Author: Reehuis, Edgar
Title: Guiding evolutionary search towards innovative solutions
Issue Date: 2013-12-17

1

Introduction

In *optimization*, the goal is to find an input for which a given mathematical function produces the *optimal* output value. The optimal output value is either the minimal or maximal value that the function can return, a choice that is specified in advance. A real-world engineering design task, such as shaping an airfoil, can be viewed as a *quality-optimization problem*: The input is the proposed design and the output its assessed quality. A design is then referred to as a *solution* for the optimization problem.

For optimizing a design task in an automated way, the input has to be re-formulated in such a way, i.e., *encoded*, that it is possible to automatically vary and adjust proposed solutions. A solution is, for instance, represented by a sequence of numbers, from which the actual design is retrieved using a *decoding* scheme. Furthermore, the quality of a design has to be quantified into a numerical score as well. Ideally, an automated pipeline is set up in which the numerical input is fed to a *simulator* that models the represented design's real-world performance in a simulated environment, and approximates its actual quality score. The simulator, accepting solutions as input and producing a quality score as output, effectively takes on the role of a *quality function*.

Given this numerical representation of solutions and the available automated quality assessment, an *optimization algorithm* can be applied that, without human interaction, aims to find an optimal solution to the design task by systematically varying the numerical input. Note that this optimality is restricted by the freedom in varying the designs that is allowed through the used encoding, and relates strictly to the simulated environment and the used quality quantification.

Evolutionary Algorithms (EAs) are methods that approach optimizing a function by using a *population* of, e.g., numerically-represented, solutions that *evolve* towards an optimal configuration. Unlike classical, *deterministic* optimization algorithms, the *stochastic* EAs rely on randomness to drive the optimization process. The to-be-optimized function provides feedback on *solution fitness*, which is used to select the most-fit solutions and thereby steer the random sampling. While the population-based EAs are relatively slow with respect to the required number of function evaluations in approaching an optimum, their moderate robustness to deceptive characterisics of the function landscape and flexibility in applying them using varying types of problem formulations make that they are frequently considered for real-world optimization tasks.

Optimization methods in general, and EAs specifically, are able to assist in uncovering *innovative solutions*, roughly defined as solutions involving unconventional new ways in obtaining good performance. By automatically varying solutions, a different path may be embarked on towards high quality than would be considered by a human engineer guided by his or her education. Nevertheless, even in using a population of solutions, an optimization method most-likely converges to only one optimized solution at a time. Due to their stochastic nature, however, different runs of an EA *are* potentially able to produce alternative solutions.

Exploration-inducing schemes exist that operate on top of optimization algorithms and make them actively target multiple high-performing solutions. We aim for exploration to be performed in a sequential, path-based fashion, allowing for application to scenarios in which optimization needs to be started from a single, fixed initial solution. An additional exploration criterion is to be integrated into the evolutionary search, next to quality, such that multiple high-quality solutions can be found in a single run, assisted by a memory of solutions encountered earlier during the search. Not playing a role in the active exploration, quality optimization forms an important, but supportive part of the overall process, which is therefore termed "search" instead of "optimization".

1.1 · Research Goals and Contribution of this Thesis

The main goal of this work is to present a method, running on top of an underlying optimization algorithm, that increases its likelihood of uncovering innovative solutions. As such, a way of measuring the capability of a method to do so will be presented, based on a definition of innovativeness of solutions in design optimization tasks.

Different exploration criteria are described in literature that are based on a knowledge base of solutions encountered earlier in the search. In this work, these are aimed to be used for inducing sequential exploration in design spaces, adjusting them where necessary. Furthermore, the different exploration criteria are brought together in a uniform naming and notational scheme, to facilitate implementation and comparison with this aim in mind.

After formulating the exploration criteria, an appropriate way of how to integrate such a criterion in the search, next to quality, is to be presented. Lastly, the developed methods are tested on a real-world design optimization task, showing the steps required in doing so, and for verification of their ability of assisting in finding innovative solutions.

1.2 · Thesis Outline

Based on the goals described above, this thesis takes the reader stepwise from a definition of innovativeness of solutions to application of a scheme for assisting in finding innovative solutions to a real-world airfoil optimization task:

- Chapter 2 introduces underlying techniques that will be applied in this work or that will be used as inspiration, discussing Evolutionary Algorithms, *approximation models*, and *surrogate-assisted search*;

- Chapter 3 discusses innovativeness in design optimization, relating it to *novelty* and *interestingness* of solutions. Furthermore, a scheme for measuring performance of a method aimed at assisting in finding innovative solutions is laid out, published on before in [**Reehuis et al., 2013b**] and [**Reehuis et al., 2013c**];

- Chapter 4 formulates criteria for inducing exploration in design spaces, reusing the ideas of novelty and interestingness. An initial comparison of the different criteria is performed on an artificial setup, to select the most-promising measures that will be considered for integration in quality-based search. Parts of this chapter were published before in [**Reehuis et al., 2011**], [**Reehuis et al., 2013a**], [**Reehuis et al., 2013b**], and [**Reehuis et al., 2013c**];

- Chapter 5 discusses different ways of integrating an exploration criterion in quality-based search. An initial comparison is performed on an artificial setup involving, like the real-world airfoil optimization task, the requirement of starting from a single fixed solution, and has large areas with *infeasible*, i.e., intolerable, solutions. This chapter continues on earlier work in [Reehuis et al., 2011], [Reehuis et al., 2013b], and [Reehuis et al., 2013c];

- Chapter 6: Two variants of a method for assisting in finding innovative solutions are applied to a real-world optimization task, described separately in [Reehuis et al., 2013b] and in [Reehuis et al., 2013c], and their performance is statistically compared;

- Chapter 7 summarizes the thesis and provides future research directions.