



T-DominO

Exploring Multiple Criteria with Quality-Diversity and the Tournament Dominance Objective

Adam Gaier¹(✉), James Stoddart², Lorenzo Villaggi², and Peter J. Bentley^{1,3}

¹ Autodesk Research, Bonn, Germany
adam.gaier@autodesk.com

² Autodesk Research, New York, USA

³ University College London, London, UK

Abstract. Real-world design problems are a messy combination of constraints, objectives, and features. Exploring these problem spaces can be defined as a Multi-Criteria Exploration (MCX) problem, whose goals are to produce a set of diverse solutions with high performance across many objectives, while avoiding low performance across any objectives. Quality-Diversity algorithms produce the needed design variation, but typically consider only a single objective. We present a new ranking, T-DominO, specifically designed to handle multiple objectives in MCX problems. T-DominO ranks individuals relative to other solutions in the archive, favoring individuals with balanced performance over those which excel at a few objectives at the cost of the others. Keeping only a single balanced solution in each MAP-Elites bin maintains the visual accessibility of the archive – a strong asset for design exploration. We illustrate our approach on a set of easily understood benchmarks, and showcase its potential in a many-objective real-world architecture case study.

Keywords: Quality-diversity · Generative design · Multi-objective

1 Introduction

Architecture projects must balance a dizzying array of objectives: daylight, views, noise, wind, cost, open spaces, carbon footprint, and ease of construction, to name a few – along with less easily optimized subjective considerations like aesthetics and comfort. In generative design (GD), where algorithms aid design exploration by producing candidate designs, the desired result is not a single solution, but a variety of high performing options [4]. A variety of options is required because the problem has more than one objective, which means there may be many possible solutions. Perhaps more importantly, a varied choice highlights design concepts to stakeholders and decision makers who then select and modify them according to messy human compromises.




Fig. 1. Calculating the Tournament Dominance Objective (T-DominO)

Though it may resemble multi-objective optimization (MOO) [10], the problem this work focuses on for GD and similar domains is different. We define the problem as a Multi-Criteria Exploration (MCX) problem, whose goals are to:

1. Produce a catalog of diverse solutions
2. with high performance across many objectives
3. while avoiding low performance across any objectives

MCX can be considered an exploratory form of MOO, just as Quality-Diversity (QD) [8,39], is an exploratory form of single-objective optimization. In contrast to MOO, in MCX we do not strive for either uniform coverage of the Pareto front, nor precise proximity to it. Uniform coverage of the front implies coverage of the extremes of the objective space – where solutions earn their place in the Pareto front by dominating on only a subset of objectives. These solutions are uninteresting for MCX, as solutions which disregard user preferences by ignoring some objectives are not useful in practice. Proximity to the front is also less important for exploration – the goal is to generate starting points, not end points. Generated solutions are rarely used without modification, reducing the effort of finding the precise front to an expensive distraction.

The QD approach seems, at first, ideal for solving MCX problems. QD algorithms provide a way of explicitly searching for diversity as defined at a high level by users. Whereas MOO strives for a maximum spread in the objective space, QD searches for spread in a user-defined ‘feature’¹ space. As opposed to objectives, these features correspond to different *ways* of solving the problem, not quantities to be minimized. In architecture the number of buildings in a building complex or the distance between them can be explicitly explored with QD in a way that is not possible with MOO.

But QD is designed to explore several features, not to optimize multiple objectives. MAP-Elites [7,33], the most widely used QD algorithm, divides the feature space and searches for the best solution in each partition. The result is a set of optimized designs organized by high level features; the performance of this collection can then be viewed as heat map projected on to the feature space, illuminating the relationship between features and performance.

¹ Also referred to in the QD literature as a behavior, descriptor, outcome, or measure.

Recent work has proposed combining MOO and QD by computing a Pareto front in every partition [38]. In the MCX case this ‘all of the above’ approach is not satisfying – MAP-Elites’ intuitive way of organizing, summarizing, and presenting the results depends on finding a single best solution for each partition. An alternate method of reconciling the approaches must be found.

In this work we leverage the insight that the diversity in objective space produced by MOO mechanisms such as crowding distance [12], or reference vectors [11] are unneeded when diversity is enforced by QD. In QD, users can choose the type of variety to explore, and trade-offs in objectives will naturally arise from those choices. When exploration of the objective space is no longer prioritized, non-domination – which favors the extremes of the front to the same degree as the center – ceases to be the most desirable attribute of a solution. In MCX, balanced solutions which perform well on all objectives are preferred. Our second insight is that the ‘balance’ of a solution can be defined in relation to a population, and that the solutions contained in the MAP-Elites archive can act as that population.

Our approach extends MAP-Elites to the exploration of problems with multiple objectives by introducing the Tournament Dominance Objective (T-DominO), which ranks individuals in a population according to an approximation of their distance to the center of the Pareto front. T-DominO awards poor scores to non-dominated solutions at the extremes of the objective space – those which excel at one objective while doing very poorly at the other – while those at the center of the front receive the highest scores (Fig. 1).

Optimizing MAP-Elites according to T-DominO provides an elegant approach to tackle MCX problems. A simple alternate ranking causes minimal disruption to the core algorithmic machinery while allowing MAP-Elites to discover varied design concepts which balance multiple objectives. Approaches which assume a single objective and single solution in each bin, such as CMA-ME [14], can still be used. Crucially, T-DominO tackles multiple objectives without sacrificing MAP-Elites’ intuitive visualization and analysis of solutions, features, and their interaction with objectives – the true goal of the algorithm when used for design.

2 Background

2.1 Generative Design

In design and architecture, experiments with human-machine collaboration are common [2, 4, 18, 26, 35, 41]. Recent work has demonstrated the viability of GD in real-world applications, from office retrofits [35], large scale trade-shows [36], to neighborhood scale planning [34]. The number of conflicting constraints, preferences, and objectives in these projects makes ‘solving’ them an ill-defined and impossible task. Optimization tools are typically used at the *beginning* of the design process rather than the end. Optimization algorithms are not used solve problems, but to *explore* them [4, 31].

The purpose of GD is less optimization and more communication. Search algorithms are used to understand the possibilities and potential of a problem

space. Objectives serve as proxies for preferences, goals, and features of interest that are often difficult or impossible to define mathematically. These objectives signify criteria of a good design or ways of counter-balancing those criteria to prevent extreme solutions which are not aligned with designers intent.

Results are then filtered and categorized in an effort to find qualitatively different design concepts. Typically designs are judged visually first, and only once a set of interesting and varied designs identified is their performance examined. This can be a clumsy process, and for GD to have real success accessibility must be a consideration at every step of the process, including optimization. An intuitive GD approach would not only find a set of solutions which balance performance over several objectives, but explicitly search for the high level diversity that sets solutions apart from each other. This is the goal of MCX.

2.2 Exploration and Optimization with Non-objective Criteria

MOO approaches strive to produce a set of non-dominated solutions that is diverse in the objective space, and as near to the Pareto optimal front as possible [37], but the diversity that interests designers is often not in objective space. Other qualities can be induced with ‘helper’ objectives in a process known as multiobjectivization [25, 27, 30]. Helper objectives can optimize quantities unrelated to performance, such as the type of cross sections in a structural frame [21], or the similarity to previous solutions [32], but are still performing minimization. Maximizing or minimizing the number of buildings on a site makes little sense, but understanding the effect of the number of buildings is a valuable insight.

QD approaches such as MAP-Elites [7, 33] search for solutions along a continuum of user-defined features, making them ideal for exploration. MAP-Elites has been used for design exploration in domains such as aerodynamics [15–17, 23, 24], and game design [1, 5, 19, 20], but has been restricted to consideration of a single objective. MAP-Elites operates by first discretizing the feature space into bins, collectively known as a map or archive. Each bin contains a single solution and its corresponding fitness value. New solutions are created by selecting and varying solutions from the map. These new solutions are evaluated and two values produced: a performance measure and a set of coordinates in the feature space. These coordinates indicate the bin to which the solution belongs. The solution is placed in the bin if it is empty, or if the candidate solution has higher performance than the current occupant of the bin, it replaces it.

The elitist nature of MAP-Elites, with only one solution per bin, puts it at odds with the idea of the Pareto front. A concurrent work [38] bridges this gap by introducing a Pareto front in each bin, and replacing fitness tournaments with non-domination. Though this technique is able to find a large set of Pareto fronts, it sacrifices the elegant method of communicating the results. Rather than viewing individual designs and correlations between features and objectives, we are left with a mass of summary statistics – useful for MOO, but not for MCX. In our work we maintain the the elitist nature of MAP-Elites, and instead replace the Pareto front with an alternate formulation of multi-objective performance.




Fig. 2. Using anchor points to calculate T-DominO.

3 Method

When tackling MCX problems, our interest lies in finding solutions which perform well on all objectives in each region of a QD feature space. The Tournament Dominance Objective (T-DominO), introduced here, ranks solutions according to an approximation of their distance to the center of the front, with the most balanced solutions ranking highest. This approximation is calculated through a series of tournaments between a solution and a set of existing points in objective space, or *anchor points* (Fig. 2). An individual is compared to each anchor point on a single objective, and for every anchor point with a lesser or equal objective value one point is awarded. This count is made for every objective, and these counts multiplied.

The T-DominO score of solution with objective values x , compared to a set of anchor points with objective values A is more precisely defined as:

$$\text{T-DominO}(x, A) = \prod_{n=1}^{objs} \sum_{m=1}^{anchors} f(x_n, A_{mn}), f(x, a) = \begin{cases} 1, & x \geq a \\ 0, & \text{else} \end{cases} \quad (1)$$

where *objs* and *anchors* is the number of objectives and anchor points.²

The integration of T-DominO into MAP-Elites can be summarized as follows:

1. a new individual, based on its feature coordinates, is assigned a bin.
2. a set of anchor points are collected from the history of elites in that bin, the k neighboring bins, and the new individual itself.
3. the T-DominO of the current elite and the challenger are computed based on these anchor points.
4. if the challenger has a greater T-DominO score it replaces the current elite, and the objective values of the replaced elite are stored in the bin to serve as a future anchor point.

² Or in python: `numpy.prod(numpy.sum(objs >= anchors,axis=0))`.

When individuals in a population are ranked according to T-DominO the result is a ranking from the center of the front outwards (Fig. 2, right). This ranking allows the combination of multiple objectives into a single score which rewards solutions with the highest balanced performance, without the need for penalty functions or arbitrary weighting of objectives.

T-DominO is based around comparisons to anchor points, and the MAP-Elites archive provides a ready source. Existing elites in the archive can be used as a sampling of the objective space, and act as anchor points distributed across both objective and feature space. Selection pressure toward improved T-DominO scores creates high performing solutions which in turn act as anchor points, in a virtuous cycle that leads to ever higher performance.

However, when selection pressure is organized around *only* the current population, cycling can occur. In some circumstances a challenger solution which is better on one objective can replace the current elite, which can in turn be replaced by the original. To prevent this behavior and ensure progress towards better solutions we track the objective values of the previous elites. While each bin continues to contain only a single elite, the objective values of previous elites are maintained to act as anchor points, preventing cycling and creating further refinement of the T-DominO landscape. A simple FIFO buffer of the objective values of a handful of past elites is sufficient.

Neighboring partitions typically have similar performance potential, and so are necessary for creating more fine-grained landscapes, but bins with solutions that dominate all or none of the solutions in the a bin provide no signal to inform selection pressure, and so we can safely limit the anchor points to those contained in the k nearest neighboring bins.

T-Domino allows the simultaneous optimization of several objectives with a single measure, relying on QD diversity mechanisms to prevent convergence on a single point. The output of T-Domino MAP-Elites is ideal for MCX – an archive with a single best balanced solution in each bin. Having a single solution in each bin allows the effects of solution features on that balance to be easily understood and visualized and so contribute to the understanding of the underlying problem, such as the correlation of features and objectives. The creation of a library of designs organized by high level features of the users choosing provides an ideal set of starting points for further refinement.

4 Benchmarks

4.1 Setup

We validate the expected behavior of MAP-Elites with T-DominO on a series of established multi-objective benchmark problems. The purpose of these tests is to validate our claims that T-Domino will:

1. Discover high performing, if not optimal, solutions
2. Produce balanced solutions whose performance does not come at the cost of large trade-offs in a subset of objectives

Benchmark Functions

RastriginMOO. To judge the performance of T-DominO on Multi-Objective QD problems, we test on a version of RastriginMOO as introduced in [38]. The Rastrigin function is a classic optimization benchmark, often used to test QD algorithms because it contains many local minima [6,14]. Here it is converted into a multiobjective benchmark by optimizing a pair of Rastrigin functions with shifted centers. We use a 10-D version with constants added so that every discovered bin has a positive effect on the aggregate QD Score. These objectives can be explicitly defined as:

$$\begin{cases} f_1(\mathbf{x}) = 200 - \left(\sum_{i=1}^n [(x_i - \lambda_1)^2 - 10 \cos(2\pi(x_i - \lambda_1))] \right) \\ f_2(\mathbf{x}) = 200 - \left(\sum_{i=1}^n [(x_i - \lambda_2)^2 - 10 \cos(2\pi(x_i - \lambda_2))] \right) \end{cases} \quad (2)$$

where $\lambda_1 = 0.0$ and $\lambda_2 = 2.2$ for f_1 and f_2 . All parameters are limited to the range $[-2, 2]$, with the feature space defined by the first two parameters.

ZDT3. When spread across the objective space is desired, objectives themselves could be used as features. This use case is demonstrated with the ZDT3 benchmark, a 30 variable problem from the ZDT MOO benchmark problem suite [42] whose hallmark is a set of disconnected Pareto-optimal fronts, and whose first parameter is value of the first objective. Parameter ranges span 0–1 with the first two parameters used as features, enforcing a spread of solutions across the range of the first objective.

DTLZ3. To illustrate T-DominO’s bias toward balanced solutions we analyze its performance on DTLZ3, a many-objective benchmark with a tunable number of objectives and variables[13]. We test with 10 parameters and 5 objectives, with the 6th and 7th parameters use as features.³

Baseline Approaches

ME Single. MAP-Elites [33] optimizing only a single objective is used to establish an upper and lower bound of performance we can expect from MAP-Elites. Blind to the second objective we can expect it to find the top performing solutions for the first. Equally important, the exploration of all bins without regard to the performance on the second objective establishes a floor for performance – the performance we could expect for having any solution in the bin.

ME Sum. We compare using the T-Domino objective with MAP-Elites [33] using the most naive way of combining multiple objective – simply adding them. Our benchmarks all have well-scaled objectives, but this is typically not the case. To simulate this difficulty we use a weighted sum, with each additional objective values increased by an order of magnitude (e.g. $\times 1$, $\times 10$, $\times 100\dots$).

³ The first n parameters are explicitly linked to the first n objectives as in ZDT3 – later parameters are used to avoid explicitly exploring the objective space.

NSGA-II. NSGA-II [12] is used as a benchmark for conventional multi-objective optimization without feature space exploration, reaching near the Pareto front on these simple benchmarks. Though it is not our goal to compete with MOO algorithms, they provide a useful metric to contextualize the difference between exploratory approaches and pure optimizers.

Settings. In all MAP-Elites approaches the feature space is partitioned a 20×20 grid, with 2 CMA-ME improvement emitters [14] performing optimization. T-Domino was computed using the neighbors from 4 bins away, using a history of the 10 most recent elites in each bin. Hyperparameters for NSGA-II were kept comparable, a population of 400 matched the 400 bins of the MAP-Elites grids, with the same number of new solutions generated per generation for the same number of generations. A standard implementation of NSGA-II from the PyMoo library [3] is used, as well as the library’s formulations for the ZDT3 and DTLZ benchmarks whose the exact formulation is included in the online supplemental. The PyRibs [40] library was used as a basis for all MAP-Elites experiments, with T-DominO implemented as a specialized archive type. All experiments were replicated 30 times, additional plots are provided in the Supplemental.

4.2 Result

Figure 3 illustrates the explored regions of objective and feature space in a single run. Using the NSGA-II solutions to outline the true Pareto front we can see where each MAP-Elites approach concentrates. In the RastriginMOO case, though each version of MAP-Elites explores identical areas of the feature space, the range of possible values in objective space is large. T-Domino produces solutions in the middle of the front, with solutions that strike a balance between the two objectives. In the ZDT3 case we see that by explicitly exploring one of the objectives we can force spread over the objective space, and provide high performing solutions in the other objective.

With more than two objectives the balance seeking property of T-DominO becomes even more pronounced. Parallel coordinate plots (Fig. 3, top right), which plot each solution as a line with one vertex per objective, make clear the differing selection pressure of T-DominO and non-dominated sorting. In contrast to the spiky lines denoting high performance on some objectives and low performance on others, T-DominO’s solutions form a flat band of even performance.

The difference of balance is critical when approaching MCX problems. To spread across the five dimensional front, solutions found by NSGA-II must span many areas with solutions that perform poorly on some objectives. If we divide the range of objective values found by NSGA-II into quartiles, only 25% of the solutions found by NSGA-II perform over the bottom quartile on all objectives. If all of these objectives are valued by the user, that means that three quarters of solutions may be discarded immediately – and this will only worsen as the number of objectives grows. In contrast, when T-DominO MAP-Elites’ results

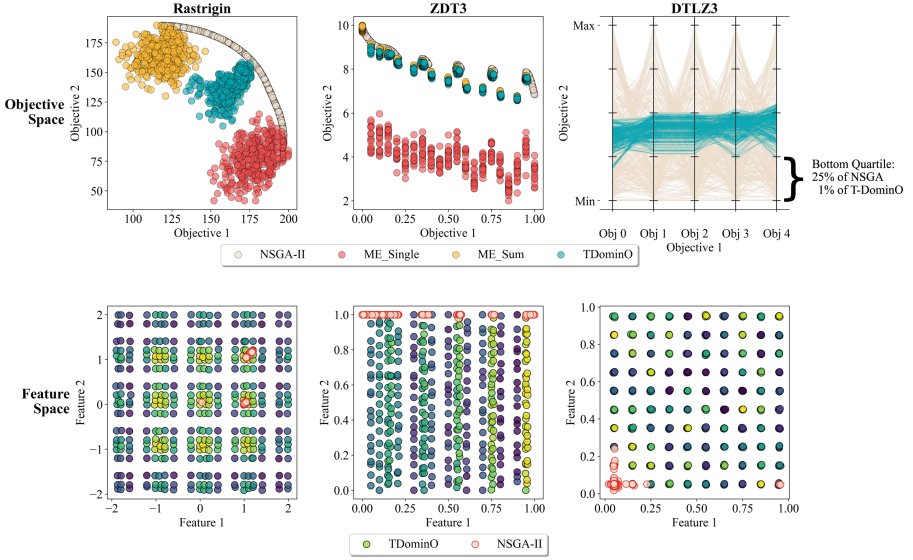


Fig. 3. Benchmark results. *Top:* Objective space as explored by each approach. *Bottom:* Feature space as explored by T-DominO MAP-Elites and NSGA-II. T-DominO points are colored by T-DominO score compute with the entire archive as anchor points.

are judged on the same scale, 99% of the solutions found by T-DominO MAP-Elites perform over the bottom quartile on all objectives.

Visualizing the distribution of found solutions in feature space (Fig. 3, bottom) gives a stark illustration of the main motivation for using a QD approach. The solutions produced by NSGA-II cluster in a tiny portion of the feature space. This region may be Pareto optimal, but QD gives us the ability to explore areas of our choosing.

5 Case Study

5.1 Setup

As a study of the applicability of T-DominO MAP-Elites to MCX problems we explore its use in optimizing building layouts for real-world residential complex. Solutions are produced using wave function collapse (WFC) [22], a popular tool for tile-based procedural content generation in games. WFC is a constraint satisfaction approach which extracts local patterns from a small set of samples, and transforms them into a set of local constraints. The constraints drive generation,




Fig. 4. Objectives, Features, and Constraints of building layout study. Shaded regions indicate portion of building site which cannot be built on. For details on computation of each metric see the online supplemental.

ensuring that every local patch of the output also exists in the set of input examples. We adapt the implementation here [28, 29]. Constrained generation systems like WFC are particularly appropriate for the semi-constrained design systems often used in residential building, such as modular or prefabricated units, and do not require the extensive curated datasets of valid designs.

WFC, though constrained, is a stochastic process. At every iteration a tile is ‘collapsed’, or fixed, with the type chosen stochastically from a list of valid tiles, and new constraints applied to its neighbors. To make this encoding more amenable to optimization we have introduced an evolvable genotype of tiles which are fixed at the beginning of this collapsing process. 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 remainder of the tiles are generated stochastically with WFC. See the supplemental material for set of used tiles and example designs.

The resulting designs are evaluated according to 4 objectives, 1 constraint, and explored along 3 features (illustrated in Fig. 4, more details in the online supplemental). The constraint was handled in a tournament fashion as in [9] – in any tournament where one solution follows the constraint and the other does not, the solution which follows the constraint wins regardless.

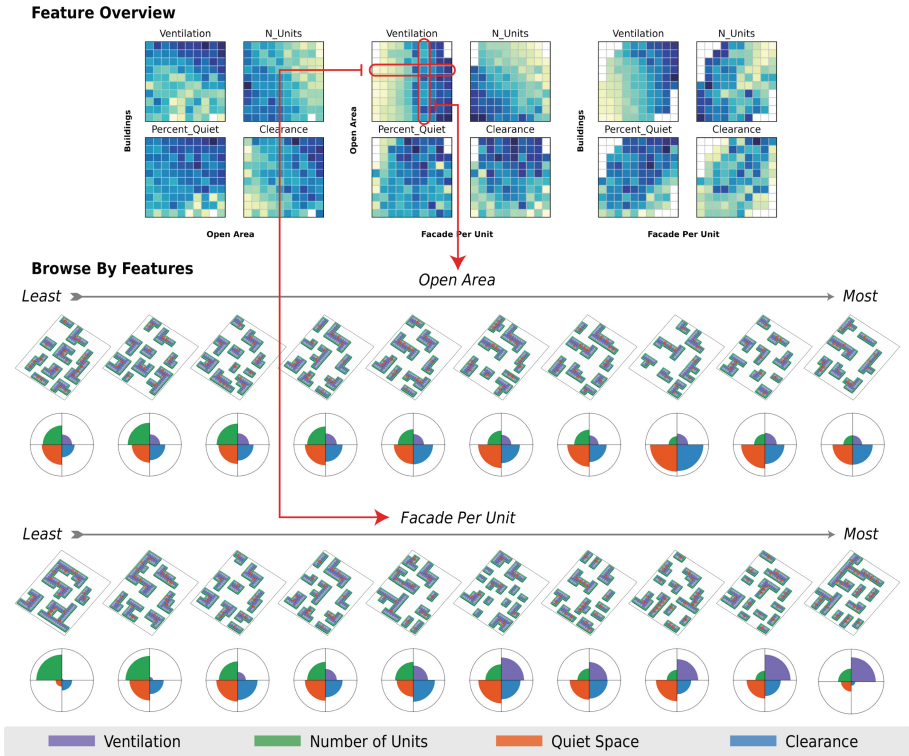


Fig. 5. Exploring building layouts generated with T-DominO MAP-Elites. Top: 2D views of 3D feature space, solutions in groups of four are identical, colored in each view by each objective (darker is better). Bottom: A walk through designs that vary along one feature dimension, with accompanying objective values. Petal plots are scaled to the final min/max objective value found in the archive. (Color figure online)

5.2 Result

Once a set of designs has been produced, we can extend on MAP-Elites' intuitive way of optimizing, organizing, and displaying solutions to multiple objectives. To better understand our 3D feature space we 'flatten' it into a set of 2D views by creating a set of new 2D archives with the desired feature axes and inserting all of the solutions from the 3D archive, forcing competition based on only two features. The result is one map for each pair of the three features (Fig. 5, top).

Each of these views can in turn be split into one map for each objective, collectively allowing correlations along feature axes and objectives to all be seen at a glance. Obvious relationships such as an increase in open area resulting in more units are clear, along with less appreciated connections: with fewer buildings ventilation is worse – unless buildings have longer facades.

Clear organized grids of solutions open up many avenues for intuitive navigation of the produced solutions. Here we show one possibility, browsing rows

or columns of designs. An area of the map can be selected, and the individual designs displayed along with their performance across objectives. Drilling down on a subset we can see the qualitative differences between large and small open areas (Fig. 5, Center), and the kinds of layouts they each represent. A combined view allows us to see large difference in objective values that may not have been apparent from a qualitative glance: even with the same amount of open area we can see there is huge amount of variation in the number of units that can fit on a site, that increasing this amount of units comes at the cost of natural ventilation, and that qualitatively this trade-off is between a few thick buildings, or several small ones (Fig. 5, Bottom).

6 Discussion

In this work we have defined a new class of problem, the MCX problem, tailored specifically to the needs of generative design. A chief aim of the generative design is to spark ideas and explore concepts, and results are typically explored by browsing designs not objectives. The measure space of MAP-Elites provides an intuitive way of creating and exploring sets of solutions with varied and understandable high level features.

T-DominO allows MAP-Elites to maintain these visual and organizational capabilities in the complex multi-criteria scenarios where they are most useful. Keeping a single solution in each bin rather than a front is about more than computational cost, it is about maintaining visual accessibility. Having a single solution in each bin simplifies browsing and selecting interesting designs. When objectives and features are correlated, the possible objectives values for each feature combination is constrained to a range – so though balanced solutions are found, the larger pattern of objective/feature relations are still clear.

T-DominO allows us to optimize for multiple objectives in a QD setting without making any other fundamental changes to the algorithm. Simple to implement, without adding any appreciable computational burden, T-DominO can be easily integrated into existing approaches. By leaving its elitist character untouched, T-DominO allows MAP-Elites to handle multiple objectives while maintaining its core visualization and presentation strengths. Equipping MAP-Elites with T-DominO allows us to generate diverse sets of well-rounded high performing solutions, creating a powerful tool for tackling MCX problems.

Acknowledgements. 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.

Supplemental Material and Code. Supplemental material and code available at: https://github.com/agaier/tdomino_ppsn.

References

1. 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), pp. 1–8. IEEE (2019)
2. Arieff, A.: New forms that function better. *Technol. Rev.* **116**(5), 94–98 (2013). TECHNOL REV 1 MAIN ST, 13 FLR, CAMBRIDGE, MA 02142 USA
3. Blank, J., Deb, K.: pymoo: multi-objective optimization in python. *IEEE Access* **8**, 89497–89509 (2020)
4. 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)
5. Charity, M., Green, M.C., Khalifa, A., Togelius, J.: Mech-elites: illuminating the mechanic space of gvg-ai. In: International Conference on the Foundations of Digital Games, pp. 1–10 (2020)
6. Cully, A.: Multi-emitter map-elites: improving quality, diversity and data efficiency with heterogeneous sets of emitters. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 84–92 (2021)
7. Cully, A., Clune, J., Tarapore, D., Mouret, J.B.: Robots that can adapt like animals. *Nature* **521**(7553), 503–507 (2015)
8. Cully, A., Demiris, Y.: Quality and diversity optimization: a unifying modular framework. *IEEE Trans. Evol. Comput.* **22**(2), 245–259 (2017)
9. Deb, K.: An efficient constraint handling method for genetic algorithms. *Comput. Methods Appl. Mech. Eng.* **186**(2–4), 311–338 (2000)
10. Deb, K.: Multi-objective optimization. In: Search methodologies, pp. 403–449. Springer, Boston (2014). https://doi.org/10.1007/978-1-4614-6940-7_15
11. Deb, K., Jain, H.: An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE Trans. Evol. Comput.* **18**(4), 577–601 (2013)
12. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Trans. Evol. Comput.* **6**(2), 182–197 (2002)
13. Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable multi-objective optimization test problems. In: Proceedings of the 2002 Congress on Evolutionary Computation, CEC 2002 (Cat. No. 02TH8600), vol. 1, pp. 825–830. IEEE (2002)
14. Fontaine, M.C., Togelius, J., Nikolaidis, S., Hoover, A.K.: Covariance matrix adaptation for the rapid illumination of behavior space. In: Proceedings of the 2020 Genetic And Evolutionary Computation Conference, pp. 94–102 (2020)
15. Gaier, A., Asteroth, A., Mouret, J.B.: Aerodynamic design exploration through surrogate-assisted illumination. In: 18th AIAA/ISSMO Multidisciplinary Analysis And Optimization Conference, p. 3330 (2017)
16. 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)
17. Gaier, A., Asteroth, A., Mouret, J.B.: Data-efficient design exploration through surrogate-assisted illumination. *Evol. Comput.* **26**(3), 381–410 (2018)
18. Gerber, D.J., Lin, S.H., Pan, B., Solmaz, A.S.: 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)

19. González-Duque, M., Palm, R.B., Ha, D., Risi, S.: Finding game levels with the right difficulty in a few trials through intelligent trial-and-error. In: 2020 IEEE Conference on Games (CoG), pp. 503–510. IEEE (2020)
20. Gravina, D., Khalifa, A., Liapis, A., Togelius, J., Yannakakis, G.N.: Procedural content generation through quality diversity. In: 2019 IEEE Conference on Games (CoG), pp. 1–8. IEEE (2019)
21. Greiner, D., Emperador, J.M., Winter, G., Galván, B.: Improving computational mechanics optimum design using helper objectives: an application in frame bar structures. In: Obayashi, S., Deb, K., Poloni, C., Hiroyasu, T., Murata, T. (eds.) EMO 2007. LNCS, vol. 4403, pp. 575–589. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70928-2_44
22. Gumin, M.: Bitmap and tilemap generation from a single example by collapsing a wave function. GitHub (2016)
23. Hagg, A., Asteroth, A., Bäck, T.: Prototype discovery using quality-diversity. In: Auger, A., Fonseca, C.M., Lourenço, N., Machado, P., Paquete, L., Whitley, D. (eds.) PPSN 2018. LNCS, vol. 11101, pp. 500–511. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-99253-2_40
24. Hagg, A., Wilde, D., Asteroth, A., Bäck, T.: Designing air flow with surrogate-assisted phenotypic niching. In: Bäck, T., et al. (eds.) PPSN 2020. LNCS, vol. 12269, pp. 140–153. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58112-1_10
25. Handl, J., Lovell, S.C., Knowles, J.: Multiobjectivization by decomposition of scalar cost functions. In: Rudolph, G., Jansen, T., Beume, N., Lucas, S., Poloni, C. (eds.) PPSN 2008. LNCS, vol. 5199, pp. 31–40. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87700-4_4
26. Holzer, D., Hough, R., Burry, M.: Parametric design and structural optimisation for early design exploration. *Int. J. Archit. Comput.* **5**(4), 625–643 (2007)
27. Jensen, M.T.: Helper-objectives: Using multi-objective evolutionary algorithms for single-objective optimisation. *J. Math. Modell. Algo.* **3**(4), 323–347 (2004)
28. Karth, I.: wfc2019f (2021). <https://github.com/ikarth/wfc-2019f>
29. Karth, I., Smith, A.M.: Addressing the fundamental tension of pcgml with discriminative learning. In: Proceedings of the 14th International Conference on the Foundations of Digital Games, pp. 1–9 (2019)
30. Knowles, J.D., Watson, R.A., Corne, D.W.: Reducing local optima in single-objective problems by multi-objectivization. In: Zitzler, E., Thiele, L., Deb, K., Coello Coello, C.A., Corne, D. (eds.) EMO 2001. LNCS, vol. 1993, pp. 269–283. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44719-9_19
31. Matejka, J., Glueck, M., Bradner, E., Hashemi, A., Grossman, T., Fitzmaurice, G.: Dream lens: exploration and visualization of large-scale generative design datasets. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, pp. 1–12 (2018)
32. Mouret, J.B.: Novelty-based multiobjectivization. In: New horizons in evolutionary robotics, pp. 139–154. Springer, Berlin (2011). https://doi.org/10.1007/978-3-642-18272-3_10
33. Mouret, J.B., Clune, J.: Illuminating search spaces by mapping elites. arXiv preprint [arXiv:1504.04909](https://arxiv.org/abs/1504.04909) (2015)
34. Nagy, D., Villaggi, L., Benjamin, D.: Generative urban design: integration of financial and energy design goals in a generative design workflow for residential neighborhood layout. In: Symposium on Simulation for Architecture and Urban Design (2018)

35. Nagy, D., et al.: Project discover: an application of generative design for architectural space planning. In: Proceedings of the Symposium on Simulation for Architecture and Urban Design, p. 7. Society for Computer Simulation International (2017)
36. Nagy, D., Villaggi, L., Zhao, D., Benjamin, D.: Beyond heuristics: a novel design space model for generative space planning in architecture (2017)
37. Panichella, A.: An adaptive evolutionary algorithm based on non-euclidean geometry for many-objective optimization. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 595–603 (2019)
38. Pierrot, T., Richard, G., Beguir, K., Cully, A.: Multi-objective quality diversity optimization. arXiv preprint [arXiv:2202.03057](https://arxiv.org/abs/2202.03057) (2022)
39. Pugh, J.K., Soros, L.B., Stanley, K.O.: Quality diversity: a new frontier for evolutionary computation. *Frontiers Robot. AI* **3**, 40 (2016)
40. Tjanaka, B., Fontaine, M.C., Zhang, Y., Sommerer, S., Dennler, N., Nikolaidis, S.: pyribs: a bare-bones python library for quality diversity optimization (2021). <https://github.com/icaros-usc/pyribs>
41. Turrin, M., Von Buelow, P., Stouffs, R.: Design explorations of performance driven geometry in architectural design using parametric modeling and genetic algorithms. *Adv. Eng. Inform.* **25**(4), 656–675 (2011)
42. Zitzler, E., Deb, K., Thiele, L.: Comparison of multiobjective evolutionary algorithms: empirical results. *Evol. Comput.* **8**(2), 173–195 (2000)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

