# Multi-objective Representation Setups for Deformation-based Design Optimization

Andreas Richter[1], Jascha Achenbach[1], Stefan Menzel[2] and Mario Botsch[1]

[1] Computer Graphics Group, Bielefeld University, Germany
[2] Honda Research Institute Europe, Germany

**Abstract.** The increase of complexity in virtual product design requires high-quality optimization algorithms capable to find the global parameter solution for a given problem. The representation, which defines the encoding of the design and the mapping from parameter space to design space, is a key aspect for the performance of the optimization process. To initialize representations for a high performing optimization we utilize the concept of *evolvability*. Our interpretation of this concept consists of three performance criteria, namely *variability*, *regularity*, and *improvement potential*, where regularity and improvement potential characterize conflicting goals. In this article we address the generation of initial representation setups trading off between these two conflicting criteria for design optimization. We analyze Pareto-optimal compromises for deformation representations with radial basis functions in two test scenarios: fitting of 1D height fields and fitting of 3D face scans. We use the Pareto front as a ground-truth to show the feasibility of a single-objective optimization targeting one preference-based trade-off. Based on the results of both optimization approaches we propose two heuristic methods, Lloyd sampling and orthogonal least squares sampling, targeting representations with high regularity and high improvement potential at the two ends of the Pareto front. Thereby, we overcome the time consuming process of an evolutionary optimization to set up high-performing representations for these two cases.

## 1 Introduction

The increasing complexity in modern industrial design processes requires advanced optimization methods to efficiently come up with novel and high-quality solutions for successful business. In automotive product design, our target application, concurrent development processes are applied to deal with different requirements, e.g., from physical domains such as aerodynamic or structural performance criteria, from manufacturing process layout, or from design features specified by current customer demands. Moreover, since these requirements change over time, an efficient development process needs to cope with dynamic environments to allow a high degree of flexibility.

Biologically-inspired population-based evolutionary optimization algorithms are designed to handle these demands [13]. The careful construction of the representation, the encoding of a design, is one of the most important aspects for the success of an optimization process because the representation defines the solution space and determines how efficient the optimizer can explore it. Although a human designer can manually set up representations for design optimization, e.g. [19], the setup might not be optimal such that even minor changes of the setup could highly improve the performance of the optimization process. Therefore, we target to develop automatic procedures to optimize the initial representation along with the designer's input or preference, which is typically given by information on the expected importance of certain design regions for the current optimization task.

Based on the concept of evolvability [17] we proposed a mathematical model for evaluating the quality of deformation representation setups in [16]. The quality of each setup is numerically quantified by three criteria, namely *variability*, *regularity*, and *improvement potential*. However, since regularity and improvement potential are conflicting targets [16] a multi-objective analysis is required to finally provide the designer the possibility to choose one trade-off setup from the set of Pareto solutions according to her/his preference.

In this article our focus is on the initial generation of optimal representation setups addressing these two conflicting targets, which we evaluate for the same test scenarios as in [16]: 1D function approximation and 3D template fitting, which both are based on RBF deformation representations. For these RBF deformations we are particularly interested in the optimal distribution of the RBF centers and their efficient computation. To this end we first perform a multi-objective optimization of center distributions to analyze the trade-off between regularity and improvement potential. Although this time-consuming optimization might be infeasible for real-word applications, it results in ground-truth solutions that we use to evaluate a weighted single-objective optimization of the center distribution. We show that such a single-objective optimization is feasible in our application and speeds up the optimization process for one particular weight, which is set according to a designer's preference. The insight gained from the multi-objective optimization furthermore allows us to derive two heuristic approaches for rapidly generating center distributions on both ends of the Pareto front, i.e., aiming solely for regularity or improvement potential, respectively. Being based on Lloyd sampling [10] or orthogonal least squares [4], both methods generate high-quality center distributions within minutes in contrast to the single-objective optimization, which runs for hours.

In Section 2 we discuss state-of-the-art approaches for setting up deformation representations and motivate our approach. In Section 3 we give the technical details for RBF deformations as the representation of our choice. In Section 4 we describe our model of evolvability that we use to evaluate these representations. This yields the basis to perform and analyze a multi-objective optimization in two test scenarios in Section 5. The Pareto front is the ground-truth for a preference-based single-objective optimization in Section 6. Moreover, the Pareto-optimal solutions motivate heuristics, Lloyd and OLS, which we discuss in Section 7.

## 2 Related Work

In shape optimization based on deformation methods designer-driven approaches are typically applied to set up initial deformation representations. The designer defines target regions where the design has to be varied/optimized and places control points adapted to these regions. For example, in [19] a control grid for free-form deformation (FFD) is manually constructed and handle regions for deformations with radial basis functions (RBF) are manually set up. For basic automated representation setups commercial tools provide a uniform distribution of control points, e.g., a glider optimization [5].

Originally, deformation representations are employed in scattered data approximation, e.g., for approximating a target shape. In [3] the control points of non-uniform rational B-splines are optimized by a gradient-based method to improve the approximation quality of a wing. In [22, 27] a uniform setup of a control grid is refined in sensitive regions, i.e parameters are added, resulting in an improved approximation. Amoignon [1, 2] tackles the problem that uniform control grids for FFD might have empty grid cells. Instead of adjusting the grid to the design he deforms the design (e.g., wings) to completely fill out the grid. To obtain RBF setups that are adapted to a target, different basis functions are iteratively evaluated and selected at fixed locations [24] or their location is being optimized [7, 14]. All these approaches are specialized to set up control points for approximating one fixed target. Thereby, they neglect numerical properties of the deformation setup which are important for, e.g., the convergence speed of an evolutionary optimization.

The representation setup of adaptive B-splines for an evolutionary design process is targeted in [15, 26]. The optimization alternates between approximation of a shape and adaptation of the representation. To test whether an adjusted representation is beneficial for the optimization this process is performed for a few iterations. Thereby the performance of a representation is measured by the objective function of the actual optimization task. In [20] the representation is optimized implicitly by adding its parameters to the approximation problem. The criterion for a high-quality representation purely depends on the target of the optimization omitting further aspects of this process like convergence speed.

In contrast, we utilize quality criteria based on the concept of evolvability to evaluate representations. We include an objective-independent criterion to address the convergence speed as well as target information or human knowledge. Based on this model we set up high-quality deformation representations for evolutionary design optimization. In the next section we give the technical details for these deformations.

## 3 RBF Representations

In a shape optimization scenario, for instance in automotive product design, the design model to be optimized (the phenotype) is typically represented by a surface polygon mesh, where the $n$ mesh vertices $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{R}^3$ represent points on the surface, which are connected by polygonal faces (usually triangles or quads). The vertex positions $\mathbf{x}_i$ could in theory be used as optimization parameters in an evolutionary optimization. However, for non-trivial models the complexity of the model easily exceeds one million vertices, thus making the direct optimization of vertex positions intractable.

Even for highly complex shapes the actual *deformations* applied during optimization are typically rather simple, low-frequency functions, which can therefore be controlled by a small number of parameters. Hence we choose as representation a deformation function $\mathbf{u}(\mathbf{x})$, which maps deformation parameters (genotypes) to shape variations (phenotypes), which are then evaluated by a fitness function. Both free-form deformation (FFD) and radial basis functions (RBFs) have been successfully employed in design optimization [19]. In this paper we focus on RBF deformations, since their kernel-based setup is more flexible than lattice-based FFD representations.

The initial design $(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ is deformed into a shape variant $(\mathbf{x}'_1, \ldots, \mathbf{x}'_n)$ by adding the displacement $\mathbf{u}(\mathbf{x}_i)$ to each vertex $\mathbf{x}_i$ (Figure 1), which for RBF deformation has the form

$$\mathbf{u}(\mathbf{x}) \;=\; \sum_{j=1}^{m} \mathbf{w}_j \, \varphi(\|\mathbf{c}_j - \mathbf{x}\|) \;=:\; \sum_{j=1}^{m} \mathbf{w}_j \, \varphi_j(\mathbf{x}) \quad . \tag{1}$$

Here, $\varphi_j(\mathbf{x}) = \varphi(\|\mathbf{c}_j - \mathbf{x}\|)$ denotes the $j$-th scalar-valued radial basis function, which is centered at $\mathbf{c}_j \in \mathbb{R}^3$ and weighted by the coefficient $\mathbf{w}_j \in \mathbb{R}^3$.

The choice of the kernel function $\varphi \colon \mathbb{R} \to \mathbb{R}$ has a significant influence on the resulting deformation and its computation complexity [18]. In this paper we employ and analyze globally-supported triharmonic thin-plate splines, $\varphi_{tri}$, as well as compactly-supported Wendland functions, $\varphi_W$, with support radii $s$ varying from rather local to more global [25]:

$$\varphi_{tri}(r) = \begin{cases} r^2 \log(r) & \text{for 2D domains,} \\ r^3 & \text{for 3D domains.} \end{cases}$$

$$\varphi_W(r) = \begin{cases} \left(1 - \frac{r}{s}\right)^4 \left(\frac{4r}{s} + 1\right) & \text{for } r < s, \\ 0 & \text{otherwise.} \end{cases}$$

The RBF deformation (and thus the deformed shape) is linear in the RBF weights $\mathbf{w}_j$. If we write the initial and deformed shapes as $(n \times 3)$-matrices $\mathbf{X} = (\mathbf{x}_1^{\mathrm{T}}, \ldots, \mathbf{x}_n^{\mathrm{T}})^{\mathrm{T}}$ and $\mathbf{X}' = (\mathbf{x}_1'^{\mathrm{T}}, \ldots, \mathbf{x}_n'^{\mathrm{T}})^{\mathrm{T}}$, respectively, we can write the shape deformation in matrix notation

$$\mathbf{X}' = \mathbf{X} + \mathbf{\Phi}\mathbf{W} \tag{2}$$

using an $(n \times m)$ RBF matrix $(\mathbf{\Phi})_{i,j} = \varphi_j(\mathbf{x}_i)$ and the RBF weights $\mathbf{W} = (\mathbf{w}_1^{\mathrm{T}}, \ldots, \mathbf{w}_m^{\mathrm{T}})^{\mathrm{T}} \in \mathbb{R}^{m \times 3}$.

In the above setting, the deformation $\mathbf{u}$ is controlled by manipulating the RBF weights $\mathbf{w}_j$, which we call *indirect manipulation.* However, it has been

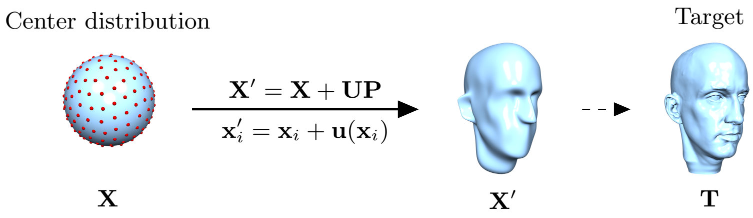
Center distribution
Target
$\mathbf{X}' = \mathbf{X} + \mathbf{UP}$
$\mathbf{x}_i' = \mathbf{x}_i + \mathbf{u}(\mathbf{x}_i)$
$\mathbf{X}$
$\mathbf{X}'$
$\mathbf{T}$

**Fig. 1.** The RBF deformation $\mathbf{u}$ transforms the initial mesh $\mathbf{X}$ to $\mathbf{X}'$ by translating each vertex $\mathbf{x}_i$ of $\mathbf{X}$ by the displacement $\mathbf{u}(\mathbf{x}_i)$. The distribution of the RBF centers (red dots) is crucial for a high-performing fit of the mesh $\mathbf{X}$ to the target $\mathbf{T}$.

shown in the context of free-form deformation that so-called *direct manipulation* is more intuitive for the human designer [8] as well as more efficient in an evolutionary optimization [12], due to the more direct and stronger causal relation between optimization parameters and the resulting shape deformation. In the RBF setting, a direct manipulation is controlled by specifying the displacement $\mathbf{d}_j$ for each center position $\mathbf{c}_j$, and then solving a linear system for the weights $\mathbf{w}_j$ that meet these interpolation constraints:

$$\mathbf{W} = \mathbf{\Psi}^{-1}\,\mathbf{D}\,, \tag{3}$$

with $\mathbf{D} = (\mathbf{d}_1^{\mathrm{T}}, \ldots, \mathbf{d}_m^{\mathrm{T}})^{\mathrm{T}} \in \mathbb{R}^{m\times 3}$ and $(\mathbf{\Psi})_{i,j} = \varphi_j(\mathbf{c}_i) \in \mathbb{R}^{m\times m}$. Combining equations (2) and (3) leads to the matrix representation of direct RBF deformation:

$$\mathbf{X}' = \mathbf{X} + \mathbf{\Phi}\mathbf{\Psi}^{-1}\,\mathbf{D}\ \ . \tag{4}$$

Note that both *indirect manipulation* (2) and *direct manipulation* (4) can be written as a linear deformation operator

$$\mathbf{X}' = \mathbf{X} + \mathbf{UP}\,, \tag{5}$$

using a deformation matrix $\mathbf{U}$, being either $\mathbf{\Phi}$ or $\mathbf{\Phi}\mathbf{\Psi}^{-1}$, and deformation parameters $\mathbf{P}$, being either $\mathbf{W}$ or $\mathbf{D}$. The deformation matrix $\mathbf{U}$ (which is the deformation setup) depends on the employed kernel and the center distribution. These two aspects define the realizable deformations and thereby the performance of a design optimization process. Evaluating and optimizing different deformation setups, different kernels, and different center distributions, allows us to initialize a high performing design optimization. The concept of evolvability reveals quality criteria that we discuss in the next section.

## 4 Evolvability for Linear Deformations

The biological concept of *evolvability* is a very promising approach to measure the *expected* performance of evolutionary processes [23]. In [17] we gathered, categorized, and extensively discussed this concept not only in the biological context, but also in the context of technical engineering. In agreement with Sterelny [21], we understand evolvability as a combination of three major attributes: *variability*, *regularity*, and *improvement potential.* Based on this classification we proposed a mathematical model to quantify the quality of linear deformation representations $\mathbf{U}$ in design optimization. In this section, we give a short summary of our model of evolvability and refer to [16] for details.

*Variability* $V(\mathbf{U})$ measures the potential of a deformation setup to explore the design space and we define it as

$$V(\mathbf{U}) = \frac{\operatorname{rank}(\mathbf{U})}{n}\,, \tag{6}$$

where $n$ is the number of vertices of a design and $\operatorname{rank}(\mathbf{U})$ denotes the rank of the matrix $\mathbf{U}$ [6]. We showed that variability is independent of the manipulation type, being either indirect or direct (equation (2) or (4)). Furthermore, for a fixed number of RBF centers it is independent of the center distribution as long as the deformation matrix has maximal rank, which is the case if centers do not coincide. In our test scenarios we assume a fixed number of centers, which results in constant variability, which thereby does not represent a conflicting interest to regularity and improvement potential.

We define *regularity* $R(\mathbf{U})$ as

$$R(\mathbf{U}) = \kappa^{-1}(\mathbf{U}) = \frac{\sigma_{\min}}{\sigma_{\max}}\,, \tag{7}$$

where $\kappa$ is the condition number of a matrix and $\sigma_{\min}$ and $\sigma_{\max}$ are its minimal and maximal singular value [6]. The regularity of a deformation setup characterizes the expected convergence speed of an evolutionary optimization. This criterion is an interpretation of the concept of robustness. Robust representations aim to prevent infeasible designs and thereby speed up the optimization process. Our regularity measure addresses the convergence directly. Furthermore, it corresponds to the concept of *causality*, because the condition number $\kappa$ characterizes the causal relation between genotype (parameter space) and phenotype (design space). In [16] we show that RBF deformation setups with a local kernel (a kernel with a small support radius) have higher regularity than setups with a global kernel and that direct manipulation has higher regularity than indirect manipulation.

We define *improvement potential* $P(\mathbf{U})$ as the potential of a representation to improve the fitness of a design. From a local point of view the most beneficial variation of a design is according to the (estimated) fitness gradient $\mathbf{g}$. Thus, we measure the improvement potential $P(\mathbf{U})$ as the approximation error to this gradient, which leads to:

$$P(\mathbf{U}) = 1 - \left\|(\mathbf{I} - \mathbf{U}\mathbf{U}^{+})\mathbf{g}\right\|_2^2\,, \tag{8}$$

with $\mathbf{U}^{+}$ being the Pseudo-inverse of $\mathbf{U}$ [6]. Because for complex design optimization applications the calculation of the fitness gradient is infeasible, designer

**Fig. 2.** Example of a (uniform) RBF center setup for function approximation. The center distribution on the initial plane is to be optimized to allow a high quality fit of the right test function (compare [16]).

knowledge and experience offer valuable insight to derive an estimated gradient. Improvement potential is independent of the manipulation type (indirect or direct), as is variability. Furthermore, we showed that global kernels have better improvement potential than local ones thus lead to solutions with higher fitness.

Our experiments in [16] revealed that regularity and improvement potential are conflicting targets, because local kernels show high regularity but low improvement potential, whereas global kernels show low regularity but high improvement potential. These two conflicting criteria motivate the multi-objective optimization in the next section to gain insight into possible trade-off setups. Such a multi-objective optimization is very time-consuming and hence infeasible for most applications. However, it results in ground truth solutions to evaluate a more efficient weighted single-objective optimization, which we propose and discuss in Section 6. To further increase efficiency we analyze heuristics for generating setups in Section 7, which are more robust against local optima where a single-objective optimization might easily get stuck.

## 5 Pareto Analysis

In this section we show the results of the multi-objective optimization of deformation setups towards the two conflicting targets regularity and improvement potential. Due to the enormous computational costs of automotive design optimization we analyze two simpler test scenarios instead, namely 1D function approximation and 3D template fitting, according to our test framework in [16].

### 5.1 Test Scenario: Height Field Approximation

In this scenario we fit an initial plane (Figure 2, left) to a target height field (Figure 2, right) by minimizing the approximation error (see [16] for the details). Instead of performing the actual fit (as in [16]), our goal in this paper is to find a well-performing *deformation setup*. We employ RBF deformation and construct the matrix $\mathbf{U}$ according to either indirect or direct manipulation, equation (2) or (4). To cover a variety of kernel types, from rather local to global, we employ compact Wendland kernels with support radii $s$ of 0.25 and 0.5, and global triharmonic kernels. Given the type of manipulation, the kernel, and the

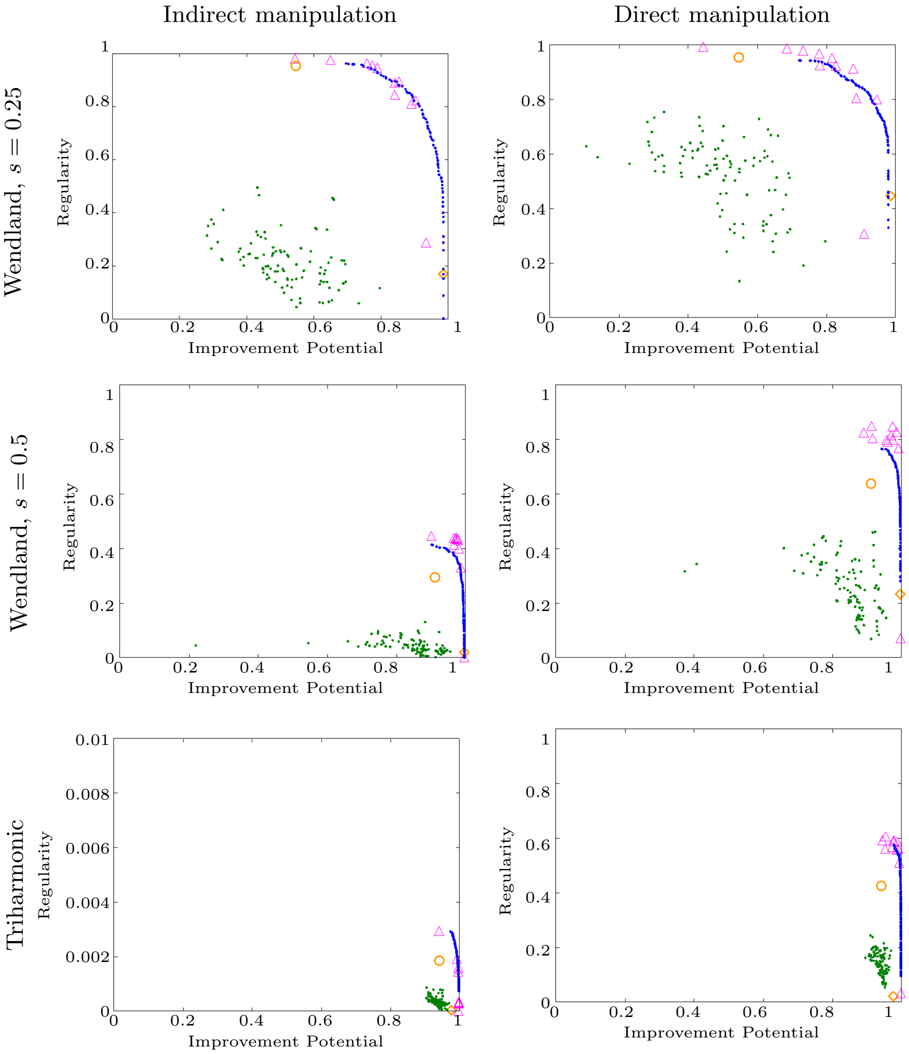

**Fig. 3.** The Pareto front (blue dots) and the initial random population (green dots) for the function approximation scenario. The magenta triangles are the results of the weighted single-objective optimization described in Section 6. Heuristic setups generated with Lloyd (orange circle) or OLS (orange diamond), discussed in Section 7, result in very regular setups or setups with a very good improvement potential.

support radius, the optimal center distribution with respect to the conflicting targets regularity and improvement potential, equations (7) and (8), is the goal of a multi-objective optimization. Because we analyze distributions with a fixed number of centers, $m = 25$, variability is constant thus not included in the

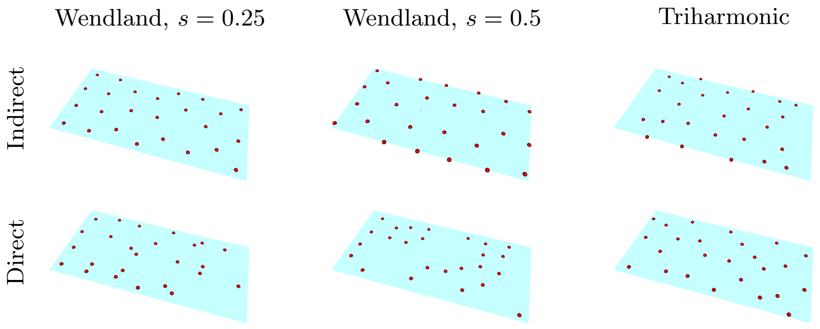

**Fig. 4.** Optimized center distributions towards regularity are uniform for indirect manipulation (top) but tends to be unintuitive for direct manipulation (bottom).

optimization. Each center has two coordinates, thus we solve a 50-dimensional optimization problem.

We realize the multi-objective optimization with the NSGA2 algorithm of the shark 2.3 library [9] with the following settings: 100 individuals, tournament selection, polynomial mutation rate with a probability of 1/50, crossover with a probability of 0.9, and 25000 iterations. We initialize the algorithm with randomized center distributions and restrict the centers to the initial plane ($[0,1] \times [0,2]$) during the optimization. With these settings one optimization run took approximately 2 days.

In Figure 3 we plot the resulting Pareto front as blue dots for the three tested kernels with indirect and direct manipulation, respectively. The green dots are the values of the initial population. The tests indicate a smooth well-shaped Pareto front. For the local Wendland kernel the front almost reaches the optimal value of 1 for regularity and improvement potential, respectively. Note that the very low regularity values of the triharmonic kernel for indirect manipulation goes along with our results in [16] and theoretical results in [25].

Especially the center distributions maximizing either regularity or improvement potential, respectively, are interesting because they can be computed through a single-objective optimization. For *indirect* manipulation we obtain uniformly distributed centers (Figure 4, top) resulting in maximal regularity, in agreement with theoretical results [25]. In contrast, the center distributions leading to optimal regularity for *direct* manipulation are unpredictable (Figure 4, bottom), which shows the advantages of an automatic procedure for distributing centers in contrast to a purely designer-driven approach.

Center distributions with maximal improvement potential are adapted to the target height field for the compact Wendland kernels (Figure 5). The distribution is denser in regions which have to be deformed more. In contrast, centers for the global triharmonic kernel are not placed in these regions (Figure 5, right), which is unintuitive for a designer. This again emphasizes the demand for an automatic construction of setups instead of a purely designer driven approach.

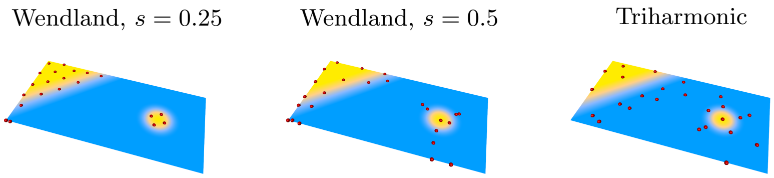

Wendland, $s = 0.25$ Wendland, $s = 0.5$ Triharmonic

**Fig. 5.** Target-adapted center distributions with optimal improvement potential. The compact kernels are mainly placed in regions with locally high fitting error (yellow), whereas triharmonic kernels are placed less intuitive (blue).

### 5.2 Test Scenario: 3D Template Fitting

In the second test scenario we deform an initial sphere to closely fit the point cloud of a given face scan (see Figure 1 and [16]). Like in the height field approximation scenario we intend to *set up* an optimal center distribution rather than performing the fitting. However, distributing centers for template fitting is more complex because the sphere and the scan are embedded in 3D such that each of the 25 centers consists of 3 coordinates, resulting in 75 parameters to be optimized. We choose the initial distributions randomly on the initial sphere, restrict the search domain to its bounding box $[-1, 1]^3$, and choose support radii of 0.5 and 1 for the Wendland kernels (since the initial domain is larger than in the function approximation scenario). Apart from the mutation rate, which we set to 1/75 according to the 75 parameters, we perform the multi-objective optimization with identical settings as in the function approximation scenario. In Figure 6 we plot the Pareto front for the three kernel types with direct and indirect manipulation, respectively. These plots are qualitatively equivalent to the plots of the function approximation scenario, compare to Figure 3.

The multi-objective optimization in both test scenarios, height field approximation and template fitting, runs up to 2 days, hence it is infeasible for most real-world applications. Instead of computing the whole Pareto front we are rather interested in one particular setup trading off regularity and improvement potential according to our preference. Therefore, we employ a weighted single-objective optimization in the next section and utilize the Pareto front as a ground truth to test if this optimization is able to converge towards the front.

## 6 Weighted Single-objective Optimization

The runtime of 2 days of a multi-objective optimization in our tests motivates alternative optimization approaches. Instead of computing the whole Pareto front the designer guides the construction of trade-off setups between regularity and improvement potential by setting a preference $\lambda \in [0, 1]$ based on his expertise.

By weighting equation (7) and (8) we define an objective function $f$ for a preference-based single-objective optimization:

$$f_\lambda(\mathbf{U}) = \lambda R(\mathbf{U}) + (1 - \lambda) P(\mathbf{U}) \ . \tag{9}$$

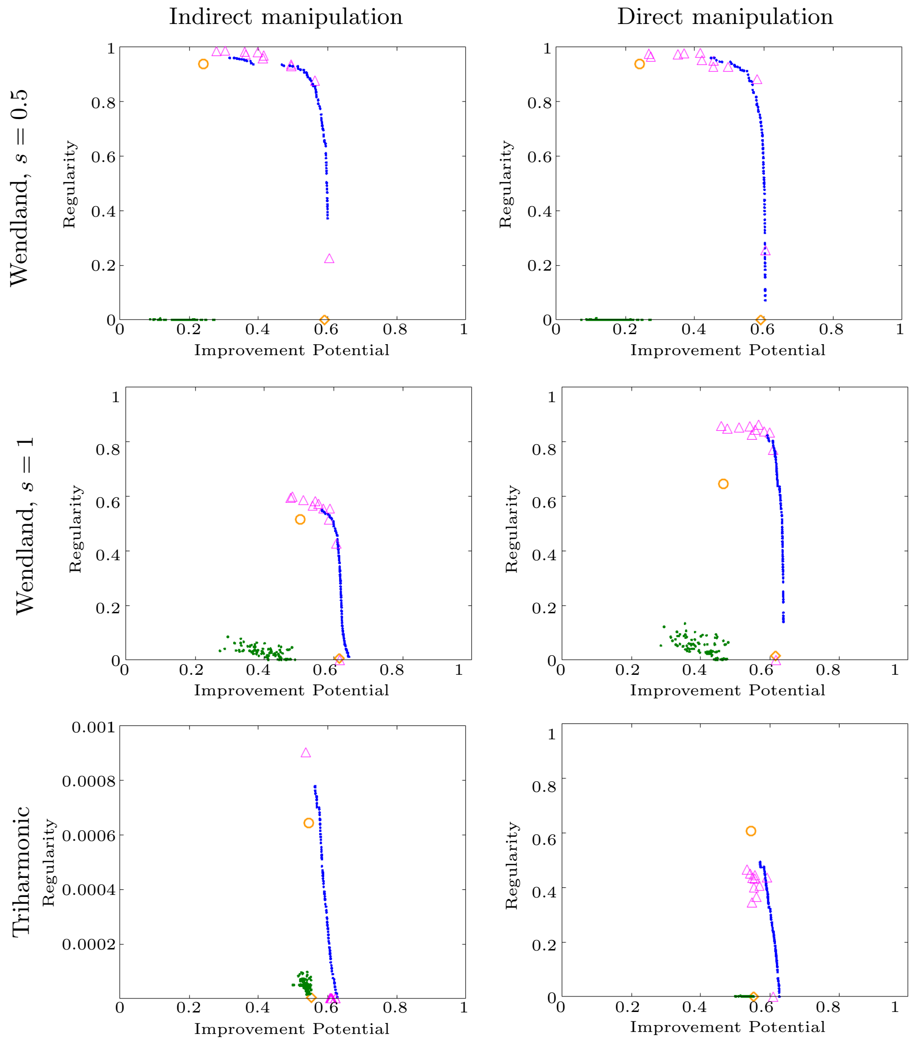

**Fig. 6.** The Pareto front (blue dots) and the initial random population (green dots) for the template fitting scenario. The magenta triangles are the results of the weighted single-objective optimization described in Section 6. Heuristic setups, described in Section 7, generated with Lloyd (orange circle) or OLS (orange diamond) result in very regular setups or setups with a very good improvement potential.

Because such a single-objective optimization might not converge to the Pareto front, we analyze this in the following. As an optimization algorithm we choose a (25,100)-CMA-ES of the shark 2.3 library [9], we choose the preferences $\lambda$ to be 0, 0.1, 0.2, . . . , 1 for equation (9) and run the optimization for 1000 generations. The optimization of a setup for one preference took approximately 2 hours.

The results of the single-objective optimization in Figure 3 and Figure 6 are depicted with the magenta triangles. The clustering of solutions, e.g., Figure 3 and Figure 6 middle, shows that uniformly distributed preferences $\lambda$ do not result in uniformly distributed solutions along the Pareto front. Therefore, a designer has to set the preference carefully. The single-objective optimization converges towards the Pareto front and even performs slightly better because of its focus in one preferred direction, except for the triharmonic kernel in the template fitting scenario, where the optimizer gets stuck in local optima (Figure 6, bottom right). This shows the feasibility of such an optimization for scenarios where a designer is interested in a setup for one particular preference. The runtime of 2 hours and these local optima motivate efficient heuristics to distribute centers.

## 7 Heuristic Setup Strategies

Heuristic methods aim to generate good center distributions in a robust and efficient manner. They are analytically and geometrically motivated but lack the guarantee to be Pareto-optimal. In our test scenarios a single-objective optimization still runs for hours and might get stuck in local optima. Because we expect these drawbacks become worse for more complex scenarios, e.g., with a more complex initial design or a larger amount of parameters, we propose and analyze a geometry-motivated approach for very regular setups and an analytically motivated approach for setups with high improvement potential.

Pareto-optimal center distributions targeting regularity for indirect manipulation are uniform distributions in all our tests (Figure 4, top). Hence, we apply Lloyd sampling which is also known as k-means clustering (see [10] and [11], respectively, for algorithmic details), which result in uniform center distributions similar to the Pareto-optimal solutions (compare Figure 2, left, and Figure 4, top). Comparing the regularity score of the resulting setup to the Pareto front (see Figure 3, Figure 6, orange circles) reveals that the Lloyd sampling is close to the front for local Wendland kernels ($s = 0.25$ for the plane or $s = 0.5$ for the sphere). Even for direct manipulation uniform Lloyd sampling results in good regularity. For the triharmonic kernel in the template fitting scenario the heuristic even out-performs the multi- and single-objective optimization (Figure 6, bottom right). According to equation (7) regularity is the ratio of the smallest to the largest singular value of the deformation matrix. For indirect RBF manipulation this singular value is bounded by the *separation distance*, which measures the minimal distance between any pairs of centers [25]. The uniform Lloyd sampling by construction has a good separation distance and thus results in good regularity. This sampling performs better than any tested random distributions, performs better than the evolutionary optimization in one test, is robust to local optima (Figure 6, bottom right), and fast to set up (one minute).

In [16] we motivated improvement potential (equation (8)) by solving the approximation problem $\mathbf{g} = \mathbf{U}\mathbf{p}$ for an estimated fitness gradient $\mathbf{g}$, the deformation matrix $\mathbf{U}$, and the deformation parameters $\mathbf{p} = (p_1, \ldots, p_m)$. Each parameter $p_j$ is the coefficient for $\varphi_j(\mathbf{x})$, which corresponds to a column $\mathbf{U}_j$ of

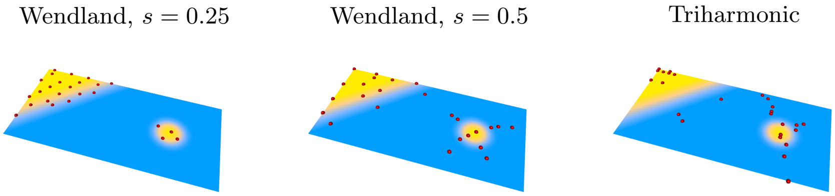

**Fig. 7.** Heuristic OLS setups with high improvement potential are adapted to the target. Wendland kernels are placed in regions with locally high fitting error (yellow) rather than in already optimal ones (blue).

the deformation matrix and to a center $\mathbf{c}_j$ for indirect manipulation. The orthogonal least squares method (OLS, detailed description and algorithm in [4, 7]) determines the influence of each parameter to minimize the approximation error to the estimated gradient in a greedy manner. According to their influence OLS ranks the parameters, thereby ranks the centers, and we select the most important ones. We initialize OLS with a large set of centers as candidates ($30 \times 30$ in 2D or $30 \times 30 \times 30$ in 3D) on a uniform grid and greedily select the best 25 ones. We cannot apply this procedure for direct manipulation because the interpolation matrix $\mathbf{\Psi}^{-1}$ in equation (4) disbands the correspondence between parameters and centers. Since direct manipulation and indirect manipulation result in equal improvement potential for identical center distributions, we simply apply this algorithm for indirect manipulation and switch to direct manipulation afterwards. For the function approximation scenario we show that target-adapted setups for the Wendland kernels in Figure 7 (left) are similar to the Pareto-optimal solutions in Figure 5. But for the global triharmonic kernels OLS results in an unintuitive center placement (Figure 7 right). Nonetheless, the OLS setups are close to the Pareto front or even hit it in both test scenarios (see Figure 3, Figure 6, orange diamonds) for the compact Wendland kernels. In conjunction with the small computation time of 1 minute, OLS is very efficient.

## 8 Summary and Future Work

The initial representation setup is crucial for the performance of an evolutionary optimization process. We analyzed the generation of RBF deformation setups for evolutionary design optimization for two test scenarios. The concept of evolvability reveals powerful criteria for setups, namely variability, regularity, and improvement potential, to measure the expected performance of a setup. Regularity and improvement potential are conflicting targets, which we therefore analyze with a multi-objective optimization. As downside this optimization process has a runtime of 2 days for our comparatively simple test scenarios.

In real-world applications we rather aim for one optimal deformation setup with respect to a user-specified preference between regularity and improvement potential. We demonstrated the feasibility of such a weighted single-objective optimization. For some tests the quality of the deformation setup is even better than the Pareto-optimal ones. This process is much faster, but it still runs for 2 hours for our simple problems. Furthermore, the single-objective optimization gets stuck in local optima in some of our tests.

In order to further improve computational performance and robustness we proposed and analyzed heuristics to generate setups. The regular setups constructed by Lloyd sampling are close to the Pareto front for local Wendland kernels. Even for direct manipulation where we lack the geometrical motivation the regularity of the setup is significantly better than a random initialization. The Lloyd sampling even out-performs the evolutionary solutions in one example. Center distributions constructed with orthogonal least squares have high improvement potential and are on or very close to the Pareto front in all tests. Both methods reduce the computational effort from 2 hours to 1 minute.

For future work we blend between the heuristics, Lloyd and OLS, according to our preference. Moreover, we analyze the generalization of our methods to alternative deformations like free-from deformation.

## Acknowledgments

Andreas Richter gratefully acknowledges the financial support from Honda Research Institute Europe (HRI-EU). Mario Botsch is supported by the Cluster of Excellence Cognitive Interaction Technology "CITEC" (EXC 277) at Bielefeld University, funded by the German Research Foundation (DFG).

## References

1. Amoignon, O., Hradil, J., Navratil, J.: A numerical study of adaptive FFD in aerodynamic shape optimization. In: Proceedings of 52nd Aerospace Sciences Meeting (2014)
2. Amoignon, O., Navrátil, J., Hradil, J.: Study of parameterizations in the project CEDESA. In: Proceedings of 52nd Aerospace Sciences Meeting (2014)
3. Becker, G., Schäfer, M., Jameson, A.: An advanced NURBS fitting procedure for post-processing of grid-based shape optimizations. In: Proceedings of 49th Aerospace Sciences Meeting (2011)
4. Chen, S., Billings, S.A., Luo, W.: Orthogonal least squares methods and their application to non-linear system identification. International Journal of control 50(5), 1873–1896 (1989)
5. Costa, E., Biancolini, M.E., Groth, C., Cella, U., Veble, G., Andrejasic, M.: RBF-based aerodynamic optimization of an industrial glider. In: Proceedings of International CAE Conference (2014)
6. Golub, G.H., Van Loan, C.F.: Matrix computations. Johns Hopkins University Press (2012)
7. Gomm, J.B., Yu, D.L.: Selecting radial basis function network centers with recursive orthogonal least squares training. IEEE Transactions on Neural networks 11(2), 306–314 (2000)
8. Hsu, W.M., Hughes, J.F., Kaufman, H.: Direct manipulation of free-form deformations. In: Proceedings of ACM SIGGRAPH. pp. 177–184 (1992)

9. Igel, C., Heidrich-Meisner, V., Glasmachers, T.: Shark. The Journal of Machine Learning Research 9, 993–996 (2008)
10. Lloyd, S.P.: Least squares quantization in PCM. IEEE Transactions on Information Theory 28(2), 129–137 (1982)
11. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: Proceedings of the fifth Berkeley symposium on mathematical statistics and probability. pp. 281–297 (1967)
12. Menzel, S., Olhofer, M., Sendhoff, B.: Direct manipulation of free form deformation in evolutionary design optimisation. In: Proceedings of the International Conference on Parallel Problem Solving From Nature. pp. 352–361 (2006)
13. Mina, A.A., Braha, D., Bar-Yam, Y.: Complex Engineered Systems: Science Meets Technology, chap. Complex Engineered Systems: A New Paradigm, pp. 1–21. Springer (2006)
14. Ohtake, Y., Belyaev, A., Seidel, H.P.: 3D scattered data approximation with adaptive compactly supported radial basis functions. In: Proceedings of IEEE International Conference on Shape Modeling Applications. pp. 31–39 (2004)
15. Olhofer, M., Jin, Y., Sendhoff, B.: Adaptive encoding for aerodynamic shape optimization using evolution strategies. In: Proceedings of IEEE Congress on Evolutionary Computation. pp. 576–583 (2001)
16. Richter, A., Achenbach, J., Menzel, S., Botsch, M.: Evolvability as a quality criterion for linear deformation representations in evolutionary optimization. In: Proceedings of IEEE Congress on Evolutionary Computation. pp. 901–910 (2016)
17. Richter, A., Botsch, M., Menzel, S.: Evolvability of representations in complex system engineering: a survey. In: Proceedings of IEEE Congress on Evolutionary Computation. pp. 1327–1335 (2015)
18. Sieger, D., Gaulik, S., Achenbach, J., Menzel, S., Botsch, M.: Constrained space deformation techniques for design optimization. Computer Aided Design 73, 40–51 (2016)
19. Sieger, D., Menzel, S., Botsch, M.: A comprehensive comparison of shape deformation methods in evolutionary design optimization. In: Proceedings of the International Conference on Engineering Optimization (2012)
20. Simões, L.F., Izzo, D., Haasdijk, E., Eiben, A.E.: Self-adaptive genotype-phenotype maps: neural networks as a meta-representation. In: International Conference on Parallel Problem Solving from Nature. pp. 110–119 (2014)
21. Sterelny, K.: Modularity in development and evolution, chap. Symbiosis, Evolvability and Modularity, pp. 490–516. University of Chicago Press (2004)
22. Vuong, A.V., Giannelli, C., Jüttler, B., Simeon, B.: A hierarchical approach to adaptive local refinement in isogeometric analysis. Computer Methods in Applied Mechanics and Engineering 200(49), 3554–3567 (2011)
23. Wagner, G.P., Altenberg, L.: Perspectives: Complex adaptations and the evolution of evolvability. Evolution 50(3), 967–976 (1996)
24. Webb, A.R., Shannon, S.: Shape-adaptive radial basis functions. IEEE Transactions on Neural Networks 9(6), 1155–1166 (1998)
25. Wendland, H.: Scattered data approximation. Cambridge University Press (2004)
26. Yang, Z., Sendhoff, B., Tang, K., Yao, X.: Target shape design optimization by evolving B-splines with cooperative coevolution. Applied Soft Computing 48, 672–682 (2016)
27. Zheng, J., Wang, Y., Seah, H.S.: Adaptive T-spline surface fitting to z-map models. In: Proceedings of the 3rd international conference on Computer graphics and interactive techniques in Australasia and South East Asia. pp. 405–411 (2005)