Evolvability as a Quality Criterion for Linear Deformation Representations in Evolutionary Optimization

Andreas Richter*, Jascha Achenbach*, Stefan Menzel† and Mario Botsch*
* Computer Graphics Group, Bielefeld University, Germany
†Honda Research Institute Europe, Germany

***Abstract*—Industrial product design is characterized by increasing complexity due to the high number of involved parameters, objectives, and boundary conditions, all typically changing over time. Population-based evolutionary design optimization targets to solve these kinds of application problems, offering efficient algorithms striving for high-quality solutions. An important factor in the optimization setup is the representation, which defines the encoding of the design and the mapping from parameter space to design space. Being able to numerically quantify the quality of different representation settings would strengthen the optimal choice of encoding. Motivated by the biological concept of evolvability, we propose three criteria, namely *variability*, *regularity*, and *improvement potential*, to evaluate linear deformation representations for their use in shape optimization problems. The first aspect characterizes the exploration potential of the design space, the second measures the expected convergence speed, and the third determines the expected improvement of the quality of a design. We propose and experimentally analyze mathematical definitions for each of the three criteria. We demonstrate the successful application of our model to two evolutionary optimization scenarios: fitting of 1D height fields and fitting of 3D face scans, both based on RBF deformations. Due to the general character of our definition we expect the transferability of our concepts to alternative deformation methods.**

I. 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 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 [1]. For these algorithms the encoding of the problem, i.e., the representation of the design component, is one of the core success factors. A careful design of the representation can affect the computation time more than parameter tuning. Thus, deciding on the type of representation and selecting optimization parameters among the representation is a crucial step which needs to be thoroughly evaluated before the optimization is initialized. To maximize the positive effect of the representation a quantification of the expected influence of different representation settings on the optimization process is required, which allows us to compare between them and to select the most promising ones. To quantify the potential of a representation in this context we analyze *evolvability*, which is known to describe such a criterion in biological evolution [2].

Automotive aerodynamic design targets the optimization of a given vehicle geometry for improved flow characteristics, e.g., to minimize drag for higher fuel efficiency. For these application scenarios shape morphing methods have shown two strong advantages [3]: First, these methods are capable of representing designs using a small number of parameters while still offering a high degree of geometric flexibility, and they can deal with complicated shape features such as holes or edges. Second, computational grids, which are typically required to determine the aerodynamic performance using finite element/volume methods, are modified simultaneously with the design update step, thereby avoiding costly (manual) remeshing processes. Among these shape morphing algorithms are state-of-the-art space deformation algorithms [4], such as deformations based on radial basis functions (RBF) or free-form deformation (FFD). In Figure 1 a typical automotive evolutionary optimization is depicted, which starts with a base design used to initialize the deformation method. For initialization a surface and an application-dependent computational mesh are represented by RBF or FFD parameters, which are determined by a careful manual representation setup. By repeating the steps of parameter recombination and mutation, followed by design variation, evaluation, and selection, an optimal design is generated.

As stated above, the initial representation setup is crucial for the success of the optimization process, hence we have to determine its quality. In the present paper we define such quality criteria for deformation representations based on the biological concept of evolvability and evaluate it in two design optimization test scenarios. First, we discuss existing technical evolvability concepts and related approaches in Section II. In Section III we discuss RBF deformations as an example of linear deformation methods, before we propose the evolvability definitions in Section IV. We analyze this concept first in a 1D function approximation scenario with RBF deformations (Section V). Afterwards we perform a more complex 3D template fitting test in Section VI.

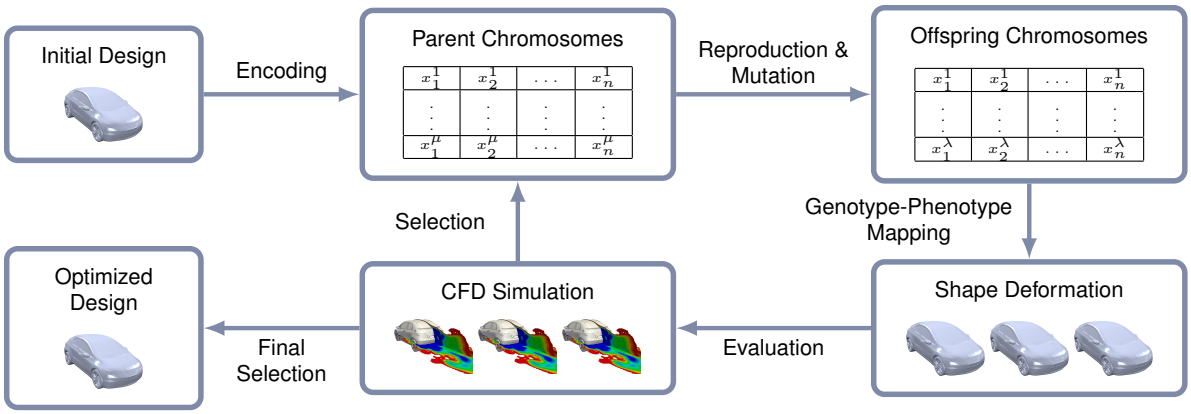

Fig. 1. Overview of an evolutionary design optimization process: After the initial design is encoded into genotype parameters, the mutation and reproduction steps produce different offspring genotypes, which lead to different design variations (phenotypes) through the genotype–phenotype mapping (the deformation representation). The new designs are evaluated through the fitness function and the most promising designs are selected. This process is repeated until convergence.

II. Related Work

It is well known from the literature that a guiding principle for the evolution success is *evolvability* [2, 5]. However, there exist a wide variety of definitions for the term evolvability, which we have gathered, categorized, and extensively discussed in [6]. For instance in the context of engineering systems, evolvability is considered as an evolved quality specified as "the ability of the configuration space (in this case, the space of genotypes and phenotypes) to produce an endless supply of viable configurations with remarkably few obvious dead-end" [1], or "considered in the sense of the capacity of a system to produce favorable phenotypic variations of a design within a moderate number of generations while avoiding non-feasible mutations" [7]. Along this line of thinking and in agreement with Sterelny [8], we understand evolvability as a combination of three major attributes, which promise a mathematically motivated quantification of encoding setups in the context of geometric representations: *variability*, *regularity*, and *improvement potential*. In the following we review and discuss these criteria, following our analysis in [6].

One aspect of a representation is to ensure the exploration of the design space—to produce a variety of designs on the one hand and to achieve specialized, optimal ones on the other hand. *Variability* characterizes this design space exploration potential of a representation. The biological concepts mainly evaluate the ratio between phenotype (design) variation and genotype (parameter) variation [9, 10]. In [7] this idea is transferred to a shape matching optimization using free-form deformation. Their variability definition is based on global information, which is derived from the analysis of the whole design space. This was possible in their configuration, since the design space was bounded by a box around the initial design, but the concept is not suitable for problems with unbounded parameter spaces—which we aim at. Local approaches analyze the genotype neighborhood of a design [11, 12], and consider the distance between the corresponding phenotypes as a measure for variability, aiming for a widely spread phenotype neighborhood. Transferring this idea to continuous genotype spaces would look for representations (genotype–phenotype mappings) with large gradient magnitudes, which is known to cause numerical problems. These limitations motivate the need for an improved variability measure.

Regularity is understood as a fitness-independent quality criterion and is interpreted in different ways. It can be regarded as an anti-outlaw condition [8] to prevent infeasible offspring. In design optimization self-intersections are such infeasibilities, which are reduced in [7], but regularity is not quantified there. Preventing or reducing infeasible designs (before costly fitness evaluation) speeds up the optimization process. Inspired by this property, we directly interpret regularity as a criterion to measure the optimization/convergence speed. Typically, optimization algorithms try to identify a trend in the fitness landscape, a direction in which the fitness improves the most. This is significantly influenced by the representation. A local (strong causal) representation maps small changes in the genotype to small phenotype variations, thereby preserving the local neighborhood structure. This makes it easier for the optimization to discover a promising direction in the fitness landscape, in contrast to non-local (weak causal) representations. Causality [13] and locality [14] are two concepts that are related to the convergence speed of evolutionary algorithms, too. But they are typically associated to the mutation operator (e.g., [15–18]), and are defined for discrete genotypes and phenotypes only. We propose a definition of regularity that generalizes these concepts to continuous design optimization problems.

In the biological context several approaches identify evolvability itself with adaptation potential or adaptation speed of a population to an environment, e.g. [19]. But this definition is rather imprecise, since a population is called adapted as soon as beneficial traits occur more often. Referring to technical optimization problems, where the goal is to improve an initial solution, the third aspect of evolvability is *improvement potential*. This criterion is investigated for optimization in varying [20, 21] as well as in static [22] fitness landscapes. Estimating the improvement potential of a representation is difficult due to, first, a lack of knowledge of the technical application, where information about the fitness landscape may not be available at the beginning of an optimization. Second, it can be computationally expensive, especially for industrial applications, if the local improvement potential is tested during the optimization by additional data sampling and (surrogate) modeling steps. These drawbacks motivate the need for novel formulation of improvement potential in design optimization.

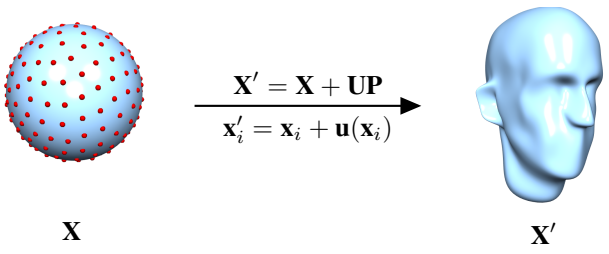

Fig. 2. The linear RBF deformation $\mathbf{u}$ transforms the initial model $\mathbf{X}$ to $\mathbf{X}'$ by translating each vertex $\mathbf{x}_i$ of $\mathbf{X}$ by the displacement $\mathbf{u}(\mathbf{x}_i)$. The red dots depict the centers $\mathbf{c}_j$ of the radial basis functions.

III. Linear Deformation 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 1 million vertices, thus making the direct optimization of vertex position intractable.

However, even for highly complex shapes the actual shape deformations applied during optimization are rather simple, low-frequency functions, which can therefore be controlled by a small number of parameters. The representation is therefore a deformation function $\mathbf{u}(\mathbf{x})$, which maps deformation parameters (genotypes) to shape variations (phenotypes), which are then evaluated by the fitness function (see Figure 1). Both free-form deformation (FFD) and radial basis functions (RBFs) have been successfully employed in design optimization [4]. In this paper we focus on RBF deformations, since their kernel-based setup is less constrained 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 to each $\mathbf{x}_i$ the displacement $\mathbf{u}(\mathbf{x}_i)$, which for RBF deformation has the form

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

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 $\mathbf{w}_j \in \mathbb{R}^3$ (see Figure 2).

The choice of the kernel function $\varphi\colon \mathbb{R} \to \mathbb{R}$ has a significant influence on the resulting deformation and the computation complexity [23]. In this paper we employ and analyze globally-supported triharmonic thin-plate splines, φ_{tri}, as well as compactly-supported Wendland functions, φ_W, with support radii s varying from rather local to more global:

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

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

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

$$\mathbf{X}' = \mathbf{X} + \mathbf{\Phi W} \qquad (2)$$

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 $\mathbf{u}$ would be controlled by manipulating the RBF weights $\mathbf{w}_j$. However, it has been shown in the context of free-form deformation that so-called *direct manipulation* is more intuitive for the human designer [24] as well as more efficient in an evolutionary optimization [25], 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 the interpolation constraints:

$$\mathbf{W} = \mathbf{\Psi}^{-1} \mathbf{D} \,, \qquad (3)$$

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

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

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

$$\mathbf{X}' = \mathbf{X} + \mathbf{UP} \,, \qquad (5)$$

using a deformation matrix $\mathbf{U}$ (being $\mathbf{\Phi}$ or $\mathbf{\Phi \Psi}^{-1}$) and deformation parameters $\mathbf{P}$ (being $\mathbf{W}$ or $\mathbf{D}$). From a mathematical point of view the matrix $\mathbf{\Psi}^{-1}$ can also be considered as a "preconditioner" that might speed up the optimization process.

The linear *deformation representation* now consists of this (constant) matrix $\mathbf{U}$, which is initially set-up by selecting the kernel function φ and placing the RBF centers $\mathbf{c}_j$. During the optimization process the algorithm selects different deformation parameters (genotypes) $\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \ldots$, which are mapped to shape variations (phenotypes) $\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \ldots$, which are successively evaluated by the fitness function $f(\mathbf{X}^{(k)})$. Figure 3 depicts the shape optimization process, where an initial plane model is deformed using direct RBF manipulation in order to closely fit a given target height field.

Our goal in this paper is to quantitatively rate the quality or the potential of a linear deformation representation setup, in order to estimate how well it will perform in an evolutionary optimization, and then to pick the most promising setup. Note that besides direct and indirect RBF deformation, all linear deformation representations can be written in the form of (5), such as, for instance, direct [24] and indirect free-form deformation [26] or linear thin shell models [27]. In the following section, we will therefore formulate evolvability-inspired quality criteria based solely on properties of the deformation matrix $\mathbf{U}$ to allow for possible generalization.

IV. Evolvability for Linear Deformations

As described in Section I, we model quality criteria for linear deformation representations in design optimization by utilizing the biological concept of evolvability. To this end, we propose mathematical formulations for the three evolvability criteria *variability*, *regularity*, and *improvement potential*. These criteria are designed to depend on the deformation matrix $\mathbf{U}$ only, and should therefore generalize beyond RBF deformations. To simplify the notation and derivation, we first assume that the displacement function $\mathbf{u}(\mathbf{x})$ is *scalar-valued*, such that the RBF coefficients $w_j \in \mathbb{R}$, hence $\mathbf{p} = (w_1, \ldots, w_m) \in \mathbb{R}^m$, and consequently $\mathbf{u} = \mathbf{U}\mathbf{p} \in \mathbb{R}^n$. Since the deformation matrix $\mathbf{U}$ is identical for the 1D and 3D deformation, this simplification does not change the resulting formulation of the evolvability criteria.

A. Variability

Variability is meant to quantify the potential for exploring the phenotype space—independent of the possibly complex objective function—by varying the genotype parameters $\mathbf{p}$ and mapping them to phenotype variations $\mathbf{u} = \mathbf{U}\mathbf{p}$. A representation has maximal variability if it can control the "1D-displacement" $u(\mathbf{x}_i)$ for each vertex $\mathbf{x}_i \in \mathbf{X}$ independently. This, however, would require an intractable number $m = n$ of optimization parameters, and therefore a much smaller number $m \ll n$ is typically chosen in practice.

Consequently, not every desired shape variation $\bar{\mathbf{u}} = (\bar{u}_1, \ldots, \bar{u}_n)$ can be represented as $\mathbf{U}\mathbf{p}$. The variability criterion will estimate how well a given arbitrary displacement $\bar{\mathbf{u}}$ can be approximated as $\mathbf{U}\mathbf{p}$, by averaging the approximation error $\|\bar{\mathbf{u}} - \mathbf{U}\mathbf{p}\|$ over all possible deformations $\bar{\mathbf{u}}$.

For a given deformation $\bar{\mathbf{u}}$, the optimal parameters $\mathbf{p}$, corresponding to the least squares approximation, can be computed through the normal equations [28]:

$$\min_{\mathbf{p}} \|\mathbf{U}\mathbf{p} - \bar{\mathbf{u}}\|^2 \quad \Leftrightarrow \quad \mathbf{U}^{\mathrm{T}}\mathbf{U}\mathbf{p} = \mathbf{U}^{\mathrm{T}}\bar{\mathbf{u}} \quad \Leftrightarrow \quad \mathbf{p} = \mathbf{U}^{+}\bar{\mathbf{u}}$$

with $\mathbf{U}^{+} = \left(\mathbf{U}^{\mathrm{T}}\mathbf{U}\right)^{-1}\mathbf{U}^{\mathrm{T}}$ being the pseudo-inverse of $\mathbf{U}$. The best-approximating deformation therefore is $\mathbf{U}\mathbf{U}^{+}\bar{\mathbf{u}}$, and its approximation error is

$$\left\|\bar{\mathbf{u}} - \mathbf{U}\mathbf{U}^{+}\bar{\mathbf{u}}\right\| \;=\; \left\|\left(\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right)\bar{\mathbf{u}}\right\| \,. \tag{6}$$

To analyze how well the design space can be explored by varying the initial design $\mathbf{X}$, we average the approximation error (6) over *all* possible variations $\bar{\mathbf{u}}$. Since a scaling of $\bar{\mathbf{u}}$ will lead to the same scaling of the approximation $\mathbf{U}\mathbf{U}^{+}\bar{\mathbf{u}}$, we restrict to unit-length variations $\|\bar{\mathbf{u}}\| = 1$, i.e., all variations $\bar{\mathbf{u}}$ from the n-dimensional unit-sphere $\mathbb{S}^n$:

$$\begin{aligned} & \frac{1}{|\mathbb{S}^n|}\int_{\bar{\mathbf{u}}\in\mathbb{S}^n} \left\|\left(\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right)\bar{\mathbf{u}}\right\|^2 \mathrm{d}\bar{\mathbf{u}} \\ \leq\; & \frac{1}{|\mathbb{S}^n|}\int_{\bar{\mathbf{u}}\in\mathbb{S}^n} \left\|\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right\|_{\mathrm{F}}^2 \left\|\bar{\mathbf{u}}\right\|^2 \mathrm{d}\bar{\mathbf{u}} \\ =\; & \left\|\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right\|_{\mathrm{F}}^2 \;=\; n - \operatorname{rank}(\mathbf{U}) \,, \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The last derivation step exploits the fact that $\mathbf{U}\mathbf{U}^{+}$ equals an n-dimensional diagonal matrix $\mathbf{I}_k$ with ones on the first k diagonal elements and zeroes afterwards, for k being the rank of $\mathbf{U}$.

Hence, the rank of the deformation operator $\mathbf{U}$ characterizes the potential for design space exploration. We therefore define the variability V based on the rank of $\mathbf{U}$, but normalize it by phenotype dimension in order to scale the values to $[0, 1]$:

$$V(\mathbf{U}) \;:=\; \frac{\operatorname{rank}(\mathbf{U})}{n} \,. \tag{7}$$

According to this definition, a representation with good variability should have maximum rank m, i.e., all deformation/optimization parameters should be truly independent. The maximum theoretical variability $V(\mathbf{U}) = 1$ is achieved for $m = n$, but typically $V(\mathbf{U}) \ll 1$ due to $m \ll n$.

B. Regularity

Regularity should characterize the numerical stability of the deformation process on the one hand, and the expected speed of convergence on the other hand. A function is considered numerically stable if a small (relative) change in the input leads to a small (relative) change in the output [29]. We want to characterize the stability of the mapping $\mathbf{p} \mapsto \mathbf{U}\mathbf{p}$, i.e., we analyze the variation of the displacement $\mathbf{u} = \mathbf{U}\mathbf{p}$ due to a change in the parameters $\mathbf{p}$. In a numerically stable deformation representation each deformation parameter p_j should have about the same influence on the resulting deformation $\mathbf{U}\mathbf{p}$. More formally, a change in genotype $\mathbf{p} \mapsto \mathbf{p} + \boldsymbol{\delta}$ leads to a corresponding change in phenotype $\mathbf{U}(\mathbf{p} + \boldsymbol{\delta}) - \mathbf{U}\mathbf{p} = \mathbf{U}\boldsymbol{\delta}$. For a stable representation the *amount* of phenotype change $\|\mathbf{U}\boldsymbol{\delta}\|$ should depend on the *amount* of genotype change $\|\boldsymbol{\delta}\|$ only—and not on the direction of $\boldsymbol{\delta}$ in genotype space.

In our case of a linear function, the numerical stability is measured by the condition number of the involved matrix:

$$\kappa(\mathbf{U}) = \frac{\sigma_{\max}}{\sigma_{\min}} \,,$$

where $\sigma_{\max}$ and $\sigma_{\min}$ denote the smallest and largest singular value of the matrix $\mathbf{U}$. In fact, the condition number measures the ratio of maximum to minimum phenotype variation for unit-norm genotype changes: $\|\mathbf{U}\boldsymbol{\delta}_{\max}\| \,/\, \|\mathbf{U}\boldsymbol{\delta}_{\min}\|$, where $\boldsymbol{\delta}_{\max}$ and $\boldsymbol{\delta}_{\min}$ correspond to the maximum/minimum right singular vectors of the matrix $\mathbf{U}$ [28].

It might be tempting to incorporate the fitness function f into the regularity criterion, for instance by analyzing the numerical stability of the mapping $\mathbf{p} \mapsto f(\mathbf{x} + \mathbf{U}\mathbf{p})$, but we cannot assume analytic knowledge of the fitness function. It is known, however, that basic evolutionary optimizations perform better if all genotype parameters have a similar influence on the phenotype variation. A few dominant parameters might slow down the optimization or otherwise require sophisticated adaptation techniques.

For evolutionary processes, *locality* is known to characterize the convergence speed. Mappings with strong locality preserve the neighborhood of a genotype in the phenotype, which is beneficial for evolutionary optimization [13, 18]. This is formulated in Rothlauf's definition of locality for discrete genotypes/phenotypes [14]:

$$\sum_{d^g_{x,y} = d^g_{\min}} \left| d^p_{x,y} - d^p_{\min} \right| \,,$$

where $d^g_{x,y}$ defines the distance between the discrete genotypes x and y, $d^p_{x,y}$ the distance between their discrete phenotypes,

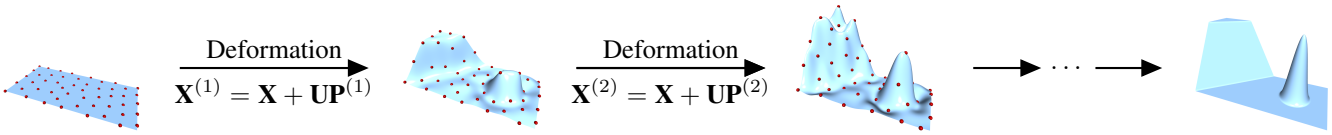

Fig. 3. The 1D function approximation scenario: The initial plane $\mathbf{X}$ is deformed via a linear deformation RBF $\mathbf{UP}$. The handle displacements $\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \ldots$ (red dots) of the direct RBF manipulation result in the deformed meshes $\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \ldots$ aiming to improve the fitting accuracy to a target height field.

and $d^g_{\min}$ and $d^p_{\min}$ the minimal values of these distances. Assuming genotype variations of equal amount ($d^g_{x,y} = d^g_{\min}$) the optimal/minimal locality is achieved if the corresponding phenotype variations also have the same amount ($d^p_{x,y} = d^p_{\min}$). Using the condition number as a regularity criterion can therefore also be understood as a generalization of the concept of locality from *discrete* genotypes/phenotypes to *continuous* optimizations.

Thus, we define regularity based on the condition number, but use the inverse in order to bound the values to $[0, 1]$, with 0 being the worst value and 1 the optimal value:

$$R(\mathbf{U}) \;:=\; \frac{1}{\kappa(\mathbf{U})} = \frac{\sigma_{min}}{\sigma_{max}}\,. \tag{8}$$

C. Improvement Potential

While a high variability allows the optimization to eventually find beneficial shape variations, and a high regularity suggests that it does so rather efficiently, both criteria cannot guarantee that the optimization performs well for the specific problem at hand—since both variability and regularity are agnostic of the fitness function. Assuming some (approximate) knowledge of the fitness function, the criterion *improvement potential* is meant to estimate the potential of the representation for improving the design w.r.t. the fitness function.

In an optimization process some regions of the phenotype might already be close to optimal, while other parts still have to be improved further. A successful representation should then allow for and promote these particular required shape variations. Analyzing whether the representation allows the optimizer to push the design towards beneficial configurations (larger fitness value) requires knowledge of ascent directions in genotype space.

The direction that locally improves the fitness function the most is its gradient with respect to genotype parameters, i.e., $\nabla_{\mathbf{p}} f(\mathbf{x} + \mathbf{Up}) := \partial f(\mathbf{x} + \mathbf{Up})/\partial \mathbf{p}$. Since the analytic fitness gradient is not known in most cases, we assume that at least an approximate gradient direction $\mathbf{g}$ is available, for instance through learning from previous optimization runs or adjoint optimization approaches.

Given the improvement direction $\mathbf{g}$, we measure how well the representation can approximate it as $\mathbf{Up}$. We proceed similar to Section IV-A, assume $\|\mathbf{g}\| = 1$, and find the least squares approximation error to be $\left\|\left(\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right)\mathbf{g}\right\|^2$. Since $\mathbf{g}$ is normalized, this error is bounded from above by 1. To have the criterion in the intuitive range $[0, 1]$, with 1 being the optimal value, we define the improvement potential $P(\mathbf{U})$ (for a given unit-norm improvement direction $\mathbf{g}$) as

$$P(\mathbf{U}) \;:=\; 1 - \left\|\left(\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right)\mathbf{g}\right\|^2\,. \tag{9}$$

The straightforward generalization to vector-valued deformation functions will be given in Section VI. After having defined quantitative formulations for the three evolvability criteria *variability*, *regularity*, and *improvement potential*, we will in the following evaluate how well these criteria are able to predict the quality and performance of different representation setups for 1D function approximation (Section V) and 3D template matching (Section VI).

V. Test Scenario: 1D Function Approximation

Empirically evaluating our evolvability criteria within an automotive design optimization scenario is impractical, since (1) the complicated fitness function inevitably requires a computationally expensive and thus slow optimization process, and (2) no ground truth in the form of an analytically known global optimum of the fitness function exists.

As a first evaluation scenario we have therefore chosen a simple least-squares function approximation problem. Starting from a plane discretized by a regular grid of 150×150 vertices $\mathbf{x}_i = (x_i, y_i)^{\mathrm{T}}$, we use a RBF function $u\colon \mathbb{R}^2 \to \mathbb{R}$ to approximate a given scalar height field, see Figure 3. The two test functions to be approximated are a simple sine wave

$$s(x, y) \;=\; \sin(\pi \cdot (x + y))^2$$

with $(x, y) \in [0, 1]^2$ and a more complex function used by Giannelli et al. [30]:

$$s(x, y) \;=\; \begin{cases} 0.5\cos(4\pi \cdot q^{0.5}) + 0.5 & q < \frac{1}{16}\,, \\ 2(y - x) & 0 < y - x < 0.5\,, \\ 1 & 0.5y - x\,, \end{cases}$$

with $(x, y) \in [0, 2] \times [0, 1]$ and $q = (x - 1.5)^2 + (y - 0.5)^2$, see Figure 4. We discretize the target functions using the same 150×150 points $\mathbf{x}_i$ and define scalar height values $s_i := s(\mathbf{x}_i)$.

The fitness function $f(\mathbf{p})$ measures the least squares approximation error between the target values s_i and the current "design" $u_{\mathbf{p}}(\mathbf{x}_i)$, and has to be minimized with respect to the deformation parameters $\mathbf{p} = (p_1, \ldots, p_m)$:

$$f(\mathbf{p}) \;=\; \sum_{i=1}^{n} \left(u_{\mathbf{p}}(\mathbf{x}_i) - s_i\right)^2 \;=\; \|\mathbf{Up} - \mathbf{s}\|^2 \;\to \min\,,$$

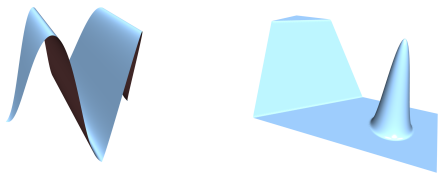

Fig. 4. Scalar test functions: Sine wave (left), Giannelli function [30] (right).

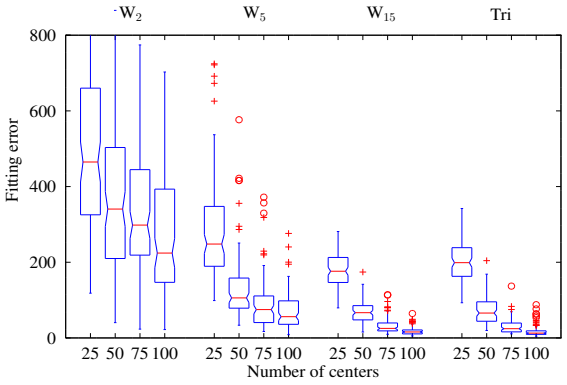

	W_2	W_5	W_{15}	Tri
Giannelli	.37 (0)	.68 (0)	.88 (0)	.88 (0)
Sine	.31 (0)	.43 (0)	.79 (0)	.93 (0)

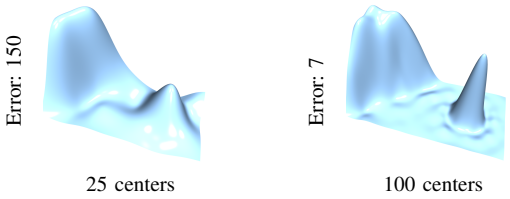

Fig. 5. Variability results for 1D function approximation. The table shows the Spearman's correlation (and p-values) between variability and the fitting error. For the widest Wendland kernel W_{15} and for the triharmonic kernel the correlation is very strong. The box-plot (left) visualizes the trend for the Giannelli test function, showing that a higher variability results in a better fit. The pictures (right) show two fitting examples for the triharmonic kernel.

with $\mathbf{s} = (s_1, \ldots, s_n)^{\mathrm{T}}$. The deformation matrix $\mathbf{U}$ is defined according to either indirect or direct RBF method, as shown in equations (2) and (4), respectively.

In this particularly simple optimization scenario the global optimum can directly be computed as in (1), leading to $\mathbf{p} = \mathbf{U}^{+}\mathbf{s}$ and $\mathbf{u} = \mathbf{U}\mathbf{U}^{+}\mathbf{s}$. We will use this solution as a reference when evaluating the evolvability criteria, since this allows us to compare the evolutionary solution to the analytic one.

In order to experimentally evaluate the evolvability criteria defined in Section IV, we generate a large variation of representation setups using different kernel types, different numbers of kernels, different support radii, and direct or indirect manipulation. We employ both the global triharmonic kernel and compact (local) Wendland kernels of varying support radius. To test the regularity and the improvement potential criterion we randomly generate 100 different center setups for each kernel function, each setup consisting of 25 centers. For the variability analysis we expand the tests to 50, 75, and 100 centers per setup. In the case of the compact Wendland kernels we set the support radii s (identical for all centers), such that each point $\mathbf{x}_i$ of the design mesh is overlapped (and hence can be varied) by at least l RBF kernels. We chose l to be 2, 5, and 15, such that we can distinguish between more local ($l = 2$) and more global ($l = 15$) setups. We denote the triharmonic kernel with *Tri* and the different Wendland kernels with W_l in the plots and tables. Analyzing these 4 kernels for both direct and indirect manipulation on 100 random setups leads to 800 deformation representations to be evaluated in total.

We then analyze how well our evolvability criteria predict the actual quality/performance of the representation by computing the Spearman correlation (using R [31]). This test analyzes the monotonic correlation between two quantities and calculates its significance (p-value). The intervals $[0, 0.2[$, $[0.2, 0.4[$, $[0.4, 0.6[$, $[0.6, 0.8[$, and $[0.8, 1]$ are classified as very weak, weak, moderate, strong, and very strong correlation [32]. We interpret p-values smaller than .01 as significant and round p-values smaller than 10^{-4} to 0. Note that our evolvability criteria should be maximized (with 1 being the optimal value), while the properties they characterize (e.g., fitting error, numbers of iterations) should be minimized. Hence, large negative correlations are better, but in order to simplify interpreting the numbers using the above intervals, we negate all reported correlation coefficients, such that 1 is the optimal value.

A. Results: Variability

Variability characterizes the potential for design space exploration, and is measured as the (normalized) rank of the deformation matrix $\mathbf{U}$, see (7). Representations with a high variability are expected to result in more accurate fits. A closer look at the deformation matrices of indirect manipulation ($\mathbf{\Phi}$) and direct manipulation ($\mathbf{\Phi}\mathbf{\Psi}^{-1}$), see (2) and (4), reveals that both matrices have the same rank, which is the dimension of the range of $\mathbf{\Phi}$. For the same setup of RBF centers, direct and indirect manipulation thus have identical variability.

In all our experiments the deformation matrices generated from random center setups had full rank m, such that the variability depends on the numbers of kernels, but not on their type and placement (as long as kernels do not coincide). We therefore increase the variability by adding more RBF centers. Since we characterize the *potential* for accurate fits, instead of the actual result of an evolutionary process, we compute the optimal fitting errors using the analytic solution and thereby rule out any negative effects of a randomized search.

The table in Figure 5 shows the correlation (and its p-values in brackets) between fitting accuracy and variability. The correlation is significant for all basis functions, very strong for triharmonic kernel and the widest Wendland kernel (W_{15}), and weak for the local ones. The box-plot, visualizing this trend for the Giannelli test (Figure 5, left), depicts a widely spread fitting error for local kernels (W_2), which decreases with increasing support radius (W_{15}). That is because local basis function with a small support are very sensitive to the random center distribution whereas global kernels are more robust. For instance, the chance to randomly place centers in already optimal regions, such they do not influence the fitting result, is much higher for local basis functions. Hence, we can state that global kernels are typically more accurate. The results for the sine function are very similar, such that we omit their visualization.

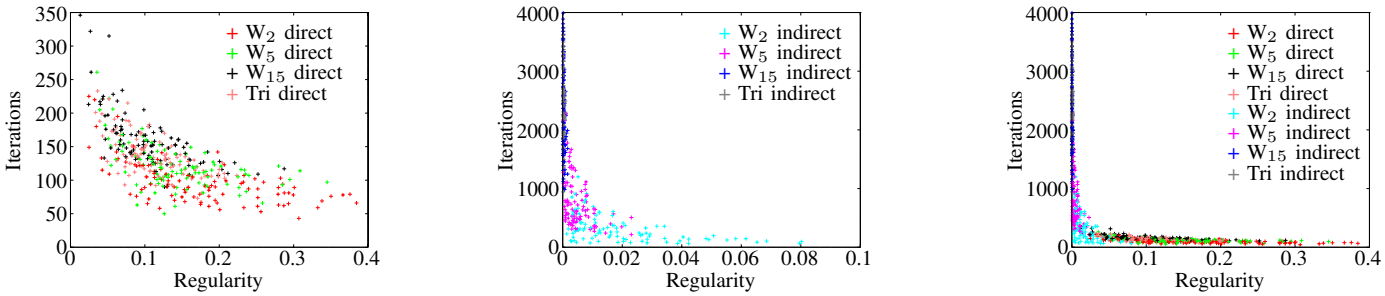

Fig. 6. Relation between regularity and convergence speed (#iterations until convergence) for the Giannelli test function. Left: direct manipulation, center: indirect manipulation, right: comparison of direct and indirect manipulation. It can be observed that a higher regularity results in a faster convergence, and that direct manipulation and local kernels are preferable in terms of convergence speed.

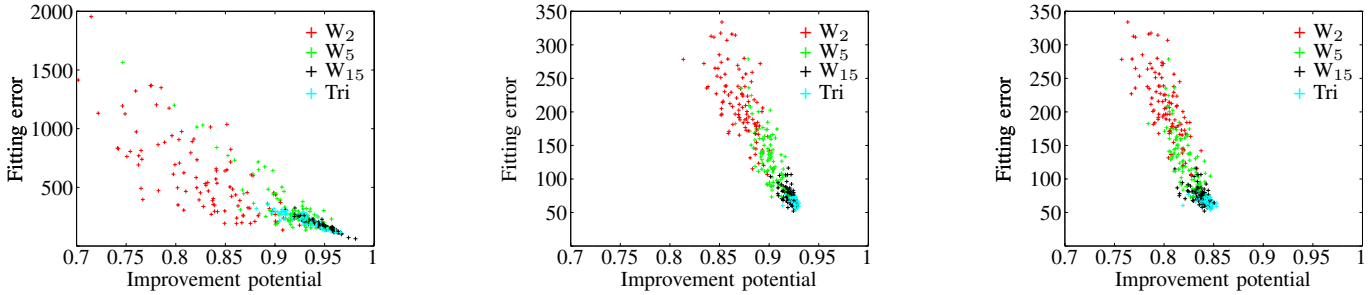

Fig. 7. Relation between the improvement potential and the fitness value (approximation error) scenario for the Giannelli function (left: $\lambda = .5$) and for template fitting (middle: $\lambda = .5$, right: $\lambda = 1$). The plots clearly demonstrate the validity of our improvement potential criterion.

B. Results: Regularity

Regularity is meant to estimate the expected convergence speed, and is computed as the inverse condition number $\kappa(\mathbf{U})$, see (8). As evolutionary optimization we use the (1,10)-CMA-ES (covariance matrix adaptation evolution strategy) of the shark3.0 library [33], since we want to use this algorithm in the automotive design optimization later on. We measure convergence speed by counting the number of iterations until the optimization converges. We consider the algorithm converged as soon as the optimizer reaches a fitness value that is within a 5% tolerance of the true analytic solution.

Figure 6 plots the regularity criterion against the number of iterations until convergence for the different kernel functions on the example of the Giannelli test function. The results for the sine test function are very similar, such that we omit the plot. Table I summarizes the Spearman's correlations of the regularity criterion and the required numbers of iterations for all employed kernel types. Only 3 out of the 16 tests are not significant and are marked in red.

TABLE I. SPEARMAN'S CORRELATION (AND P-VALUES) BETWEEN REGULARITY AND CONVERGENCE SPEED FOR THE 1D FUNCTION APPROXIMATION PROBLEM.

	Tri	W_{15}	W_5	W_2	all
Giannelli					
direct	.63 (0)	.71 (0)	.51 (0)	.60 (0)	.91 (0)
indirect	.22 (.027)	.39 (0)	.44 (0)	.53 (0)	
Sine					
direct	.80 (0)	.64 (0)	.82 (0)	.58 (0)	.87 (0)
indirect	.2 (.05)	.13 (.19)	.51 (0)	.60 (0)	

For direct manipulation we observe a strong correlation between regularity and convergence speed for all basis functions. This correlation becomes weaker, but is still moderate, for the significant indirect manipulation results. When comparing indirect and direct manipulation, Figure 6 reveals that in all our experiments direct manipulation setups have a better regularity and converge faster than the indirect ones, which is in line with the results of [25].

The relation of our regularity criterion with Rothlauf's locality measure [14] motivates the analysis of Wendland kernels with different support radii. It can be observed in Figure 6 that more local kernels converge faster than more global ones, which is also hinted at by their better regularity values. In Figure 6 (right) all kernels (direct and indirect manipulation) are plotted together, showing a strong correlation between regularity and convergence speed for a range of different methods, which is also confirmed by the Spearman's coefficients in Table I (column 6).

Overall it can be stated that our regularity criterion is a good indicator for the convergence speed of a deformation representation. Our direct-vs-indirect and local-vs-global results are in agreement with results known from the literature, stating that w.r.t. convergence speed direct manipulation and local kernels should be preferred.

C. Results: Improvement Potential

The improvement potential estimates how much a given representation can potentially improve the fitness value, based on how accurate a given *approximate* gradient $\mathbf{g}$ can be reproduced, as described in (9). We expect deformation setups that can approximate the direction $\mathbf{g}$ well to result in solutions with a better fitness value. To emulate this approximate knowledge

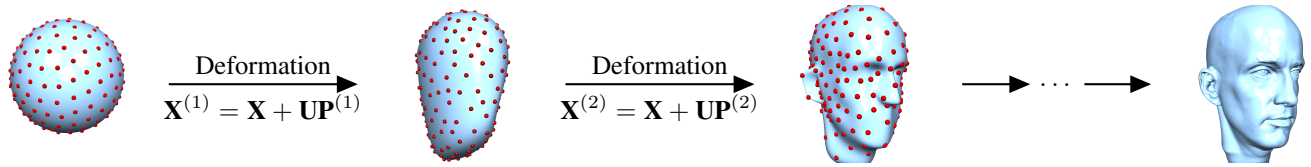

Fig. 8. The 3D template fitting scenario: An initial sphere $\mathbf{X}$ is deformed via a linear deformation function $\mathbf{UP}$. The handle displacement $\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \ldots$ (red dots) results in the deformed meshes $\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \ldots$ aiming to improve the fitting accuracy to a target scan.

of a beneficial direction, we define $\mathbf{g}$ to be a smoothed version of the true analytic gradient $\nabla_{\mathbf{p}} f$:

$$\tilde{\mathbf{g}} \;=\; \lambda \frac{\nabla_{\mathbf{p}} f}{\|\nabla_{\mathbf{p}} f\|} + (1-\lambda)\frac{\mathbf{1}}{\|\mathbf{1}\|}\,, \qquad \mathbf{g} \;=\; \frac{\tilde{\mathbf{g}}}{\|\tilde{\mathbf{g}}\|}\,,$$

where $\mathbf{1} = (1, \ldots, 1)$ is a vector with n ones. The blending parameter λ models the "reliability" of the approximate gradient $\mathbf{g}$, which matches the exact gradient for $\lambda = 1$ and does not contain any fitness-related knowledge for $\lambda = 0$. We conduct our experiments with $\lambda = 0.5$ and $\lambda = 0.75$.

A first observation is that for a fixed center placement the improvement potential is identical for direct and indirect manipulation (analogously to the variability criterion), such that we only have to consider different kernel types, support radii, and kernel placements.

Table II shows the Spearman's correlation between our improvement potential and the final fitting accuracy. Since we want to measure the *potential* for improvement, we measure the fitting accuracy by the distance to the (known) analytic solution. Since the (quadratic) function approximation problem can be solved by setting the fitness gradient to zero, the correlation is 1 if a precise gradient is known ($\lambda = 1$), which is not shown. For a small distortion ($\lambda = .75$) the correlation between our criterion and the resulting fitting error still is close to optimal. Even tests with a larger distortion ($\lambda = .5$) result in a strong correlation for the Giannelli function. Although the correlation values for the sine test function are worse, all results are still significant for both test functions. Figure 7 plots the improvement potential against the resulting fitting error. The left image depicts the very strong correlation for the Giannelli test function with the different kernel functions. It can be observed that the triharmonic kernel and the Wendland kernel with largest support approximate the estimated gradient direction better and also result in a better solution. A similar trend can be observed for the sine test function as well.

Overall, these results clearly demonstrate that the idea of estimating the improvement potential by approximating an (approximate) fitness gradient works well even for imprecise gradient information, which might very well be available in practical real-world applications.

TABLE II. SPEARMAN'S CORRELATION (AND P-VALUES) BETWEEN IMPROVEMENT POTENTIAL AND FITTING ACCURACY FOR THE 1D FUNCTION APPROXIMATION PROBLEM.

	Tri	W_{15}	W_5	W_2	all
Giannelli test					
$\lambda = .75$	.99 (0)	.99 (0)	.98 (0)	.99 (0)	.99 (0)
$\lambda = .5$	.91 (0)	.95 (0)	.71 (0)	.75 (0)	.89 (0)
Sine test					
$\lambda = .75$	1.00 (0)	.97 (0)	.96 (0)	.98 (0)	.98 (0)
$\lambda = .5$	.99 (0)	.76 (0)	.32 (0)	.34 ($5{\cdot}10^{-4}$)	.86 (0)

VI. TEST SCENARIO: 3D TEMPLATE FITTING

In the previous 1D function approximation scenario we were able to analytically compute the global optimum, which we exploited for the analysis of the evolvability criteria. The 3D template fitting scenario described in this section is considerably more complex, since the fitness function has many local minima and the global solution is not known. The goal is to fit a given triangle mesh $\mathbf{X}$ (a sphere in our experiments) to a target face scan $\mathbf{T}$ using an RBF deformation $\mathbf{u}\colon \mathbb{R}^3 \to \mathbb{R}^3$, as depicted in Figure 8. The sphere model is discretized using $n \approx 10k$ vertices, the scan consists of $m \approx 12k$ points.

In contrast to the 1D function approximation problem, where each point $(x, y, 0)^{\mathrm{T}}$ on the plane corresponds to a point $(x, y, s(x, y))^{\mathrm{T}}$ on the height field, there is no such one-to-one correspondence between the vertices of the sphere and points of the scan. Hence, the fitness function, which measures the approximation error between the two models, computes distances between any point on the (deformed) sphere and its closest point of the scan, and vice versa. These closest points are denoted by

$$\mathbf{c}_{\mathbf{T}}(\mathbf{x}_i) = \arg\min_{\mathbf{t}_j \in \mathbf{T}} \|\mathbf{x}_i - \mathbf{t}_j\|\,, \quad \mathbf{c}_{\mathbf{X}}(\mathbf{t}_j) = \arg\min_{\mathbf{x}_i \in \mathbf{X}} \|\mathbf{t}_j - \mathbf{x}_i\|\,.$$

These closest-point-pairs vary during the iterative optimization depending on the deformed sphere $\mathbf{X}^{(k)}$ at iteration k. This yields the (non-static) fitness function

$$f(\mathbf{P}) \;=\; \frac{1}{n}\sum_{i=1}^{n} \left\| \mathbf{x}_i + \mathbf{u}_{\mathbf{P}}(\mathbf{x}_i) - \mathbf{c}_{\mathbf{T}}\left(\mathbf{x}_i^{(k)}\right) \right\|^2 + \frac{1}{m}\sum_{j=1}^{m} \|\mathbf{t}_j - \mathbf{c}_{\mathbf{X}^{(k)}}(\mathbf{t}_j)\|^2\,, \tag{10}$$

where $\mathbf{x}_i^{(k)}$ denotes the i-th vertex of the mesh $\mathbf{X}^{(k)}$. Note that for the sake of a simpler formulation, we do not incorporate any regularization terms, which otherwise should be used to prevent over-fitting and yield a higher surface quality, as discussed, e.g., in [34].

In the evolutionary setting, we again employ the CMA-ES of the shark library [33] to minimize the fitness function (10). Since we cannot compute the analytic solution as a ground truth reference for analyzing the evolvability criteria, we fall back to the gradient-based Gauss-Newton approach of [34] for computing the reference solution.

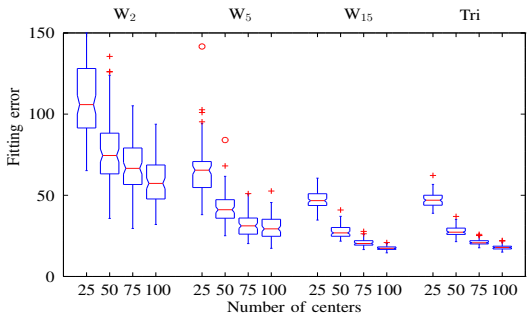

	W_2	W_5	W_{15}	Tri
Template fit	.65 (0)	.77 (0)	.95 (0)	.95 (0)

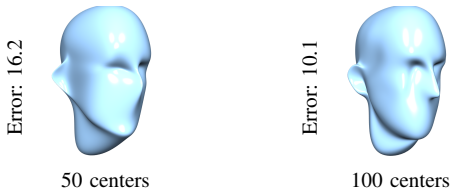

Fig. 9. Variability results for the 3D template fitting. For the widest Wendland kernel W_{15} and for the triharmonic kernel the correlation is strongest.

The analysis of the *variability* criterion confirms the findings of the previous section: Since the RBF matrices have full rank, adding more RBF kernels directly increases the variability, which correlates (very) strongly with the fitting accuracy, see Figure 9.

Analyzing the relation of our *regularity* criterion and convergence speed reveals significant correlation in all but one case (indirect manipulation, triharmonic kernel), see Table III. Plotting regularity against numbers of required iterations (Figure 10) shows the same trend as in the 1D function approximation scenario, which confirms that our regularity formulation characterizes convergence speed even in more complex scenarios.

We have to slightly adjust the formulation of the improvement potential, since the 3D deformation leads to an $(n \times 3)$-dimensional fitness gradient, such that we replace the vector norm of equation (9) by the Frobenius norm. Given a normalized *approximate* gradient $\|\mathbf{G}\|_F = 1$, the improvement potential is

$$P(\mathbf{U}) := 1 - \left\| \left(\mathbf{I} - \mathbf{U}\mathbf{U}^{+}\right) \mathbf{G} \right\|_F^2 . \qquad (11)$$

Since the closest-point-correspondences change during the optimization procedure, even the exact gradient of the initial state $\mathbf{X}$ will be inaccurate after the first few iterations. We therefore use this exact initial gradient ($\lambda = 1$), as well as distorted version ($\lambda = 0.5$ and $\lambda = 0.75$) to evaluate the improvement potential. The correlation between the improvement potential and the final fitting accuracy is significant for all kernel types. For the Wendland kernel the correlation is strong, whereas the correlation is weaker for the triharmonic kernel, see Table IV. We can again observe a strong correlation when we analyze the test samples of all kernels (see Figure 7, center and right). These results demonstrate again that the approximate gradient information $\mathbf{G}$ can be rather inaccurate and will still lead to a valid prediction of the eventual fitness.

TABLE III. SPEARMAN'S CORRELATION (AND P-VALUES) BETWEEN REGULARITY AND CONVERGENCE SPEED FOR THE 3D TEMPLATE FITTING.

Template fit	Tri	W_{15}	W_5	W_2	all
direct	.86 (0)	.71 (0)	.56 (0)	.47 (0)	.87 (0)
indirect	.11 (.3)	.42 (0)	.48 (0)	.34 (10^{-4})	

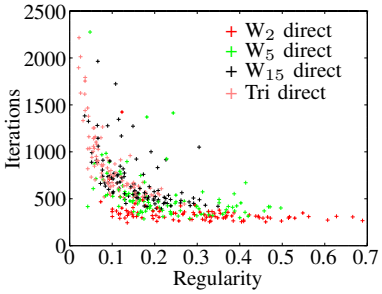

Fig. 10. The correlation between regularity and convergence speed for the 3D template fitting is equivalent to the results of the 1D function approximation case. We plot the results of direct manipulation only, since the poor convergence of the indirect triharmonic configuration (200k iterations) would otherwise dominate the plot.

TABLE IV. CORRELATIONS (AND P-VALUES) BETWEEN IMPROVEMENT POTENTIAL AND FITTING ACCURACY FOR 3D TEMPLATE FITTING.

Template fit	Tri	W_{15}	W_5	W_2	all
$\lambda = 1$	.38 (10^{-4})	.61 (0)	.74 (0)	.71 (0)	.93 (0)
$\lambda = .75$	.38 (10^{-4})	.60 (0)	.74 (0)	.74 (0)	.91 (0)
$\lambda = .5$	.38 (10^{-4})	.57 (0)	.67 (0)	.66 (0)	.90 (0)

The analysis of the 3D template fitting scenario confirms the results of the simpler 1D function approximation setting: Our formulations of variability, regularity, and improvement potential indeed represent reliable criteria for evaluating the quality of representation setups. The regularity experiments again reveal that direct manipulation performs better than indirect manipulation. Regarding variability and improvement potential, there is no difference between the two manipulation approaches, which thus motivates the use of direct manipulation for future evolutionary optimizations. Also in agreement with the 1D function approximation results, local kernels converge faster, whereas global kernels result in more accurate approximations, which leads to (interesting) conflicting goals when setting up deformation representations.

VII. SUMMARY AND FUTURE WORK

A smart representation design tremendously supports the efficiency of industrial product optimization. Inspired by the biological concept of evolvability we propose a mathematical model for quantifying the quality of linear deformation representations. Our formulation is based on the three characteristics variability, regularity, and improvement potential.

Variability covers the design space exploration potential independently of the fitness or objective function. Its definition is based on the rank of the deformation matrix, and our results confirm that increasing the operator's variability improves the quality of the solution.

Regularity, which we define as the inverse of the condition number of the deformation matrix, characterizes the expected convergence speed of the optimization algorithm. Our experiments show a significant and strong correlation between regularity and convergence speed, and reveal that direct manipulation and local RBF kernels have the highest regularity and converge the fastest.

Improvement potential depends on the ability to reproduce an approximate fitness gradient, implying a numerical quantification of most beneficial local variations. Our experiments support that already a coarse gradient information yields significant correlations between improvement potential and fitting accuracy. Best results are achieved by global RBF kernels.

Based on these findings, two research directions seem most promising: First, since variability, regularity, and improvement potential are partly contradicting, a multi-objective optimization should be used to reveal the Pareto front of RBF representation setups. Second, we will analyze how well our evolvability criteria generalize to different kinds of linear deformation methods, such as free-form 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] A. A. Mina, D. Braha, and Y. Bar-Yam, *Complex Engineered Systems: Science Meets Technology*. Springer, 2006, ch. Complex Engineered Systems: A New Paradigm, pp. 1–21.

[2] G. P. Wagner and L. Altenberg, "Perspectives: Complex adaptations and the evolution of evolvability," *Evolution*, vol. 50, no. 3, pp. 967–976, 1996.

[3] D. Sieger, S. Menzel, and M. Botsch, "A comprehensive comparison of shape deformation methods in evolutionary design optimization," in *Proceedings of the International Conference on Engineering Optimization*, 2012.

[4] D. Sieger, S. Menzel, and M. Botsch, "On shape deformation techniques for simulation-based design optimization," in *New Challenges in Grid Generation and Adaptivity for Scientific Computing*. Springer, 2015, pp. 281–303.

[5] M. Kirschner and J. Gerhart, "Evolvability," *Proceedings of the National Academy of Sciences*, vol. 95, no. 15, pp. 8420–8427, 1998.

[6] A. Richter, M. Botsch, and S. Menzel, "Evolvability of representations in complex system engineering: a survey," in *Proceedings of IEEE Congress on Evolutionary Computation*, 2015, pp. 1327–1335.

[7] H. Lehmann and S. Menzel, "Evolvability as concept for the optimal design of free-form deformation control volumes," in *Proceedings of IEEE Congress on Evolutionary Computation*, 2012, pp. 1–8.

[8] K. Sterelny, *Modularity in development and evolution*. University of Chicago Press, 2004, ch. Symbiosis, Evolvability and Modularity, pp. 490–516.

[9] A. Wagner, "Robustness and evolvability: a paradox resolved," *Proceedings of the Royal Society B: Biological Sciences*, vol. 275, no. 1630, pp. 91–100, 2008.

[10] J. L. Payne, J. H. Moore, and A. Wagner, "Robustness, evolvability, and the logic of genetic regulation," *Artificial Life*, vol. 20, no. 1, pp. 111–126, 2014.

[11] J. Lehman and K. O. Stanley, "Abandoning objectives: Evolution through the search for novelty alone," *Evolutionary Computation*, vol. 19, no. 2, pp. 189–223, 2011.

[12] D. Shorten and G. Nitschke, "How evolvable is novelty search?" in *Proceedings of the IEEE International Conference on Evolvable Systems*, 2014, pp. 125–132.

[13] T. Weise, R. Chiong, and K. Tang, "Evolutionary optimization: Pitfalls and booby traps," *Journal of Computer Science and Technology*, vol. 27, no. 5, pp. 907–936, 2012.

[14] F. Rothlauf, *Representations for Genetic and Evolutionary Algorithms*. Springer, 2006.

[15] B. Sendhoff, M. Kreutz, and W. Von Seelen, "A condition for the genotype-phenotype mapping: Causality," in *Proceedings of the International Conference on Genetic Algorithms*, 1997, pp. 73–80.

[16] E. Galván-López, J. McDermott, M. ONeill, and A. Brabazon, "Defining locality as a problem difficulty measure in genetic programming," *Genetic Programming and Evolvable Machines*, vol. 12, no. 4, pp. 365–401, 2011.

[17] T. Seaton, J. F. Miller, and T. Clarke, "An ecological approach to measuring locality in linear genotype to phenotype maps," in *Proceedings of the European Converence on Genetic Programming*, 2012, pp. 170–181.

[18] A. Thorhauer and F. Rothlauf, "On the locality of standard search operators in grammatical evolution," in *Proceedings of the International Conference on Parallel Problem Solving From Nature*, 2014, pp. 465–475.

[19] M. Pigliucci, "Is evolvability evolvable?" *Nature Reviews Genetics*, vol. 9, no. 1, pp. 75–82, 2008.

[20] M. Bercachi, P. Collard, M. Clergue, and S. Verel, "Do not choose representation just change: An experimental study in states based EA," in *Proceedings of Annual Conference on Genetic and Evolutionary Computation*, 2009, pp. 1799–1800.

[21] D. Tarapore and J.-B. Mouret, "Comparing the evolvability of generative encoding schemes," in *Artificial Life: Proceedings of the International Conference on the Synthesis and Simulation of Living Systems*, 2014, pp. 55–62.

[22] J. Reisinger and R. Miikkulainen, "Acquiring evolvability through adaptive representations," in *Proceedings of Annual Conference on Genetic and Evolutionary Computation*, 2007, pp. 1045–1052.

[23] D. Sieger, S. Gaulik, J. Achenbach, S. Menzel, and M. Botsch, "Constrained space deformation techniques for design optimization," *Computer Aided Design*, vol. 73, pp. 40–51, 2016.

[24] W. M. Hsu, J. F. Hughes, and H. Kaufman, "Direct manipulation of free-form deformations," in *Proceedings of ACM SIGGRAPH*, 1992, pp. 177–184.

[25] S. Menzel, M. Olhofer, and B. Sendhoff, "Direct manipulation of free form deformation in evolutionary design optimisation," in *Proceedings of the International Conference on Parallel Problem Solving From Nature*, 2006, pp. 352–361.

[26] T. W. Sederberg and S. R. Parry, "Free-form deformation of solid geometric models," in *Proceedings of ACM SIGGRAPH*, 1986, pp. 151–159.

[27] M. Botsch and O. Sorkine, "On linear variational surface deformation methods," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 1, pp. 213–30, 2008.

[28] G. H. Golub and C. F. Van Loan, *Matrix computations*. Johns Hopkins University Press, 2012.

[29] L. N. Trefethen and I. David Bau, *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.

[30] C. Giannelli, B. Jüttler, and H. Speleers, "THB-splines: The truncated basis for hierarchical splines," *Computer Aided Geometric Design*, vol. 29, no. 7, pp. 485–498, 2012.

[31] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2015. [Online]. Available: https://www.R-project.org

[32] D. I. Weir, "Spearman's correlation," 2015. [Online]. Available: http://www.statstutor.ac.uk/resources/uploaded/spearmans.pdf

[33] A. Auger and N. Hansen, "Turorial: CMA-ES: Evolution strategies and covariance matrix adaptation," in *Proceedings of the Conference on Genetic and Evolutionary Computation*, 2011, pp. 991–1010.

[34] J. Achenbach, E. Zell, and M. Botsch, "Accurate face reconstruction through anisotropic fitting and eye correction," in *Proceedings of Vision, Modeling & Visualization*, 2015, pp. 1–8.