Towards Directed Open-Ended Search by a Novelty Guided Evolution Strategy

Lars Gräning, Nikola Aulig, Markus Olhofer

2010

Preprint:

This is an accepted article published in Parallel Problem Solving from Nature - PPSN XI. The final authenticated version is available online at: https://doi.org/[DOI not available]

Towards Directed Open-ended Search by a Novelty Guided Evolution Strategy

Lars Graening[1], Nikola Aulig[2], and Markus Olhofer[1]

[1] Honda Research Institute Europe GmbH,
Carl-Legien-Straße 30, 63073 Offenbach/Main, Germany
{lars.graening,markus.olhofer}@honda-ri.de
[2] Technical University Darmstadt
Karolinenplatz 5, 64289 Darmstadt
nikola.aulig@web.de

Abstract. In the conceptional phases of design optimization tasks it is required to find new innovative solutions to a given problem. Although evolutionary algorithms are suitable methods to this problem, the search of a wide range of the solution space in order to identify novel concepts is mainly driven by random processes and is therefore a demanding task, especially for high dimensional problems. To improve the exploration of the design space additional criteria are proposed in the presented work which do not evaluate solely the quality of a solution but give an estimation of the probability to find alternative optima. To realize these criteria, concepts of novelty and interestingness are employed. Experiments on test functions show that these novelty guided evolution strategies identify multiple optima and demonstrate a switching between states of exploration and exploitation. With this we are able to provide first steps towards an alternative search algorithm for multi-modal functions and the search during conceptual design phases.

Key words: Evolutionary algorithm, open-endedness, interestingness, multi-objective optimization, novelty detection, prediction error, niching

1 Introduction

Evolutionary algorithms have various properties which make them suited to solve complex real world problems. One of these properties is the ability to identify multiple solutions in multi-modal quality functions. The importance of this property lies in the fact that very often conceptually different solutions can be found in real world applications which offer alternative realizations for the design of a system. The selection of the best suited solution can only be done by experts in the related field due to the complexity of the overall problem or due to the non-technical nature of the criteria, for example aesthetic arguments or the necessary distinctiveness to other available solutions used in other products.
In order to enhance this behavior various improvements of the algorithms are described in the literature. A common way is to strengthen the 'exploratory' behavior by increasing the mutation rate. Although this requires a low dimensional

search space in order to find solutions within an acceptable number of generations, successful examples of this approach can be found in the field termed *creative evolutionary search* [1] and in the early phases of design optimisation in which new concepts for a solution have to be identified on simplified models.
Another strategy is followed by *Niching Algorithms* which identify multiple optimal solutions by maintaining diversity within a population [2–4]. The simplest niching approach is fitness sharing where the fitness of individuals is reduced if they are located close together within a niching radius. Various improvements of these ideas are available like the dynamic niche sharing which uses a dynamic peak identification (DPI) algorithm to recognize the forming of niches and the fitness is shared among individuals within one niche. The dynamic niching algorithm [2] introduces mating restrictions instead of changing the fitness. While these methods require an a-priori estimation of the size of the niches, the derandomized CMA-ES implements an adaptive niche radius [4]. Although, niching algorithms provides the potential for identifying several distinct optimal solutions, the possible number of optima which can be discovered has to be specified in advance and remains limited by the population size.
An alternative way to guide the search towards new and innovative solutions is the integration of human creativity into the process [5, 6] in which a human user is involved in the generation of variations or alternatively in the selection process. Although this process turned out to be very powerful it is limited to problems for which a solution can be found within a small number of evaluations, in which a human operator is able to judge the solutions by their intuition or knowledge of the process.
In this work we outline an alternative method for the determination of optima in a multi-modal fitness landscape which is fundamentally different to existing methods. We propose to guide the search by an additional criterion which directly relates to new and unexplored areas of the search space. This criterion is based on novelty or interestingness measures. Although the concept of novelty exists in the subjective perceptions of individuals and is generally difficult to describe, various attempts to define the concept can be found in different fields of science like psychology, active learning or evolutionary robotics which allow to formulate measures suitable for a numerical calculation.
Silberschatz and Tuzhilin's [7] for example state that interestingness depends on the user who is examining a pattern. They point out that something that is interesting for one user might not be interesting for another one. Schmidhuber [8] argues that if something is too unexpected it appears random and is no longer interesting. Along this line, Saunders [9] refers to the Wundt curve and writes: "... the most interesting experiences are those that are similar-yet-different to those that have been experienced previously". Based on these works it becomes already obvious that novelty as well as interestingness can only be evaluated based past experience. In the field of active learning, Risi et all. [10] evaluates novelty simply by measuring the similarity to existing solutions stored in an archive. Similar to the work of Risi, Lehman at all. [11] defined novelty through sparseness that is evaluated based on already generated solutions. In the domain of developmental

robotics, motivated by the concept of intrinsic motivation, Oudeyer and Kaplan [12] provide a comprehensive summary on alternative techniques for quantifying interestingness and novelty. Most of the different attempts share the idea that a model that builds up a compact representation of the search space is used to produce an indicator for novelty or interestingness and allow the identification of parameter regimes which should be sampled. In this work a mechanism for the detection of novelty or interestingness is utilized to actively guide the search to alternative optima in a multi-modal search problem using evolutionary algorithms. The novelty measure provides an additional criteria besides the original quality function. The target is to guide the search by the newly added criteria towards currently unexplored regions of the search space and additionally to start new exploration phases after temporary convergence of the population. It is demonstrated on simple test functions that a combination of both criteria allows the algorithm to guide the population towards alternative local optima after the localization and convergence of the population to a formerly identified optimum. In the next section we describe the novelty metrics used here and their integration into the algorithm in more detail. In section 3 first experimental results are presented where the proposed algorithm is compared to the niching and standard evolutionary algorithm. The paper concludes with a discussion on the results and an outline for future work in section 4.

2 Novelty Guided Evolution Strategy

2.1 Overall Framework

The schematic view of the proposed algorithm is depicted in Fig. 1. After the reproduction, recombination and mutation of the parent population, the fitness is assigned to each individual of the produced offspring population. A second criteria is introduced evaluating the individual's novelty based on a model of the quality function which is adapted by newly evaluated individuals. The selection incorporates at least two criteria, the novelty and the quality function.

2.2 Novelty Evaluation and World Model Adaptation

An individual is said to be novel if it does not meet the expectations derived from the accumulated knowledge about the search space. Therefore, the implemented novelty metric is defined as follows. If the expected quality value, estimated by the world model, differs from the calculated one, a high novelty value is assigned to the individual. As depicted in Fig. 2 the implemented novelty metric is calculated as,

$$\varepsilon(\boldsymbol{x}_i^t) = |f(\boldsymbol{x}_i^t) - \hat{f}(\boldsymbol{x}_i^t)|, \; i \in [1 \dots \lambda] \tag{1}$$

where λ defines the size of the offspring population, t the current generation, $\boldsymbol{x}_i^t$ is the design vector and $f(\boldsymbol{x}_i^t)$, $\hat{f}(\boldsymbol{x}_i^t)$ defines the calculated and the predicted quality function value respectively. $\varepsilon(\boldsymbol{x}_i^t)$ reflects the prediction error of the world model and is used for the quantification of the novelty. This formulation of the

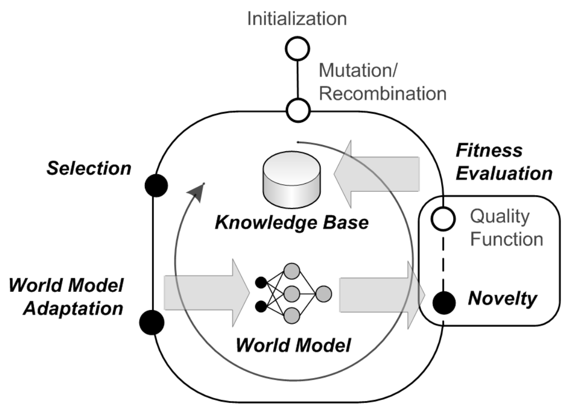

Fig. 1. Flowchart of the novelty guided optimization framework.

novelty metric equals the *Predictive novelty motivation* concept described by Oudeyer and Kaplan [12]. Solutions with a high prediction error assign a high novelty value to the individual.
The world model is adopted to predict the quality function value by means of estimating $\hat{f}(\boldsymbol{x}_i^t)$. Since multilayer feed-forward neural networks have successfully been employed as universal function approximators they are implemented for the world model. As already shown by Bishop [13], the neural network model is well suited to estimate the novelty of a solution. Given a pre-defined model structure, the network is updated in each generation $t-1$ using RProp [14], a variation of the back-propagation algorithm, together with cross-validation to prevent over-fitting. Data from generation $t-\gamma$ to generation $t-1$ that is added to the knowledge base during evolution is used for training. The parameter γ controls whether more global or localized model of the search space is generated.

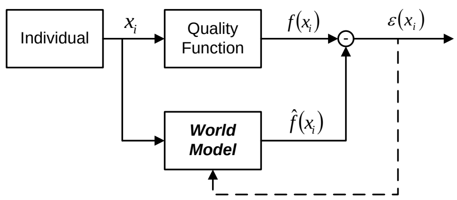

Fig. 2. Estimation of the prediction error which defines the novelty of a solution. The world model builds a compact representation of the search space and is adopted to predict the target function value.

2.3 Selection Strategy

During the fitness evaluation a fitness vector is assigned to each offspring,

$$\boldsymbol{f}(\boldsymbol{x}_i^t) = [f(\boldsymbol{x}_i^t),\ \varepsilon(\boldsymbol{x}_i^t)]^T, \tag{2}$$

containing the actual quality value and the prediction error estimating the novelty. It has to be noted that the additional criteria is dynamic in the sense that each time the world model is updated the estimated novelty value changes for one and the same solution. The multi objective optimisation problem can be transformed into a single objective optimisation by a *linear weighted aggregation*. This method leaves us with the problem of choosing an adequate weight. A high weight on the novelty objective would result in an extensive explorative behavior while a high weight on the actual quality function would result in an intensive exploitation of a single optimal solution. The desired behavior of the evolution strategy is a process that identifies successively several optima in regions with high fitness values but which does not exploit only one optimal solution. Aside the linear weighted aggregation, we employ a Pareto optimal selection criterion. The implemented strategy is derived from the crowded tournament selection suggested in the NSGA-II algorithm [15].

3 Experimental Results

The following experiments target the study of the basic characteristic of the *novelty guided ES*. The proposed evolution strategy is compared to three existing strategies, namely *standard ES*, *dynamic niche sharing* and *open-ended ES*.

3.1 Characteristics of the novelty guided evolution strategy

In the first experiments, the behavior of the introduced novelty guided ES is studied on a two dimensional multi-modal test function in which the design vector covers two variables, $\boldsymbol{x}_i^t = [x_{1i}^t,\ x_{2i}^t]^T$ or in short $\boldsymbol{x} = [x_1,\ x_2]^T$. The test function is constructed by a superposition of $N_G = 6$ 2D Gauss functions and is mathematically defined as follows:

$$f(\boldsymbol{x}) = -\sum_{j=1}^{N_G} e^{-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu}_j)\boldsymbol{\Sigma}^{-1}\cdot(\boldsymbol{x}-\boldsymbol{\mu}_j)+1}, \tag{3}$$

where $\boldsymbol{\mu}_j$ is the center and $\boldsymbol{\Sigma}^{-1}$ the covariance matrix of the Gauss kernel. Each center $\boldsymbol{\mu}_j$ of the different Gauss functions defines approximately the location of one local optima. The Gauss kernels are inverted, simply to transfer the task from a maximization into a minimization task. The resulting quality function is depicted in Fig. 3 a). All runs were calculated for $N = 500$ generations with a parent population size of $\mu = 20$ and an offspring population size of $\lambda = 100$. The standard evolution strategy is a (μ, λ) strategy with global step size

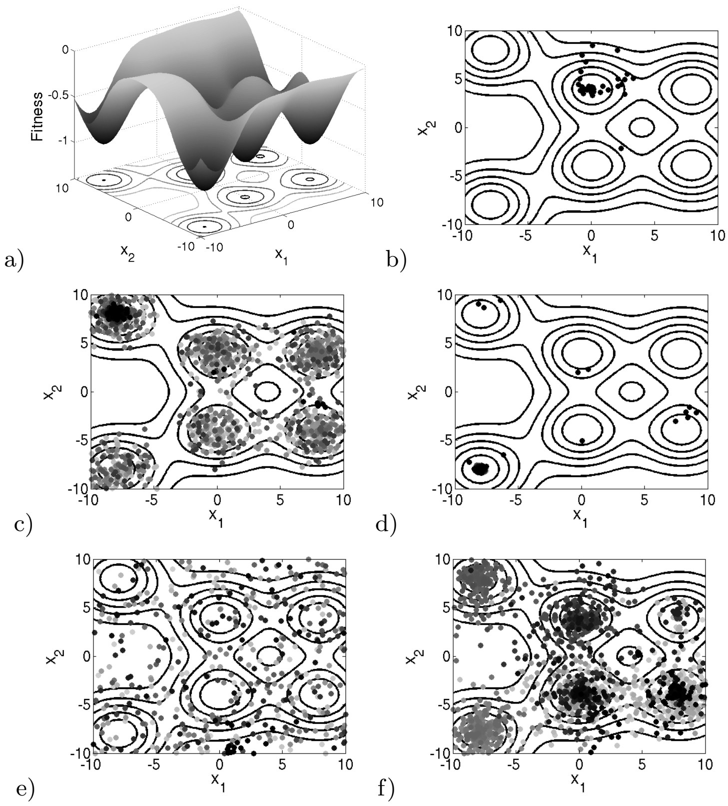

Fig. 3. a) the test function with 6 optima, b) results of the standard ES, c) niching with optimal niching radius, d) niching with an improper niching radius, e) open-ended evolution, targeting the generation of novel designs only and f) the results of the novelty guided ES.

control. Recombination as well as mutation are applied to produce the offspring population. For the niching algorithm a niching radius ρ is defined according to Shir [3]. It has to be noted that the calculation of an adequate niching radius requires knowledge about the number of optima, which is usually not available. The world model that is needed for the calculation of the novelty metric is realized using a multi-layer network with 10 hidden neurons and with sigmoidal activation function. The data of the offspring population from 5 generations is used for the training of the network weights. The dominance based ranking with crowding distance is used as selection operator in the novelty guided ES. The results of the different experiments are summarized in Fig. 3. For each experiment, the contour plot of the fitness landscape together with the generated solutions is shown. The fill color of the dots indicate the generation number at which the solution has been produced. Dark indicates early and bright late generations. Fig 3 b) shows the result of a standard ES. The algorithm converges

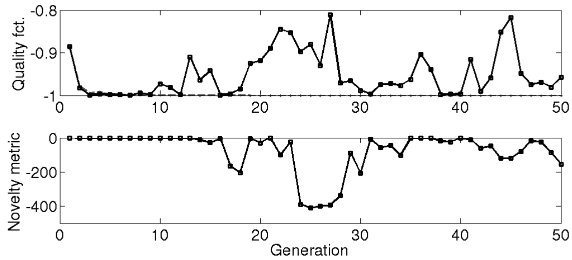

Fig. 4. Progress of quality and novelty in novelty guided ES. The selection of solutions with high degree of novelty allows to escape from optimal solutions.

directly (within about 20 generations) into the next local optimum. After that the algorithm is converged. Fig. 3 c) and d) present the result of the dynamic niching algorithm. While in c) an optimal niching radius has been used, in d) the niching radius is over-estimated (twice the optimal radius) what easily happens if no knowledge about the number of optima is given. In the case of a correct estimation of the niching radius the individuals distribute well between all the optima. If the niching radius is wrong or the population size too small, the niching algorithm might get stuck in a limited number of optimal solutions. In Fig. 3 e) results of a novelty driven evolutionary search are shown in which the search is only based on the novelty criterion neglecting information given by the quality function (open ended evolution). The algorithm does not exploit one of the six optima and diverges towards the boundaries of the search space as expected. Thus, a pure novelty driven optimization is quite inefficient and should be used for the exploration of the search space only. As can be seen from Fig. 3 f) the proposed novelty guided ES is able to locate all optimal solutions. Compared to the niching algorithm the optima are not exploited in parallel but rather sequentially. This sequential exploitation of the optima comes from the interplay between the two objectives, the quality function and the novelty metric. Fig. 4 shows the development of the quality and novelty value of the best offspring in the first 50 generations. In early generations the algorithm starts to exploit a nearby optimal solution exactly as it is done in the standard ES. After about 10 generations the influence of the novelty measure on the selection increases. Novel but worse solutions are selected. This allows the optimizer to escape from one optimum and exploit another one. This interplay between quality and novelty repeats until the algorithm is stopped. Since, a local model is used here, the algorithm visits optima multiple times due to a limited memory of the model.

3.2 Comparative study on a high dimensional test function

To carry out experiments on higher dimensional quality functions the multivariate Gauss kernel is used for the construction of an n dimensional multi-modal test function. Instead of the superposition of the Gauss kernels the max operator is applied to prevent the shift of the optima from the Gauss center. The n

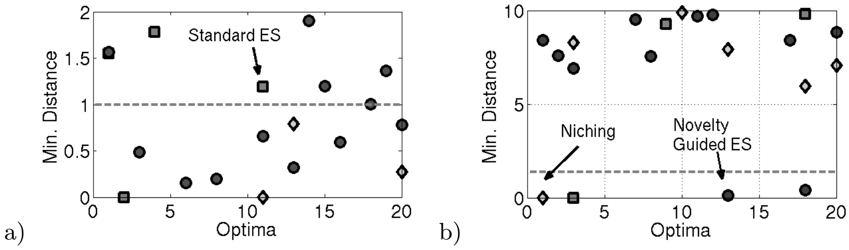

Fig. 5. Illustration of the minimal distance of the produced individuals to each of the 20 optimal solutions (x-axis). Only the closest solutions are shown. a) shows the results on the 5 dimensional and b) on the 10 dimensional fitness landscape.

dimensional test function is defined as follows:

$$f(\boldsymbol{x}) = - \max_{j \in [1 \ldots N_G]} \{ e^{-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu}_j)\boldsymbol{\Sigma}^{-1}\cdot(\boldsymbol{x}-\boldsymbol{\mu}_j)+1} \}. \tag{4}$$

An $n = 5$ and an $n = 10$ dimensional variant of $f(\boldsymbol{x})$ have been constructed. The number of Gauss kernels and thus the number of optima is set to $N_G = 20$. The Gauss centers are distributed randomly in the search space but remain the same for the different strategies. The minimal distance of all generated individuals to each optima is used to evaluate the strategies. If the distance runs below a threshold of $\tau = 1.0$, an optima is classified as being identified. The novelty guided ES has been compared to niching and standard ES. Since, the open-ended ES does not tend to exploit any of the optimal solutions it is skipped from the comparison in these experiments. Related to the preceding experimental setup, the number of hidden neurons of the world model has been increased to 25. For the niching algorithm, the correct number of optima has been used to estimate the radius ρ. Each optimization has been performed 5 times with different random seeds in order to retrieve a first idea on the reliability of the different strategies.

The results of the experiments are summarized in Fig. 5 and Tab. 1. Fig. 5 visualizes the evaluation of a typical run a) on the 5 and b) on the 10 dimensional quality function. The index of the optima is mapped onto the x-axis while the y-axis shows the distance of the closest solution to each optima. The dotted line indicates the threshold applied for counting the number of approached local optimal solutions. In Tab. 1 the mean $\overline{g}$ and the variance σ^2 of the number of ap-

a)

ES	$\overline{g}$	σ^2
Standard ES	1.4	0.3
Niching	2.2	0.7
Novelty Guided ES	5.6	1.3

b)

ES	$\overline{g}$	σ^2
Standard ES	1.0	0.0
Niching	1.0	0.0
Novelty Guided ES	1.2	0.2

Table 1. Summary of the results on the a) 5 and b) 10 dimensional multi-modal Gauss function. $\overline{g}$, σ^2 are mean and variance of the number of approached optimal solutions.

proached optimal solutions over 5 runs is summarized. Again, it can be observed that the standard ES exploits one single optimal solution only independent of the search space dimension. The distance to the remaining 19 optima remains large. Concerning the number of approached optima, the novelty guided ES outperformes niching on the 5 dimensional test function. The proposed strategy approaches from minimal 4 to about 7 out of 20 optima, while niching reaches only about 3 optima at maximum. As can be seen from Tab. 1 b), none of the strategies is performing well on the 10 dimensional test function. However, the novelty guided ES is at least able to approach 2 optima in one out of the five optimization runs.

4 Discussion

The experiments on the test functions show the general feasibility of the proposed method. Evolutionary Strategies allow, combined with the concept of novelty measures, the determination of multiple optima on a multi-modal quality function. In contrast to other algorithms, novelty guided evolution strategies allow a sequential process of alternating phases of exploration and exploitation on multi-modal quality functions in which the exploration is guided by novelty or interestingness measures instead of randomly sampling the search space. In the presented initial experiments the additional criteria, which is introduced to guide the search to alternative solutions, is based on a novelty measure, purely relying on the prediction error of a model. The generation of purely novel solutions is usually a simple task solved easily by e.g. generating sufficiently large mutations in a unconstrained search space. Preferably, the new direction should be an estimation of the most likely area for new optima or at least an area from which to sample in order to increase the chance to determine a new optima. In this sense the utilization of a novelty measure cannot be the final answer which was already stated in [12]. In order to determine useful search directions, measures of interestingness are required which guide the search towards areas which are interesting in the sense that knowledge about the design space is generated in order to finally determine areas with high probability of high fitness values. Therefore the evaluation of measures based on the learning rate of a model will be the next step to tackle more realistic problems for example in the field of aerodynamic design in which areas of high noise or even chaotic parameter regimes are expected. In general, models of the quality functions are necessary to determine the novelty or the interestingness of design areas. Assuming that approximation models are generally more simple than the original quality function all approximation models can only realize local models, valid in a limited area of the design space. Therefore a second step in our future efforts in the development of the algorithm is the integration and the adaptation of model ensembles in which each single model represent a different area of the global search space. Utilizing model ensembles also avoids the oscillation of the optimization process between to optima which can be observed in the presented results. The reason is that the sampling of one optimum results in an adaptation of the model in a way that old information in the model is removed. Keeping an ensemble of models allows us to avoid the overwriting of former acquired information of the global search space. In this sense the presented work has to be seen as a starting point for the research of novelty guided evolution strategies.

References

1. Bentley, P.J., Corne, D.W., eds.: Creative Evolutionary Systems. Morgan Kaufmann (2001)
2. Shir, O.M., Bäck, T.: Niching in evolution strategies. In: GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation, New York, NY, USA, ACM (2005) 915–916
3. Shir, O.M.: Dynamic niching in evolution strategies with covariance matrix adaptation. In: In Proceedings of the 2005 Congress on Evolutionary Computation CEC-2005, Piscataway, IEEE Press (2005) 2584–2591
4. Shir, O.M., Bäck., T.: Niche radius adaptation in the cma-es niching algorithm. In: Parallel Problem Solving from Nature (PPSN IX), Reykjavik, Iceland (2006)
5. Herdy, M.: Evolution strategies with subjective selection. In: PPSN IV: Proceedings of the 4th International Conference on Parallel Problem Solving from Nature, London, UK, Springer-Verlag (1996) 22–31
6. Takagi, H.: Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. In: Proceedings of the IEEE. Volume 89. (2001) 1275–1290
7. Silberschatz, A., Tuzhilin, A.: What makes patterns interesting in knowledge discovery systems. IEEE Transactions on Knowledge and Data Engineering **8** (1996) 970–974
8. Schmidhuber, J.: What's interesting? Idsia-35-97, IDSIA, Switzerland (1997)
9. Saunders, R.: Curious Design Agents and Artificial Creativity. PhD thesis, Faculty of Architecture, The University of Sydney (2001)
10. Risi, S., Vanderbleek, S.D., Hughes, C.E., Stanley, K.O.: How novelty search escapes the deceptive trap of learning to learn. In: GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation, New York, NY, USA, ACM (2009) 153–160
11. Lehman, J., Stanley, K.O.: Exploiting open-endedness to solve problems through the search for novelty. In: Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI), Cambrige MA, MIT Press (2008)
12. Oudeyer, P.Y., Kaplan, F.: What is intrinsic motivation? a typology of computational approaches. Frontiers in Neurorobotics (2007)
13. Bishop, C.M.: Novelty detection and neural network validation. In: Proc. IEE Conference on Vision and Image Signal Processing. (1994) 217–222
14. Igel, C., Husken, M.: Improving the rprop learning algorithm. In: Second International Symposium on Neural Computation. (2000) 115–121
15. Deb, K., Agrawal, S., Pratap, A., Meyarivan, T.: A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In: Parallel Problem Solving from Nature, Springer (2000) 849–858