Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Novelty-guided Restarts for Diverse Solutions in Optimization of Airfoils

Edgar Reehuis, Markus Olhofer, Bernhard Sendhoff, Thomas Bäck

2013

Preprint:

This is an accepted article published in Evolve 2013. The final authenticated version is available online at: https://doi.org/[DOI not available]

Emmerich et al. (Eds.): EVOLVE 2013, July 10-13, 2013
A Bridge between Probability, Set Oriented Optimization, and Evolutionary Computation
©LIACS, Leiden University, nl: ISBN 978-2-87971-118-8, ISSN 2222-9434

Novelty-Guided Restart for Diverse Solutions in Optimization of Airfoils

Edgar Reehuis[a,b], Markus Olhofer[b], Bernhard Sendhoff[b], and Thomas Bäck[a]

[a]Natural Computing Group
LIACS, Leiden University
Niels Bohrweg 1, 2333 CA Leiden, The Netherlands
{ereehuis,baeck}@liacs.nl
http://natcomp.liacs.nl

[b]Honda Research Institute Europe GmbH
Carl-Legien-Straße 30, 63073 Offenbach/Main, Germany
{markus.olhofer,bernhard.sendhoff}@honda-ri.de
http://www.honda-ri.de

Abstract: A restart scheme is proposed that alternates between optimization on quality and novelty with respect to a domain-specific distance measure, aimed at finding a diverse set of high-quality airfoil designs. Compared to the combined outcome of multiple standard optimization runs, the result set of the restart scheme is more diverse and it manages to find solutions of strictly higher quality, while requiring only half of the budget of quality evaluations.

1 Introduction

In design optimization of airfoils one is typically faced with high-dimensional search spaces that are highly constrained by geometric and manufacturing requirements. In certain problem instances, the simulator used for modeling the aerodynamic properties of candidate geometries will only return trustworthy results if the geometry under consideration is located well in *feasible* space, that is, not violating any such constraint. The search therefore needs to be started from a basic solution known to be feasible. As the mesh grid used in simulating is distributed according to this basic geometry, at each stage of the search the candidate solutions have to resemble the basic shape to a critical extent for the simulator to produce meaningful results, or produce results at all.
When trying to find *innovative* solutions in an automated way, we are hindered by the fact that the search needs to be started from basic feasible solutions. Though applying a stochastic optimization method, multiple instances of the search will converge to largely similar optimized solutions. In order to uncover innovative solutions, we have to escape the basin of attraction of the solution that optimization typically converges to when started from a certain basic solution. Moving sufficiently far away, guided by some exploration scheme, has to be done stepwise, as we need to check that we are still in feasible space to be able to start quality-based optimization again.
In our definition, an innovative solution is of better, or tolerably less quality than the *state-of-the-art* in the application domain. Tolerably less means still sufficient for allowing the solution to be adopted. Furthermore, an innovative solution is *novel* compared to the state-of-the-art, i.e., different from it according to a *domain-specific distance measure*. Here the state-of-the-art is the set of the *highest*-quality solutions, generally accepted as such, in the application domain. As it is not straightforward to comprehensively compile this reference set of state-of-the-art solutions, while commonly based on implicit assumptions ex-

isting in the mind of the human engineers, we intend not to search for innovative solutions directly but for a *diverse* set of solutions that are novel when comparing them to each other using the domain-specific distance measure.
Taking inspiration from the work of Cuccu et al. [1], we propose a restart scheme that alternates between a fitness optimization phase and a novelty optimization phase: A standard fitness optimization is carried out until some stopping condition is met, after which the novelty with respect to all solutions sampled in the fitness optimization is maximized, until a second stopping condition is met. The fitness optimization is then restarted from the found solution maximizing novelty, the idea is that this allows for restarting the fitness optimization outside of the basins of attraction of the earlier converged-to optima.
In searching for diverse solutions, *niching* [6] is a commonly used approach, in which candidate solutions are grouped into different niches during optimization based on their domain-specific resemblance. Typically multiple niches are maintained in parallel, requiring the number of niches to be indicated in advance, and requiring an approach of exploring for new niches until this number is met. The alternating scheme can be seen as a form of sequential niching, omitting a number of niches to be indicated, and with emphasis on finding new niches through the repeated novelty maximization phases.
Resuming, we aim for an optimization method to produce a diverse set of high-quality solutions, motivated by the potential of innovative solutions being present amongst these, ultimately to be determined by a human domain-expert. Therefore, to measure performance of a search method an approach is required for scoring its result set by reflecting both the quality of the found solutions and the diversity between the solutions in this set. For each solution we determine the diversity of the *level set* approximation within the result set that the solution's quality score gives rise to. Each such "solution level set" contains the solution itself and all other solutions that have quality equal to or better than the solution's quality score. Requiring a *set diversity measure* that applies the domain-specific distance measure (we use *Solow-Polasky* [7]), we get for each solution a pair of values consisting of its quality value and the diversity score of the level set approximation based on its quality value, which we can use to obtain a two-dimensional *Pareto front* that illustrates the performance of the method.

2 Diversity in Two-dimensional Optimization of Airfoils

To lay out the different possible sources for determining novelty of solutions in the considered problem class of two-dimensional optimization of airfoils, we sketch the path from solution vector to fitness value. For a solution vector that represents a certain airfoil geometry the aerodynamic behavior of that geometry is approximated by *computational fluid dynamics* (CFD) simulation. The simulated *flow field* around the geometry is numerically abstracted into a fitness value, schematically

$$\overset{\text{genotype}}{\mathbf{x}} \xrightarrow[NURBS]{\text{decoding}} \overset{\text{phenotype}}{P(\mathbf{x})} \xrightarrow[CFD]{\text{simulation}} \overset{\text{sim. phenotype}}{\underset{\textit{flow field}}{S(P(\mathbf{x}))}} \xrightarrow{\text{analysis}} \overset{\text{properties}}{\left.\begin{array}{l}\mathrm{pr}_1(S(P(\mathbf{x})))\\ \cdots\\ \mathrm{pr}_k(S(P(\mathbf{x})))\end{array}\right\}} \xrightarrow[heuristics]{\text{aggregation}} \overset{\text{fitness}}{f(\mathbf{x}).} \tag{1}$$

The genotype $\mathbf{x}$, consisting of n input parameters x_i, represents an airfoil geometry that is obtained through interpreting the input parameters as *non-uniform rational B-spline* (NURBS) *control points*, resulting in the phenotype of the solution, denoted $P(\mathbf{x})$. CFD simulation is used to generate a flow field visualization around the phenotype which we address by *simulated phenotype* of $\mathbf{x}$, denoted $S(P(\mathbf{x}))$. From this flow field

approximation, scores for pre-defined properties $\mathrm{pr}_j(S(P(\mathbf{x})))$ are determined, which are then aggregated into a single fitness value according to some heuristic weight vector $\mathbf{w}$,

$$f(\mathbf{x}) = \sum_j w_j \mathrm{pr}_j(S(P(\mathbf{x}))). \tag{2}$$

From the different available characterizations we determine novelty based on the two-dimensional phenotypes, as these allow for a relatively simple domain-specific distance measure. The two phenotypes are sampled equidistantly along the spline, and each sample is again divided in two sets of points, one set representing the *suction* side (side featuring higher air speed) and one set representing the *pressure* side (side featuring lower air speed) per airfoil. The *summed-up Hausdorff distance* [4] is then applied on the suction sets and the pressure sets. It is defined as

$$\mathrm{d}_{\Sigma\mathrm{H}}(A, B) = \sum_{\mathbf{a}\in A} \min_{\mathbf{b}\in B} \|\mathbf{a} - \mathbf{b}\| + \sum_{\mathbf{b}\in B} \min_{\mathbf{a}\in A} \|\mathbf{a} - \mathbf{b}\|, \tag{3}$$

that is, for each two-dimensional point $\mathbf{a} = (a_1, a_2) \in A$ calculate the distance to the closest point $\mathbf{b} \in B$, and vice versa, and sum all these distances up. Then, to determine the difference between the two solutions $\mathbf{x}_1, \mathbf{x}_2$, we add up the summed-up Hausdorff distance between the suction side samples and the distance between the pressure side samples,

$$\begin{aligned} \mathrm{d}_{\mathrm{airfoil}}(\mathbf{x}_1, \mathbf{x}_2) \;=\; & \mathrm{d}_{\Sigma\mathrm{H}}(\mathrm{suction}(P(\mathbf{x}_1)),\ \mathrm{suction}(P(\mathbf{x}_2))) \quad + \\ & \mathrm{d}_{\Sigma\mathrm{H}}(\mathrm{pressure}(P(\mathbf{x}_1)),\ \mathrm{pressure}(P(\mathbf{x}_2)))\ . \end{aligned} \tag{4}$$

With the domain-specific distance measure having been defined, if we now want to express the diversity in a set of solutions S, a set diversity measure is required that applies the domain-specific distance measure to pairs of solutions. A possible measure is Solow-Polasky [7], defined as

$$\mathrm{D}_{\mathrm{SP}}(S) = \sum_{i,j} b_{ij}, \tag{5}$$

with

$$\boldsymbol{B} = \boldsymbol{A}^{-1},\ a_{ij} = \exp(-\theta \cdot \mathrm{d}_{\mathrm{airfoil}}(\mathbf{x}_i, \mathbf{x}_j)),\ 1 \le i, j \le |S|,$$

that is, calculate the sum of the elements b_{ij} of $\boldsymbol{B}$, which is the inverse of $\boldsymbol{A}$. The elements a_{ij} of $\boldsymbol{A}$ depend on the difference between all pairs of solutions $\mathbf{x}_i, \mathbf{x}_j$ from set S, and on θ that should be set to some positive value and scales the relation between $\mathrm{d}_{\mathrm{airfoil}}$ and D_{SP} [7]. We use $\theta = 1$. Decreasing θ results in lower $\mathrm{D}_{\mathrm{SP}}(S)$ values, decreasing the influence of $\mathrm{d}_{\mathrm{airfoil}}$ on the score. Regardless of the chosen θ, $\mathrm{D}_{\mathrm{SP}}(S)$ always results in a real number in the interval $[1, |S|]$.

3 Novelty-guided Restarts for Diverse Solutions

Cuccu et al. [1] present a restart scheme aimed at leaving optima after having converged there, by restarting the search from the most novel point in an archive of solutions that were found to be highly novel upon sampling. We propose a similar scheme that performs optimization on fitness until some convergence criterion is met, and then optimizes on novelty, starting from the converged-to solution until the search ends up in infeasible space, i.e., all solutions in the current population have become infeasible. Following, the fitness optimization phase is restarted from the most novel solution found in the last novelty optimization

phase. The alternating scheme is run on top of the *CMA-ES* [2], and between restarts its covariance matrix is reinitialized, resetting the stepsize to the problem-specific initial value.
Cuccu et al. calculate novelty as the average difference to the closest k solutions currently in a partial archive of earlier sampled solutions and the current population (*sparseness*) [3, 5]. In this work a different novelty expression is used, namely *uniqueness* [5]. A comparison is made to all solutions sampled before during fitness optimization and the difference to the single closest solution is adopted as novelty score. The idea is that all these solutions are in the basin of attraction of an earlier converged-to optimum, and therefore are to be avoided.
The alternating scheme could be viewed as a way of incrementally populating a level set based on implicitly promoting diversity through stepwise maximizing novelty. Formally, given a fitness function $f : X \to \mathbb{R}$ and a threshold fitness value f_{thres}, we define a level set as

$$L_{f_{\text{thres}}} = \begin{cases} L^{-}_{f_{\text{thres}}} = \{\mathbf{x} \in X | f(\mathbf{x}) \leq f_{\text{thres}}\} & \text{if } f \text{ is to be minimized} \\ L^{+}_{f_{\text{thres}}} = \{\mathbf{x} \in X | f(\mathbf{x}) \geq f_{\text{thres}}\} & \text{if } f \text{ is to be maximized} \end{cases}, \tag{6}$$

that is, the level set $L_{f_{\text{thres}}}$ contains all solutions from domain X with fitness equal or better than f_{thres}. The novelty-guided restarts would serve as a variation operator, stepwise adding a new local optimum to the level set approximation that consists of the converged-to solutions in the fitness optimization phase and that has the fitness of the lowest-quality solution in this set as its f_{thres}.

4 Experiments and Discussion

Two-dimensional optimization is performed on the *outlet guide vanes* (OGVs) for the *low-pressure compressor* in a small turbofan engine, involving 16 control points for adjusting the two-dimensional shape of the blade, thus 32 decision variables. OGVs are *stator blades*, that is, non-moving blades intended for straightening the air flow, arranged around a central axis in a cylindrical fashion. Simulating aerodynamic behavior is done using a deterministic in-house *Navier-Stokes flow solver*. The fitness function is calculated as a to be minimized weighted sum of seven properties, see Eq. (1) and (2), namely the resulting deviation angle of the air flow, the occurring pressure loss, the extent of violation of four minimal thickness constraints along the blade, and a penalty depending on the extent to which the solver has not converged. We apply the alternating fitness/novelty scheme on top of a (12,24)-CMA-ES, which is allowed to run for 4000 generations (96k evaluations by the simulator). For comparison, we run multiple normal fitness-based optimizations using the same (12,24)-CMA-ES, for 1000 generations each (24k evaluations). Multiple of such runs can jointly be seen as a single run that has been restarted a number of times. A problem-specific initial stepsize σ_{init} of 1×10^{-3} is used.
Within the alternating scheme we stop the fitness phase when we assume that it has converged, considered as such when the current stepsize is smaller than $0.5 \times \sigma_{\text{init}}$. In the novelty phase the stopping condition is met when we assume that we are in infeasible space: Either all candidates have become strictly infeasible (a geometry is disconnected, i.e., contains loops, or the simulator does not converge for it), or the average fitness value of the best individual over the last 5 generations exceeds 1×10^3 (at this problem-specific quality value we surmise that the shape is maximally distorted for still allowing to find a way back to a high-performing solution). Finally, in determining novelty using the summed-up Hausdorff distance, the geometries are sampled in 200 points along the spline. Code was implemented using Shark Machine Learning Library v2.3.3.
In Fig. 1 the outcome of a single run of the alternating scheme is plotted against 8 standard fitness optimization runs combined (time to complete per run, using 24 processor cores per run, approximately

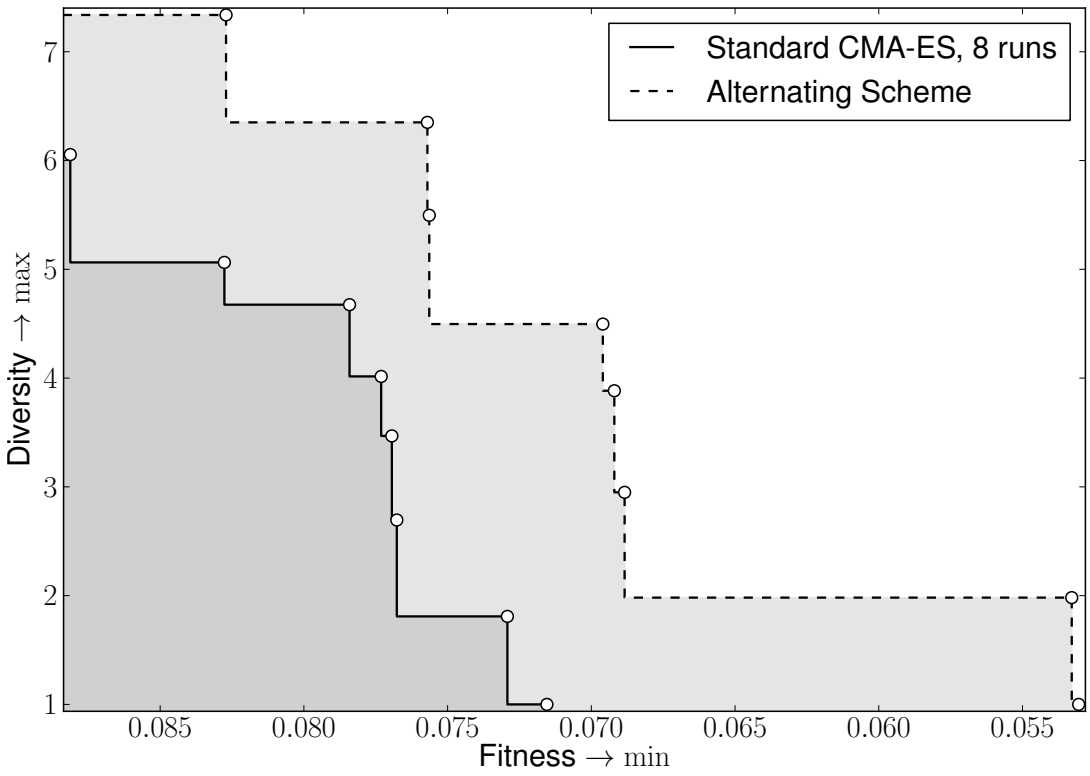

Figure 1: **Results.** For each solution in the final set of a single run of the alternating fitness/novelty scheme, and of 8 standard runs combined, the fitness and induced diversity scores are plotted (see Section 1). The alternating scheme results in a larger dominated surface in fitness/diversity space, while using only 96k evaluations vs. jointly 192k for the standard runs. The alternating scheme finds a set of 8 solutions that is more diverse than the combined set of 8 standard runs (7.33 vs. 6.05), *and* finds 5 solutions strictly better with respect to fitness than the standard runs' best (0.0530–0.0696 vs. 0.0715).

two weeks and 30 hours, respectively). The alternating fitness/novelty scheme finds a set of solutions that is more diverse than the combined outcome of the 8 standard runs and even manages to find solutions that are of better quality while using only half of the evaluations. For future work alternative domain-specific distance measures deriving from air speed along the airfoils and flow field visualizations are considered.

References

[1] Cuccu, G., Gomez, F., Glasmachers, T.: Novelty-Based Restarts for Evolution Strategies. In: 2011 IEEE Congress on Evolutionary Computation, CEC 2011. pp. 158–163. IEEE Computer Society (2011)

[2] Hansen, N.: The CMA Evolution Strategy: A Comparing Review. In: Towards a New Evolutionary Computation, Studies in Fuzziness and Soft Computing, vol. 192, pp. 75–102. Springer (2006)

[3] Lehman, J., Stanley, K.: Abandoning Objectives: Evolution through the Search for Novelty Alone. Evolutionary Computation 19(2), 189–223 (2011)

[4] Olhofer, M., Arima, T., Sonoda, T., Fischer, M., Sendhoff, B.: Aerodynamic shape optimisation using evolutionary strategies. In: Parmee, I., Hajela, P. (eds.) Optimization in Industry. pp. 83–94. Springer (2002)

[5] Reehuis, E., Olhofer, M., Emmerich, M., Sendhoff, B., Bäck, T.: Novelty and Interestingness Measures for Design-space Exploration. In: 15th Annual Conference on Genetic and Evolutionary Computation, GECCO'13. ACM Press (2013)

[6] Shir, O., Emmerich, M., Bäck, T.: Adaptive Niche Radii and Niche Shapes Approaches for Niching with the CMA-ES. Evolutionary Computation 18(1), 97–126 (2010)

[7] Ulrich, T., Thiele, L.: Maximizing Population Diversity in Single-Objective Optimization. In: N. Krasnogor et al. (ed.) 13th Annual Conference on Genetic and Evolutionary Computation, GECCO'11. pp. 641–648. ACM Press (2011)