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

3

Innovativeness in Design Optimization

Finding *innovative solutions* is an objective that holds generally for design optimization tasks, as when it is not aimed for directly, it can at least be stated to be valued in general. To be able to assist human engineers in finding innovative solutions using automated search, what exactly constitutes an innovative solution has to be determined.

[Garcia and Calantone, 2002] state in their survey on innovation terminology that, while often used interchangeably, it is important to make a distinction between the constructs *innovation* and *innovativeness* (i.e., the property of being innovative): An innovative solution becomes an innovation once it has passed through production and marketing phases and is diffused into the marketplace. Factors accounting for a swift diffusion of such an invention into society are addressed in [Rogers, 2010]. Furthermore, [Garcia and Calantone, 2002] stress to clearly distinguish between innovativeness relating to *products* and *organizational* innovativeness: Organizational innovativeness involves the *tendency* of an organization to develop or adopt innovative products, i.e., lead to or participate in an innovation, while the innovativeness of a product that a firm markets or adopts *cannot* be regarded as a measure of organizational innovativeness. In this thesis, innovativeness is addressed from its viewpoint of being a property of technological products.

[Deb and Srinivasan, 2006] describe the practice of using automated search for uncovering innovative solutions. Their approach, termed *innovization* (e.g., [Deb and Srinivasan, 2006, Deb, 2013]), is reformulating optimization problems involving a single search objective as a *multiobjective* problem with two or more conflicting search

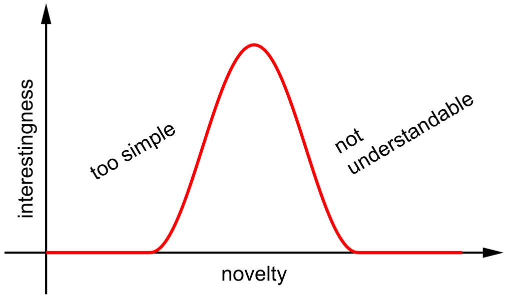

Figure 3.1 Novelty vs. Interestingness. Interestingness shown as a function of novelty. Image courtesy of [**Graziano et al., 2011**].

objectives, for instance utility and cost. In optimizing on these objectives in parallel, the aim is to best-approximate the optimal multiobjective set of solutions that all have a different trade-off in objective values: Of each such a solution cannot be stated that another solution in the set is better with respect to all objectives. This approximated *non-dominated* set is then analyzed to find common design principles shared by the different solutions in the set. Potentially, this uncovers common solution principles, new to the human engineer, that lead to high-performing designs. Thus, innovization is an approach of using automated search for gaining a deeper understanding of the optimization problem at hand, potentially offering new insights that *can* lead to innovative products.

For a definition that allows for quantifying innovativeness, we turn to the discussion on *creative* products in [**Saunders, 2002**]. Generally, creative products are recognized as being *useful* and *novel*. Usefulness has to do with being appropriate, valuable, or aesthetic, while novelty is related to originality, unexpectedness, and surprisingness. Furthermore, a creative product is likely to remain useful for some time, whereas its novelty dissipates much quicker [**Saunders, 2002**]. We take innovativeness as synonymous to creativity in this context, and thus depending on utility and the more volatile novelty.

Other work describing novelty of solutions contrast it to the actual *interestingness* of a solution, related to learning and model building. Novelty of an observation, action, or solution, is a subjective notion in the sense that it depends on *similarity* to what was encountered before, and on the *state* of the agent, approach, or model used in *determining* this similarity. Easily modeled and frequently occurring solutions quickly lose their novelty, while solutions that are subject to *noise* always remain novel

Table 3.1 Terminology. Overview of the terms that are used in defining *innovative solutions* and their likelihood of being adopted by the engineering community, to get further-developed into an innovative product.

Term	Description
Solution	Product proposal, product design.
State-of-the-art	Highest-quality solutions, generally accepted as such, in the application domain.
Innovation	The process of an innovative product getting adopted by the customer base, i.e., the end users.
Innovative, innovativeness	1) *Novel* and 2) of *better or tolerably worse quality* compared to 3) the *state-of-the-art*; synonymous to creative.
Novel, novelty	Difference with respect to a *reference set*, expressed using a *domain-specific distance measure*.
Interestingness	Increase of understanding that is given rise to, with respect to a *reference model*; corresponds to: Sufficient, but not too extreme novelty.

[**Graziano et al., 2011**]. For efficient learning or modeling of a solution space, those solutions are selected at each time step that improve the understanding or model of the solution space most: An interesting solution has sufficient (not too simple, boring) but not too extreme novelty (beyond understandable, for instance, noise) [**Schmidhuber, 1997**]. In learning, not the most-novel but the most-interesting solutions are to be selected, see Figure 3.1.

In the sections that follow, we further relate these notions to the aim of using automated search to assist in finding innovative solutions. We start by defining novelty, innovativeness, and interestingness in Section 3.1. An overview of the terminology that will be used is given in Table 3.1. Next, in Section 3.2, it is concluded that finding innovative solutions is best done indirectly through aiming for diversity. Possible approaches for efficiently searching for diverse solutions, inspired by the discussion on novelty and interestingness, are described in Section 3.3. The chapter is summarized in Section 3.4.

3.1 · Novelty, Innovativeness, and Interestingness

The aim of this thesis is to develop an automated search method that assists a human engineer in finding *innovative solutions*. Here, a *solution* is to be viewed as a *product proposal* or product design, to be further developed into an actual product by a human domain expert. Thus, *product innovativeness* is meant, related to technological products, and not *organizational* innovativeness from a business perspective (see [**Garcia and Calantone, 2002**]). An innovative product, i.e., developed

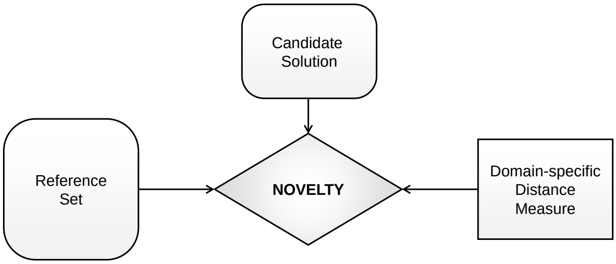

Figure 3.2 Novelty. The novelty of a candidate solution is expressed with respect to a reference set of solutions, and is quantified as the domain-specific distance to the reference solution that is closest to the candidate solution.

from an innovative solution, has the potential of leading to an *innovation*, which is roughly defined as an innovative product attaining market penetration [**Garcia and Calantone, 2002**]. A solution's innovativeness should thus quantify this potential for the product that the solution gives rise to.

To express innovativeness of a solution, *novelty* of a solution is introduced first. We define novelty as a property that is expressed relative to a *comparison* or *reference set* of solutions, by applying a *distance measure* appropriate for the application domain of the solutions. Such a distance measure is required to highlight those aspects of the resulting product that are relevant from the application perspective (e.g., depending on the application, the color of an automobile is not, or actually might be relevant). This is formalized as follows, see also Figure 3.2.

Notion 3.1 Novelty. *The novelty of a solution* $\mathbf{x}$ *is defined as its distance to solutions* $\mathbf{x}_{\text{ref}}$ *in a reference set* S_{ref}*, with respect to a domain-specific distance measure* d*. A solution is generally represented by a vector of values.*

For quantification, a *novelty measure* is required that derives from the measured domain-specific distances between the candidate solution $\mathbf{x}$ and the solutions $\mathbf{x}_{\text{ref}}$ in the reference set. We adopt the *smallest observed distance* (referred to as *uniqueness* in Chapter 4) as novelty score.

Like novelty, innovativeness is expressed with respect to a reference set of solutions. We define an innovative solution to be of better, or tolerably less quality than the solutions in this reference set. "Tolerably less" means still sufficient for adoption by a human domain expert, relative to the quality of the solutions in the reference set. A solution's innovativeness can then be quantified by its novelty with respect to the solutions in the reference set.

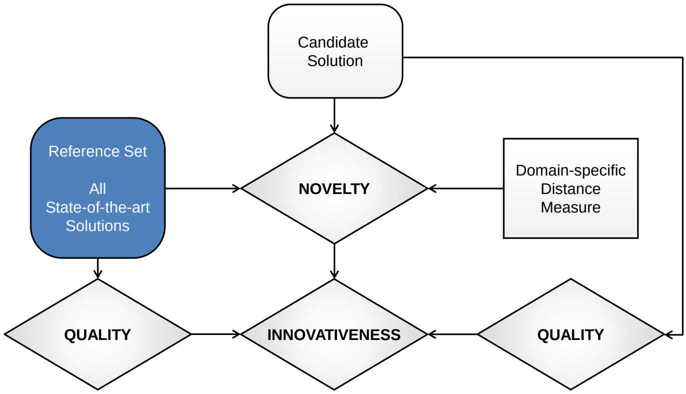

Figure 3.3 Innovativeness. To determine its innovativeness, a candidate solution is compared to the set of state-of-the-art solutions in its application domain. An innovative solution is first of all of better, or tolerably less quality than the most-resembling state-of-the-art solution. Upon meeting this requirement, the innovativeness of a solution is quantified by its novelty, which we define to be equal to the domain-specific distance to this most-resembling solution.

The *state-of-the-art* solutions are the *highest*-quality solutions, generally accepted as such, in an application domain. The meaning of a solution in this sense is that of the product design of an already established high-quality product instance. For concisely expressing innovativeness, the comparison is thus made to these state-of-the-art solutions. When speaking of innovativeness of a solution, it is clear which application domain is meant, namely, of the product that the solution represents. We will therefore simply refer to innovativeness, when innovativeness with respect to the reference set of the state-of-the-art solutions in the application domain is meant. This is formalized as follows, see also Figure 3.3.

Notion 3.2 Innovativeness. *An innovative solution* $\mathbf{x}$ *statisfies a domain-specific required quality level* f_{req}. *The quality level* f_{req} *is derived from the set of best-quality solutions* $S_{\text{state-of-the-art}}$ *in the application domain. Upon meeting this requirement, the innovativeness of a solution* $\mathbf{x}$ *is quantified by its novelty with respect to* $S_{\text{state-of-the-art}}$.

The domain-specific required quality level can be formulated as a *relative* value Δf_{req} that puts a restriction on the allowed quality difference to the most-resembling solution $\mathbf{x}_{\text{closest}}$ as follows:

$$|f(\mathbf{x}) - f(\mathbf{x}_{\text{closest}})| < \Delta f_{\text{req}}, \text{ or } f(\mathbf{x}) \text{ is better than } f(\mathbf{x}_{\text{closest}}), \tag{3.1}$$

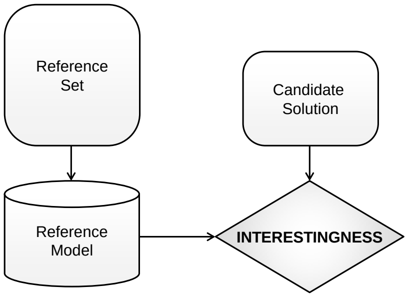

Figure 3.4 Interestingness. The interestingness of a candidate solution is determined with respect to a reference model obtained from reference solutions. Those solutions that represent new knowledge but still resemble the reference solutions (i.e., sufficiently but not extremely novel solutions) are the most-interesting ones, as they typically increase the understanding contained in the reference model most.

where

$$\mathbf{x}_{\text{closest}} = \underset{\mathbf{x}_{\text{state-of-the-art}} \in S_{\text{state-of-the-art}}}{\operatorname{argmin}} d(\mathbf{x}, \mathbf{x}_{\text{state-of-the-art}}). \tag{3.2}$$

Interestingness

By linearly depending on novelty, innovativeness does not fully express the likelihood of a solution actually being adopted and used by domain experts. This is best captured by the notion of *interestingness* [**Schmidhuber, 1997**]: To be perceived as interesting, a *concept*, understood here in the most-general sense, should be challenging by differing sufficiently from available knowledge, while not deviating so much that it cannot be understood at all. Likewise, an interesting solution should be novel, while still be in line with existing designs [**Saunders and Gero, 2001**].

Interestingness is a highly subjective, application-dependent notion: It depends on one's understanding of the application domain, built upon pre-existing knowledge and experience [**Saunders and Gero, 2001**]. We define a solution to be interesting if it greatly improves the current understanding of the application domain. A closed-form expression of interestingness is not feasible due to the subjectiveness involved. Where novelty is expressed with respect to a reference set, interestingness is related to a *reference model*. By *model*, a *representation of understanding*, i.e., of knowledge, is meant that can be used to generalize from *observed* concepts to as-yet-*unwitnessed* concepts, for instance in mathematical function approximation or as exists in the human mind. This is formalized as follows, see also Figure 3.4.

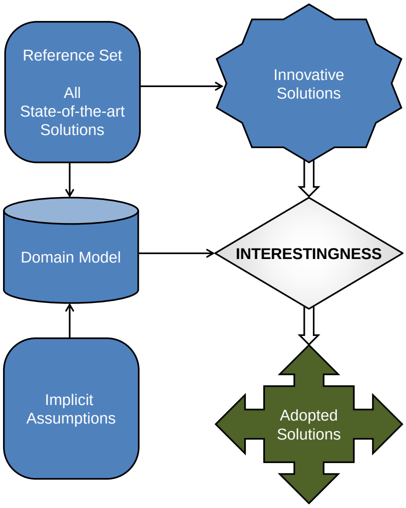

Figure 3.5 Interestingness of Innovative Solutions. Stepping beyond innovativeness, to express the likelihood of a solution actually being adopted by the engineering community, it should be sufficiently novel, but not so novel that it is not fully understood. This interestingness of a solution is expressed with respect to a model of the domain, as exists in a human domain expert.

Notion 3.3 Interestingness. *The interestingness of a solution is the increase of the understanding contained in a reference model that it leads to, after using it to refine the model. A reference model is established from a reference set of solutions.*

Evaluating the interestingness of innovative solutions requires a *domain model* based on the knowledge and understanding of the state-of-the-art solutions in the application domain (see Figure 3.5). Such a model is the comprehension existing in the mind of a human domain expert: The expert finds the candidate solution interesting based on his or her (implicit) knowledge and understanding of the established solutions in the domain. Alternatively, in principle, an artificial model can be used that approximates a certain mapping between solution characteristics, based on a set of established state-of-the-art solutions: By retraining the approximation model on a candidate solution, it is able to predict the to-be-learned mapping better, and this *learning progress* then characterizes the interestingness of a solution. Depending on the type of reference model, and even on the expert that was asked to evaluate, interestingness rankings are likely to show large variation. It is therefore important to describe the exact nature of the used reference model.

A further distinction can be made between interestingness to the *product devel-*

oper, i.e., the human domain expert, and the *product user*, i.e., prospective customers. As the aim of this work is to present methods for automated assistance in finding innovative solutions, ultimately to be developed into actual innovative products, we use the domain-expert perspective of interestingness, which thus expresses the likelihood for a product proposal found in automated search to be adopted by the human engineer and further developed by him or her into an actual product. Notwithstanding, interestingness from the customer's perspective *is* important for the commercial marketing strategy of the further-developed actual product [**Garcia and Calantone, 2002**], namely, it expresses the product's likelihood to be adopted by the potential customer base and thereby formally leading to a *technological innovation* according to the OECD (Organisation for Economic Co-operation and Development) definition [**Freeman, 1991, Garcia and Calantone, 2002**]: *'innovation' is an iterative process initiated by* [...] *a technology-based invention* [...] *striving for the commercial success of the invention*. We take "technology-based invention" synonymous to "innovative product". A theory further explaining the speed of the customer-base's adoption of innovative products is *diffusion of innovations* [**Rogers, 2010**], originally posed in 1962.

Discussion

Recapitulating, our view on innovative solutions is that they are novel and satisfy the minimal-quality requirement for adoption by the engineering community, both derived from the reference set of state-of-the-art solutions in the application domain. The most-innovative solution, however, is not likely to be the most-interesting for an engineer. For instance, when solution quality is based on simulation and the engineer cannot fully understand the simulated physical behavior that the solution gives rise to, he or she will not use it. An argument motivating this is that solutions found through simulation are prone to modeling and/or approximation errors. Hence, an innovative solution is only interesting to an engineer if its innovative properties are "graspable" for him or her. This is analogous to how children learn and play: They are attracted to things that are challenging, while at the same time understandable with respect to their current knowledge and experience.

Not the most-innovative solution, but the most-interesting sufficient-quality solution has the best chance of being selected and further-developed into an actual product. Whereas the most-innovative solution (highest novelty) *in theory* has the best potential of leading to an innovation, i.e., the process of the innovative product that it is further developed into actually being adopted by the end users, it is the most-interesting innovative solution (sufficient, but not too extreme novelty) that has the best potential of leading to an innovation *in practice*, as it is the most-likely to be adopted by the engineering community and developed into a product.

Notwithstanding, given the intractable nature of interestingness, for finding renewing but not too extreme solutions, the search should be steered towards solutions with maximum innovativeness. The view on interestingness is that it is evaluated in a post-processing step, using a model that is external to the search, which is the general approach of a human domain expert judging produced results after optimization.

We ask ourselves whether an expert could find a solution interesting that was not evaluated as being innovative in our approach. The expert has to agree on the used domain-specific distance measure and on the reference set of state-of-the-art solutions, so he or she should only find solutions interesting that are viewed as novel based on these, and if not, this indicates that distance measure and reference set are to be adjusted. It is possible that low-quality solutions not meeting minimum quality requirements are perceived as interesting, but these then have such optimal novelty (i.e., are located in the appropriate range for inducing learning progress) that the quality tolerance is (implicitly) stretched.

3.2 · Search Objective: Towards Innovative Solutions through Diversity

Searching for innovative solutions requires defining an appropriate domain-specific distance measure that differentiates between solutions. Innovativeness furthermore requires comparing candidate solutions to all established state-of-the-art solutions in the application domain. On the one hand, compiling this comprehensive reference set is hardly realizable, also taking, e.g., historical solutions that were abandoned into account. On the other hand, it is difficult to formulate established, e.g., patented, solutions in such a way that they can be interpreted automatically through putting them in the right notation.

We thus aim for innovative solutions, but it is difficult to express the true innovative value of solutions in an automated way. For certain applications, it could be argued that novelty is expressed with reasonable accuracy through identifying a relatively small number (e.g, 4–5) of baseline instances of established solutions. Alternatively, we drop the requirement of formulating reference solutions altogether. Instead of searching for innovative solutions directly, we aim for solutions that are of *high quality* and

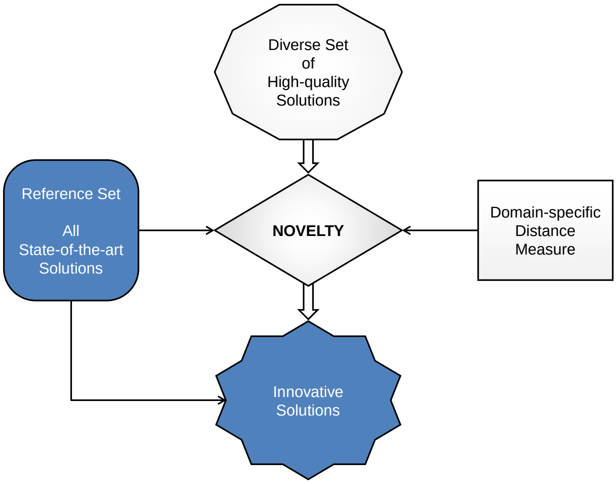

Figure 3.6 Towards Innovative Solutions through Diversity. A diverse set of high-quality solutions potentially contains solutions that are innovative, that is, novel with respect to all established high-quality solutions in the application domain and of better or sufficient quality compared to this state-of-the-art. Diverse means that the solutions are different from each other according to the same domain-specific distance measure as used in novelty determination.

are *diverse* with respect to each other, see Figure 3.6. The idea behind this is that a diverse set of high-quality solutions potentially contains innovative solutions. This chance increases with higher diversity in the set of solutions found.

For determining diversity, defining the domain-specific distance measure is still required, now for comparing candidate solutions with each other. This is formalized as follows, see also Figure 3.7.

Notion 3.4 Diversity. *The diversity of a set S of solutions is expressed by a set diversity measure* D *that applies a domain-specific distance measure* d *to all pairs of solutions* $\mathbf{x}, \mathbf{x}' \in S$.

3.2.1 · Measuring Performance

Resuming, we aim for a search method to produce a diverse set of high-quality solutions, motivated by the potential of innovative solutions being present among them. Therefore, to measure performance of a search method, an approach is required for scoring its result set by reflecting both the quality of the solutions found and the diversity between the solutions in this set.

Postponing a possibly required specification of a minimally-acceptable quality

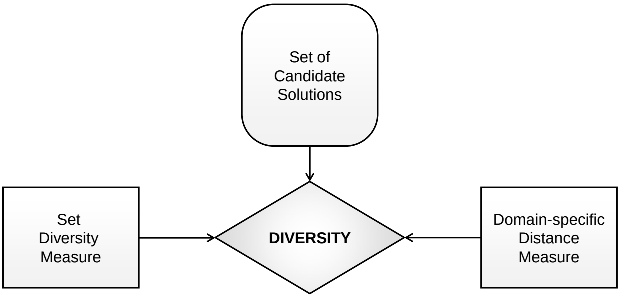

Figure 3.7 Diversity. The diversity of a set of candidate solutions is obtained using a set diversity measure that applies a domain-specific distance measure to all pairs of solutions from the set of candidate solutions.

level, we determine all *level-set* approximations possible in a method's result set, governed naturally by the quality values of the solutions contained in the result set. For each such choice of minimal quality level, the diversity is evaluated of the resulting level-set approximation. Each "solution level set" contains the solution itself and all other found solutions that have quality equal to or better than the solution's quality score. Formally, given a quality function $f\colon X \to \mathbb{R}$ and a threshold quality value f_{thres}, we define a level set as

$$L_{f_{\text{thres}}} = \begin{cases} L^-_{f_{\text{thres}}} = \{\mathbf{x} \in X | f(\mathbf{x}) \leq f_{\text{thres}}\} & \text{if } f \text{ is to be minimized} \\ L^+_{f_{\text{thres}}} = \{\mathbf{x} \in X | f(\mathbf{x}) \geq f_{\text{thres}}\} & \text{if } f \text{ is to be maximized} \end{cases}, \tag{3.3}$$

that is, the level set $L_{f_{\text{thres}}}$ contains all solutions from domain X with quality equal to or better than f_{thres}.

Set Diversity Measure

Each solution thus gets assigned a diversity score through its quality value serving as f_{thres} for a level-set approximation within the result set. To determine the diversity in a set of solutions S, a *set diversity measure* $\mathrm{D}(S)$ is required that applies the domain-specific distance measure to all pairs of solutions. Following [Ulrich and Thiele, 2011], we assume a symmetric domain-specific distance measure,

$$\mathrm{d}(\mathbf{x},\mathbf{x}') = \mathrm{d}(\mathbf{x}',\mathbf{x}) \text{ for any pair of solutions } \mathbf{x},\mathbf{x}'. \tag{3.4}$$

In [Solow and Polasky, 1994], three requirements are given for a set diversity measure, structured in [Ulrich et al., 2010, Ulrich and Thiele, 2011]:

1. *Monotonicity in variety*: Diversity increases under adding additional unique

solutions, that is, when a solution is added that is not a duplicate of a solution already present in the set;

2. *Twinning*: Diversity is constant under adding a solution that is a duplicate of a solution already present in the set;

3. *Monotonicity in distance*: Diversity increases under more dissimilarity between pairs of solutions, that is, for equal-size sets S, S', it holds that

$$\mathrm{D}(S') \geq \mathrm{D}(S) \iff \mathrm{d}(\mathbf{x}'_i, \mathbf{x}'_j) \geq \mathrm{d}(\mathbf{x}_i, \mathbf{x}_j),\ 1 \leq i, j \leq |S|. \tag{3.5}$$

For obtaining a set diversity measure $\mathrm{D}(S)$ with the properties listed above, [Solow and Polasky, 1994] start from the domain-specific distances between each pair of solutions $\mathbf{x}_i, \mathbf{x}_j \in S$ and transform these into similarity scores between 0 and 1,

$$\mathrm{sim}_{\mathrm{d},\theta}(i, j, S) = \exp(-\theta \cdot \mathrm{d}(\mathbf{x}_i, \mathbf{x}_j)),\ \theta > 0. \tag{3.6}$$

Here, θ scales the "reverse" non-linear transformation, a smaller value for θ results in higher similarity scores with less spread. The similarity scores are put in a $|S|$-by-$|S|$ matrix $\boldsymbol{A}$,

$$a_{ij} = \mathrm{sim}_{\mathrm{d},\theta}(i, j, S),\ 1 \leq i, j \leq |S|. \tag{3.7}$$

The set diversity measure termed *Solow–Polasky* [Ulrich and Thiele, 2011] then follows by taking the sum of the elements of the inverse of $\boldsymbol{A}$,

$$\mathrm{D}_{\mathrm{SP}}(S) = \sum_{i,j} b_{ij},\ \boldsymbol{B} = \boldsymbol{A}^{-1}. \tag{3.8}$$

$\mathrm{D}_{\mathrm{SP}}(S)$ results in a real number in the interval $[1, |S|]$: When $\mathrm{sim}_{\mathrm{d},\theta}(i, j, S) = 0$ for all $i \neq j$ (all solutions are completely dissimilar), then $\boldsymbol{B}$ (as well as $\boldsymbol{A}$) is the identity matrix and $\mathrm{D}_{\mathrm{SP}}(S)$ is equal to $|S|$; if $\mathrm{sim}_{\mathrm{d},\theta}(i, j, S)$ approaches 1 for all i, j (all solutions are almost perfectly identical), then $\mathrm{D}_{\mathrm{SP}}(S)$ also approaches 1 [Solow and Polasky, 1994].

Pareto Front Interpretation

For each solution we thus get a pair of values: Its quality value and the diversity score of the level-set approximation based on its quality value. We illustrate a method's performance from these pairs of values by interpreting them as a *Pareto front* approximation in quality/diversity space, see Figure 3.8.

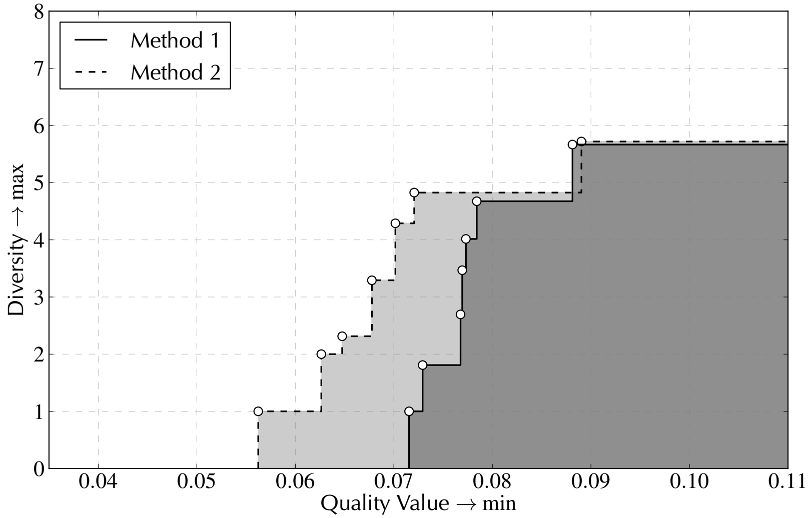

Figure 3.8 Performance Comparison. Per method, of each possible level set in its result set, the value-pair of the quality threshold and diversity score is plotted, governed by the quality values of the solutions in the result set. Interpreting these pairs jointly as a Pareto-front approximation, the dominated surface in quality/diversity space is visualized. Using a reference point of $(0.11, 0)$, method 1 has an obtained hypervolume of 0.184 and method 2 of 0.235.

Allowing for easier comparison, in a further step, the Pareto front is abstracted into a singular value expressing the dominated surface or *hypervolume* [**Zitzler and Thiele, 1998**] in the quality/diversity space that it covers with respect to a *reference point*. The reference point should represent worse than possible quality and diversity scores to make sure that all possible quality/diversity pairs contribute to the hypervolume.

3.3 · Efficiently Finding Diverse High-quality Solutions

Starting from the population-based optimization paradigm of Evolutionary Algorithms (see Chapter 2), the aim is to define a search approach that is able to efficiently produce diverse high-quality solutions. Taking inspiration from the discussion on innovativeness, we envision using an *exploration criterion*, next to the quality criterion, based on *online* novelty or interestingness to help steer the search, see Figure 3.9 and Figure 3.10. The idea is to make the search for diverse solutions more efficient by directing it into novel or interesting areas, and away from areas that are “known” already, based on a memory of *earlier sampled solutions* in the optimization run. Recently, the approach of attaining diversity by rewarding online novelty was also described in [**Lehman et al., 2013**].

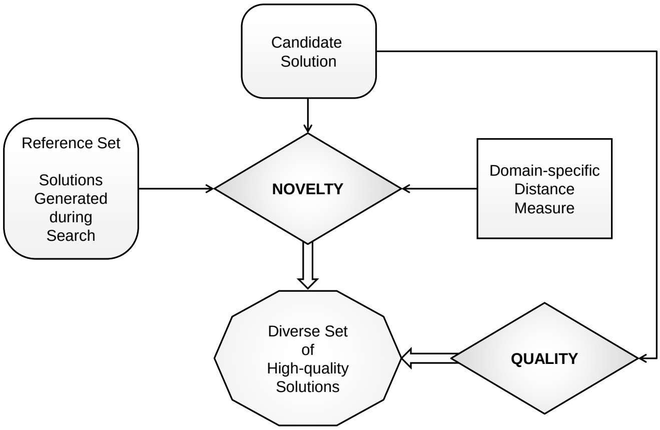

Figure 3.9 Using Novelty for Finding Diverse Solutions. A candidate solution is compared to solutions generated earlier during the search, using a domain-specific distance measure. The resulting novelty value serves as additional criterion in quality-based search, intended for improving exploration in order to assist in finding a diverse set of high-quality solutions.

Importantly, we are re-using ideas that are used for expressing innovativeness, but applied in this way, they do not provide the actual innovativeness of solutions found. In the online setup, the comparison is done to solutions found earlier in the search only, whereas actual innovativeness is derived from the reference set of all state-of-the-art solutions in the application domain.

It is straightforward to see how stepwise selection for novelty promotes domain-specific distance between found solutions, and thus the diversity of the entire set of found solutions with respect to the same domain-specific distance measure. Selecting on interestingness, on the other hand, induces a pattern of exploration that is optimal for obtaining an accurate model of the search space by promoting the learning progress that candidate solutions give rise to. The question is whether this exploration pattern implicitly helps in finding solutions that are diverse with respect to a distance measure that is unrelated to the model, as compared to selecting on novelty that *is* related to the distance measure. Presumably, selecting on interestingness makes the exploration more efficient by moving quickly through areas where the information density is low, thereby helping to find different high-quality solutions.

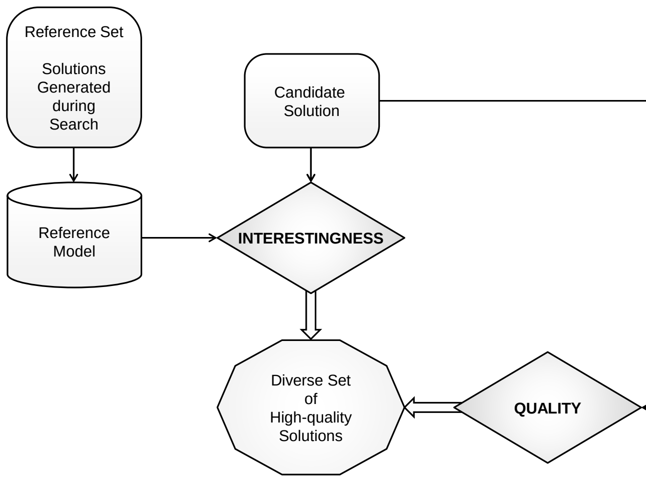

Figure 3.10 Using Interestingness for Finding Diverse Solutions. A reference model is derived from solutions generated earlier in the search. Each candidate solution gets an interestingness score based on the improvement of the model that it leads to. Using interestingness as additional criterion should move the search into areas of the search space with higher information density, thereby implicitly assisting quality-based search in finding a diverse set of high-quality solutions.

3.4 · Summary

We define *innovative solutions* as being *novel* and of tolerable quality, both with respect to the *reference set* of all *state-of-the-art* solutions in the application domain. Novelty is the difference to the closest solution from this set with respect to a *domain-specific distance measure*.

As evaluating novelty and innovativeness requires composing the comprehensive set of state-of-the-art solutions, we instead adopt finding a diverse set of high-quality solutions as search objective. A set of high-quality solutions, different from each other with respect to the domain-specific distance measure mentioned above, potentially contains innovative solutions.

The likeliness of an innovative solution actually being adopted by a human engineer is formalized as its *interestingness*. Interestingness is related to the *learning progress* that a solution leads to with respect to the domain model that exists in the mind of the human domain expert. As in learning, a solution should be sufficiently novel but not too extreme for it to give rise to maximum learning progress. Nevertheless, as interestingness is a highly subjective notion, the aim of the search remains finding innovative solutions (thus with maximum novelty), through optimizing on diversity and quality.

Lastly, deriving from the discussion on innovativeness, the idea is presented of using novelty and interestingness within the search to efficiently find diverse high-quality solutions. Instead of the set of state-of-the-art solutions, a *dynamic* reference set of solutions generated earlier during the search is used to obtain *exploration criteria* that can be used to steer the search into novel or interesting areas of the search space.

The formulation of such novelty and interestingness-based exploration criteria is worked out in Chapter 4, and in Chapter 5, the appropriate introduction of an exploration criterion into quality-based search is addressed. Chapter 6 describes the application of the developed methods to an airfoil optimization task, including the definition of a domain-specific distance measure.