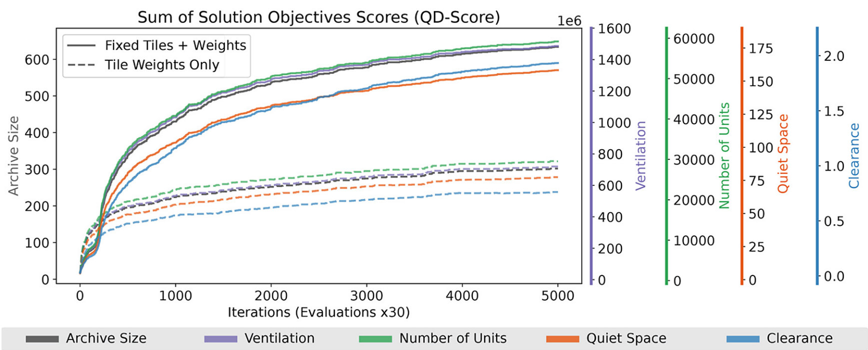

Fig. 7. Sum of objective values of best designs found in every feature region by MAP-Elites. Fixing tiles (solid line) as part of the optimization process dramatically improves the diversity of solutions found (Archive Size) as well as the performance of the solutions on the target objectives vs. using only global weights (dotted line).

Optimization Settings. At each generation 30 new individuals were created by mutating parent individuals with at 50% probability of adding a tile and a 50% probability of removing a tile. One run consists of 5000 generations. One full run takes approximately 8h on a 32 core workstation. The feature space was divided into a 10x10x10 grid, to create a collection of solutions, or archive, of up to 1000 solutions. Features were explored between the ranges of [1.8–3.6], [300–500], [4–14] for the Façade Per Unit, Open Area, and Number of Buildings respectively.

3 Findings and Discussion

Our approach generates a high performing set of apartment layouts which vary along the provided features -- illuminating the relationship between these features and performance. Viewing the performance of designs organized by these features we see that layouts with fewer large buildings tend towards poor natural ventilation —- an effect

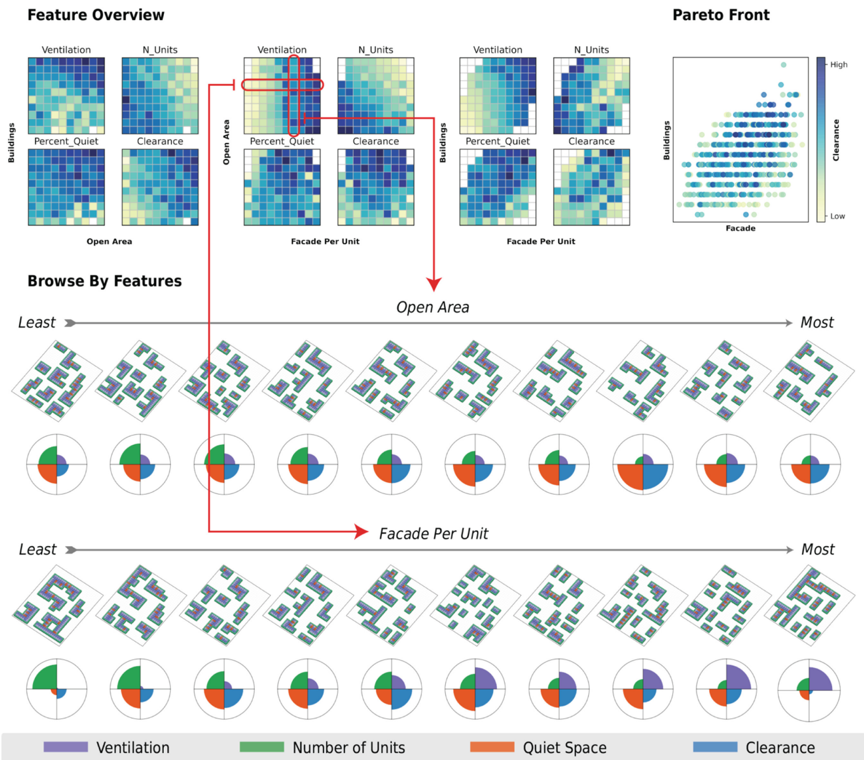

Fig. 8. Example visualization approaches made possible by the feature-centered Quality-Diversity approach. Top Left: Top solutions are organized by feature combinations, with each feature region represented as a bin colored by performance on each objective. Relationships are visible at a glance, allowing rapid identification of promising regions. Top Right: The same overview presented with a single objective as a Pareto Front. Bottom: Browsing designs by features. The highlighted feature regions are explored as a walk-through feature space, allowing designers to browse designs in an intuitively structured way.

that can be remedied with longer, more convoluted facades. Conversely, layouts with few large buildings and larger open areas interspersed across the site tend to have less noise. These insights are easily identifiable with the map-based visualization approaches (Fig. 7).

Our method of encoding and optimizing WFC-based solutions makes this type of exploration possible. Optimization of tile weights guides the direction of WFC and iteratively fixing tiles provides further control -- resulting in improvements in the quality of solutions produced while accelerating optimization by an order of magnitude. The fixed tile approach adds an intuitive method of steering optimization. Fixed tiles can be manually or parametrically placed to guide design outcomes around constraints like stair locations, courtyards, or existing structures (Fig. 8).

Fig. 9. An example of 3D feature mapping for intuitive navigation of design spaces. Here we show how additional attributes can be extracted (such as number of open spaces and area of individual open spaces) to further aid the design space navigation.

High-level features are valuable to designers for decision-making, but difficult to integrate into multi-objective frameworks -- QD allows the design space to be viewed through the lens of these features (Fig. 9). WFC allows design intent to be communicated through concrete, visual examples as well as promoting a multi-scalar approach to design where internal layouts, building footprint and site organization happen all simultaneously. Both advances flow from the same principle: for generative design to be useful, it must be intuitive.

Acknowledgments. The authors would like to thank Renaud Danhaive, Jeffrey Landes, and the entire Spacemaker team for their invaluable site analysis tool and expertise as well as Mark Davis and David Benjamin for their guidance and support.

References

1. Bradner, E., Iorio, F., Davis, M.: Parameters Tell the Design Story: Ideation and Abstraction in Design Optimization. In: Proceedings of the Symposium on Simulation for Architecture & Urban Design, vol. 26 Society for Computer Simulation International (2014)
2. Gerber, D.J., et al.: Design optioneering: multi-disciplinary design optimization through parameterization, domain integration and automation of a genetic algorithm. In: Proceedings of the 2012 Symposium on Simulation for Architecture and Urban Design, pp. 1–8 (2012)
3. Druot, T., et al.: Multi-objective optimization of aircrafts family at conceptual design stage. In: Inverse Problems, Design and Optimization Symposium. Albi, France, vol. 29, pp. 773–79 (2013)
4. Cully, A., Demiris, Y.: Quality and diversity optimization: a unifying modular framework. IEEE Trans. Evol. Comput. **22**(2), 245–259 (2017)
5. Mouret, J.B., Clune, J.: Illuminating Search Spaces by Mapping Elite. ArXiv Preprint ArXiv:1504.04909 (2015)
6. Cully, A., Clune, J., Tarapore, D., Mouret, J.B.: Robots that can adapt like animals. Nature **521**(7553), 503–507 (2015)
7. Lehman, J., Stanley, K.O.: Evolving a diversity of virtual creatures through novelty search and local competition. In: Proceedings of the 13th annual conference on Genetic and evolutionary computation, pp. 211–218 (2011)
8. Gaier, A., Asteroth, A., Mouret, J.B.: Data-efficient exploration, optimization, and modeling of diverse designs through surrogate-assisted illumination. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 99–106 (2017)
9. Gaier, A., Asteroth, A., Mouret, J.B.: Data-efficient design exploration through surrogate-assisted illumination. Evol. Comput. **26**(3), 381–410 (2018)
10. Khalifa, A., Lee, S., Nealen, A., Togelius, J.: Talakat: Bullet hell generation through constrained map-elites. In: Proceedings of The Genetic and Evolutionary Computation Conference, pp. 1047–1054 (2018)
11. Alvarez, A., Dahlskog, S., Font, J., Togelius, J.: Empowering quality diversity in dungeon design with interactive constrained map-elites. In: 2019 IEEE Conference on Games (CoG) IEEE, pp. 1–8 (2019)
12. Fontaine, M.C., et al.: Illuminating mario scenes in the latent space of a generative adversarial network. arXiv preprint arXiv:2007.05674 (2020)
13. Nagy, D., Villaggi, L., Zhao, D., Benjamin, D.: Beyond heuristics: a novel design space model for generative space planning in architecture (2017)

14. Gumin, M.: Wave Function Collapse. https://github.com/mxgmn/WaveFunctionCollapse (2016)
15. Karth, I., Smith, A.M.: WaveFunctionCollapse is constraint solving in the wild. In: Proceedings of the 12th International Conference on the Foundations of Digital Games, pp. 1–10 (2017)
16. Merrell, P.C.: Model synthesis. Ph.D. Dissertation. University of North Carolina at Chapel Hill (2009)
17. Lin, B., Jabi, W., Diao, R.: Urban space simulation based on wave function collapse and convolutional neural network. In: Proceedings of the 11th Annual Symposium on Simulation for Architecture and Urban Design, pp. 1–8 (2020)
18. Mintrone, A., Erioli, A.: Training spaces – fostering machine sensibility for spatial assemblages through wave function collapse and reinforcement learning. In: proceedings of the 39th eCAADe Conference, Volume 1, University of Novi Sad, Novi Sad, Serbia, pp. 8-10 (2021)
19. Chasioti, E.: Gameplay with encoded architectural tilesets: A computational framework for building massing design using the Wave Function Collapse algorithm. BARC0141: Built Environment Dissertation (2020)
20. Sandhu, A., Chen, Z., McCoy, J.: Enhancing wave function collapse with design-level constraints. In: Proceedings of the 14th International Conference on the Foundations of Digital Games (FDG '19: The Fourteenth International Conference on the Foundations of Digital Games, San Luis Obispo California USA: ACM), pp. 1–9 (2019)
21. i-Tree, Califronia Urban Forests, Urban Ecos and Cal Fire (2020). https://www.itreetools.org/support/resources-overview/i-tree-methods-and-files/new-carbon-equations-and-methods-2020
22. Gaier, A. Stoddart J., Villaggi L., Bentley, P.J.: T-Domino: Exploring Multiple Criteria with Quality Diversity and Tournament Dominance Objective. In: Parallel Problem Solving From Nature. Springer (2022). https://doi.org/10.1007/978-3-031-14721-0_19