Multi-Objective Hyperparameter Tuning and Feature Selection using Filter Ensembles

Martin Binder*

Ludwig-Maximilians-Universität München
Munich, Germany
martin.binder@stat.uni-muenchen.de

Janek Thomas

Fraunhofer Institute for Integrated Circuits IIS
Erlangen, Germany
janek.thomas@scs.fraunhofer.de

Julia Moosbauer*

Ludwig-Maximilians-Universität München
Munich, Germany
julia.moosbauer@stat.uni-muenchen.de

Bernd Bischl

Ludwig-Maximilians-Universität München
Munich, Germany
bernd.bischl@stat.uni-muenchen.de

ABSTRACT

Both feature selection and hyperparameter tuning are key tasks in machine learning. Hyperparameter tuning is often useful to increase model performance, while feature selection is undertaken to attain sparse models. Sparsity may yield better model interpretability and lower cost of data acquisition, data handling and model inference. While sparsity may have a beneficial or detrimental effect on predictive performance, a small drop in performance may be acceptable in return for a substantial gain in sparseness. We therefore treat feature selection as a multi-objective optimization task. We perform hyperparameter tuning and feature selection simultaneously because the choice of features of a model may influence what hyperparameters perform well.

We present, benchmark, and compare two different approaches for multi-objective joint hyperparameter optimization and feature selection: The first uses multi-objective model-based optimization. The second is an evolutionary NSGA-II-based wrapper approach to feature selection which incorporates specialized sampling, mutation and recombination operators. Both methods make use of parameterized filter ensembles.

While model-based optimization needs fewer objective evaluations to achieve good performance, it incurs computational overhead compared to the NSGA-II, so the preferred choice depends on the cost of evaluating a model on given data.

CCS CONCEPTS

• **Theory of computation** → **Evolutionary algorithms**; Non-convex optimization; *Mixed discrete-continuous optimization*; • **Computing methodologies** → *Supervised learning*; **Feature selection**;

KEYWORDS

Feature Selection, Hyperparameter Optimization, Multiobjective Optimization, Evolutionary Algorithms, Model-based Optimization

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
GECCO '20, July 8–12, 2020, Cancún, Mexico
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7128-5/20/07.
<https://doi.org/10.1145/3377930.3389815>

ACM Reference Format:

Martin Binder*, Julia Moosbauer*, Janek Thomas, and Bernd Bischl. 2020. Multi-Objective Hyperparameter Tuning and Feature Selection using Filter Ensembles. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancún, Mexico. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3377930.3389815>

1 INTRODUCTION

Machine learning models often need to satisfy multiple objectives simultaneously to accommodate the nature of a practical setting. Usually the main goal is predictive performance. Especially on large and complex datasets this necessitates highly nonlinear algorithms, which have hyperparameters that need to be chosen carefully. *Hyperparameter optimization* poses a substantial challenge in machine learning: Besides few model-specific methods [16, 30], there are no general analytic representations of model performance w.r.t. hyperparameter settings. Performance therefore needs to be estimated using test-set evaluation or cross-validation. Hyperparameter optimization is therefore an expensive black-box optimization problem.

Besides predictive performance, model sparsity is frequently another desirable objective. According to Guyon and Elisseeff [17], sparser models help with interpretability, i.e. a better understanding of the underlying process that generated the data. In addition to that, predictions can be made faster and more cost-effectively. Sparser models may even have better predictive performance, since they regularize against overfitting.

The process of *feature selection* aims to select a small subset of relevant features while still constructing models with sufficient or even optimal predictive performance. There are two distinct model-agnostic approaches to feature selection [17]: Filters and wrappers. *Filters* use proxy measures to rank features by their estimated explanatory power, independently of the learning algorithm being employed. These include information theoretic measures, correlation measures, distance measures or consistency measures [8]. In contrast, *wrappers* [28] optimize model test-set performance directly over the space of feature subsets. Because every feature subset evaluation requires either one or multiple model fits, exhaustive search is usually infeasible, and a black-box discrete optimization search strategy is necessary. Commonly used are simple greedy methods like forward or backward search. More advanced methods like evolutionary algorithms can improve upon this [37]. Because

*These authors contributed equally to this work.

they directly optimize learner performance, wrappers often yield better results [37].

Feature selection is often considered as a single-objective task. Sometimes the feature selection step is only used to optimize performance [28]. However, in many applications it is desirable to forego a small drop in performance for a substantial gain in sparseness. This leads to a natural treatment of the feature selection problem as a *multi-objective optimization problem*: Maximize predictive performance while minimizing the number of features selected. Feature selection methods may aggregate model performance and number of features into a single objective function through a penalization term [37]. However, this implies a trade-off between performance and sparsity must be specified a-priori, which may be difficult.

Multi-objective optimization methods try to find a *set* of solutions that represent different trade-offs between the goals, enabling the user to consider the possible alternatives and to choose a fitting solution a-posteriori. This is beneficial for feature selection [36].

Hyperparameter optimization and feature selection are often performed in separate steps. We argue that jointly optimizing over the combined spaces of hyperparameters and feature subsets is beneficial and appropriate: The optimal choice of hyperparameter configuration is very likely to depend on the specific features that are included and vice-versa. Also, it is likely more computationally efficient to explore the joint spaces simultaneously. When using the wrapper approach this combination is not trivial: The exponentially large binary search space of selected features now has to be fused with the mixed numeric-categorical space of hyperparameters.

We present and adapt model-agnostic holistic approaches for both aspects discussed above: multi-objective and joint optimization of hyperparameters and feature sets. To guide the search, our methods make use of combinations of feature filter scores which are combined in a *filter ensemble*.

Our approaches can be considered *hybrid* filter-wrappers: they are fundamentally “wrapper”-based, because they optimize model-performance, but also make use of filters.

2 RELATED WORK AND CONTRIBUTIONS

In recent literature, *Bayesian optimization* (“BO”), also referred to as *model-based optimization*, has become a popular method for hyperparameter tuning of learning algorithms [33], often outperforming simple grid search or random search. Originally, BO generally relied on a Gaussian process (GP) model, such as in the Efficient Global Optimization (EGO) algorithm [24]. However, the GP does not typically scale well to high dimensions and large numbers of data points. Random forests have been proposed as an alternative surrogate model, as used in the popular *SMAC* hyperparameter optimization method [22].

Automated machine learning (AutoML) deals with the configuration and optimization of complete machine learning pipelines, often encompassing data pre-processing, ML models, ensembling, and possibly post-processing steps. AutoML may encompass choosing among the many different possibilities of what method to use at each stage of the pipeline, optimizing the hyperparameters for these methods, and combining them in ways that yield well-performing ensembles. Feature selection by filter or wrapper methods is an

important pre-processing step already part of some AutoML frameworks. Consequently, AutoML has the potential to jointly optimize hyperparameters and included features. *autosklearn* [14] for example integrates a filter-based feature selector, parameterized by a filter measure and a percentage indicating the fraction of highest-ranked features to be included. Bayesian optimization is used to find the optimal filter (among a set of possible filters) and the best feature selection rate. The AutoML framework *TPOT* [31] uses an evolutionary algorithm to search over different machine learning pipelines, and also provides different feature selection strategies in its search space. TPOT can perform multi-objective optimization with regard to model performance and model complexity, measured as the number of processing steps, but it does not take model sparsity into account.

However, even though hyperparameter optimization and feature selection are both present in some AutoML frameworks, it is not their goal to find a good trade-off between predictive performance and sparseness. Feature selection is merely used to improve predictive performance, without considering the preference for sparse models in light of better interpretability or other benefits. In fact, these frameworks often tend to produce considerably complex models. They may even introduce additional features through feature engineering in pursuit of increasing predictive performance as the only goal. The results are often large and heterogeneous ensembles that are hard to interpret and deploy [32].

Evolutionary algorithms are especially well suited for multi-objective optimization. According to the survey by Xue et al. [37], *genetic algorithms* (GAs) are among the most commonly applied techniques for multi-objective feature selection. Inspired by natural evolution, GAs apply recombination and mutation operators to iteratively improve the population of solution candidates. Through techniques like non-dominated sorting [9, 25], genetic algorithms have become a powerful tool for multi-objective optimization. Several GA-based methods have been proposed for the task of wrapper-based feature selection [13, 18, 35].

There have been first investigations on simultaneous multi-objective hyperparameter optimization and feature selection using GAs: Bouraoui et al. [4] proposed an SVM-wrapper approach based on the NSGA-II, using a shared representation of the feature configuration and algorithm hyperparameters. While their method of combining the search spaces is similar to our GA-based approach, Bouraoui et al. [4] do not use specialized initialization and mutation operators that we have found to be necessary for good performance¹. Another limitation is that their approach is not model agnostic but limited to SVMs only.

Numerous methods have been proposed for *feature filtering*, and there are known ways of combining filters into *filter ensembles*. These may rely on applying the same filter on diverse datasets (“homogeneous approach”), or different filters on the same dataset (“heterogeneous approach”), and there are different ways of aggregating the filter rankings [3]. We extend the heterogeneous approach into a hybrid filter-wrapper selection method by optimizing a parameterized ranking aggregator.

¹See the Ablation Study in the Supplement, where their approach, “Variant (1)”, is outperformed by all our methods on every dataset and learning algorithm with no exception.

There has been a lot of research on feature selection and hyperparameter optimization. However, we found that there is no *general* algorithm or framework to perform model-agnostic simultaneous hyperparameter optimization and feature selection for predictive performance and model sparsity in a multi-objective fashion. We therefore choose to tackle this problem from two directions. First we adapt a standard evolutionary multi-objective optimization method, the NSGA-II, for the particular problem of feature selection and hyperparameter tuning. We contrast this with a classical hyperparameter optimization algorithm, based on BO, extended to perform feature selection.

Our main contributions are:

- (1) We adapt the NSGA-II to the problem at hand by introducing specialized sampling and mutation operators that make use of filter ensembles to enhance optimization performance.
- (2) We investigate how multi-objective Bayesian optimization (MOBO) can be used for this problem and propose an effective method that uses a filter ensemble for feature selection.
- (3) By conducting a benchmark of these approaches on a variety of tasks for different machine learning algorithms, we provide a comparison between our approaches and show how these algorithms perform compared to suitable baselines.

3 PROBLEM STATEMENT

We consider the machine learning problem of a given feature space $\mathcal{X}$ of vectors with p components (“features”), an arbitrary outcome space $\mathcal{Y}$ (for example $\{-1, 1\}$ for a binary classification task), and a performance measure $L : \mathbb{R}^g \times \mathcal{Y} \rightarrow \mathbb{R}$ measuring the quality of predictions (in $\mathbb{R}^g$) given ground truth values. g is 1 for regression tasks, and equal to the number possible outcome classes for classification. Data samples $\mathcal{D} = \left\{ \left(\mathbf{x}^{(1)}, y^{(1)} \right), \dots, \left(\mathbf{x}^{(n)}, y^{(n)} \right) \right\} \in (\mathcal{X} \times \mathcal{Y})^n$ are assumed to be n i.i.d. realizations of random variables (X, Y) which follow a joint distribution $\mathbb{P}_{XY}$. A dataset is thus characterized by n , the number of samples available, and p , the number of *features*.

Let $\mathcal{A}_{\lambda, s}$ be a learning algorithm that takes the given dataset $\mathcal{D}$ and constructs a model $f : \mathcal{X} \rightarrow \mathbb{R}^g$. The learning algorithm is parameterized by the *feature configuration vector* $s \in \{0, 1\}^p$, where $s_j = 1$ denotes that feature j is included in the model. The sparseness of the resulting model is thus determined by the *Hamming weight* of s , i.e. the number of components of s that are 1. *Hyperparameters* of the learning algorithm are summarized in a vector $\lambda \in \Lambda$. Λ is a (possibly mixed) bounded space and may contain numeric, integer, and categorical values².

In general we are trying to construct a model $f = \mathcal{A}_{\lambda, s}(\mathcal{D})$ which minimizes the *generalization error* $\text{GE}[f] = \mathbb{E}_{\mathbb{P}_{XY}} [L(f(\mathbf{x}), y)]$. However, the generalization error can only be estimated using in-sample data $\widehat{\text{GE}}[\mathcal{A}_{\lambda, s}, \mathcal{D}]$ through a resampling technique such as cross-validation.

This setup gives rise to the bi-objective hyperparameter optimization and feature selection problem:

$$\min_{\lambda \in \Lambda, s \in \{0, 1\}^p} \left(\widehat{\text{GE}}[\mathcal{A}_{\lambda, s}, \mathcal{D}], \sum_{i=1}^p c_i s_i \right).$$

²Hierarchical dependencies of hyperparameters, e.g. kernel hyperparameters that appear only if a specific kernel is chosen, are a common extension but not considered in this work.

The setup regards estimated generalization error as one objective and the cost of features considered as another.

It is possible, and a trivial extension of our method, to consider arbitrary costs c_i for each feature i . However, in our benchmarks we limit ourselves to equal costs $c_i = 1/p$. The resulting measure corresponds to the *fraction of selected features*, $\text{ffrac} = \frac{1}{p} \sum_i s_i$, ranging from 0 to 1.

4 MULTI-OBJECTIVE HYPERPARAMETER TUNING AND FEATURE SELECTION: TWO APPROACHES

Our two methods are based on two popular multi-objective optimization methods, adapted for the particular task of hyperparameter tuning and feature selection: A model-based approach, and an approach based on an evolutionary algorithm. Both approaches make use of feature filters to accelerate the search.

4.1 Feature Filters and Filter Ensembles

The number of possible feature configurations s is exponential in p , so a large number of performance evaluations would be necessary to explore the performance space. Therefore, both our optimization approaches make use of *feature filters*. These are methods that heuristically score individual features according to their apparent relevance for the outcome variable. Because there are various methods to estimate this relevance, we consider a collection of M feature filters, which generate scores $F^m(\mathcal{D})_j$, $m = 1, \dots, M$, $j = 1, \dots, p$ for each feature j of a training dataset $\mathcal{D}$. Different methods may generate scores on different scales, but we rank-transform and scale these scores to values ranging in equi-distant steps from least (value of 0) to most (value of 1) relevant for the outcome variable.

We propose combining multiple filter methods into *filter ensembles* similarly to Dittman et al. [12], but extending their method by using *weighted* average rank aggregation. The ensemble filter score EF_j for feature j is calculated as the weighted average according to weight vector $\mathbf{w} \in [0, 1]^M$, $\sum_i w_i = 1$:

$$\text{EF}_j(\mathbf{w}) = \sum_{m=1}^M w_m F^m(\mathcal{D})_j. \quad (1)$$

The weighting parameter $\mathbf{w}$ can be optimized, extending the filter into a hybrid filter-wrapper approach.

4.2 Bayesian Optimization Approach

Bayesian optimization (BO), also often referred to as *sequential model-based optimization*, has been successfully used for machine learning hyperparameter optimization in many applications [33]. The principle of BO is based on two steps, which are performed in turn. First, a so-called *surrogate model* is fitted to model the relationship between decision variables (e.g. hyperparameter values) and the objective value (e.g. estimated generalization performance). This model generates cheap approximations of the (generally expensive to evaluate) objective function values. In a second step, an *infill criterion* is used to find promising decision values to be evaluated on this expensive function. The infill criterion has to face a trade-off between “exploitation”—evaluating points for which the surrogate

GA-MO-FE: NSGA-II with feature ensemble mutation

GA-MO: NSGA-II with Hamming-weight preserving mutation

GA-MO-FE-NJ: NSGA-II for features; hyperparameters fixed

BO-MO-FE, BO-SO-FE, BO-MO-FE-NJ: BO with filter ensemble

BO-MO, BO-SO: BO with individual filter selection

Figure 1: Representation of individuals in different variants of our proposed optimization methods as described in Section 5.3; λ , w , and s are vectors as described in Section 4.

predicts good performance—and “exploration”—evaluating points where predictive uncertainty is high.

There are different ways of adapting BO to perform multi-objective Bayesian optimization (“MOBO”) [21], which has been applied successfully to hyperparameter tuning [19]. We chose the *Pareto* method [27]. Pareto scalarizes the different objectives through the Chebyshev norm using a random weight vector, which is sampled again for every point being proposed. Pareto has the advantage over many other MOBO methods of only requiring a single surrogate model fit for a proposed point [19], but other MOBO approaches can trivially be used as a substitute within our method.

The number of possible feature configurations is exponential in p , so a large number of evaluations would be necessary to fit an accurate surrogate model on performance values. We therefore investigate the use of prior knowledge from *feature filter* scores to reduce the dimensionality of the search space. Instead of searching over the discrete space of possible feature configuration vectors $s \in \{0, 1\}^p$, we restrict the search space to the *fraction of selected features* $\text{ffrac} \in [0, 1]$. The explicit feature configuration s results from selecting the ffrac most highly ranked features according to an individual filter (4.2.1) or according to the filter ensemble (4.2.2). This re-parametrization reduces the set of possible feature configurations being considered, and it may miss features that work well in combination but get a poor filter score because they perform poorly individually. Ultimate performance therefore depends on the quality of the filters being used, and we investigate methods that use combinations of filters to lessen the dependency on individual filter performance and therefore boost overall performance.

Two possible methods for simultaneous use of multiple feature filter methods are considered (see also Figure 1):

4.2.1 Individual filter selection. In this method, we introduce a discrete *filter selection hyperparameter* m , as well as a *feature fraction hyperparameter* $\text{ffrac} \in [0, 1]$. For each model evaluation, only the most relevant $\lceil p \cdot \text{ffrac} \rceil$ features, according to filter with

index m , are included in the model. This approach is similar to the one taken in auto-sklearn [14], although there the feature selection problem was not considered as a multi-objective problem.

4.2.2 Filter ensemble selection. This method uses the filter ensemble as shown in Equation 1 and introduces the vector w , as well as the aforementioned feature fraction ffrac , as hyperparameters. The most relevant $\lceil p \cdot \text{ffrac} \rceil$ features, according to $\text{EF}_j(w)$, are included in the model.

4.3 Evolutionary Approach

The *Nondominated Sorting Genetic Algorithm II* (NSGA-II) [10] is an evolutionary multi-objective algorithm that uses nondominated-sorting to preferably select individuals close to the Pareto-front of the problem. It iterates through generations of each a *reproduction*, a *crossover*, a *mutation*, and a *survival* step that generate the population of the next generation.

GAs often represent individuals as vectors of binary, discrete, or continuous values, depending on the optimization problem. Because hyperparameters generally have various types, we use the Cartesian products of operators that operate in different ways on the various types, following Li et al. [29]. This means e.g. that numeric hyperparameters undergo Gaussian mutation, while categorical hyperparameters undergo uniform mutation etc. Table 4 in the supplement summarizes the chosen recombination and mutation operators for the respective hyperparameter types. For the hyperparameter mutations, we use self-adapting step sizes and mutation probabilities as suggested by Li et al. [29].

The feature configuration vector parameter s plays a special role, because s maps to the objective of the fraction of selected features, which is being optimized, in a straightforward manner. We therefore consider the initialization and mutation of s in detail.

4.3.1 Geometric initialization. A naive approach for feature configuration initialization would be Bernoulli-sampling of each feature selection bit s_i individually, possibly biased towards a low expected number of selected features to favor relatively sparse solutions [1]. However, this gives rise to a binomial distribution of the number of selected features $\sum_i s_i$ with standard deviation $\sim O(\sqrt{p})$. For even moderately large values of p , this fails to cover the objective space evenly along the dimension of the selected feature fraction. Figure 6 in the supplement illustrates this. We propose to sample values of s such that the sum of selected features covers the whole feature fraction objective. Therefore, we elect to use a truncated geometric distribution of number of included features to encode our preference for sparse models. This is achieved by sampling the desired number of included features S as a truncated geometrically distributed random integer between 0 and p , and then uniformly sampling from all vectors s that satisfy $\sum_i s_i = S$.

The method introduces the success probability of the geometric distribution as a configuration parameter. It can be set by the user to encode a relative preference for sparsity. We chose to use an empirically determined value by fitting decision trees on 100 random subsets of 90% of the data set and determining the average number of distinct split variables.

4.3.2 Filter-ensemble based initialization. It is possible to enhance the geometric initialization by including prior knowledge

gained from feature filter methods. The goal is to select the features that have high filter scores with larger probability than the ones with lower ranking, while still having an approximately geometric distribution over the number of total features selected. We make the initial distribution of bit j dependent on the filter ensemble value as described in equation 1, with $\mathbf{w}$ uniformly randomly sampled from the simplex $\mathbf{w} \in [0, 1]^M$, $\sum_{m=1}^M w_m = 1$. For each individual to be initialized, we first sample S as in 4.3.1. Each bit j is then sampled from a Bernoulli-distribution with parameter:

$$\pi_B(\mathbf{w}, S) = \frac{\text{EF}_j(\mathbf{w})(S+1)}{\text{EF}_j(\mathbf{w})S + (1 - \text{EF}_j(\mathbf{w}))(p - S) + 1}. \quad (2)$$

4.3.3 Hamming-weight preserving mutation. Performing random bit-flip mutation on the feature selection vector entails similar problems to Bernoulli-initialization: a bit-flip with probability $\pi_{\text{Mut}_{\text{bin}}}$ is equivalent to erasing a bit with probability $2\pi_{\text{Mut}_{\text{bin}}}$ and sampling it anew from a $\frac{1}{2}$ -Bernoulli distribution. This biases the mutation result towards the $\frac{1}{2}$ point of the feature fraction objective. Instead, we choose to approximately preserve the Hamming weight. This is done by sampling erased bits from a Bernoulli-distribution with parameter $\pi_B = (S+1)/(p+2)$, where $S = \sum_i s_i$ is the Hamming weight of the original vector.

4.3.4 Filter-ensemble based mutation. Just as for initialization, mutation can be made dependent on feature filter ensemble ranks to preferentially include more relevant features. As in the Hamming-weight preserving mutation, each bit is erased with a probability of $2\pi_{\text{Mut}_{\text{bin}}}$ and then drawn from a Bernoulli-distribution. The parameter is as in equation 2 for filter-based initialization, with S the Hamming weight of the original vector, and $\mathbf{w}$ a weighting vector. This weighting vector itself becomes part of the search space to achieve self-adaption, as described in Li et al. [29], which results in $\mathbf{w}$ being optimized during the NSGA-II run. This way, the Hamming-weight is (approximately) preserved.

4.4 Implementation and Reproducibility

All proposed methods are implemented and publicly accessible through the `mosmafs` R package published on CRAN³, using the `mlr` [2] machine learning framework as backend. For full reproducibility of the results we publish the code used to perform the benchmark experiments on Github⁴.

5 EXPERIMENTS

We conduct experiments to answer the following research questions empirically:

- (1) *Evolutionary vs. Bayesian Optimization:* How do the proposed methods—the evolutionary and the Bayesian optimization approach—perform relative to each other?
- (2) *Effect of Filter Ensembles:* Do the methods benefit from using filter ensembles?
- (3) *Multi-Objective vs. Single-Objective:* Does multi-objective optimization find much sparser solutions without a major loss in predictive performance compared to single-objective optimization?

³<https://CRAN.R-project.org/package=mosmafs>

⁴<https://github.com/compstat-lmu/mosmafs/tree/master>

Table 1: Description of the datasets being used. n denotes the number of observations and p the total number of features. The class ratio gives the proportional size of the smaller outcome class. The dataset id (*did*) is the unique identifier for the dataset on the OpenML platform [34].

name	n	p	class ratio	n/p	did
wdbc	569	30	0.37	18.97	1510
ionosphere	351	33	0.36	10.64	59
sonar	208	60	0.47	3.47	40
hill-valley	1212	100	0.50	12.12	1479
tecator	240	124	0.43	1.94	851
semeion	319	256	0.50	1.25	41973
madeline	3140	259	0.50	12.12	41144
lsvt	126	307	0.33	0.41	1484
madelon	2600	500	0.50	5.20	1485
isolet	600	617	0.50	0.97	41966
cnae-9	240	282	0.50	0.85	41967
arcene	200	9961	0.44	0.02	1458
AP_Breast_Colon	630	10935	0.45	0.06	1145
AP_Colon_Kidney	546	10935	0.48	0.05	1137

- (4) *Simultaneous Hyperparameter Tuning and Feature Selection:* Is it beneficial to perform hyperparameter optimization and feature selection *simultaneously* compared to performing the tasks sequentially?

5.1 Benchmark Datasets

Because our implementation is based off the `mlr` framework, it can work with a wide variety of regression and classification tasks, even in conditions of missing values or class imbalance. However, we conduct our experiments on a set of well-behaved problems to limit the scope of the experiments necessary to show effectiveness.

We consider real-world binary classification tasks that are publicly accessible through the OpenML platform [34] (see Table 1). To eliminate algorithmic factors that might influence the result (like class imbalance correction, feature encoding, or handling of missing values), we included datasets fulfilling the following criteria: (roughly) balanced outcome classes, a purely numeric feature space, and no missing values. Further, datasets have been chosen to represent a large variety in terms of dimensionality ($30 \leq p \leq 10937$)⁵ and in terms of instances per dimension ($0.02 \leq n/p \leq 18.97$).

5.2 Learning Algorithms and their Hyperparameters

We consider three different classifiers that are tuned by the optimization algorithms proposed in Section 4: The support vector machine classifier (SVM) [6] with Gaussian kernel, extreme gradient boosting (xgboost) [5] and the kernelized k-nearest-neighbor classifier (kkn) [39]. We decided to consider these learners because they represent three very distinct learning paradigms, and because they are generally regarded responsive to hyperparameter tuning. The hyperparameter spaces that are being tuned over are presented

⁵To limit the computational resources necessary for our experiments we did not investigate even larger datasets.

Table 2: Hyperparameter spaces over which tuning was performed for the three learning algorithms *xgboost*, *support vector machine* and *kernelized k-Nearest-Neighbors* respectively. Ranges marked with * were tuned on a logarithmic scale. For the large datasets *madeline*, *madelon*, *arcene*, *AP_Breast_Colon* and *AP_Colon_Kidney*, the *xgboost nrounds* parameter was set fixed to 2000 and early stopping after 10 rounds was enabled.

SVM		xgboost	
kernel	rbfdot	nrounds	$\{1, 2, \dots, 2000\}$
sigma	$[2^{-10}, 2^{10}]^*$	eta	$[0.01, 0.2]$
C	$[2^{-10}, 2^{10}]^*$	gamma	$[2^{-7}, 2^6]^*$
knn		max_depth	$\{3, \dots, 20\}$
		colsample_bytree	$[0.5, 1]$
k	$\{1, 2, \dots, 50\}$	colsample_bylevel	$[0.5, 1]$
distance	$[1, 100]$	lambda	$[2^{-10}, 2^{10}]^*$
kernel	{rectangular, optimal, triangular, biweight}	alpha	$[2^{-10}, 2^{10}]^*$
		subsample	$[0.5, 1]$

in Table 2. They also react differently to high dimensionality: knn is vulnerable to the “curse of dimensionality”, the SVM is a regularized modeling algorithm, and the xgboost algorithm does implicit feature selection because it is a tree-based learner.

5.3 Algorithms

We perform hyperparameter optimization and feature selection for the learners described in Section 5.2. The datasets used are presented in Section 5.1. We use different configurations of our methods for comparisons described in the following. See also Figure 1 for a schematic representation of individual configuration vectors used in each method.

To answer research question 1, comparing the evolutionary and Bayesian optimization methods, we consider:

GA-MO-FE The NSGA2 with filter ensemble based initialization (as in 4.3.2) and mutation (as in 4.3.4).

BO-MO-FE Multi-objective Bayesian optimization with filter ensemble selection as described in 4.2.2.

To study research question 2, i.e. the contribution of filter ensembles to the methods, we compare both methods to these algorithms without filter ensembles:

GA-MO The NSGA2 with filter ensemble based initialization (as in 4.3.2) but *without* filter-based mutation, instead only using Hamming-weight preserving mutation as in 4.3.3.

BO-MO Multi-objective Bayesian optimization with individual filter selection, as described in 4.2.1.

We illuminate research question 3 by considering feature selection and hyperparameter optimization as a single-objective task, run with the following methods:

BO-SO Single-objective Bayesian optimization with individual filter selection (see 4.2.1)—this is a method similar to auto-sklearn’s approach [14] which we consider to be a state-of-the-art approach.

BO-SO-FE Single-objective Bayesian optimization with filter ensemble selection (see 4.2.2), which is our strongest single objective baseline.

To address research question 4, we consider variants of our algorithms that perform feature selection and hyperparameter optimization in separate steps. We construct two algorithms to approximate the most straightforward simplifications to non-joint optimization. Note that the GA-approach has a focus on feature selection, while the BO-approach is more directed at hyperparameter tuning.

GA-MO-FE-NJ We initially optimize hyperparameters for the learning algorithms on each dataset by running standard single objective BO on the *full* datasets without performing feature selection, through 500 model evaluations. The resulting hyperparameter values are then fixed, while our full NSGA2-variant (using both filter ensemble initialization and mutation) only performs multi-objective feature selection.

BO-MO-FE-NJ To optimize hyperparameters and filter weights we use the BO-SO-FE method, and then evaluate the final model with the fixed hyperparameters. We iterate over equally spaced frac values and evaluate the model with all filters. This constructs a set of models with different trade-offs between sparsity and model performance, from which we construct a Pareto-set.

The methods have access to $M = 5$ filter methods⁶: Information Gain, Random Forest Feature Importance [23], Joint Mutual Information (JMI) [38], Minimal Conditional Mutual Information Maximization (CMIM) [15], Area Under the Curve (AUC).

All NSGA2 variants use $\mu = 80$ as population and $\lambda = 15$ as offspring size which Khan and Baig [26] found to perform well in a feature selection setting. We choose an overall per-individual mutation probability of 0.3 and an overall per-pair-of-individuals crossover probability of 0.7.

Our Bayesian optimization methods use a random forest as surrogate to model the mixed discrete and continuous hyperparameter space similar to the state-of-the-art hyperparameter optimization toolbox SMAC [22]. The infill criterion used is LCB [7]. In each iteration, a batch of 15 configurations is proposed in parallel as described in Horn et al. [20].

5.4 Evaluation

We measure the performance of resulting models by their mean misclassification error ($\text{mmce} = \frac{1}{n} \sum_i \mathbb{I}_{y^{(i)} \neq \hat{y}^{(i)}}$) on a validation set with ground truth values $y^{(i)}$ and model predictions $\hat{y}^{(i)}$.

To get an estimate of the optimization performance that is unbiased by potential overtuning, we performed *nested resampling*: During the whole optimization run, each optimization algorithm is only allowed to assess model performance $\text{mmce}_{\text{optim}}$ through (*inner*) cross-validation (CV) on a optimization set $\mathcal{D}_{\text{optim}}$. The final performance of the algorithm is reported as the performance of

⁶The methods were chosen by running many filters on a range of datasets and performing hierarchical clustering on the differences of their feature rankings; see the Supplement for more details.

the solution candidates trained on $\mathcal{D}_{\text{optim}}$ and evaluated on a test set $\mathcal{D}_{\text{test}}$. This procedure is repeated 10 times on *outer* CV folds $(\mathcal{D}_{\text{optim}}^{(k)}, \mathcal{D}_{\text{test}}^{(k)}), k = 1, \dots, 10$.

As a proper multi-objective performance measure we consider the *dominated hypervolume* [40] (domHV) with reference point $W = (1, 1)$, which corresponds to the worst possible values w.r.t. the two objectives. To prevent overtuning effects from skewing our results, we proceed as follows:

- (1) We let each optimization algorithm report its Pareto set, taking into account only $\mathcal{D}_{\text{optim}}$. This corresponds to returning non-dominated individuals w.r.t. (mmce_{optim}, frac). GA-based methods take only candidates from their current generation, while BO-based methods get the Pareto set of all candidates seen so far.
- (2) We calculate the *generalization* dominated hypervolume domHV_{gen}, i.e. the hypervolume that is dominated by this Pareto set of candidate configurations w.r.t. the (mmce_{test}, frac) evaluated on $\mathcal{D}_{\text{test}}$.⁷

For each of the experiments, we allow 2000 model evaluations, with one evaluation corresponding to computing a single inner 10-fold CV on $\mathcal{D}_{\text{optim}}$. This corresponds to approximately 130 iterations for both the model-based and the evolutionary approach. We test performance differences for significance using the critical difference test of Demšar [11], which is a non-parametric test for comparison of different methods on different datasets.

6 RESULTS

6.1 Evolutionary vs. Bayesian Optimization, Effect of Filter Ensembles

The global rank analysis as well as a critical difference test [11] presented in Figure 2 show that our most advanced BO-based method, BO-MO-FE, significantly outperforms all GA methods. The optimization trace seems to indicate that including the feature ensemble also confers an advantage, possibly because linear combinations of filters can perform better than any individual filter. This difference is not statistically significant after 2000 evaluations, however.

Figure 5 shows an additional aspect to take into consideration: overall runtime. BO-MO-FE usually has a slight performance advantage over GA-MO-FE in absolute terms, but comes with computational overhead that may have to be considered if individual model performance evaluations are relatively fast. The overhead comes mostly from surrogate model fitting and infill optimization, with runtime dependent on the number of hyperparameters, but independent of dataset size. The overhead becomes less relevant when optimizing large datasets or slow models.

6.2 Multi-Objective vs. Single-Objective

The goal of our multi-objective methods is not to optimize predictive performance by itself, and instead to explore the optimal trade-offs between performance and sparseness. It is still interesting to look at the best performing model configurations being found by the multi-objective methods, and to compare them to the best models found

⁷Note that domHV_{gen} is not the same as the dominated hypervolume of a population on the test set. Instead, only the individuals that are non-dominated according to $\mathcal{D}_{\text{optim}}$ are used to calculate their dominated hypervolume on $\mathcal{D}_{\text{test}}$.

Table 3: domHV_{gen} after 2000 evaluations of simultaneous hyperparameter tuning / feature selection methods compared to corresponding non-simultaneous methods. Best results for GA and BO in bold. Values averaged over datasets, see supplement for results by dataset.

Learner	BO-MO-FE	BO-MO-FE-NJ	GA-MO-FE	GA-MO-FE-NJ
kknn	0.9217	0.8670	0.9108	0.9134
SVM	0.9331	0.8719	0.9241	0.8892
xgboost	0.9164	0.8940	0.9165	0.9121

by single-objective methods, as done in Figure 4. Here, again, the BO methods outperform GA methods. Except for the *kknn* learning algorithm, the MO-method is on par with the SO methods.

The BO-SO method employs a feature filtering step, and the models it chooses will often be sparse to some degree, for regularization. Figure 3 compares the BO-SO method to the multi-objective methods in terms of both sparseness and predictive performance: For each learning algorithm, dataset, and CV fold of BO-SO, the best performing models from BO-MO-FE and GA-MO-FE were chosen that are *at least* as sparse. Note that the *length* of arrows in this figure depends mainly on the size of the population returned by the methods and is not directly of interest. Instead the *slope* should be considered, with steeper arrows pointing to the right indicating a cheaper trade for sparseness vs. performance (and arrows pointing to the left indicating results outperforming BO-SO). The plot shows that both MO methods often found much sparser models than the SO baseline while giving up very little in predictive performance.

6.3 Simultaneous Hyperparameter Tuning and Feature Selection

Table 3 shows the domHV_{gen} of our joint optimization methods, compared to their non-joint correspondents. In most cases the joint optimization confers a considerable performance advantage, especially compared to the BO-SO-FE-NJ method. We assume that the advantage of joint over non-joint methods depends on the interaction strength of hyperparameter performance and sparseness.

7 CONCLUSION

Our results show that both evolutionary approaches and model-based approaches can efficiently perform model-agnostic multi-objective optimization to simultaneously tune hyperparameters and select features. We have also shown that performing these tasks simultaneously has an advantage over running them separately.

Our experiments suggest that using multi-objective model-based optimization to tune feature inclusion based on single filter works well, although a parameterized filter ensemble performs better.

Our other proposed method based on an NSGA-II enhanced with specialized initialization and mutation sampling seems to perform almost as well, although it does not reach the level of the best Bayesian optimization based approach. The advantage of the NSGA-II, however, is that it does not introduce as much computational overhead, which can make up a considerable part of overall runtime if model performance evaluations themselves are cheap.

Figure 2: Comparison of different methods. Top: the results of a global rank analysis based on $\text{domHV}_{\text{gen}}$ to compare the proposed methods BO-MO(-FE) and GA-MO(-FE), i.e. both the BO and GA with and without filter ensemble. Ranks are computed per dataset and algorithm (ties are ranked by their average rank) and then averaged. Higher values are better. Bottom: Non-parametric critical difference test [11] performed after the full budget of 2000 evaluations at significance level $\alpha = 0.05$, based on $\text{domHV}_{\text{gen}}$ rank as above. The BO-MO-FE method statistically significantly outperforms all GA methods, although the absolute difference is small (see Figure 5).

Figure 3: Comparison of multi-objective and single-objective methods applied to the SVM learning algorithm: Performance $\text{mmce}_{\text{test}}$ and the fraction of included features frac found after 2000 evaluations by baseline BO-SO (tail end of arrows), BO-MO-FE (head of solid arrows) and GA-MO-FE (head of dashed arrows). Each dataset (Table 1) is shown, values are averaged over 10 outer CV runs. Choice of individuals for MO methods described in Section 6.2.

Figure 4: Comparison of multi-objective and single-objective methods: $\text{mmce}_{\text{test}}$ of the best $\text{mmce}_{\text{optim}}$ individual found until a particular evaluation, averaged over 10 outer CV folds and over all datasets. Lower values are better. Note that $\text{mmce}_{\text{test}}$ is not monotonically improving because individuals with better $\text{mmce}_{\text{optim}}$ may be found that perform worse on the test set than previous individuals.

Figure 5: Performance ($\text{domHV}_{\text{gen}}$) and runtime of GA-MO-FE (tail end of arrow) and BO-MO-FE (head of arrow) on each dataset (Table 1) and learning algorithm (Table 2), averaged over 10 outer CV runs. BO-MO-FE has moderately improved performance over GA-MO-FE, but at a cost in runtime.

Our final recommendation is therefore to use the Bayesian optimization approach in combination with parameterized Filter ensembles if model evaluations are expensive, if computational resources are cheap, and if it is important to get configurations that perform very close to optimal, given their sparseness. The NSGA-II approach is suitable if model evaluations are cheap and if marginal degradation of performance are acceptable.

ACKNOWLEDGMENTS

This work has been partially funded by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A. The authors of this work take full responsibilities for its content.

This work was supported by the Bavarian Ministry for Economic Affairs, Infrastructure, Transport and Technology through the Center for Analytics-Data-Applications (ADA-Center) within the framework of “BAYERN DIGITAL II”.

REFERENCES

- [1] Bernd Bischl, Igor Vatolkin, and Mike Preuss. Selecting small audio feature sets in music classification by means of asymmetric mutation. In *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature*, volume 6238, pages 314–323, 2010. ISBN 3642158439.
- [2] Bernd Bischl, Michel Lang, Jakob Bossek, Leonard Judt, Jakob Richter, Tobias Kuehn, and Erich Studerus. mlr: Machine Learning in R. *Journal of Machine Learning Research*, 17(170):1–5, 2014.
- [3] Amparo Bolón-Canedo, Verónica; Alonso-Betanzos. Ensembles for feature selection: A review and future trends. *Information Fusion*, 11 2018. doi: 10.1016/j.inffus.2018.11.008.
- [4] Amal Bouraoui, Salma Jamoussi, and Yassine BenAyed. A multi-objective genetic algorithm for simultaneous model and feature selection for support vector machines. *Artificial Intelligence Review*, 50(2):261–281, 2018. ISSN 15737462.
- [5] Tianqi Chen and Carlos Guestrin. *XGBoost: A Scalable Tree Boosting System*. ACM, 2016. ISBN 9781450342322.
- [6] Corinna Cortes and Vladimir Vapnik. Support-Vector Networks. *Machine Learning*, 20(3):273–297, 1995. ISSN 15730565.
- [7] Dennis D Cox and Susan John. A statistical method for global optimization. In *[Proceedings] 1992 IEEE International Conference on Systems, Man, and Cybernetics*, pages 1241–1246. IEEE, 1992.
- [8] M. Dash and H Liu. Feature selection for classification. *Intelligent Data Analysis*, 1(3):131–156, 1997. ISSN 1088467X.
- [9] Kalyanmoy Deb and R B Agrawal. Simulated binary crossover for continuous search space. *Complex Systems*, 9(2):115–148, 1994. ISSN 08912513.
- [10] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. ISSN 1089778X.
- [11] Janez Demsar. Statistical comparisons of classifiers over multiple data sets. *J. Mach. Learn. Res.*, 7:1–30, December 2006. ISSN 1532-4435.
- [12] David J. Dittman, Taghi M. Khoshgoftaar, Randall Wald, and Amri Napolitano. Comparing two new gene selection ensemble approaches with the commonly-used approach. In *2012 11th International Conference on Machine Learning and Applications*. IEEE, December 2012. doi: 10.1109/icmla.2012.175. URL <https://doi.org/10.1109/icmla.2012.175>.
- [13] C. Emmanouilidis, A. Hunter, and J. Macintyre. A multiobjective evolutionary setting for feature selection and a commonality-based crossover operator. In *Proceedings of the 2000 Congress on Evolutionary Computation. CEC00 (Cat. No.00TH8512)*, volume 1, pages 309–316. IEEE, 2000. ISBN 0-7803-6375-2.
- [14] Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter. Efficient and Robust Automated Machine Learning, 2015. URL <http://papers.nips.cc/paper/5872-efficient-and-robust-automated-machine-learning>.
- [15] Francois Fleuret. Fast Binary Feature Selection with Conditional Mutual Information. *Journal of Machine Learning Research*, 5:1531–1555, 2004.
- [16] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, Proceedings of Machine Learning Research. PMLR, 2018.
- [17] Isabelle Guyon and Andre Elisseeff. An Introduction to Variable and Feature Selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003. ISSN 00032670.
- [18] Tarek M. Hamdani, Jin-Myung Won, Adel M. Alimi, and Fakhri Karray. Multi-objective Feature Selection with NSGA II. In *Adaptive and Natural Computing Algorithms*, pages 240–247. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.
- [19] Daniel Horn and Bernd Bischl. Multi-objective parameter configuration of machine learning algorithms using model-based optimization. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8. IEEE, 2016.
- [20] Daniel Horn, Tobias Wagner, Dirk Biemann, Claus Weihs, and Bernd Bischl. Model-based multi-objective optimization: Taxonomy, multi-point proposal, toolbox and benchmark. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9018, pages 64–78, 2015. ISBN 9783319159331.
- [21] Daniel Horn, Tobias Wagner, Dirk Biemann, Claus Weihs, and Bernd Bischl. Model-based multi-objective optimization: Taxonomy, multi-point proposal, toolbox and benchmark. In *Lecture Notes in Computer Science*, pages 64–78. Springer International Publishing, 2015. doi: 10.1007/978-3-319-15934-8_5. URL https://doi.org/10.1007/978-3-319-15934-8_5.
- [22] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential Model-Based Optimization for General Algorithm Configuration. In *International conference on learning and intelligent optimization*. Springer, 2011. URL <http://www.cs.ubc.ca/labs/beta/Projects/SMAC/papers/11-LION5-SMAC.pdf>.
- [23] Hemant Ishwaran, Udaya B Kogalur, Eiran Z Gorodeski, Andy J Minn, and Michael S Lauer. High-dimensional variable selection for survival data. *Journal of the American Statistical Association*, 105(489):205–217, 2010. ISSN 01621459.
- [24] Donald R Jones, Matthias Schonlau, and William J Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13: 455–492, 1998. URL <http://www.ressources-actuarielles.net/EXT/ISFA/1226.nsf/0/f84f7ac703bf5862c12576d8002f5259/{%protect%T1%textdollar}FILE/Jones98.pdf>.
- [25] T. Meyarivan K. Deb, A. Pratap, S. Agarwal. A fast and elitist multi-objective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation*. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [26] A. Khan and A. R. Baig. Multi-objective feature subset selection using non-dominated sorting genetic algorithm. *Journal of Applied Research and Technology*, 13(1):145–159, 2015. ISSN 16656423.
- [27] J. Knowles. ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, feb 2006. ISSN 1089-778X. doi: 10.1109/TEVC.2005.851274. URL <http://ieeexplore.ieee.org/document/1583627/>.
- [28] Ron Kohavi and George H. John. Wrappers for feature subset selection. *Artificial Intelligence*, 97(1-2):273–324, 2002. ISSN 00043702.
- [29] Rui Li, Michael T.M. Emmerich, Jeroen Eggermont, Thomas Bäck, M. Schütz, J. Dijkstra, and J. H.C. Reiber. Mixed integer evolution strategies for parameter optimization. *Evolutionary Computation*, 21(1):29–64, 2013. ISSN 10636560.
- [30] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018.
- [31] Randal S. Olson and Jason H. Moore. TPOT: A Tree-Based Pipeline Optimization Tool for Automating Machine Learning. pages 151–160. Springer, Cham, 2019. doi: 10.1007/978-3-030-05318-5_8. URL http://link.springer.com/10.1007/978-3-030-05318-5_8.
- [32] Florian Pfisterer, Stefan Coors, Janek Thomas, and Bernd Bischl. Multi-objective automatic machine learning with autogxboostmc. *arXiv preprint arXiv:1908.10796*, 2019.
- [33] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical Bayesian Optimization of Machine Learning Algorithms. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2951–2959, 2012.
- [34] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014. ISSN 19310145. URL <http://dl.acm.org/citation.cfm?doid=2641190.2641198>.
- [35] Kashif Waqas, Rauf Baig, and Shahid Ali. Feature subset selection using multi-objective genetic algorithms. In *IEEE 13th International Multitopic Conference*, pages 1–6, 2009. ISBN 9781424448722.
- [36] Bing Xue, Mengjie Zhang, and Will N Browne. Particle swarm optimization for feature selection in classification: A multi-objective approach. *IEEE Transactions on Cybernetics*, 43(6):1656–1671, 2013. ISSN 21682267.
- [37] Bing Xue, Mengjie Zhang, Will N. Browne, and Xin Yao. A Survey on Evolutionary Computation Approaches to Feature Selection. *IEEE Transactions on Evolutionary Computation*, 20(4):606–626, 2016. ISSN 1089778X.
- [38] H Yang. Data visualization and feature selection: New algorithms for nongaussian data. *Advances in Neural Information Processing Systems*, pages 687–693, 1999.
- [39] Kai Yu, Liang Ji, and Xuegong Zhang. Kernel Nearest-Neighbor Algorithm. *Neural Processing Letters*, 15(2):147–156, 2002. ISSN 13704621. doi: 10.1023/A:1015244902967. URL <http://link.springer.com/10.1023/A:1015244902967>.
- [40] Eckart Zitzler and Lothar Thiele. Multiobjective Optimization Using Evolutionary Algorithms — a Comparative Case Study. In Agoston E. Eiben, Thomas Bäck, Marc Schoenauer, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature — PPSN V*, pages 292–301, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg. ISBN 978-3-540-49672-4.