

Constraint-handling for Optimization with Support Vector Surrogate Models

A Novel Decoder Approach

Jörg Bremer and Michael Sonnenschein

Department of Computing Science, University of Oldenburg, Uhlhornsweg, Oldenburg, Germany

Keywords: Smart Grid, Soft Computing, Evolutionary Optimization, Constraint Modeling, SVDD.

Abstract: A new application for support vector machines is their use for meta-modeling feasible regions in constrained optimization problems. We here describe a solution for the still unsolved problem of a standardized integration of such models into (evolutionary) optimization algorithms with the help of a new decoder based approach. This goal is achieved by constructing a mapping function that maps the whole unconstrained domain of a given problem to the region of feasible solutions with the help of the the support vector model. The applicability to real world problems is demonstrated using the load balancing problem from the smart grid domain.

1 INTRODUCTION

A popular class of commonly used heuristics for solving hard optimization problems is known as evolutionary search methods. These methods usually work with candidate solutions that encode each parameter within an allowed interval between a lower and an upper limit and try to improve them within these bounds. Thus, all solutions are defined in a d -dimensional hypercube. Nevertheless, due to additional constraints, not all of these solutions are usually feasible. Effectively solving real world optimization problems often suffers from the presence of constraints that have to be obeyed when exploring alternative solutions. Evolutionary algorithms have been widely noticed due to their potential for solving complex (discontinuous or non differentiable) numerical functions. However, a full success in the field of nonlinear programming problems is still missing, because constraints have not been addressed for integration in a systematic way (Michalewicz and Schoenauer, 1996; Kramer, 2010).

For constraint-handling techniques, the general constrained continuous nonlinear programming (NLP) problem is often used as problem formulation:

Find $x \in \mathbb{R}^d$ that optimizes $f(x)$ subject to a set of constraints:

$$\begin{aligned} \text{equalities: } & g_i(x) = 0; & 1 \leq i \leq m \\ \text{inequalities: } & g_j(x) \leq 0; & 1 \leq j \leq n. \end{aligned} \quad (1)$$

Real world problems often additionally face nonlinear constraints or such that are not given in an explicit

formulation. One example for a not explicitly given constraint is a simulation model that devalues given solutions as not feasible judged by simulation runs. We are going to focus (but not restrict ourselves) on the latter type.

In general, the set of constraints defines an arbitrary shaped region within the search space (the hypercube defined by parameter bounds) that contains all feasible solutions. Taking into account nonlinear constraints, the NLP is generally intractable (Michalewicz and Schoenauer, 1996). Evolutionary Algorithms approximately solve non linear optimization very efficiently. Nevertheless, surprisingly low effort has been put in the integration of constraint handling techniques (cf. (Kramer, 2010)). Standard constraint-handling techniques are for example the introduction of a penalty for infeasible solutions, the separation of objectives and constraints to transforms a given optimization problem into an unconstrained many-objective one, or decoder approaches to give an algorithm hints on how to construct feasible solutions by imposing a relationship between feasibility and decoder solution.

At the same time, support vector machines and related approaches have been shown to have excellent performance when trained as classifiers for multiple purposes, especially real world problems. As a use case related to describing the region where some given data resides in, (Tax and Duin, 1999) developed the support vector domain description (SVDD) as a one-class support vector classification approach

that is capable of learning the region that is defined by some given training data and has therefore been harnessed for example by (Bremer et al., 2011) as a model for the feasible region. This model so far only allows for afterwards checking the feasibility of an already given solution.

What we will now add to these two worlds is a new decoder approach for integrating constraints that are modeled by a support vector classifier into evolutionary optimization in a standardized way. The basic idea is to construct a mapping from the original, unconstrained domain of the problem (the hypercube) to the feasible space. In this way, the problem is transferred into an unconstrained one. All that we need for constructing this mapping is the set of support vectors together with the associated weights from the model.

After a brief review of constraint handling techniques and black-box modelling with support vector approaches, we introduce the underlying model and describe the construction of our mapping approach in detail. We present results from several test scenarios with artificial optimization problems and conclude with results from applying our method to the load balancing problem in smart grid scenarios.

2 RELATED WORK

Several techniques for handling constraints are known. Nevertheless, many are concerned with special cases of NLP or require priori knowledge of the problem structure for proper adaption (Michalewicz and Schoenauer, 1996). We will briefly discuss some prominent representatives of such techniques. A good overview can for instance be found in (Coello Coello, 2002) or, more recently, in (Kramer, 2010).

2.1 Constraint-handling

We will start with an overview on some currently used conventional constraint handling techniques followed by an introduction to the upcoming black-box approach of meta-modelling constraints.

Penalty. A widely and long since used approach for constraint handling is the introduction of a penalty into the objective function that devalues all solutions that violate some constraint. In this way, the problem is transformed into an unconstrained one. Most commonly used are exterior penalties that draw outside solutions towards the feasible region in contrast to interior ones that keep solutions inside, but require to start with a feasible solution (Coello Coello, 2002).

Separation of Objectives and Constraints. Constraints or aggregations of constraints may be treated

as separate objectives. This leads to a transformation into a (unconstrained) many objective problem. Such approaches have some computational disadvantages from determining Pareto optimality or may lack the ability (in the case of a disjoint region) to escape a sub-region (Coello Coello, 2002). Moreover, a functional description of constraints must be known here in advance. Such description is not readily available when using surrogate models that hide original relations and only model the original behaviour.

Solution Repair. Some combinatorial optimization problem allow for an easy repair of infeasible solutions. In this case, it has been shown that repairing infeasible solutions often outperforms other approaches (Liepins and Vose, 1990). This approach is closely related to the decoder based approaches.

Decoder. In order to give an algorithm hints on how to construct a solution, sometimes so called decoders impose a relationship between feasibility and decoder solutions. For example, (Koziel and Michalewicz, 1999) proposed a homomorphous mapping between an n-dimensional hyper cube and the feasible region in order to transform the problem into a topological equivalent one that is easier to handle, although with a need for extra parameters that have to be found empirically and with some extra computational efforts. In contrast, we will see later how a similar approach can be automatically derived from a given support vector description. Earlier approaches e.g. used Riemannian mapping (Kim, 1998).

2.2 Meta-modelling of Constraints

A relatively new constraint handling technique is the use of meta-models for black-box optimization scenarios with no explicitly given constraint boundaries. Such a model allows for efficiently checking feasibility and may allow for a search for the constraint boundary between a feasible and an infeasible solution in case a repair of a mutation is needed.

Various classification or regression methods might be harnessed for creating such models for the boundary (Kramer, 2010). There are two main reasons for using such an approach:

1. Substituting computational costs for constraint evaluation by a comparatively easy check through the model in case of computational hard constraint evaluation (e.g. in technical simulations).
2. Efficient communication in distributed environments and for spreading one's own set of alternative solutions to other actors in a distributed optimization scenario (e.g. in multi agent systems).

An example from the smart grid (the prospective future electricity grid) domain for the latter case that has been realized by an SVDD approach can be found in (Bremer et al., 2011). Besides, this domain serves as example for scenarios with (at least partly) unknown functional relationships of the constraints. The feasible region can sometimes only be derived with lacking full knowledge on hidden variables or intrinsic relations that determine the operability of a electric device and therewith the feasible region. This is due to the fact that the feasibility of a solution is derived from a technical simulation of the energy resources in order to gain a set of alternative feasible solutions (Bremer et al., 2010). The authors therefore have their model learned by SVDD from a set of operable (feasible) examples; (Blank et al., 2011) used a two-class SVM for learning operation point and bias (regarding allowed voltage and current bands) of a line in a power grid for easier classifying a grid state as feasible or not. In both cases, at the time of searching for the optimum, the only available information is the model, i.e. a set of support vectors and associated weights. Every information about the original constraints is no longer available in such scenarios.

2.3 Load Balancing in Smart Grids

For our real world use case we will briefly review load balancing approaches. Within the framework of today's (centralized) operation planning for power stations, different heuristics are harnessed. Short-term scheduling of different generators assigns (in its classical interpretation) discrete-time-varying production levels to energy generators for a given planning horizon (Pereira et al., 2008). It is known to be an NP-hard problem (Guan et al., 2003). Determining an exact global optimum is, in any case, not possible in practice until ex post due to uncertainties and forecast errors. Additionally, it is hard to exchange operational constraints in case of a changed setting (e.g. a new composition of energy resources) of the generation system.

Coordinating a pool of distributed generators and consumers with the intend to provide a certain aggregated load schedule for active power has some objective similarities to controlling a virtual power plant (VPP). Within the smart grid domain the volatile character of such a group of distributed energy resources (DER) has additionally to be taken into account. On an abstract level, approaches for controlling groups of distributed devices can be roughly divided into centralized and distributed scheduling algorithms.

Centralized approaches have long time dominated the discussion (Tröschel and Appelrath, 2009) and

are discussed in the context of static pools of DER with drawbacks and restrictions regarding scalability and particularly flexibility. Recently, distributed approaches gained more and more importance. Different works proposed hierarchical and decentralized architectures based on multi-agent systems and market based computing (Kok et al., 2008). Newer approaches try to establish self-organization between actors within the grid (Kamper and Esser, 2009; Mihailescu et al., 2011; Ramchurn et al., 2011).

In load balancing scenarios, a scheduling algorithm (centralized or distributed) must know for each participating energy resource which load schedules are actually operable (satisfy all constraints) and which are not. Each energy resource has to restrict its possible operations due to several constraints. These can be distinguished into hard constraints (usually technically rooted, e.g. minimum and/or maximum power input or output) and soft constraints (often economically or ecologically rooted, e.g. personal preferences like noise pollution in the evening).

When determining an optimal partition of the schedule for load distribution, an alternative schedule is taken from each DER's search space of individual operable schedules (individual feasible region) in order to assemble a desired aggregate load schedule. The feasible region of each DER might be efficiently encoded by SVDD. The model of each DER is so far submitted to a scheduling instance and used for checking feasibility of solutions. The question for a direct integration of these models in any optimization algorithm has as yet not been answered.

3 MAPPING ALGORITHM

We will now describe the integration of a SVDD based black-box model into an arbitrary evolutionary optimization algorithm with proper and effective constraint handling. The integration will not rely on any knowledge on the set of constraints as these are hidden by the model. We here propose handling the constraints in a different way. When learning the region that contains all feasible solutions by harnessing SVDD, the topological traits of this region are captured and encoded in a set of support vectors and a weight vector. We will harness this implicitly inherent information to transform the original parameter hypercube to resemble this region.

3.1 SVDD-Model for Feasible Regions

As a prerequisite for our mapping, we assume that the feasible region of an optimization problem has been

encoded by SVDD as e.g. described in (Bremer et al., 2011). We will briefly describe this approach before deriving our new utilization method. Given a set of data samples $x_i \in \mathcal{X}$, the inherent structure of the region where the data resides in is derived as follows: After mapping the data to a high dimensional feature space, the smallest images enclosing sphere is determined. When mapping back the sphere to data space, its pre-image forms a contour (not necessarily connected) enclosing the data sample.

This task is achieved by determining a mapping $\Phi : \mathcal{X} \subset \mathbb{R}^d \rightarrow \mathcal{H}, x \mapsto \Phi(x)$ such that all data from a sample from a region \mathcal{X} is mapped to a minimal hypersphere in some high-dimensional space \mathcal{H} . The minimal sphere with radius R and center a in \mathcal{H} that encloses $\{\Phi(x_i)\}_N$ can be derived from minimizing $\|\Phi(x_i) - a\|^2 \leq R^2 + \xi_i$ with $\|\cdot\|$ as the Euclidean norm and slack variables $\xi_i \geq 0$ for soft constraints.

After introducing Lagrangian multipliers and further relaxing to the Wolfe dual form, the well known Mercer's theorem (cf. e.g. (Schölkopf et al., 1999)) may be used for calculating dot products in \mathcal{H} by means of a kernel in data space: $\Phi(x_i) \cdot \Phi(x_j) = k(x_i, x_j)$. In order to gain a more smooth adaption, it is known to be advantageous to use a Gaussian kernel: $k_G(x_i, x_j) = e^{-\frac{1}{2\sigma^2} \|x_i - x_j\|^2}$ (Ben-Hur et al., 2001). Putting it all together, the equation that has to be maximized in order to determine the desired sphere is:

$$W(\beta) = \sum_i k(x_i, x_i) \beta_i - \sum_{i,j} \beta_i \beta_j k(x_i, x_j). \quad (2)$$

With $k = k_G$ we get two main results: the center $a = \sum_i \beta_i \Phi(x_i)$ of the sphere in terms of an expansion into \mathcal{H} and a function $R : \mathbb{R}^d \rightarrow \mathbb{R}$ that allows to determine the distance of the image of an arbitrary point from $a \in \mathcal{H}$, calculated in \mathbb{R}^d by:

$$R^2(x) = 1 - 2 \sum_i \beta_i k_G(x_i, x) + \sum_{i,j} \beta_i \beta_j k_G(x_i, x_j). \quad (3)$$

Because all support vectors are mapped right onto the surface of the sphere, the sphere radius R_S can be easily determined by the distance of an arbitrary support vector to the center a . Thus the feasible region can now be modeled as $\mathcal{F} = \{x \in \mathbb{R}^d | R(x) \leq R_S\} \approx \mathcal{X}$.

So far, such models have for example been used for efficiently communicating the feasible region of controllable energy resources in smart grid scenarios (Bremer et al., 2011). Only the comparably small set of support vectors together with a reduced version of vector β that contains non zero weight values (denoted w) for the support vectors has to be submitted. The model might then be used as a black-box that abstracts from any explicitly given form of constraints and allows for an easy and efficient decision on whether a given solution is feasible or not.

Moreover, as the radius functions Eq. 3 maps to \mathbb{R} , it allows for a conclusion about how far away a solution is from feasibility. Nevertheless, a systematic constraint-handling during optimization is not induced in this way. In the following, we will develop a more sophisticated way of integrating such SVDD surrogate models into optimization.

3.2 The Decoder Approach

Let \mathcal{F} denote the feasible region within the parameter domain of some given optimization problem bounded by an associated set of constraints. We do not make any assumptions on the constraints. It is known, that pre-processing the data by scaling it to $[0, 1]^d$ leads to better adaption (Juszczak et al., 2002). According to (Bremer et al., 2011), some energy domain problems require a rescaling of the domain to $[0, 1]^d$ for easier handling, too. For this reason, we consider optimization problems with scaled domains and denote with $\mathcal{F}_{[0,1]}$ the likewise scaled region of feasible solutions. We want to construct a mapping

$$\gamma : [0, 1]^d \rightarrow \mathcal{F}_{[0,1]} \subseteq [0, 1]^d; x \mapsto \gamma(x) \quad (4)$$

that is able to map the unit hypercube $[0, 1]^d$ onto the d-dimensional region of feasible solutions. We achieve this mapping as a composition of three functions: $\gamma = \Phi_\ell^{-1} \circ \Gamma_a \circ \hat{\Phi}_\ell$. Instead of trying to find a direct mapping to $\mathcal{F}_{[0,1]}$ we go through the kernel space. The commutative diagram (Eq. 5) sketches the idea. We start with an arbitrary point $x \in [0, 1]^d$ from the unconstrained d-dimensional hypercube and map it to an ℓ -dimensional manifold in kernel space that is spanned by the images of the ℓ support vectors. After drawing the mapped point to the sphere in order to pull it into the image of the feasible region, we search the pre-image of the modified image to get a point from $\mathcal{F}_{[0,1]}$.

$$\begin{array}{ccc} x \in [0, 1]^d & \xrightarrow{\hat{\Phi}_\ell} & \hat{\Psi}_x \in \mathcal{H}^{(\ell)} \\ \downarrow \gamma & & \downarrow \Gamma_a \\ x^* \in \mathcal{F}_{[0,1]} \subseteq [0, 1]^d & \xleftarrow{\Phi_\ell^{-1}} & \tilde{\Psi}_x \in \mathcal{H}^{(\ell)} \end{array} \quad (5)$$

3.2.1 Mapping to the SV Induced Subspace $\mathcal{H}^{(\ell)}$ with an Empirical Kernel Map

We will now have a closer look onto the respective steps of this procedure. Let

$$\begin{aligned} \Phi_\ell : \mathbb{R}^d &\rightarrow \mathbb{R}^\ell, \\ x &\mapsto k(\cdot, x)|_{\{s_1, \dots, s_\ell\}} \\ &= (k(s_1, x), \dots, k(s_\ell, x)) \end{aligned} \quad (6)$$

be the empirical kernel map w.r.t. the set of support vectors $\{s_1, \dots, s_\ell\}$. If Φ_ℓ is modified to

$$\hat{\Phi}_\ell : \mathbb{R}^d \rightarrow \mathcal{H}^{(\ell)}, \quad (7)$$

$$x \mapsto K^{-\frac{1}{2}}(k(s_1, x), \dots, k(s_\ell, x))$$

with $K_{ij} = k(s_i, s_j)$: the kernel Gram Matrix, then Eq. 7 maps points x, y from input space to \mathbb{R}^ℓ , such that $k(x, y) = \hat{\Phi}_\ell(x) \cdot \hat{\Phi}_\ell(y)$ (cf. (Schölkopf et al., 1999)).

With $\hat{\Phi}_\ell$ we are able to map arbitrary points from $[0, 1]^d$ to some ℓ -dimensional space $\mathcal{H}^{(\ell)}$ that contains a lower dimensional projection of the sphere. Again, points from $\mathcal{F}_{[0,1]}$ are mapped into or onto the projected sphere, outside points go outside the sphere and must be moved in $\mathcal{H}^{(\ell)}$ towards the center in order to draw them into the image of the feasible region.

3.2.2 Re-adjustment in Kernel Space

In general, in kernel space \mathcal{H} the image of the region is represented as a hypersphere \mathcal{S} with center a and radius $R_\mathcal{S}$ (Eq. 3). Points outside this hypersphere are not images of points from \mathcal{X} , i.e. in our case, points from $\mathcal{F}_{[0,1]}$ are mapped (by Φ) into the sphere or onto its surface (support vectors), points from outside $\mathcal{F}_{[0,1]}$ are mapped outside the sphere. Actually, using a Gaussian kernel, Φ maps each point into a n -dimensional manifold (with sample size n) embedded into infinite dimensional \mathcal{H} . In principle, the same holds true for a lower dimensional embedding spanned by ℓ mapped support vectors and the ℓ -dimensional projection of the hypersphere therein.

We now want to pull points from outside the feasible region into that region. As we do have rather a description of the image of the region, we draw images of outside points into the image of the region, i.e. into the hypersphere; precisely into its ℓ -dimensional projection. For this purpose we use

$$\tilde{\Psi}_x = \Gamma_a(\hat{\Psi}_x) = \hat{\Psi}_x + \mu \cdot (a - \hat{\Psi}_x) \cdot \frac{R_x - R_\mathcal{S}}{R_x} \quad (8)$$

to transform the image $\hat{\Psi}_x$ produced in step 1) into $\tilde{\Psi}_x \in \hat{\Phi}_\ell(\mathcal{F}_{[0,1]})$ by drawing $\hat{\Psi}_x$ into the sphere. Alternatively, the simpler version

$$\tilde{\Psi}_x = a + \frac{(\hat{\Psi}_x - a) \cdot R_\mathcal{S}}{R_x} \quad (9)$$

may be used for drawing $\hat{\Psi}_x$ just onto the sphere but then without having to estimate parameter $\mu \in [1, R_x]$. Parameter μ allows us to control how far a point is drawn into the sphere ($\mu = 1$ is equivalent to eq. 9, $\mu = R_x$ draws each point onto the center). In this way, each image is re-adjusted proportional to the original distance from the sphere and drawn into the direction of the center.

Points from the interior are also moved under mapping gamma in order to compensate for additional points coming from the exterior. In this way, the whole unit hypercube is literally squeezed to the form of the feasible region without a too large increasing of the density at the boundary. Though, if the feasible region is very small compared with the hypercube, density at the boundary increases (depending on the choice of μ). On the other hand, the likelihood of an optimum being at the boundary increases likewise. So, this might be a desired effect.

After this procedure we have $\tilde{\Psi}_x$ which is the image of a point from $\mathcal{F}_{[0,1]}$ in terms of a modified weight vector \tilde{w}^{Γ_a} .

3.2.3 Finding an Approximate Pre-image

As a last step, we will have to find the pre-image of $\tilde{\Psi}_x$ in order to finally get the wanted mapping to $\mathcal{F}_{[0,1]}$. A major problem in determining the pre-image of a point from kernel space is that not every point from the span of Φ is the image of a mapped data point (Schölkopf et al., 1999). As we use a Gaussian kernel, none of our points from kernel space can be related to an exact pre-image except for trivial expansions with only one term (Kwok and Tsang, 2004). For this reason, we will look for an approximate pre-image whose image lies closest to the given image using an iterative procedure after (Mika et al., 1999). In our case (Gaussian kernel), we iterate x^* to find the point closest to the pre-image and define approximation Φ_ℓ^{-1} by equation

$$x_{n+1}^* = \frac{\sum_{i=1}^{\ell} (\tilde{w}_i^{\Gamma_a} e^{-\|s_i - x_n^*\|^2 / 2\sigma^2} s_i)}{\sum_{i=1}^{\ell} (\tilde{w}_i^{\Gamma_a} e^{-\|s_i - x_n^*\|^2 / 2\sigma^2})}. \quad (10)$$

As an initial guess for x_0^* we take the original point x and iterate it towards $\mathcal{F}_{[0,1]}$. As this procedure is sensitive to the choice of the starting point, it is important to have x as a fixed starting point in order to ensure determinism of the mapping. Empirically, x has showed up to be a useful guess.

Finally, we have achieved our goal to map an arbitrary point from $[0, 1]^d$ into the region of feasible solutions described merely by a given set of support vectors and associated weights: x_n^* is the sought after image under mapping γ of x that lies in $\mathcal{F}_{[0,1]}$.

We will now use this method for a decoder approach that transforms constrained optimization problems into unconstrained ones by automatically constructing mapping γ from a SVDD model of the feasible region that has been learned from a set of feasible example solutions.




Figure 1: 1(a): Sample from a artificial double banana shaped region. 1(b): Re-sampling the feasible region by mapping random points from $[0, 1]^2$. 1(c): marked optima of the Six-hump camel back objective function withing the used domain.

4 EXPERIMENTS

We present evaluation results of our decoder method, starting with several theoretical test cases and results from the smart grid load balancing problem.

4.1 General Test Cases

We started testing our method with several artificially constrained optimization problems and test functions. We consider optimization problems as described in section 1 and use the above described procedure as constraint handling technique, i.e. we transform problem Eq. 1 into an unconstrained optimization problem by applying mapping γ :

$$\text{optimize } f(\gamma(x)), \text{ s.t. } x \in [0, 1]^d. \quad (11)$$

Of course, the restriction to the unit hypercube still entails a box constraint, but as these are easily handled by almost all algorithm implementations, this is not a serious obstacle. In the case of evolutionary algorithms, the method can be easily applied by defining the neighborhood in $[0, 1]^d$ and search the whole unit hypercube, but evaluate a solution x at the position $\gamma(x)$. Note that mapping $\gamma(x)$ generates a feasible solution regardless of the choice of x . Therefore optimization might always start with an arbitrary (randomly chosen) $x \in [0, 1]^d$ without having to find a feasible start solution first.

For some first tests, we generated random samples from toy regions and used them as training sets for SVDD. The retrieved support vectors and weights are taken as a model for feasible region $\mathcal{F}_{[0,1]}$ as it would be determined by constraints in real world problems. Figure 1(a) shows an example with a 2-dimensional double banana set. With these models, we constructed our mapping γ . As a first test, a set of 1 million equally distributed points has been randomly picked from $[0, 1]^2$ and mapped. Figure 1(b) shows the result with mapped points depicted in green.

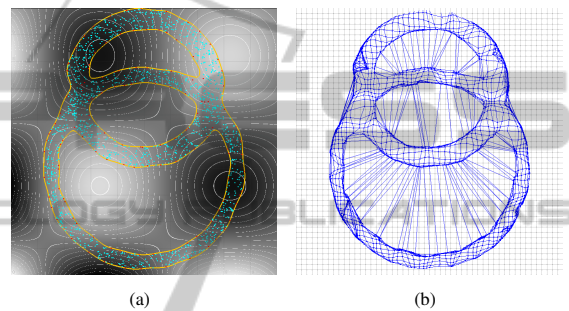


Figure 2: 2(a): Randomly generated double ring data set as representation of a second toy region. Again, the orange line denotes the learned boundary that encloses the feasible space. 2(b): Mapping a mesh with grid size 0.02 into the learned region of a double ring data set.

Next, we applied standard particle swarm optimization (PSO) (Kennedy and Eberhart, 1995) and standard artificial bee colony (ABC) optimization (Karaboga and Basturk, 2007) in order to find optima of several standard test objective functions. For this purpose, both algorithms have been equipped with mapping γ in order to evaluate solutions at the corresponding position inside $\mathcal{F}_{[0,1]}$, while the topology of the neighbourhood that both algorithms operate on is defined as the whole unit hypercube. Figure 1(c) shows an example of found optima for the above sketched setting (both succeeded equally good). In this case, the well known Six-hump camel back function

$$f_{CB}(x_1, x_2) = (4 - 2.1x_1^2 + \frac{x_1^4}{3})x_1^2 + x_1x_2 + (-4 + 4x_2^2)x_2^2 \quad (12)$$

has been used with the domain $-1.9 \leq x_1 \leq 1.9$, $-1.1 \leq x_2 \leq 1.1$ scaled to $[0, 1]^2$.

As a second test case, double ring data sets (Figure 2(a)) have been generated. The contour plot in the background shows as objective function the well

known Shubert function (Michalewicz, 1996). For the depicted configuration, different almost equally good local optima are situated near different distant positions at the boundary of the feasible region. Nevertheless, all algorithms equipped with mapping γ succeeded in finding optima inside (or at the boundary of) the feasible region.

As a next step, we compared how fast a solution converges with a mapped objective function. We compared the performance, i.e. the speed of convergence, with exterior penalty approaches. Such a constraint handling approach entails additional penalty values to solutions outside the feasible region. According to (Richardson et al., 1989), a penalty function that reflects the distances from the feasible region, is supposed to lead to better performance. Therefore, we have chosen the distance function of the SVDD (Eq. 3) as the penalty that attracts an outside solution to the feasible region. As we do not have any information on the original constraints, it is not in general possible to model penalties based on the number of or based on any individual constraint.




Figure 3: Speed of convergence for solving the scenario depicted in Figure 2 by a PSO. Top: PSO with mapping γ ; bottom: PSO with exterior penalty relative to the distance from the sphere center.

Figure 3 shows the result of 1 million test runs on a scenario with double rings and Shubert as objective. The PSO algorithm converges faster with mapping than with penalty. The whole swarm operates completely inside the feasible region from the beginning when using a mapped objective function while retaining normal swarm behaviour in $[0, 1]^d$. The same holds true for the ACO case.

Tables 1 to 3 show further results. Table 1 focusses on the population size of swarm based approaches. Table 2 shows further results (the lower the better) on various combinations of algorithms and

Table 1: PSO and ABC and their absolute fitnesses (lower is better) after n iterations for the Shubert function as test function and double banana as constraint region.

ALGORITHM	n	PENALTY	MAPPING
PSO	5	70.912 ± 89.714	-13.360 ± 0.256
	10	24.734 ± 68.949	-13.435 ± 0.278
	25	5.923 ± 47.904	-13.446 ± 0.28
	50	0.013 ± 41.368	-13.477 ± 0.286
ABC	5	67.189 ± 89.783	-13.310 ± 0.17
	10	22.128 ± 64.272	-13.367 ± 0.137
	25	-2.897 ± 35.965	-13.496 ± 0.175
	50	-10.236 ± 15.765	-13.596 ± 0.153

Table 2: Results for different objective functions and Algorithms. A double banana set has been used for all objectives. ABC: 87.98 and PSO: 83.22 percent were valid solutions for the penalty case; mapping: 100 percent.

OBJECTIVE	ALG.	PENALTY	MAPPING
SHUBERT	PSO 5	-1.63 ± 38.95	-13.3 ± 0.41
SHUBERT	ABC 5	-11.59 ± 8.35	-13.6 ± 0.35
SHUBERT	PSO 50	-13.99 ± 0.06	-13.87 ± 0.13
SHUBERT	ABC 50	-13.93 ± 0.08	-13.89 ± 0.01
BRANIN	PSO 5	135.39 ± 80.35	36.36 ± 0.91
BRANIN	ABC 5	46.91 ± 33.78	36.15 ± 0.19
ZAKHAROV	PSO 5	1.54 ± 14.9	0.39 ± 0.08
ZAKHAROV	ABC 5	0.51 ± 3.9	0.38 ± 0.0
SHEKELS FOXHOLES	PSO 5	12.68 ± 0.01	12.67 ± 0.01
SHEKELS FOXHOLES	ABC 5	12.67 ± 0.01	12.67 ± 0.0
BOHACHEVSKY 2	PSO 5	0.93 ± 11.21	0.26 ± 0.1
BOHACHEVSKY 2	ABC 5	0.4 ± 2.76	0.24 ± 0.01
BOHACHEVSKY 2	PSO 50	0.24 ± 0.0	0.24 ± 0.0
BOHACHEVSKY 2	ABC 50	0.24 ± 0.01	0.24 ± 0.0
HIMMELBLAU	PSO 5	179.84 ± 30.05	137.6 ± 1.96
HIMMELBLAU	ABC 5	145.08 ± 14.38	136.03 ± 1.1

Table 3: Results for different objective functions and Algorithms. This time, the double ring has been used for all objectives. ABC: 91.71 and PSO: 98.4 percent were valid for the penalty case; mapping 100 percent.

OBJECTIVE	ALG.	PENALTY	MAPPING
SHUBERT	PSO 5	-13.61 ± 9.02	-14.07 ± 0.75
SHUBERT	ABC 5	-14.14 ± 0.49	-14.25 ± 0.22
SHUBERT	PSO 50	-14.43 ± 0.0	-14.43 ± 0.01
SHUBERT	ABC 50	-14.43 ± 0.0	-14.42 ± 0.01
BRANIN	PSO 5	43.61 ± 40.24	33.14 ± 0.17
BRANIN	ABC 5	33.56 ± 1.73	33.12 ± 0.02
ZAKHAROV	PSO 5	0.18 ± 0.15	0.14 ± 0.01
ZAKHAROV	ABC 5	0.18 ± 0.08	0.14 ± 0.0
SHEKELS FOXHOLES	PSO 5	12.67 ± 0.01	12.67 ± 0.0
SHEKELS FOXHOLES	ABC 5	12.67 ± 0.0	12.67 ± 0.0
BOHACHEVSKY 2	PSO 5	0.43 ± 0.22	0.32 ± 0.05
BOHACHEVSKY 2	ABC 5	0.38 ± 0.11	0.28 ± 0.02
BOHACHEVSKY 2	PSO 50	0.27 ± 0.0	0.27 ± 0.0
BOHACHEVSKY 2	ABC 50	0.28 ± 0.01	0.27 ± 0.0
HIMMELBLAU	PSO 5	127.11 ± 19.04	121.84 ± 0.31
HIMMELBLAU	ABC 5	123.1 ± 1.69	121.74 ± 0.06

further objective functions for the case of the double banana dataset. Table 3 shows the respective results for the ring data set. Stated are respectively absolute achieved fitnesses, thus the ratio between mapping and penalty is to be compared. A major drawback of penalty approaches is the fact that it not always converges to a feasible solution. Whereas the mapping method always served feasible solutions, the penalty approach failed in up to 17 percent of the test runs.

Nevertheless, the inaccuracy inherent in the model from learning the region still remains an inaccuracy for mapping. But, the same holds true for all approaches that are based on such a surrogate model, including the penalty approach. Although, we made the observation that γ performs better at sharp edges than the decision boundary Eq. 3 (cf. Figure 1(b)).

Figure 2(b) shows the result of mapping a regular mesh from $[0, 1]^2$ onto the double rings. The mapped mesh shows how points from different parts of the feasible region become neighbours under the γ by bypassing the infeasible region inside the rings.




Figure 4: A typical result for a higher dimensional test case; dashed: penalty approach.

Figure 4 shows some results for test runs on an 8-dimensional problem with a stretched ellipse as feasible region and Himmelblau (Himmelblau, 1972) as objective function. We compared artificial bee algorithms with 5 individuals for the mapping case and 200 individuals for the penalty case. Nevertheless, the mapping approach performs better and some penalty runs still converged to a infeasible solution, showing the superiority of the mapping approach.

4.2 The Smart Grid Use Case

Next, we applied our method to the following real world problem from the smart grid domain: An individual schedule has to be determined for each member of a pool of micro-co-generation (CHP) plants such that the aggregated electric load schedule of all plants resembles a given (probably demanded by market) target schedule in an optimal way. For the sake of simplicity, we will consider optimality as a close as possible adaption of the aggregated (sum of individual loads) schedule to the requested one. Optimality usually refers to additional local (individual cost) as well as to global (e.g. environmental impact) objectives. When determining an optimal partition of the schedule for load distribution, exactly one alternative schedule is taken from each generators search space of individual operable schedules in order to assemble the desired aggregate schedule.

Therefore, the optimization problem is: finding any combination of schedules (one from each DER with \mathcal{X}_i as the set of possible choices) that resembles

the target schedule l_T as close as possible, i.e. minimize the Euclidean distance between aggregated and target schedule:

$$\left\| \sum_i x_i - l_T \right\| \rightarrow \min, \text{ s.t. } x_i \in \mathcal{X}_i. \quad (13)$$

Of course, each generator has individual constraints such as time varying buffer charging, power ranges, minimum ON/ OFF times, etc. Thus, we simulated individual plants. For our simulations, we used simulation models of modulating CHP-plants (combined heat and power generator capable of varying the power level) with the following specification: Min. / max. electrical power: 1.3 / 4.7 kW, min. / max. thermal power: 4 / 12.5 kW; after shutting down, the device has to stay off for at least 2 hours.




Figure 5: Relationship of electrical and thermal power for an EcoPower CHP; modified after (Thomas, 2007).

The relationship between electrical (active) power and thermal power was modeled after Figure 5. In order to gain enough degrees of freedom for varying active power, each CHP is equipped with an 800 litre thermal buffer store. Thermal energy consumption is modeled and simulated by a model of a detached house with its several heat losses (heater is supposed to keep the indoor temperature on a constant level) and randomized warm water drawing for gaining more diversity among the devices.

For each simulated household, we implemented an agent capable of simulating the CHP (and surroundings and auxiliary devices) on a meso-scale level with energy flows among different model parts but no technical details. All simulations have so far been done with a time resolution of 15 minutes for different forecast horizons. Although, our method is indifferent about any such time constraints. We have run several test series with each CHP randomly initialized with different buffer charging levels, temperatures and water drawing profiles. The feasible spaces of individual CHP had been encoded with the SVDD approach. These support vector models have then been used for the search for optimal schedules: with a penalty approach on the one hand and with the proposed mapping on the other. Figure 6 shows a typical result. We used a co-variance matrix adaption evolution strategy (CMA-ES) approach (Hansen, 2006) for finding combinations of schedules that best resemble the dashed target schedule in the top chart (Figures

6(a) and 6(b)). Both seem to have equally good results, but, looking at individual loads (in middle) and the temperatures (bottom) reveals that the penalty approach gets easily stuck at an (at least partly) infeasible solution whereas the mapping approach succeeds with feasible solutions. This effect amplifies with the number of plants and therefore with the number of used penalties. Moreover, the mapping approach most times converges faster as Figure 6(c) shows for this specific example.

Considering the complexity, additional computational costs are entailed on solution evaluation. Step 2 growing quadratically with the number of support vectors ℓ is decisive together with the number of iterations necessary for finding the pre-image in step 3. Empirically, during our experiments, we observed for instance a mean number of 6.75 ± 0.3 for the case of the 2-dimensional double banana and 36.3 ± 26.4 for the case of a stretched 8-dimensional ellipse in order to reach convergence with 10^{-8} accuracy. Additionally, this number reduces in the course of optimization as soon as the evolution approaches feasible space. Otherwise, fewer function evaluations are necessary with our decoder approach, because we never evaluate infeasible solutions and we do not have to check feasibility during optimization. Both effects put into perspective the computational costs.

5 CONCLUSIONS

Many real world optimization problems face the effect of constraints that restrict the search space to an arbitrary (nonlinear) shaped possibly disjoint region that contains the feasible solutions. Conventional constraint handling techniques often require the set of constraints to be a priori known and are hardly applicable for black-box models of feasible regions. Although penalties may be used with such models, the task of correctly tuning the objective with these additional losses stays an error prone job due to the unknown nature of the original constraints that are no longer known at optimization time.

We proposed a new constraint handling technique for support vector modeled search spaces and demonstrated its applicability and usefulness with the help of theoretical test problems as well as for a real world optimization problem taken from the smart grid domain. The major benefit of this approach is the universal applicability for problem transformation, solution repair and standardized integration in arbitrary evolutionary algorithms by constructing a modified objective function and treating the whole unconstrained domain as valid for search. So far, we have restricted

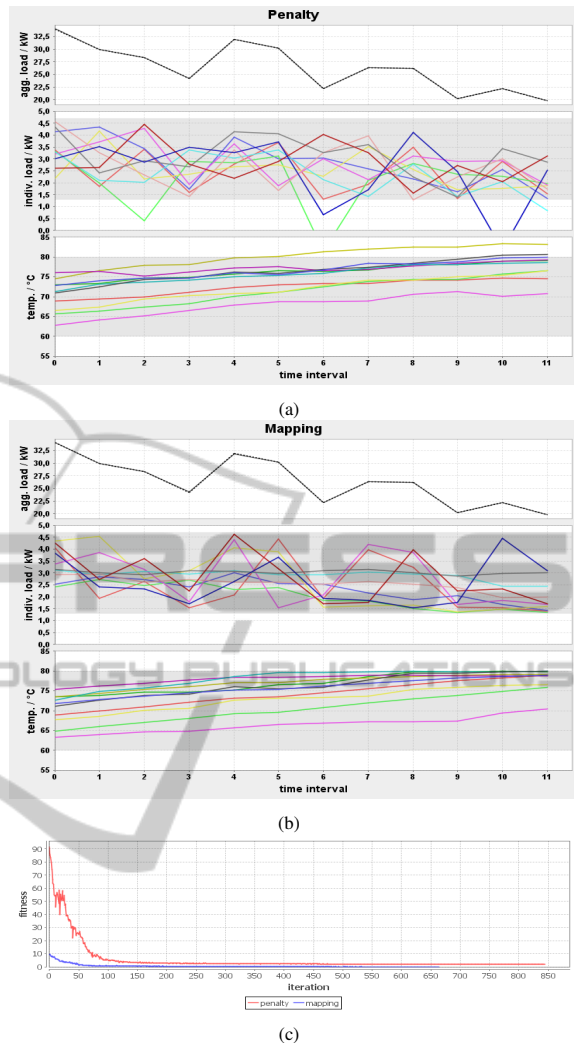


Figure 6: Results from a smart grid load balancing scenario with a standard penalty constraint-handling technique on the left (6(a)) and the mapping based approach on the right (6(b)); 6(c) compares the speed of convergence.

ourselves to problems scaled to $[0, 1]^d$. Further tests will show whether this limitation should be kept or whether arbitrary domains perform equally good.

ACKNOWLEDGEMENTS

The Lower Saxony research network 'Smart Nord' acknowledges the support of the Lower Saxony Ministry of Science and Culture through the Niedersächsisches Vorab grant programme (grant ZN 2764).

REFERENCES

- Ben-Hur, A., Siegelmann, H. T., Horn, D., and Vapnik, V. (2001). Support vector clustering. *Journal of Machine Learning Research*, 2:125–137.
- Blank, M., Gerwinn, S., Krause, O., and Lehnhoff, S. (2011). Support vector machines for an efficient representation of voltage band constraints. In *Innovative Smart Grid Technologies*. IEEE PES.
- Bremer, J., Rapp, B., and Sonnenschein, M. (2010). Support vector based encoding of distributed energy resources' feasible load spaces. In *IEEE PES Conference on Innovative Smart Grid Technologies Europe*, Chalmers Lindholmen, Gothenburg, Sweden.
- Bremer, J., Rapp, B., and Sonnenschein, M. (2011). Encoding distributed search spaces for virtual power plants. In *IEEE Symposium Series in Computational Intelligence 2011 (SSCI 2011)*, Paris, France.
- Coello Coello, C. A. (2002). Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191(11-12):1245–1287.
- Guan, X., Zhai, Q., and Papalexopoulos, A. (2003). Optimization based methods for unit commitment: Lagrangian relaxation versus general mixed integer programming. volume 2, page 1100 Vol. 2.
- Hansen, N. (2006). The CMA evolution strategy: a comparing review. In Lozano, J., Larranaga, P., Inza, I., and Bengoetxea, E., editors, *Towards a new evolutionary computation. Advances on estimation of distribution algorithms*, pages 75–102. Springer.
- Himmelblau, D. (1972). *Applied nonlinear programming*. McGraw-Hill.
- Juszcak, P., Tax, D., and Duin, R. P. W. (2002). Feature scaling in support vector data description. In Depretere, E., Belloum, A., Heijnsdijk, J., and van der Stappen, F., editors, *Proc. ASCI 2002, 8th Annual Conf. of the Advanced School for Computing and Imaging*, pages 95–102.
- Kamper, A. and Esser, A. (2009). Strategies for decentralised balancing power. In A. Lewis, S. Mostaghim, M. R., editor, *Biologically-inspired Optimisation Methods: Parallel Algorithms, Systems and Applications*, number 210 in Studies in Computational Intelligence, pages 261–289. Springer, Berlin, Heidelberg.
- Karaboga, D. and Basturk, B. (2007). A powerful and efficient algorithm for numerical function optimization: artificial bee colony (ABC) algorithm. *Journal of Global Optimization*, 39(3):459–471.
- Kennedy, J. and Eberhart, R. (1995). Particle swarm optimization. In *Neural Networks, 1995. Proceedings., IEEE International Conference on*, volume 4, pages 1942–1948 vol.4. IEEE.
- Kim, D. G. (1998). Riemann mapping based constraint handling for evolutionary search. In *SAC*, pages 379–385.
- Kok, K., Derzsi, Z., Gordijn, J., Hommelberg, M., Warmer, C., Kamphuis, R., and Akkermans, H. (2008). Agent-based electricity balancing with distributed energy resources, a multiperspective case study. *Hawaii International Conference on System Sciences*, 0:173.
- Koziel, S. and Michalewicz, Z. (1999). Evolutionary algorithms, homomorphous mappings, and constrained parameter optimization. *Evol. Comput.*, 7:19–44.
- Kramer, O. (2010). A review of constraint-handling techniques for evolution strategies. *Appl. Comp. Intell. Soft Comput.*, 2010:3:1–3:19.
- Kwok, J. and Tsang, I. (2004). The pre-image problem in kernel methods. *Neural Networks, IEEE Transactions on*, 15(6):1517–1525.
- Liepins, G. E. and Vose, M. D. (1990). Representational issues in genetic optimization. *Journal of Experimental and Theoretical Artificial Intelligence*, 2.
- Michalewicz, Z. (1996). *Genetic algorithms + data structures = evolution programs (3rd ed.)*. Springer-Verlag, London, UK.
- Michalewicz, Z. and Schoenauer, M. (1996). Evolutionary algorithms for constrained parameter optimization problems. *Evol. Comput.*, 4:1–32.
- Mihailescu, R.-C., Vasirani, M., and Ossowski, S. (2011). Dynamic coalition adaptation for efficient agent-based virtual power plants. In *Proceedings of the 9th German conference on Multiagent system technologies, MATES'11*, pages 101–112, Berlin, Heidelberg. Springer-Verlag.
- Mika, S., Schölkopf, B., Smola, A., Müller, K. R., Scholz, M., and Rätsch, G. (1999). Kernel PCA and denoising in feature spaces. In *Proceedings of the 1998 conference on Advances in neural information processing systems II*, pages 536–542, Cambridge, MA, USA. MIT Press.
- Pereira, J., Viana, A., Lucus, B., and Matos, M. (2008). A meta-heuristic approach to the unit commitment problem under network constraints. *International Journal of Energy Sector Management*, 2(3):449–467.
- Ramchurn, S. D., Vytelingum, P., Rogers, A., and Jennings, N. R. (2011). Agent-based control for decentralised demand side management in the smart grid. In *AA-MAS*, pages 5–12.
- Richardson, J. T., Palmer, M. R., Liepins, G. E., and Hilliard, M. R. (1989). Some guidelines for genetic algorithms with penalty functions. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 191–197, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Schölkopf, B., Mika, S., Burges, C., Knirsch, P., Müller, K.-R., Rätsch, G., and Smola, A. (1999). Input space vs. feature space in kernel-based methods. *IEEE Transactions on Neural Networks*, 10(5):1000–1017.
- Tax, D. M. J. and Duin, R. P. W. (1999). Data domain description using support vectors. In *ESANN*, pages 251–256.
- Thomas, B. (2007). *Mini-Blockheizkraftwerke: Grundlagen, Gerätetechnik, Betriebsdaten*. Vogel Buchverlag.
- Tröschel, M. and Appelrath, H.-J. (2009). Towards reactive scheduling for large-scale virtual power plants. In Braubach, L., van der Hoek, W., Petta, P., and Pokahr, A., editors, *MATES*, volume 5774 of *Lecture Notes in Computer Science*, pages 141–152. Springer.