



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/210793/>

Version: Accepted Version

Proceedings Paper:

Buchanan Berumen, Edgar, Hickinbotham, Simon John, Dubey, Rahul et al. (2024) A quality diversity study in EvoDevo processes for engineering design. In: 2024 IEEE Congress on Evolutionary Computation (CEC). IEEE World Congress on Computational Intelligence 2024, 30 Jun - 05 Jul 2024 IEEE, JPN.

<https://doi.org/10.1109/CEC60901.2024.10612076>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A quality diversity study in EvoDevo processes for engineering design*

*Note: Sub-titles are not captured in Xplore and should not be used

Anonymous Authors

Abstract—For a long time engineering design has relied on human engineers manually crafting and refining designs using their expertise and experience. In Bio-inspired Evolutionary Development (EvoDevo), generative algorithms are employed to investigate a broader design space that may go beyond what human engineers have considered. Previous literature has demonstrated the use of quality and diversity (QD) algorithms in evolutionary approaches to drive the process to better quality solutions. This paper provides a study to understand the effects of using QD algorithms in EvoDevo processes for engineering design. This paper also analyses the impact of using different behavioural characterisations (BC) in the performance of the quality of the solutions found. The results demonstrates that quality and diversity algorithms can find better solutions than other EAs for engineering design problems. It was also found that the characterisation of the BC is important to get the best results.

Index Terms—evodevo, generative design, structural engineering, quality diversity, neural networks

I. INTRODUCTION

Evolutionary Developmental (*EvoDevo*) algorithms add a developmental component (*Devo*) to traditional evolutionary algorithms (*Evo*). Recent work has utilised the *EvoDevo* process in the engineering context to produce sets of engineering designs. Previous work [1] demonstrated the design of bracket components using developmental or growth processes. Previous work [1] demonstrated the use of *EvoDevo* to improve the design of bridge-like structures.

In contrast to traditional *Evo* processes, the population is a set of “designers” that optimize the same starting structure, or seedling, during the growth process. The designer in this context takes the form of an artificial gene regulatory network (GRN). The structure is composed of cells, and the same copy of the GRN resides in each cell of the structure. The GRN takes as an input the state of a cell itself and changes to the cell are taken from the GRN output (s). These changes occur at every *Developmental* step. After the *Developmental* process finishes, the fitness information is feedback to the *Evo* process to evolve the GRNs. This evolved GRN can be subsequently used for a similar problem with different conditions and the growth process will create potentially a new design for the new conditions.

In previous work, the seedling was hand-designed by the user. However, recent work found that the shape of the seedlings has an influence on the performance of the *EvoDevo*

process where seedlings with relatively good performance could lead to sub-optimal designs and sub-optimal seedlings could yield subjectively better designs and this lead to the question regarding how to design the most appropriate seedling for a given class of tasks. The authors proposed three solutions and one of them was to use Quality Diversity algorithms (QD) [2] to design the seedlings.

This paper explores this same problem but from a different perspective. QD algorithms are used to evolve the final designs with the expectation that these algorithms would explore a broader space and with this avoid local optima regions. This paper also questions if the definition of Behavioural Characterisation (BC) (size and features) at different levels of abstraction has an influence on the performance of QD algorithms. In addition it explores the effectiveness of considering negative quality regions in the BC.

The motivation for using QD within an EvoDevo algorithm for design is similar to that for using QD in evolutionary robotics applications [3]–[5]: the space of possible designs is very large and highly heterogeneous, so driving diversity in the population is a potential means to overcome local optima. Thus, it is important to demonstrate that QD has a beneficial contribution to evolution and then to show how such an approach can be configured and tested. In particular, the measure of BC has many potential formulations and it needs to be shown how to select an appropriate one. Finally, diversity measures may push a significant proportion of the population into an unproductive region of the design space, and mechanisms for limiting this potential outcome need to be explored.

The null hypotheses to be tested in this paper are the following:

- 1) The definition of the BCs design has no statistically significant influence on the results.
- 2) The consideration of quality-only regions when designing BC has no significant difference when compared with regions that include sub-optimal quality regions.
- 3) The performance of QD algorithms is not statistically significantly different than a multi-objective algorithm when applied to *EvoDevo* processes.

The key contributions of this paper are two-fold. 1. This is the first implementation of QD algorithms inside *EvoDevo* processes, and with this, the process exhibits better performance for the problem given in this paper (in this case a Warren truss). 2. The demonstration and comparison of

TABLE I: Experimental parameters. * This grid size is used for the Devo BC as one of the dimensions cannot be higher than 10 due to the nature of that dimension.

Parameter	Value
Emitters	4
Grid size	50x50 (250x10*)
Sigma	0.8
Generations	100
Population	100
Devo steps	10
Sample size	15

performance with different BC when applied in *EvoDevo* and the results in this paper indicate that it is important to consider the information from the growth process itself to find better solutions.

II. EXPERIMENTAL METHODOLOGY

The Evolutionary Developmental algorithm *EvoDevo* used is similar to [1] and parameters can be found in table I. The structure explored in this paper is a Warren truss as illustrated in figure 1. The vertices in the bottom are fully constrained to prevent movement, the left vertex has a pinned support and the vertex on the right has roller support, and a load of 300N is applied in the middle bottom vertex.




Fig. 1: Seedling. The bottom vertices of the seedling are fixed, the left vertex is pinned and the right vertex is supported by a roller. The load of 300N is applied to the vertex in the middle of the structure.

The evolutionary algorithm evolves the GRNs in this case implemented as evolvable neural networks. These GRNs regulate the growth of each cell, in this case a triangle, in the structure. The structure grows from a starting structure (seedling) to the final design after a pre-determined number of growth steps. The GRN is able to change the cross-sectional area of the edges of the triangle and the locations of the vertices that are not loaded or supported. All the edges in the seedling have a starting cross-sectional area of 10mm^2 and the change range is $[2.5, 10]\text{mm}^2$. Each vertex can be moved up to 1mm in the x and y directions for each of the total 10 growth steps.

In order to test hypothesis 3, two algorithms are evaluated: multi-objective (MO), using the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [6] algorithm; and the QD algorithm Covariance Matrix Adaptation MAP-Elites (ME) [7].

For the first algorithm, MO, the two objectives to minimise are volume and deflection at the final design of the growth process. Assuming that the structure has m edges, the volume objective is the sum of the volume for each edge in the structure and is calculated using equation 1:

$$V = \sum_{m \in M} A_m L_m \quad (1)$$

where A and L represent the cross-sectional area and length respectively for member m . Assuming that there are n vertices, the deflection objective measures the maximum deflection of any vertex in the structure as shown in equation 2 when a load is applied to the structure. The magnitude of deflection is estimated using Finite Element Analysis (FEA) software [1].

$$D = \max[d_0, \dots, d_n] \quad (2)$$

For the second algorithm, ME, the quality metric and behaviour characterization need to be defined. The quality metric is the normalised equal-weighted sum of the two objectives used in MO: volume and deflection. The normalization is relative to the seedling. The behaviour characterization (BC) is a descriptor used to quantify and incentivise the diversity of solutions in QD algorithms. In this paper, four different two-dimensional BC are analysed for each layer of abstraction to test hypothesis 2: genotypic, phenotypic, developmental and performance. The BC are described next, values can be found in table II and examples of landscapes are shown in figure 2.

- *Genotypic (Gen)*: This BC captures information from the GRN in this case a feed-forward fully-connected neural network. The first dimension is the number of active connections out of 84 total connections and the second dimension is the weighted sum of all the active connections.
- *Phenotypic (Phe)*: This BC takes information from the structure itself where the BC is the mean concentration of strain energies in the x and y axis estimated by the FEA across all the edges.
- *Developmental (Devo)*: This BC concerns the growth process. The first dimension, referred to as reward, is the addition of the quality scores at each step and relative across the previous growth step. The second dimension indicates the growth step where the highest quality score was found during the Devo process.
- *Performance (Perf)*: This layer is the objectives used for MO: volume and deflection.

For the BCs, different values were chosen to test hypothesis 2 with the objective of identifying whether the range of values used as BC has an influence on the performance of the algorithm. For this, three different sets of values were evaluated for each BC with the exception of Gen and Phe

TABLE II: Behavioural characterisations (BC)

BC	Features	Dims 1	Dims 2	Dims 3
Gen	Connections number	[0, 84]	[26, 62]	-
	Aggregated weights	[-84, 84]	[-60, 60]	-
Phe	Mean-strain-x	[0, 32]	-	-
	Mean-strain-y	[0, 9]	-	-
Perf	Volume (mm^3)	[500, 2600]	[500, 1265]	[600, 1265]
	Deflection (mm)	[0.1, 1.1]	[0.1, 0.808]	[0.2, 0.808]
Devo	Aggregated reward	[-0.5, 0.5]	[0.0, 0.5]	[0, 0.45]
	Peak step fitness	[0, 9]	[0, 9]	[0, 9]

where the size cannot be reduced further without excluding positive values. More information on this and the reasoning for selecting the specific values shown in table II are described in section III-B.

The results shown in this paper are analysed with three different metrics: improvement relative to the seedling, the Pareto-agnostic hypervolume and the QD-metric. The improvement relative to seedling is the quality metric described and converted to a percentage for human readability. The hypervolume metric (HV) [8] is used to assess the Pareto front produced by each strategy where high values of HV represent larger Pareto fronts and the seedling is used as the reference point. The QD score [9] measures the total addition of elite fitness across the entire archive of solutions.

It is important to clarify that this paper follows the 'A-B' convention to describe the strategy used to evolve the designs where 'A' stands for the algorithm used and 'B' is the BC used by the algorithm. For example, an experiment using the ME algorithm with the 'Gen' BC is shown as 'ME-Gen'. In addition, every single strategy can evaluate the QD-score relative to other BC landscapes when this happens the algorithm is not mentioned and only the BC is shown (i.e. 'Gen'), this information is relevant in section III-D.

Two different statistical tests are used in this paper. The Vargha-Delaney A-test [10] was used for the CMA-ME parameter calibration of the sigma, grid size and number of emitters parameters (robustness analysis) and to identify the number of samples (15) required to minimize the effects of uncertainty and to increase confidence when analysing the results (consistency analysis). Robustness analysis and consistency analysis can be found in the supplementary material¹

The Mann-Whitney U-test [11] is used to assess the hypothesis that the samples from two groups are not mutually independent. For this, a three-score system is used where * represents $p < 0.05$, ** represent $p < 0.01$, *** represents $p < 0.001$ and p is the probability.

The code and data can be found on-line².

III. RESULTS

This section first provides a qualitative analysis of the four different BC landscapes which describes their size and shape. Then, the effects of selecting different range values for the BCs using the HV as a metric are explored and the best strategies

are highlighted. Lastly, the performance using the QD score of each strategy with each combination of BC is analysed to identify the strategy that performs best across the different BCs.

A. Qualitative analysis of the BC landscapes

The landscapes for each BC shown in figure 2 are taken from one of the samples of evolving with MO, ME-Gen, ME-Phe, ME-Devo and ME-Perf.

The Gen BC is shown in figure 2a and d. The first thing to notice is that ME-Gen explores a wider space than MO. However, ME-Gen still only explores a small portion of the available search space. In addition, it appears that MO converged to a global optima whereas this is not the case for ME-Gen where it appears there are multiple local optima exist in the quality domain. This might be because this BC is far removed from the quality of the structure, therefore ME-Gen is exploring regions of the Gen landscape that might not have a great contribution to the quality.

The shape of the Phe-BC landscape (figure 2b and e) is bell-like and an optimal gradient is located roughly at the middle of the x-axis therefore highlighting that symmetrical strain energy distributions seem to provide the better quality solutions. In addition, strain energies located closer to 0 in the y-axis provide the worst results. However, when comparing the landscapes between MO and ME-Phe, it seems that there is no great difference with the exception that ME-Phe finds better quality solutions in some regions of the landscape.

The best solutions for the Devo BC (figure 2c and f) are located in the upper-right region of the diversity map and the shape is curve-like. This region maximizes the reward and makes the growth process more efficient as the best solution is found in the last developmental step. The curve-like shape illustrates that some rewards are unreachable with a low number of developmental steps. Roughly half the area of this landscape contains negative rewards. ME-Devo seems to exhibit a thicker curve than MO meaning the former is exploring more regions of the landscape.

The Perf BC landscape (figure 2g and h) demonstrates the regions of optimality and a Pareto front can be visualized. This landscape contains negative values at the top-right corner. ME-Perf explores more regions in the landscape than MO this is more apparent in the top-right and bottom-left regions of the landscape.

In conclusion, the BC significantly affects the shape of the landscape and the distribution of quality designs. The

¹Supplementary material will be made available after paper-acceptance

²Code will be made available after paper-acceptance

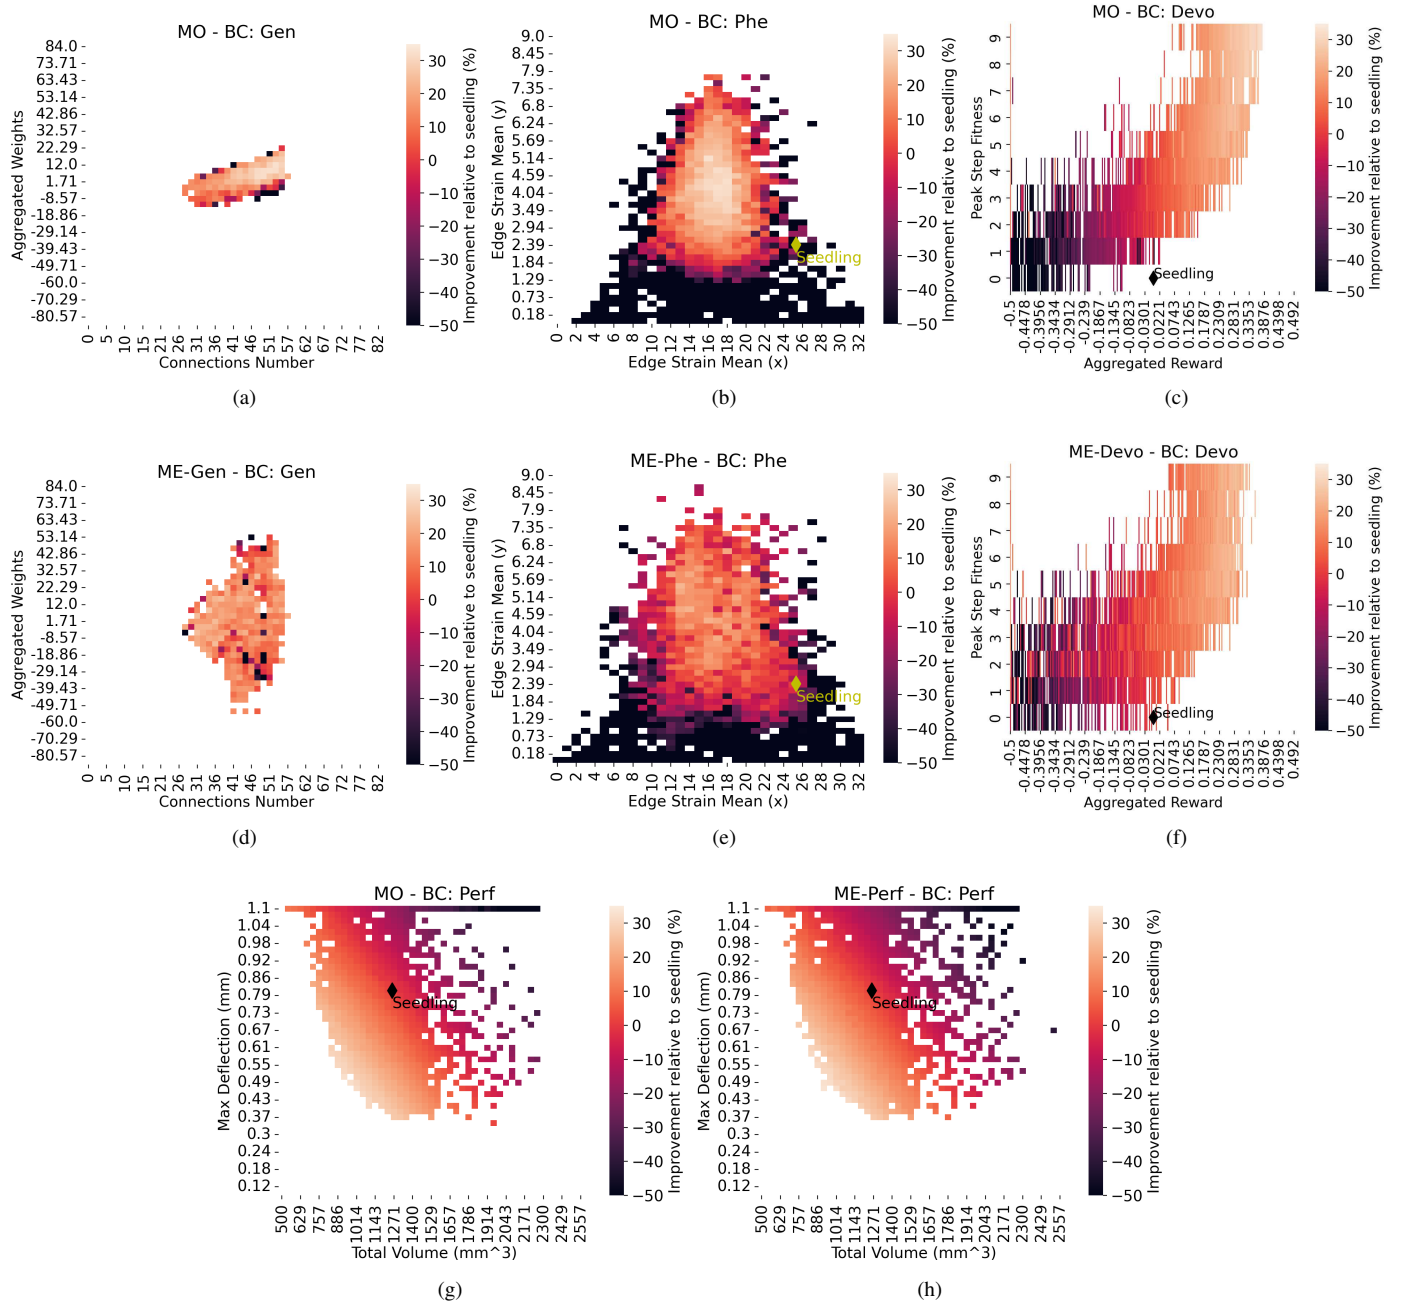


Fig. 2: Example of the four BCs landscapes: Gen, Phe, Devo and Perf. Two different samples are shown for each evolutionary strategy: MO and ME. The BC defines the shape and distribution of the quality of solutions. ME explore a wider region of the space.

ME strategy explores a wider space than the MO strategy. However, a question arises, *are the negative quality and empty regions worth considering in the BC to find the best designs when using QD algorithms?* This is explored in the next section.

B. Analysing the BC ranges

Different values for BC ranges, table II, are analyzed to identify if the negative quality regions are worth encapsulating

when designing the BC. The initial values, dims 1, were chosen for the BCs for each layer to cover most of the available landscape including positive and negative quality regions. The values for dims 2 and dims 3 were chosen in regions where positive quality designs were aggregated. It is important to highlight that the grid size is not changed, therefore if the range is reduced the granularity is increased for that region.

Figure 3a compares the QD score for each strategy at the

last generation. There is no statistical difference between ME-Gen 1 and ME-Gen 2. This demonstrates that the BC range reduction has no impact on the performance of this specific BC.

ME-Devo 1 and ME-Devo 2 are *** significantly different. This demonstrates that removing half of the BC landscape containing negative-quality solutions has a positive impact on the performance of the algorithm. There is no difference between ME-Devo 2 and ME-Devo 3.

ME-Perf 1 and ME-Perf 3 are * significantly different. The improvement is lower than ME-Devo and this might related to the proportion of negative-quality regions removed from the range for BC where the amount of negative regions is subjectively higher for ME-Devo than ME-Perf.

Figure 3a also illustrates that ME-Devo 2 and ME-Perf 2 are ** significantly different than MO, therefore this rejects hypothesis 3 and demonstrates that this QD algorithm finds better solutions for the engineering problem shown in this paper.

Figure 3b shows the convergence of the HV metric over generations for selected strategies. The convergence speed is correlated to the final HV values where ME-Perf 3 finds the best HV values.

It is therefore possible to conclude that the exclusion of negative quality regions improves the performance of the algorithm for ME-Devo and ME-Perf. This is effective up to a limit after that it becomes insignificant, hereby this rejects hypothesis 2. The next section explores the explorative effectiveness of each strategy using the QD score.

C. QD score comparison

In this section, the QD score is compared for each combination of strategy and BC landscape. This is to identify the strategy that has the highest QD scores across all the BC landscapes.

Figure 4 and table III illustrate in different ways the QD-score for each combination of strategy and BC. In table III, the median values are shown for each strategy with each BC. The blue cells highlight the best strategy for each given BC landscape. In this case, ME-Devo has the highest QD values for Gen, Phe and Devo BCs. ME-Pef has the highest QD values for the Pef BC landscapes. In other words, even though the ME-Devo is designed to explore the Devo BC landscape this strategy not only has the best QD-scores for this BC but also this strategy has the best scores for the Gen and Phe BC landscapes. This means that ME-Devo is able to explore more space and find better elites in 3 out of the 4 BCs.

The previous statement can also be visualised with figure 4. In the barplot in figure 4a, the highest values are achieved by ME-Devo in 3 out of the 4 BC landscapes and in similar way ME-Devo is able to get close to the four regions in the radar plot in figure 4b. Please note that the values of the radar plot are normalised for visualisation purposes

In conclusion, hypothesis 1 is rejected and the BC has an impact on the exploration where the ME strategy with Devo


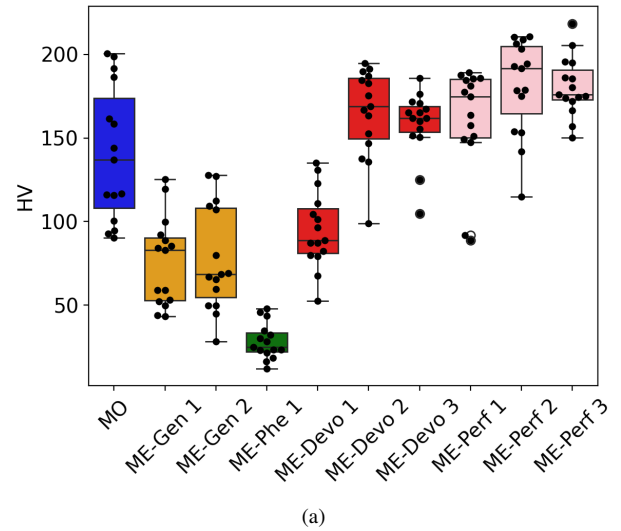


Fig. 3: Results of comparing MO and ME with different BCs using HV as metric. Figure (a) shows the HV values at the last generation and figure (b) shows the convergence of the HV over time for selected strategies (ME-Devo and ME-Perf). The ME strategy is finding better values for HV than MO for ME-Devo 3 and ME-Perf 2. Designing the BC with positive quality solutions provides better solutions for ME-Devo 3 and ME-Perf 2. Convergence is faster for ME-Devo 3 and ME-Perf 3 than MO

BC seems to provide a better exploration power across the rest of BC landscapes.

D. Growth process from evolved designs

Examples of the best quality structures found by each strategy can be visualised in figure 5. Most strategies seem to reinforce the top arch and the edges connecting to the load and this is because these edges experience the highest forces in the structure and hence the highest deflection. There are also elements of symmetry and some edges are set vertically (figure 5e).

TABLE III: Median value for each strategy measure against each BC landscape. The QD score for the best-performing strategy is highlighted in blue. A u-test is performed between the first and second best performing QD-score and the results are shown in parentheses.

		BC measured with			
		Gen	Phe	Devo	Perf
Strategy evolved with	MO	2358	-16367	17810	13121
	ME-Gen	3399	-22636	10967	2793
	ME-Phe	-6552	-40422	2443	-502
	ME-Devo	9606 (***)	-15751 (*)	24647 (***)	14902
	ME-Perf	13121	-19902	20714	19072 (***)

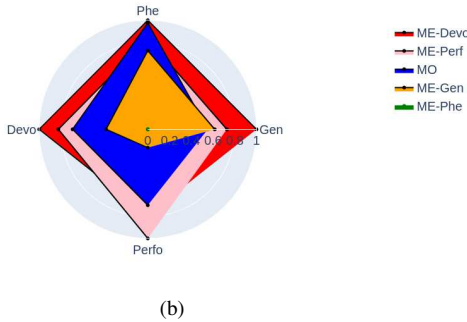
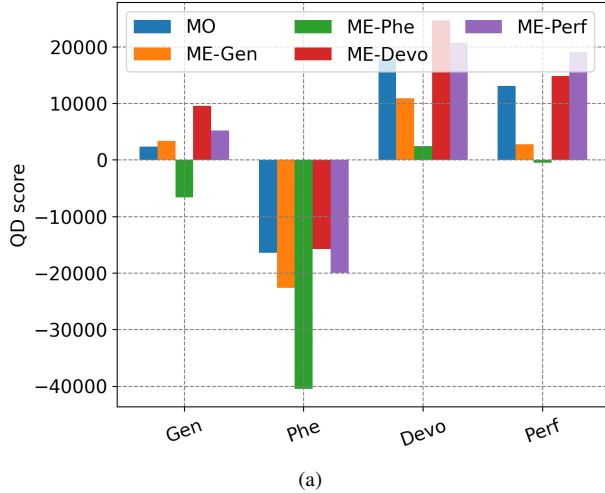


Fig. 4: Real value QD score bar plots (a) and normalised QD score radar plot (b). The ME-Devo strategy has the highest values across the Gen, Phe and Devo BCs. ME-Perf has the highest value for Perf

Figure 6 illustrates the growth of the structures shown in figure 5. Figure 6a shows the improvement at each step relative to the seedling. All strategies appear to have a consistent increase in performance. ME-Perf and ME-Phe had a miniscule drop in improvement at the last step. This does not occur in ME-Devo where one of the attributes of BC identifies in which step during the growth process the highest quality was achieved.

Figure 6b presents the improvement relative to the previous step for the growth process for each structure. As the growth from the structure matures the amount of improvement at each step diminishes. The drop in improvement at each step for

ME-Perf is not as consistent as ME-Devo.

In conclusion, the results indicate that ME-Devo might create GRNs with more efficient and consistent growth processes and this might enhance exploration as shown in section III-D.

IV. CONCLUSION

The starting conditions, that is the seedling, in this Evolutionary Development (EvoDevo) algorithm used in generative design for engineering structures can drive the process to sub-optimal solutions as shown in [1]. This paper proposed to make use of QD algorithms to prevent the EvoDevo process from stagnating in a sub-optimal solution by promoting diversity in designs. For this, four different behaviour characterisation (BC)s are explored as options for the QD algorithm. Each BC has properties from different layers in the EvoDevo process.

The main conclusions of this paper are summarised as follows: 1) The QD algorithm used in this paper (CME-ME [7]) provides better designs than NSGA-II for the problem considered (Warren truss). 2) The BC related to the growth process seems to provide better results than the rest. This might show that it is important to consider the information during the growth process to produce the best results in an EvoDevo process.

This paper also shows that the design BC should contain as much of the positive quality region for the algorithm to exhibit its best performance. Also, it is important to note that the list of BC used in this paper is not exhaustive and better BC might exist that improve the performance of the algorithms. Regardless, the question still prevails. *How to design the appropriate BC and the range values for this BC?*

For this, three options can be used from literature:

- 1) Manual calibration: similar to the work presented in this paper, the BCs can be manually crafted. However, this approach is time-consuming, and requires knowledge from the designer and the BC design might not provide the best global results.
- 2) Automatic BC definition [12], [13]: this type of algorithm uses the knowledge accumulated during the evolutionary process to construct BCs on the fly with unsupervised dimensions reduction techniques. However, this approach requires prior data to be gathered and this could be computationally expensive
- 3) Human assisted QD algorithms [14]: in this approach, a human subjectively adds metrics to the solutions according to the similarity of solutions. This approach

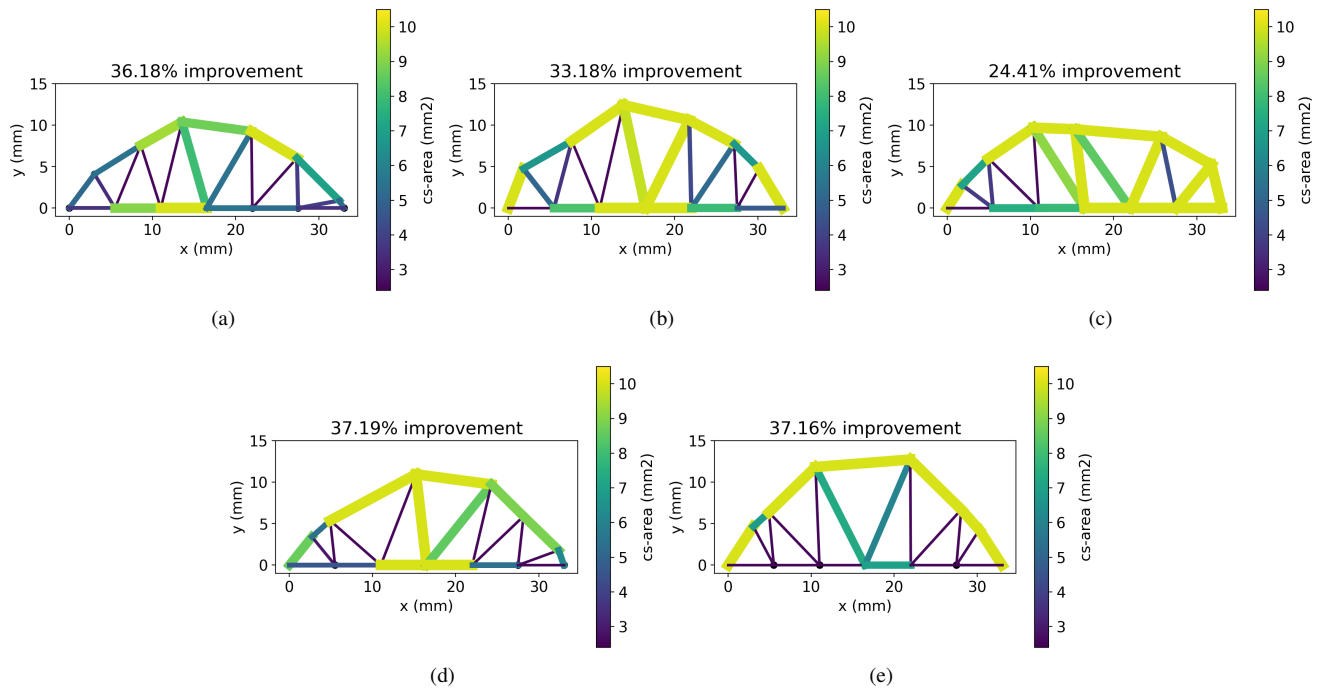


Fig. 5: Best solutions found by each strategy: MO (a), ME-Gen (b), ME-Phe (c), ME-Devo (d) and ME-Perf (e). The evolved structure exhibits the traits of reinforced arcs and edges connecting to the load.

relies on the knowledge and time of the person making the decisions.

Regardless, the evolved GRN should be generalizable for different problems. The selected QD algorithm should accommodate these needs.


Further work will explore the previous options to design BC will be analysed in the context of EvoDevo. In addition, the autonomous design of seedling will be explore to enhance the final results evolved.

ACKNOWLEDGMENT


Information was removed to anonymize the paper.

REFERENCES

- [1] A. Anonymous.
- [2] J. K. Pugh, L. B. Soros, and K. O. Stanley, "Quality diversity: A new frontier for evolutionary computation," *Frontiers in Robotics and AI*, vol. 3, p. 40, 2016.
- [3] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.
- [4] J. Nordmoen, F. Veenstra, K. O. Ellefsen, and K. Glette, "Quality and diversity in evolutionary modular robotics," in *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2020, pp. 2109–2116.
- [5] L. Grillotti and A. Cully, "Unsupervised behavior discovery with quality-diversity optimization," *IEEE Transactions on Evolutionary Computation*, vol. 26, no. 6, pp. 1539–1552, 2022.
- [6] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE transactions on evolutionary computation*, vol. 6, no. 2, pp. 182–197, 2002.
- [7] M. C. Fontaine, J. Togelius, S. Nikolaidis, and A. K. Hoover, "Covariance matrix adaptation for the rapid illumination of behavior space," in *Proceedings of the 2020 genetic and evolutionary computation conference*, 2020, pp. 94–102.
- [8] A. P. Guerreiro, C. M. Fonseca, and L. Paquete, "The hypervolume indicator: Computational problems and algorithms," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–42, 2021.
- [9] J. K. Pugh, L. B. Soros, P. A. Szerlip, and K. O. Stanley, "Confronting the challenge of quality diversity," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 2015, pp. 967–974.
- [10] A. Vargha and H. D. Delaney, "A critique and improvement of the cl common language effect size statistics of mcgraw and wong," *Journal of Educational and Behavioral Statistics*, vol. 25, no. 2, pp. 101–132, 2000.
- [11] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The annals of mathematical statistics*, pp. 50–60, 1947.
- [12] A. Cully, "Autonomous skill discovery with quality-diversity and unsupervised descriptors," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2019, pp. 81–89.
- [13] D. M. Bossens, J.-B. Mouret, and D. Tarapore, "Learning behaviour-performance maps with meta-evolution," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, 2020, pp. 49–57.
- [14] L. Ding, J. Zhang, J. Clune, L. Spector, and J. Lehman, "Quality diversity through human feedback," *arXiv preprint arXiv:2310.12103*, 2023.



(a)



(b)

Fig. 6: Growth process for the structures shown in figure 5. Figure (a) shows the progression of each structure relative to the seedling. All strategies appear to have continuous growth except by ME-Perf and ME-Phe which suffer a small drop at the last step. Figure (b) shows the improvement relative to the previous steps. The improvement diminishes over time as the growth process converges. Lines of interest are highlighted. ME-Devo appears to have the most consistent progression.