Universiteit Leiden
The Netherlands

Multi-objective mixed-integer evolutionary algorithms for building spatial design

Blom, K. van der

Citation

Blom, K. van der. (2019, December 11). *Multi-objective mixed-integer evolutionary algorithms for building spatial design*. Retrieved from https://hdl.handle.net/1887/81789

Version:	Publisher's Version
License:	Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden
Downloaded from:	https://hdl.handle.net/1887/81789

Note: To cite this publication please use the final published version (if applicable).

Cover Page

Universiteit Leiden

Leiden University Repository

The handle http://hdl.handle.net/1887/81789 holds various files of this Leiden University dissertation.

Author: Blom, K. van der
Title: Multi-objective mixed-integer evolutionary algorithms for building spatial design
Issue Date: 2019-12-11

Chapter 10

Conclusions

This thesis investigated many different aspects of multi-objective mixed-integer optimisation problems related to the built environment. In this conclusion, first the contributions of each chapter are briefly summarised, together with answers to their corresponding research questions. After that, paths to future discoveries are discussed that opened up as a result of this work.

10.1 Summary

Chapter 1 gave a high level introduction to the topic of study in this thesis.

Chapter 2 then provided a starting point to this study by introducing key concepts used throughout the thesis. Specifically, it covered the basics of optimisation for a single objective, as well as for multiple objectives. Evolutionary computation was introduced next as an approach to solve such optimisation problems. Finally, it covered the main application domain of this work: building spatial design.

Chapter 3 introduced a representation for building spatial designs in response to the first research question:

> *RQ1 How can elements of the solution space be represented?*

The introduction of the supercube representation makes it possible to encode an arbitrary number of spaces, and to control their dimensions. To be able to check the feasibility of building spatial designs, a number of constraints on the representation were introduced as well. These constraints are polynomial expressions directly formulated on the binary and continuous decision variables. Hence they can be exactly computed in a straightforward manner, and, at least in principle, can be used in equation based solvers.

Chapters 4 and 5 looked at how to manage the constraints used in the supercube in order to answer the second research question:

RQ2 How can the discovery of feasible designs be ensured?

Penalty functions that act on constraint violating solutions are shown to aid the search process in Chapter 4. They are, however, also found not to scale sufficiently for larger supercube sizes. Chapter 5 presents problem specific initialisation and mutation operators that navigate only the feasible space, and thus do not suffer from the larger infeasible regions present in larger sized supercubes.

Chapter 6 investigated the benefits of a local search strategy with the third research question in mind:

RQ3 How can local search contribute to the improvement of solutions found during global search in a multi-objective setting?

Hypervolume indicator gradient ascent multi-objective optimisation (HIGA-MO) is shown to work for a real world problem, and with numerically approximated gradients. Even so, the number of evaluations needed for numerical approximation of the considered high dimensional problem, prevents it from outperforming the considered evolutionary algorithm (SMS-EMOA-SC). This means that as a local search component HIGA-MO would likely outperform SMS-EMOA-SC at some point due to the convergence properties of the hypervolume gradient, but the number of evaluations needed would be impractical.

Chapter 7 evaluated the value of the data that is generated during optimisation based on the fourth research question:

RQ4 What can be learned about building spatial design from the optimisation process?

Based on optimisation data, key properties of high quality building spatial designs can be learned, by automatically extracting design rules from the data. These design rules can then be communicated to design experts to prove to them that the discovered solutions are trustworthy, since they typically correspond to well-known design rules used by experts.

Chapter 8 explored the possibilities of a general mixed-integer algorithm for multi-objective optimisation to answer the fifth research question:

> *RQ5 How can a generally applicable multi-objective mixed-integer algorithm be developed?*

Promising first steps are made towards a general purpose multi-objective mixed-integer evolution strategy (MOMIES). MOMIES is shown to perform well in practice, and recombination has been discovered to be more valuable in multi-objective optimisation than previously believed, although the reasons are still unclear. Step size adaptation for the multi-objective case requires more thought, as the current methods still show erratic behaviour and are not capable of reliably tracking optimal mutation step sizes.

Chapter 9 presented various applications of the developed algorithms in response to the sixth research question:

> *RQ6 How applicable are the developed algorithms to real world problems?*

In application, SMS-EMOA-SC proved to be a useful component in a hybrid algorithm that also explored which supercube configurations were most interesting to work with. Furthermore, SMS-EMOA-SC was shown to be a good optimiser, but the settings and objectives considered in the earlier academic cases do not always align with the practicalities of the real world. In the nominal discrete space of structural design optimisation MOMIES proved to be a highly effective algorithm.

In summary, this thesis has made tools available for the multi-objective optimisation of building spatial designs. To this end a representation has been introduced, as well as specialised operators for the building spatial design problem. Furthermore, the potential benefits of a memetic algorithm, where stochastic global search is combined with deterministic local search, has been explored. Practical utilisation of these methods has also been studied. Both by analysis of the optimisation data to help human designers, as well as through real world studies of the algorithms. Finally, a general algorithm, not restricted to the building spatial design problem, has also been developed.

10.2 Future Work

Although considerable progress has been made in the context of this thesis, much more work remains. Some perspectives for this future research are included here, and are briefly discussed.

New or modified representations that contain fewer infeasible or duplicate solutions than the supercube representation would be very helpful in tackling larger building designs, and in speeding up the search process in general. Chapter 3 already suggested a number of possible directions, such as the use of an integer representation instead of the current binary one. However, as also mentioned before, caution is needed when considering what is or is not a duplicate since this is largely problem dependent. With a new representation it is also likely that new operators are needed. As such, representations and operators should be considered in concert to avoid poorly compatible combinations.

Local search with the hypervolume indicator (HVI) gradient successfully improved the Pareto front approximation found during global search. It was, however, not able to outperform SMS-EMOA-SC. The primary shortcoming is found in the many evaluations needed to approximate the gradient for the considered high dimensional problem. To overcome this, methods are needed that reduce the computational costs of an algorithm using the HVI gradient. One interesting direction is the use stochastic gradient descent methods, which is currently used with success in the optimisation of weights in deep neural networks . These networks, like building spatial designs (in the supercube representation), require the adaptation of a large number of continuous variables for a given discrete structure. Another option would be the use of surrogate models that can be learned from previously obtained evaluation data. These would both allow cheap numerical approximation of the HVI gradient on the model, rather than the true function, and enable the navigation of discrete space. Another way to reduce computation costs is to move fewer points based on the HVI gradient. Only nondominated points could be considered, rather than all points, for instance.

Since differences in performance are larger between discrete subspaces than within them, median attainment curves struggle to adequately present the performance differences for algorithms that sometimes end up in one discrete subspace, and sometimes in another. New multi-objective performance measures for mixed-integer problems that take these things into account are thus a must.

Integrating the insights that can be gained from optimisation data into the optimisation process has the potential to improve, and possibly speed up the optimisation process. These insights could be used to only perform expensive evaluations for candidate solutions that look at least somewhat promising, instead of for all candidates. Another new direction in this area could be the use of learned design qualities in guiding the variation operators. By biasing the operators to directions that are likely to produce improved designs, the search could be speeded up. Naturally, the risk is that unexpected design alternatives will not be found. How to balance this bias with the needed exploratory components will therefore require careful consideration.

Finally, another promising path for future work is continue the development and analysis of MOMIES, the general purpose multi-objective mixed-integer optimisation technique that resulted from this thesis. To be a full fledged general optimiser, MOMIES will also have to include a constraint handling component capable of operating in the multi-objective mixed-integer environment. In addition, step size adaptation has to be tailored for multi-objective optimisation to improve its effectiveness. This is particularly true for the integer and nominal discrete adaptation mechanisms, since they were shown to have particularly irregular behaviour. Recombination was shown to be surprisingly effective for the multi-objective case, but it remains unclear why. Investigating this is especially interesting since it was previously shown that the value of recombination can be questioned in the multi-objective setting [101]. A deeper understanding could elucidate when it is of value to use recombination, and when to avoid it.