

Regression with Linear Factored Functions

Wendelin Böhmer and Klaus Obermayer

Neural Information Processing Group, Technische Universität Berlin,
Sekt. MAR5-6, Marchstr. 23, D-10587 Berlin, Germany
{wendelin,oby}@ni.tu-berlin.de, <http://www.ni.tu-berlin.de>

Abstract. Many applications that use empirically estimated functions face a *curse of dimensionality*, because integrals over most function classes must be approximated by sampling. This paper introduces a novel *regression*-algorithm that learns *linear factored functions* (LFF). This class of functions has structural properties that allow to analytically solve certain integrals and to calculate point-wise products. Applications like *belief propagation* and *reinforcement learning* can exploit these properties to break the curse and speed up computation. We derive a regularized greedy optimization scheme, that learns factored basis functions during training. The novel regression algorithm performs competitively to *Gaussian processes* on benchmark tasks, and the learned LFF functions are with 4-9 factored basis functions on average very compact.

Keywords: regression, factored functions, curse of dimensionality

1 Introduction

This paper introduces a novel regression-algorithm, which performs competitive to *Gaussian processes*, but yields *linear factored functions* (LFF). These have outstanding properties like analytical *point-wise products* and *marginalization*.

Regression is a well known problem, which can be solved by many non-linear architectures like *kernel methods* (Shawe-Taylor and Cristianini, 2004) or *neural networks* (Haykin, 1998). While these perform well, the estimated functions often suffer a *curse of dimensionality* in later applications. For example, computing an integral over a neural network or kernel function requires to sample the entire input space. Applications like *belief propagation* (Pearl, 1988) and *reinforcement learning* (Kaelbling et al., 1996), on the other hand, face large input spaces and require therefore efficient computations. We propose LFF for this purpose and showcase its properties in comparison to kernel functions.

1.1 Kernel regression

In the last 20 years, kernel methods like *support vector machines* (SVM, Vapnik, 1995; Boser et al., 1992) have become a de facto standard in various practical applications. This is mainly due to a sparse representation of the learned classifiers with so called *support vectors* (SV). The most popular kernel method for regression,

Gaussian processes (GP, see Rasmussen and Williams, 2006; Bishop, 2006), on the other hand, requires as many SV as training samples. Sparse versions of GP aim thus for a small subset of SV. Some select this set based on constraints similar to SVM (Vapnik, 1995; Tipping, 2001), while others try to conserve the spanned linear function space (*sparse GP*, Csató and Opper, 2002; Rasmussen and Williams, 2006). There exist also attempts to construct new SV by averaging similar training samples (e.g. Wang et al., 2012).

Well chosen SV for regression are usually not sparsely concentrated on a decision boundary as they are for SVM. In fact, many practical applications report that they are distributed uniformly in the input space (e.g. in Böhmer et al., 2013). Regression tasks restricted to a small region of the input space may tolerate this, but some applications require predictions everywhere. For example, the *value function* in reinforcement learning must be generalized to each state. The number of SV required to *represent* this function equally well in each state grows exponentially in the number of input-space dimensions, leading to Bellman’s famous curse of dimensionality (Bellman, 1957).

Kernel methods derive their effectiveness from linear optimization in a non-linear *Hilbert space* of functions. Kernel-functions parameterized by SV are the non-linear *basis functions* in this space. Due to the functional form of the kernel, this can be a very ineffective way to select basis functions. Even in relatively small input spaces, it often takes hundreds or thousands SV to approximate a function sufficiently. To alleviate the problem, one can construct complex kernels out of simple prototypes (see a recent review in Gönen and Alpaydm, 2011).

1.2 Factored basis functions

Diverging from all above arguments, this article proposes a more radical approach: to construct the non-linear basis functions directly during training, without the detour over kernel functions and support vectors. This poses two main challenges: to select a *suitable functions space* and to *regularize the optimization* properly. The former is critical, as a small set of basis functions must be able to approximate any target function, but should also be easy to compute in practice.

We propose *factored functions* $\psi_i = \prod_k \psi_i^k \in \mathcal{F}$ as basis functions for regression, and call the linear combination of m of those bases a *linear factored function* $f \in \mathcal{F}^m$ (LFF, Section 3). For example, *generalized linear models* (Nelder and Wedderburn, 1972) and *multivariate adaptive regression splines* (MARS, Friedman, 1991) are both LFF. Computation remains feasible by using *hinge functions* $\psi_i^k(x_k) = \max(0, x_k - c)$ and restricting the scope of each factored function ψ_i . In contrast, we assume the general case without restrictions to functions or scope.

Due to their structure, LFF can solve certain integrals analytically and allow very efficient computation of point-wise products and marginalization. We show that our LFF are universal function approximators and derive an appropriate *regularization* term. This regularization promotes smoothness, but also retains a high degree of variability in densely sampled regions by linking smoothness to uncertainty about the sampling distribution. Finally, we derive a novel regression algorithm for LFF based on a greedy optimization scheme.

Functions learned by this algorithm (Algorithm 1, see pages 7 and 16) are very compact (between 3 and 12 bases on standard benchmarks) and perform competitive with Gaussian processes (Section 4). The paper finishes with a discussion of the computational possibilities of LFF in potential areas of application and possible extensions to *sparse regression* with LFF (Section 5).

2 Regression

Let $\{\mathbf{x}_t \in \mathcal{X}\}_{t=1}^n$ be a set of n *input samples*, i.i.d. drawn from an input set $\mathcal{X} \subset \mathbb{R}^d$. Each so called “training sample” is *labeled* with a real number $\{y_t \in \mathbb{R}\}_{t=1}^n$. *Regression* aims to find a function $f : \mathcal{X} \rightarrow \mathbb{R}$, that predicts the labels to all (previously unseen) test samples as well as possible. Labels may be afflicted by *noise* and f must thus approximate the mean label of each sample, i.e., the function $\mu : \mathcal{X} \rightarrow \mathbb{R}$. It is important to notice that *conceptually* the noise is introduced by two (non observable) sources: noisy labels y_t and noisy samples \mathbf{x}_t . The latter will play an important role for regularization. We define the conditional distribution χ of observable samples $\mathbf{x} \in \mathcal{X}$ given the non-observable “true” samples $\mathbf{z} \in \mathcal{X}$, which are drawn by a distribution ξ . In the limit of infinite samples, the *least squares* cost-function $\mathcal{C}[f]_{\chi, \mu}$ can thus be written as

$$\lim_{n \rightarrow \infty} \inf_f \frac{1}{n} \sum_{t=1}^n \left(f(\mathbf{x}_t) - y_t \right)^2 = \inf_f \iint \xi(d\mathbf{z}) \chi(d\mathbf{x}|\mathbf{z}) \left(f(\mathbf{x}) - \mu(\mathbf{z}) \right)^2. \quad (1)$$

The cost function \mathcal{C} can never be computed *exactly*, but *approximated* using the training samples¹ and assumptions about the unknown noise distribution χ .

3 Linear factored functions

Any non-linear function can be expressed as a linear function $f(\mathbf{x}) = \mathbf{a}^\top \boldsymbol{\psi}(\mathbf{x})$, $\forall \mathbf{x} \in \mathcal{X}$, with m non-linear basis functions $\psi_i : \mathcal{X} \rightarrow \mathbb{R}$, $\forall i \in \{1 \dots, m\}$. In this section we will define *linear factored functions* (LFF), that have *factored basis functions* $\psi_i(\mathbf{x}) := \psi_i^1(x_1) \cdot \dots \cdot \psi_i^d(x_d) \in \mathcal{F}$, a regularization method for this function class and an algorithm for regression with LFF.

3.1 Function class

We define the class of linear factored functions $f \in \mathcal{F}^m$ as a linear combination (with linear parameters $\mathbf{a} \in \mathbb{R}^m$) of m factored basis functions $\psi_i : \mathcal{X} \rightarrow \mathbb{R}$ (with parameters $\{\mathbf{B}^k \in \mathbb{R}^{m_k \times m_k}\}_{k=1}^d$):

$$f(\mathbf{x}) := \mathbf{a}^\top \boldsymbol{\psi}(\mathbf{x}) := \mathbf{a}^\top \left[\prod_{k=1}^d \boldsymbol{\psi}^k(x_k) \right] := \sum_{i=1}^m a_i \prod_{k=1}^d \sum_{j=1}^{m_k} B_{ji}^k \phi_j^k(x_k). \quad (2)$$

¹ The distribution ξ of “true” samples \mathbf{z} can *not* be observed. We approximate in the following ξ with the training-sample distribution. This may be justified if the sample-noise χ is comparatively small. Although not strictly rigorous, the presented formalism helps to put the regularization derived in Proposition 2 into perspective.

LFF are formally defined in Appendix A. In short, a basis function ψ_i is the *point-wise product* of one-dimensional functions ψ_i^k in each input dimension k . These are themselves constructed as linear functions of a corresponding one-dimensional base $\{\phi_j^k\}_{j=1}^{m_k}$ over that dimension and ideally can approximate arbitrary functions². Although each factored function ψ_i is very restricted, a linear combination of them can be very powerful:

Corollary 1 *Let \mathcal{X}_k be a bounded continuous set and ϕ_j^k the j 'th Fourier base over \mathcal{X}_k . In the limit of $m_k \rightarrow \infty, \forall k \in \{1, \dots, d\}$, holds $\mathcal{F}^\infty = L^2(\mathcal{X}, \vartheta)$.*

Strictly this holds in the limit of infinitely many basis functions ψ_i , but we will show empirically that there exist close approximations with a small number m of factored functions. One can make similar statements for other bases $\{\phi_j^k\}_{j=1}^\infty$. For example, for Gaussian kernels one can show that the space \mathcal{F}^∞ is in the limit equivalent to the corresponding *reproducing kernel Hilbert space* \mathcal{H} .

LFF offer some structural advantages over other universal function approximation classes like neural networks or reproducing kernel Hilbert spaces. Firstly, the *inner product* of two LFF in $L^2(\mathcal{X}, \vartheta)$ can be computed as products of one-dimensional integrals. For some bases³, these integrals can be calculated analytically without any sampling. This could in principle break the curse of dimensionality for algorithms that have to approximate these inner products numerically. For example, input variables can be *marginalized* (integrated) out analytically (Equation 9 on Page 12). Secondly, the *point-wise product* of two LFF is a LFF as well⁴ (Equation 10 on Page 12). See Appendix A for details. These properties are very useful, for example in *belief propagation* (Pearl, 1988) and *factored reinforcement learning* (Böhmer and Obermayer, 2013).

3.2 Constraints

LFF have some degrees of freedom that can impede optimization. For example, the norm of $\psi_i \in \mathcal{F}$ does not influence function $f \in \mathcal{F}^m$, as the corresponding linear coefficients a_i can be scaled accordingly. We can therefore introduce the *constraints* $\|\psi_i\|_\vartheta = 1, \forall i$, without restriction to the function class. The factorization of inner products (see Appendix A on Page 12) allows us furthermore to rewrite the constraints as $\|\psi_i\|_\vartheta = \prod_k \|\psi_i^k\|_{\vartheta^k} = 1$. This holds as long as the product is one, which exposes another unnecessary degree of freedom. To finally make the solution unique (up to permutation), we define the constraints as $\|\psi_i^k\|_{\vartheta^k} = 1, \forall k, \forall i$. Minimizing some $\mathcal{C}[f]$ w.r.t. $f \in \mathcal{F}^m$ is thus equivalent to

$$\inf_{f \in \mathcal{F}^m} \mathcal{C}[f] \quad \text{s.t.} \quad \|\psi_i^k\|_{\vartheta^k} = 1, \quad \forall k \in \{1, \dots, d\}, \quad \forall i \in \{1, \dots, m\}. \quad (3)$$

² Examples are Fourier bases, Gaussian kernels or hinge-functions as in MARS.

³ E.g. Fourier bases for continuous, and Kronecker-delta bases for discrete variables.

⁴ One can use the trigonometric product-to-sum identities for Fourier bases or the Kronecker delta for discrete bases to construct LFF from a point-wise product without changing the underlying basis $\{\{\phi_i^k\}_{i=1}^{m_k}\}_{k=1}^d$.

The *cost function* $\mathcal{C}[f|\chi, \mu]$ of Equation 1 with the constraints in Equation 3 is equivalent to *ordinary least squares* (OLS) w.r.t. linear parameters $\mathbf{a} \in \mathbb{R}^m$. However, the optimization problem is *not* convex w.r.t. the parameter space $\{\mathbf{B}^k \in \mathbb{R}^{m_k \times m}\}_{k=1}^d$, due to the nonlinearity of products.

Instead of tackling the global optimization problem induced by Equation 3, we propose a *greedy* approximation algorithm. Here we optimize at iteration i one linear basis function $\psi_i =: g =: \prod_k g^k \in \mathcal{F}$, with $g^k(x_k) =: \mathbf{b}^{k\top} \phi^k(x_k)$, at a time, to fit the residual $\mu - f$ between the true *mean label* function $\mu \in L^2(\mathcal{X}, \vartheta)$ and the current regression estimate $f \in \mathcal{F}^{i-1}$, based on all $i - 1$ previously constructed factored basis functions $\{\psi_i\}_{i=1}^{i-1}$:

$$\inf_{g \in \mathcal{F}} \mathcal{C}[f + g|\chi, \mu] \quad \text{s.t.} \quad \|g^k\|_{\vartheta^k} = 1, \quad \forall k \in \{1, \dots, d\}. \quad (4)$$

3.3 Regularization

Regression with any powerful function class requires regularization to avoid overfitting. Examples are *weight decay* for neural networks (Haykin, 1998) or parameterized *priors* for Gaussian processes. It is, however, not immediately obvious how to regularize the parameters of a LFF and we will derive a regularization term from a Taylor approximation of the cost function in Equation 1.

We aim to enforce smooth functions, especially in those regions our knowledge is limited due to a lack of training samples. This *uncertainty* can be expressed as the *Radon-Nikodym derivative*⁵ $\frac{\vartheta}{\xi} : \mathcal{X} \rightarrow [0, \infty)$ of our factored measure ϑ (see Appendix A) w.r.t. the sampling distribution ξ . Figure 1 demonstrates at the example of a uniform distribution ϑ how $\frac{\vartheta}{\xi}$ reflects our empirical knowledge of the input space \mathcal{X} .

We use this uncertainty to modulate the *sample noise distribution* χ in Equation 1. This means that frequently sampled regions of \mathcal{X} shall yield low, while scarcely sampled regions shall yield high variance. Formally, we assume $\chi(d\mathbf{x}|\mathbf{z})$ to be a Gaussian probability measure over \mathcal{X} with mean \mathbf{z} and a *covariance matrix* $\Sigma \in \mathbb{R}^{d \times d}$, scaled by the local uncertainty in \mathbf{z} (modeled as $\frac{\vartheta}{\xi}(\mathbf{z})$):

$$\int \chi(d\mathbf{x}|\mathbf{z})(\mathbf{x} - \mathbf{z}) = \mathbf{0}, \quad \int \chi(d\mathbf{x}|\mathbf{z})(\mathbf{x} - \mathbf{z})(\mathbf{x} - \mathbf{z})^\top = \frac{\vartheta}{\xi}(\mathbf{z}) \cdot \Sigma, \quad \forall \mathbf{z} \in \mathcal{X}. \quad (5)$$

⁵ Technically we have to assume that ϑ is *absolutely continuous* in respect to ξ . For “well-behaving” distributions ϑ , like the uniform or Gaussian distributions we discuss in Appendix A, this is equivalent to the assumption that in the limit of infinite samples, each sample $\mathbf{z} \in \mathcal{X}$ will *eventually* be drawn by ξ .

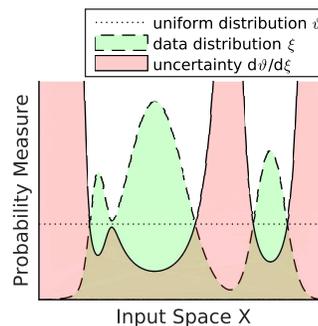


Fig. 1. We interpret the Radon-Nikodym derivative $\frac{d\vartheta}{d\xi}$ as *uncertainty measure* for our knowledge of \mathcal{X} . Regularization enforces smoothness in uncertain regions.

In the following we assume without loss of generality⁶ the matrix Σ to be diagonal, with the diagonal elements called $\sigma_k^2 := \Sigma_{kk}$.

Proposition 2 *Under the assumptions of Equation 5 and a diagonal covariance matrix Σ , the first order Taylor approximation of the cost \mathcal{C} in Equation 4 is*

$$\tilde{\mathcal{C}}[g] := \underbrace{\|g - (\mu - f)\|_{\xi}^2}_{\text{sample-noise free cost}} + \sum_{k=1}^d \sigma_k^2 \underbrace{\left\| \frac{\partial}{\partial x_k} g + \frac{\partial}{\partial x_k} f \right\|_{\vartheta}^2}_{\text{smoothness in dimension } k}. \quad (6)$$

Proof: see Appendix C on Page 14. \square

Note that the approximated cost $\tilde{\mathcal{C}}[g]$ consists of the sample-noise free cost (measured w.r.t. training distribution ξ) and d regularization terms. Each term prefers functions that are smooth⁷ in one input dimension. This enforces smoothness everywhere, but allows exceptions where enough data is available. To avoid a cluttered notation, in the following we will use the symbol $\nabla_k f := \frac{\partial}{\partial x_k} f$.

3.4 Optimization

Another advantage of cost function $\tilde{\mathcal{C}}[g]$ is that one can optimize one factor function g^k of $g(\mathbf{x}) = g^1(x_1) \cdot \dots \cdot g^d(x_d) \in \mathcal{F}$ at a time, instead of time consuming *gradient descend* over the entire parameter space of g . To be more precise:

Proposition 3 *If all but one factor function g^k are considered constant, Equation 6 has an analytical solution. If $\{\phi_j^k\}_{j=1}^{m_k}$ is a Fourier base, $\sigma_k^2 > 0$ and $\vartheta \ll \xi$, then the solution is also unique.*

Proof: see Appendix C on Page 14. \square

One can give similar guarantees for other bases, e.g. Gaussian kernels. Note that Proposition 3 does *not* state that the optimization problem has a unique solution in \mathcal{F} . Formal convergence statements are not trivial and empirically the parameters of g do not converge, but evolve around orbits of equal cost instead. However, since the optimization of *any* g^k cannot increase the cost, any sequence of improvements will converge to (and stay in) a *local minimum*. This implies a *nested* optimization approach, that is formulated in Algorithm 1 on Page 16:

- An *inner loop* that optimizes one factored basis function $g(\mathbf{x}) = g^1(x_1) \cdot \dots \cdot g^d(x_d)$ by selecting an input dimension k in each iteration and solve Equation 6 for the corresponding g^k . A detailed derivation of the optimization steps of the inner loop is given in Appendix B on Page 13. The choice of k influences

⁶ Non-diagonal covariance matrices Σ can be cast in this framework by projecting the input samples into the eigenspace of Σ (thus diagonalizing the input) and use the corresponding eigenvalues λ_k instead of the regularization parameters σ_k^2 's.

⁷ Each regularization term is measured w.r.t. the factored distribution ϑ . We also tested the algorithm without consideration of ‘‘uncertainty’’ $\frac{\vartheta}{\xi}$, i.e., by measuring each term w.r.t. ξ . As a result, regions outside the hypercube containing the training set were no longer regularized and predicted arbitrary (often extreme) values.

Algorithm 1 (abstract) – a detailed version can be found on Page 16

```

while new factored basis function can improve solution do
  initialize new basis function  $g$  as constant function
  while optimization improves cost in Equation 6 do
    for random input dimension  $k$  do
      calculate optimal solution for  $g^k$  without changing  $g^l, \forall l \neq k$ 
    end for
  end while // new basis function  $g$  has converged
  add  $g$  to set of factored basis functions and solve OLS
end while // regression has converged

```

the solution in a non-trivial way and further research is needed to build up a rationale for any meaningful decision. For the purpose of this paper, we assume k to be chosen randomly by permuting the order of updates.

The *computational complexity* of the inner loop is $\mathcal{O}(m_k^2 n + d^2 m_k m)$. Memory complexity is $\mathcal{O}(d m_k m)$, or $\mathcal{O}(d m_k n)$ with the optional cache speedup of Algorithm 1. The loop is repeated for random k until the cost-improvements of all dimensions k fall below some small ϵ .

- After convergence of the inner loop in (outer) iteration \hat{i} , the new basis function is $\psi_{\hat{i}} := g$. As the basis has changed, the linear parameters $\mathbf{a} \in \mathbb{R}^{\hat{i}}$ have to be readjusted by solving the ordinary least squares problem

$$\mathbf{a} = (\Psi \Psi^\top)^{-1} \Psi \mathbf{y}, \text{ with } \Psi_{it} := \psi_i(\mathbf{x}_t), \forall i \in \{1, \dots, \hat{i}\}, \forall t \in \{1, \dots, n\}.$$

We propose to stop the approximation when the newly found basis function $\psi_{\hat{i}}$ is no longer *linearly independent* of the current basis $\{\psi_i\}_{i=1}^{\hat{i}-1}$. This can for example be tested by comparing the *determinant* $\det(\frac{1}{n} \Psi \Psi^\top) < \epsilon$, for some very small ϵ .

4 Empirical evaluation

In this section we will evaluate the novel LFF regression Algorithm 1, printed in detail on Page 16. We will analyze its properties on low dimensional toy-data, and compare its performance with sparse and traditional Gaussian processes (GP, see Rasmussen and Williams, 2006; Bishop, 2006).

4.1 Demonstration

To showcase the novel Algorithm 1, we tested it on an artificial two-dimensional regression toy-data set. The $n = 1000$ training samples were drawn from a noisy spiral and labeled with a sinus. The variance of the Gaussian sample-noise grew with the spiral as well:

$$\mathbf{x}_t = 6 \frac{t}{n} \begin{bmatrix} \cos\left(6 \frac{t}{n} \pi\right) \\ \sin\left(6 \frac{t}{n} \pi\right) \end{bmatrix} + \mathcal{N}\left(\mathbf{0}, \frac{t^2}{4n^2} \mathbf{I}\right), \quad y_t = \sin\left(4 \frac{t}{n} \pi\right), \quad \forall t \in \{1, \dots, n\}. \quad (7)$$

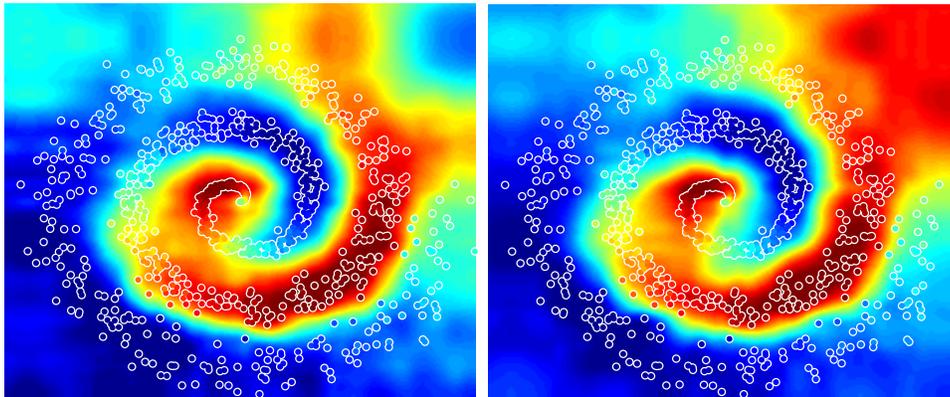


Fig. 2. Two LFF functions learned from the same 1000 training samples (white circles). The color inside a circle represents the training label. Outside the circles, the color represents the prediction of the LFF function. The differences between both functions are rooted in the randomized order in which the factor functions g^k are updated. However, the similarity of the sampled region indicates that poor initial choices can be compensated by subsequently constructed basis functions.

Figure 2 shows one training set plotted over two learned⁸ functions $f \in \mathcal{F}^m$ with $m = 21$ and $m = 24$ factored basis functions, respectively. Regularization constants were in both cases $\sigma_k^2 = 0.0005, \forall k$. The differences between the functions stem from the randomized order in which the factor functions g^k are updated. Note that the sampled regions have similar predictions. Regions with strong differences, for example the upper right corner, are never seen during training.

In all our experiments, Algorithm 1 always converged. Runtime was mainly influenced by the input dimensionality ($\mathcal{O}(d^2)$), the number of training samples ($\mathcal{O}(n)$) and the eventual number of basis functions ($\mathcal{O}(m)$). The latter was strongly correlated with approximation quality, i.e., bad approximations converged fast. Cross-validation was therefore able to find good parameters efficiently and the resulting LFF were always very similar near the training data.

4.2 Evaluation

We compared the regression performance of LFF and GP with cross-validation on five regression benchmarks from the *UCI Machine Learning Repository*⁹:

- The *concrete compressive strength* data set (*concrete*, Yeh, 1998) consists of $n = 1030$ samples with $d = 8$ dimensions describing various concrete mixture-

⁸ Here (and in the rest of the paper), each variable was encoded with 50 Fourier cosine bases. We tested other sizes as well. Few cosine bases result effectively in a low-pass filtered function, whereas every experiment with more than 20 or 30 bases behaved very similar. We tested up to $m_k = 1000$ bases and did not experience over-fitting.

⁹ <https://archive.ics.uci.edu/ml/index.html>

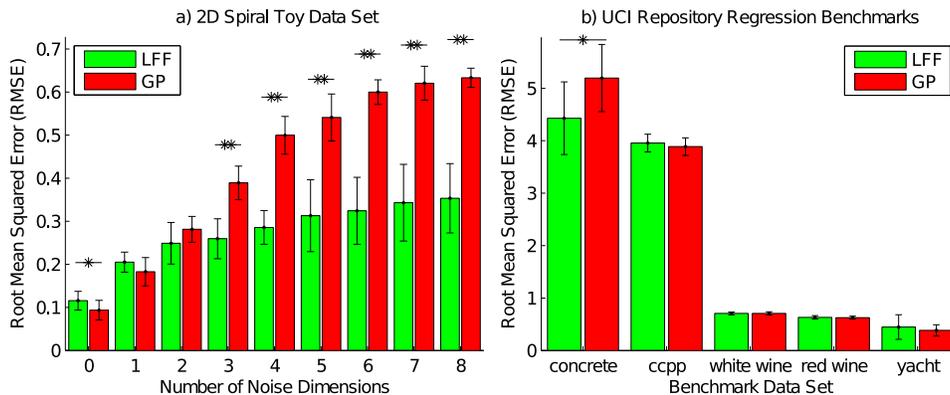


Fig. 3. Mean and standard deviation within a 10-fold cross-validation of a) the toy data set with additional independent noise input dimensions and b) all tested UCI benchmark data sets. The stars mark *significantly* different distribution of RMSE over all folds in both a *paired-sample t-test* and a *Wilcoxon signed rank test*. Significance levels are: one star $p < 0.05$, two stars $p < 0.005$.

components. The target variable is the real-valued compression strength of the mixture after it hardened.

- The *combined cycle power plant* data set (*ccpp*, Tüfekci, 2014) consists of $n = 9568$ samples with $d = 4$ dimensions describing 6 years worth of measurements from a combined gas and steam turbine. The real-valued target variable is the energy output of the system.
- The *wine quality* data set (Cortez et al., 2009) consists of two subsets with $d = 11$ dimensions each, which describe physical attributes of various white and red wines: the set contains $n = 4898$ samples of *white wine* and $n = 1599$ samples of *red wine*. The target variable is the estimated wine quality on a discrete scale from 0 to 10.
- The *yacht hydrodynamics* data set (*yacht*, Gerritsma et al., 1981) consists of $n = 308$ samples with $d = 6$ dimensions describing parameters of the *Delft yacht hull* ship-series. The real-valued target variable is the residuary resistance measured in full-scale experiments.

To demonstrate the advantage of factored basis functions, we also used the 2d-spiral toy-data set of the previous section with a varying number of additional input dimensions. Additional values were drawn i.i.d. from a Gaussian distribution and are thus independent of the target labels. As the input space \mathcal{X} grows, kernel methods will increasingly face the curse of dimensionality during training.

Every data-dimension (except the labels) have been translated and scaled to zero mean and unit-variance before training. Hyper-parameters were chosen w.r.t. the mean of a 10-fold cross-validation. LFF-regression was tested for the uniform noise-parameters $\sigma_k^2 \in \{10^{-10}, 10^{-9.75}, 10^{-9.5}, \dots, 10^{10}\}, \forall k$, i.e. for 81 different hyper-parameters. GP were tested with Gaussian kernels $\kappa(\mathbf{x}, \mathbf{y}) = \exp(-\frac{1}{2\bar{\sigma}^2} \|\mathbf{x} - \mathbf{y}\|_2^2)$ using kernel parameters $\bar{\sigma} \in \{10^{-1}, 10^{-3/4}, 10^{-1/2}, \dots, 3\}$ and prior-parameters

Table 1. 10-fold cross-validation RMSE for benchmark data sets with d dimensions and n samples, resulting in m basis functions. The cross-validation took h hours.

DATA SET	d	n	#SV	RMSE LFF	RMSE GP	m LFF	h LFF	h GP
Concrete	8	1030	927	4.429 ± 0.69	5.196 ± 0.64	4.2 ± 0.8	3.00	0.05
CCPP	4	9568	2000	3.957 ± 0.17	3.888 ± 0.17	8.8 ± 2.0	1.96	1.14
White Wine	11	4898	2000	0.707 ± 0.02	0.708 ± 0.03	4.2 ± 0.4	4.21	0.69
Red Wine	11	1599	1440	0.632 ± 0.03	0.625 ± 0.03	4.7 ± 0.7	3.25	0.13
Yacht	6	308	278	0.446 ± 0.23	0.383 ± 0.11	4.2 ± 0.6	0.43	0.005

$\beta \in \{10^{-2}, 10^{-1}, \dots, 10^{10}\}$ (see Bishop, 2006, for the definition), i.e. for 221 different hyper-parameter combinations. The number of support vectors in standard GP equals the number of training samples. As this is not feasible for larger data sets, we used the MP-MAH algorithm (Böhmer et al., 2012) to select a uniformly distributed subset of 2000 training samples for sparse GP (Rasmussen and Williams, 2006).

Figure 3a demonstrates the advantage of factored basis functions over kernel methods during training. The plot shows the *root mean squared errors*¹⁰ (RMSE) of the two dimensional spiral toy-data set with an increasing number of independent noise dimensions. GP solves the initial task better, but clearly succumbs to the curse of dimensionality, as the size of the input space \mathcal{X} grows. LFF, on the other hand, significantly overtake GP from 3 noise dimensions on, as the factored basis functions appear to be less affected by the curse. Another difference to GP is that decreasing performance automatically yields less factored basis functions (from 19.9 ± 2.18 with 0, to 6.3 ± 0.48 bases with 8 noise dimensions).

Figure 3b and Table 1 show that our LFF algorithm performs on all evaluated real-world benchmark data sets comparable to (sparse) GP. RMSE distributions over all folds were statistically indistinguishable, except for an advantage of LFF regression in the concrete compressive strength data set ($p < 0.01$ in a *t-test* and $p < 0.02$ in a *signed rank test*). As each basis function requires many iterations to converge, LFF regression runs considerably longer than standard approaches. However, LFF require between 3 and 12 factored basis functions to achieve the *same* performance as GP with 278-2000 kernel basis functions.

5 Discussion

We presented a novel algorithm for regression, which constructs factored basis functions during training. As *linear factored functions* (LFF) can in principle approximate any function in $L^2(\mathcal{X}, \vartheta)$, a regularization is necessary to avoid over-fitting.

¹⁰ RMSE are not a common performance metric for GP, which represent a *distribution* of solutions. However, RMSE reflect the objective of regression and are well suited to compare our algorithm with the *mean* of a GP.

Here we rely on a regularization scheme that has been motivated by a Taylor approximation of the least-squares cost function with (an approximation of) virtual sample-noise. RMSE performance appears comparable to Gaussian processes on real-world benchmark data sets, but the factored representation is considerably more compact and seems to be less affected by distractors.

At the moment, LFF optimization faces two challenges. (i) The optimized cost function is not convex, but the local minimum of the solution may be controlled by selecting the next factor function to optimize. For example, MARS successively adds factor functions. Generalizing this will require further research, but may also allow some performance guarantees. (ii) The large number of inner-loop iterations make the algorithm slow. This problem should be mostly solved by addressing (i), but finding a trade-off between approximation quality and runtime may also provide a less compact shortcut with similar performance.

Preliminary experiments also demonstrated the viability of LFF in a *sparse regression* approach. Sparsity refers here to a limited number of input-dimensions that affect the prediction, which can be implemented by adjusting the sample-noise parameters σ_k^2 during training for each variable \mathcal{X}_k individually. This is of particular interest, as factored functions are ideally suited to represent sparse functions and are in principle *unaffected* by the curse of dimensionality in function representation. Our approach modified the cost function to enforce LFF functions that were constant in all noise-dimensions. We did not include our results in this paper, as choosing the first updated factor functions g^k poorly resulted in basis functions that rather fitted noise than predicted labels. When we enforce sparseness, this initial mistake can afterwards no longer be rectified by other basis functions, in difference to the presented Algorithm 1. However, if this can be controlled by a sensible order in the updates, the resulting algorithm should be much faster and more robust than the presented version.

There are many application areas that may exploit the structural advantages of LLF. In *reinforcement learning* (Kaelbling et al., 1996), one can exploit the factorizing inner products to break the curse of dimensionality of the state space (Böhmer and Obermayer, 2013). Factored transition models also need to be learned from experience, which is essentially a sparse regression task. Another possible field of application are *junction trees* (for Bayesian inference, see e.g. Bishop, 2006) over continuous variables, where sparse regression may estimate the conditional probabilities. In each node one must also marginalize out variables, or calculate the point-wise product over multiple functions. Both operations can be performed analytically with LFF, the latter at the expense of more basis functions in the resulting LFF. However, one can use our framework to *compress* these functions after multiplication. This would allow junction-tree inference over mixed continuous and discrete variables.

In summary, we believe our approach to approximate functions by constructing non-linear factored basis functions (LFF) to be very promising. The presented algorithm performs comparable with Gaussian processes, but appears less sensitive to large input spaces than kernel methods. We also discussed some potential extensions for sparse regression that should improve upon that, in particular on runtime,

and gave some fields of application that would benefit greatly from the algebraic structure of LFF.

Acknowledgments The authors thank Yun Shen and the anonymous reviewers for their helpful comments. This work was funded by the *German science foundation* (DFG) within SPP 1527 *autonomous learning*.

Appendix A LFF definition and properties

Let \mathcal{X}_k denote the subset of \mathbb{R} associated with the k 'th variable of input space $\mathcal{X} \subset \mathbb{R}^d$, such that $\mathcal{X} := \mathcal{X}_1 \times \dots \times \mathcal{X}_d$. To avoid the curse of dimensionality in this space, one can integrate w.r.t. a *factored probability measure* ϑ , i.e. $\vartheta(d\mathbf{x}) = \prod_{k=1}^d \vartheta^k(dx_k)$, $\int \vartheta^k(dx_k) = 1, \forall k$. For example, ϑ^k could be uniform or Gaussian distributions over \mathcal{X}_k and the resulting ϑ would be a uniform or Gaussian distribution over the input space \mathcal{X} .

A function $g : \mathcal{X} \rightarrow \mathbb{R}$ is called a *factored function* if it can be written as a product of one-dimensional *factor functions* $g^k : \mathcal{X}_k \rightarrow \mathbb{R}$, i.e. $g(\mathbf{x}) = \prod_{k=1}^d g^k(x_k)$. We only consider factored functions g that are twice integrable w.r.t. measure ϑ , i.e. $g \in L^2(\mathcal{X}, \vartheta)$. Note that not all functions $f \in L^2(\mathcal{X}, \vartheta)$ are factored, though. Due to *Fubini's theorem* the d -dimensional inner product between two factored functions $g, g' \in L^2(\mathcal{X}, \vartheta)$ can be written as the product of d one-dimensional inner products:

$$\langle g, g' \rangle_{\vartheta} = \int \vartheta(d\mathbf{x}) g(\mathbf{x}) g'(\mathbf{x}) = \int \prod_{k=1}^d \vartheta^k(dx_k) g^k(dx_k) g'^k(dx_k) = \prod_{k=1}^d \langle g^k, g'^k \rangle_{\vartheta^k}.$$

This trick can be used to solve the integrals at the heart of many least-squares algorithms. Our aim is to *learn* factored basis functions ψ_i . To this end, let $\{\phi_j^k : \mathcal{X}_k \rightarrow \mathbb{R}\}_{j=1}^{m_k}$ be a well-chosen¹¹ (i.e. universal) basis on \mathcal{X}_k , with the space of linear combinations denoted by $\mathcal{L}_{\phi}^k := \{\mathbf{b}^{\top} \phi^k | \mathbf{b} \in \mathbb{R}^{m_k}\}$. One can thus approximate factor functions of ψ_i in \mathcal{L}_{ϕ}^k , i.e., as linear functions

$$\psi_i^k(x_k) := \sum_{j=1}^{m_k} B_{ji}^k \phi_j^k(x_k) \in \mathcal{L}_{\phi}^k, \quad \mathbf{B}^k \in \mathbb{R}^{m_k \times m}. \quad (8)$$

Let \mathcal{F} be the space of all factored basis functions ψ_i defined by the factor functions ψ_i^k above, and \mathcal{F}^m be the space of all linear combinations of those m factored basis functions (Equation 2).

Marginalization of LFF can be performed analytically with Fourier bases ϕ_j^k and uniform distribution ϑ (many other bases can be analytically solved as well):

$$\int \vartheta^l(dx_l) f(\mathbf{x}) = \sum_{i=1}^m \left(a_i \sum_{j=1}^{m_l} B_{ji}^l \underbrace{\langle \phi_j^l, 1 \rangle_{\vartheta^l}}_{\text{mean of } \phi_j^l} \right) \left[\prod_{k \neq l} \psi_i^k \right] \stackrel{\text{Fourier}}{=} \sum_{i=1}^m \underbrace{a_i B_{1i}^l}_{\text{new } a_i} \left[\prod_{k \neq l} \psi_i^k \right]. \quad (9)$$

¹¹ Examples for continuous variables \mathcal{X}_k are Fourier cosine bases $\phi_j^k(x_k) \sim \cos((j-1)\pi x_k)$, and Gaussian bases $\phi_j^k(x_k) = \exp\left(\frac{1}{2\sigma^2}(x_k - s_{kj})^2\right)$. Discrete variables may be represented with Kronecker-delta bases $\phi_j^k(x_k = i) = \delta_{ij}$.

Using the trigonometric *product-to-sum* identity $\cos(x) \cdot \cos(y) = \frac{1}{2}(\cos(x-y) + \cos(x+y))$, one can also compute the point-wise product between two LFF f and \bar{f} with cosine-Fourier base (solutions to other Fourier bases are less elegant):

$$\begin{aligned} \tilde{f}(\mathbf{x}) &:= f(\mathbf{x}) \cdot \bar{f}(\mathbf{x}) \\ &\stackrel{\text{Fourier}}{=} \sum_{i,j=1}^{m\bar{m}} \underbrace{a_i \bar{a}_j}_{\text{new } \tilde{a}_t} \prod_{k=1}^d \sum_{l=1}^{2m_k} \overbrace{\left(\frac{1}{2} \sum_{q=1}^{l-1} B_{qi}^k \bar{B}_{(l-q)j}^k + \frac{1}{2} \sum_{q=l+1}^{m_k} B_{qi}^k \bar{B}_{(q-l)j}^k \right)}^{\text{new } \tilde{B}_{it}^k} \phi_l^k(x_k), \end{aligned} \quad (10)$$

where $t := (i-1)\bar{m} + j$, and $B_{ji}^k := 0, \forall j > m_k$, for both f and \bar{f} . Note that this increases the number of basis functions $\tilde{m} = m\bar{m}$, and the number of bases $\tilde{m}_k = 2m_k$ for each respective input dimension. The latter can be counteracted by *low-pass filtering*, i.e., by setting $\tilde{B}_{ji}^k := 0, \forall j > m_k$.

Appendix B Inner loop derivation

Here we will optimize the problem in Equation 6 for one variable \mathcal{X}_k at a time, by describing the update step $g^k \leftarrow g^{/k}$. This is repeated with randomly chosen variables k , until convergence of the cost $\tilde{\mathcal{C}}[g]$, that is, until all possible updates decrease the cost less than some small ϵ .

Let in the following $\mathbf{C}^k := \langle \phi^k, \phi^{k\top} \rangle_{\vartheta^k}$ and $\dot{\mathbf{C}}^k := \langle \nabla_k \phi^k, \nabla_k \phi^{k\top} \rangle_{\vartheta^k}$ denote covariance matrices, and $\mathbf{R}_l^k := \frac{\partial}{\partial \mathbf{b}^k} \langle \nabla_l g, \nabla_l f \rangle_{\vartheta}$ denote the derivative of one regularization term. Note that for some choices of bases $\{\phi_j^k\}_{j=1}^{m_k}$, one can compute the covariance matrices analytically before the main algorithm starts, e.g. Fourier cosine bases have $C_{ij}^k = \delta_{ij}$ and $\dot{C}_{ij}^k = (i-1)^2 \pi^2 \delta_{ij}$.

The approximated cost function in Equation 6 is

$$\tilde{\mathcal{C}}[g] = \|g\|_{\xi}^2 - 2\langle g, \mu - f \rangle_{\xi} + \|\mu - f\|_{\xi}^2 + \sum_{k=1}^d \sigma_k^2 \left(\|\nabla_k g\|_{\vartheta}^2 + 2\langle \nabla_k g, \nabla_k f \rangle_{\vartheta} + \|\nabla_k f\|_{\vartheta}^2 \right).$$

The non-zero gradients of all inner products of this equation w.r.t. parameter vector $\mathbf{b}^k \in \mathbb{R}^{m_k}$ are

$$\begin{aligned} \frac{\partial}{\partial \mathbf{b}^k} \langle g, g \rangle_{\xi} &= 2 \langle \phi^k \cdot \prod_{l \neq k} g^l, \prod_{l \neq k} g^l \cdot \phi^{k\top} \rangle_{\xi} \mathbf{b}^k, \\ \frac{\partial}{\partial \mathbf{b}^k} \langle g, \mu - f \rangle_{\xi} &= \langle \phi^k \cdot \prod_{l \neq k} g^l, \mu - f \rangle_{\xi}, \\ \frac{\partial}{\partial \mathbf{b}^k} \langle \nabla_l g, \nabla_l g \rangle_{\vartheta} &= \frac{\partial}{\partial \mathbf{b}^k} \langle \nabla_l g^l, \nabla_l g^l \rangle_{\vartheta^l} \prod_{s \neq l} \overbrace{\langle g^s, g^s \rangle_{\vartheta^s}}^1 = 2 \delta_{kl} \dot{\mathbf{C}}^k \mathbf{b}^k, \\ \mathbf{R}_l^k &:= \frac{\partial}{\partial \mathbf{b}^k} \langle \nabla_l g, \nabla_l f \rangle_{\vartheta} = \begin{cases} \dot{\mathbf{C}}^k \mathbf{B}^k \left[\mathbf{a} \cdot \prod_{s \neq k} \mathbf{B}^{s\top} \mathbf{C}^s \mathbf{b}^s \right], & \text{if } k = l \\ \mathbf{C}^k \mathbf{B}^k \left[\mathbf{a} \cdot \mathbf{B}^{l\top} \dot{\mathbf{C}}^l \mathbf{b}^l \cdot \prod_{s \neq k \neq l} \mathbf{B}^{s\top} \mathbf{C}^s \mathbf{b}^s \right], & \text{if } k \neq l \end{cases} \end{aligned}$$

Setting this to zero yields the unconstrained solution g_{uc}^k ,

$$\mathbf{b}_{uc}^k = \left(\overbrace{\langle \phi^k \cdot \prod_{l \neq k} g^l, \prod_{l \neq k} g^l \cdot \phi^{k\top} \rangle_{\xi} + \sigma_k^2 \dot{\mathbf{C}}^k}^{\text{regularized covariance matrix } \tilde{\mathbf{C}}^k} \right)^{-1} \left(\langle \phi^k \cdot \prod_{l \neq k} g^l, \mu - f \rangle_{\xi} - \sum_{l=1}^d \mathbf{R}_l^k \sigma_l^2 \right). \quad (11)$$

However, these parameters do not satisfy to the constraint $\|g'^k\|_{\vartheta^k} \stackrel{!}{=} 1$, and have to be normalized:

$$\mathbf{b}'^k := \frac{\mathbf{b}_{uc}^k}{\|g_{uc}^k\|_{\vartheta^k}} = \frac{\mathbf{b}_{uc}^k}{\sqrt{\mathbf{b}_{uc}^{k\top} \mathbf{C}^k \mathbf{b}_{uc}^k}}. \quad (12)$$

The inner loop finishes when for all k the improvement¹² from g^k to g'^k drops below some very small threshold ϵ , i.e. $\tilde{\mathcal{C}}[g] - \tilde{\mathcal{C}}[g'] < \epsilon$. Using $g'^l = g^l, \forall l \neq k$, one can calculate the left hand side:

$$\begin{aligned} \tilde{\mathcal{C}}[g] - \tilde{\mathcal{C}}[g'] &= \|g\|_{\xi}^2 - \|g'\|_{\xi}^2 - 2\langle g - g', \mu - f \rangle_{\xi} \\ &\quad + \sum_{l=1}^d \sigma_l^2 \left[\underbrace{\|\nabla_l g\|_{\vartheta}^2}_{\mathbf{b}^{l\top} \mathbf{C}^l \mathbf{b}^l} - \underbrace{\|\nabla_l g'\|_{\vartheta}^2}_{\mathbf{b}'^{l\top} \mathbf{C}^l \mathbf{b}'^l} - 2 \underbrace{\langle \nabla_l g - \nabla_l g', \nabla_l f \rangle_{\vartheta}}_{(\mathbf{b}^k - \mathbf{b}'^k)^{\top} \mathbf{R}_l^k} \right] \\ &= 2\langle g - g', \mu - f \rangle_{\xi} + \mathbf{b}^{k\top} \bar{\mathbf{C}}^k \mathbf{b}^{k\top} - \mathbf{b}'^{k\top} \bar{\mathbf{C}}^k \mathbf{b}'^{k\top} - 2(\mathbf{b}^k - \mathbf{b}'^k)^{\top} \left(\sum_{l=1}^d \mathbf{R}_l^k \sigma_l^2 \right). \end{aligned} \quad (13)$$

Appendix C Proofs of the propositions

Proof of Proposition 2: The 1st order Taylor approximation of any $g, f \in L^2(\mathcal{X}, \xi\chi)$ around $\mathbf{z} \in \mathcal{X}$ is $f(\mathbf{x}) = f(\mathbf{z} + \mathbf{x} - \mathbf{z}) \approx f(\mathbf{z}) + (\mathbf{x} - \mathbf{z})^{\top} \nabla f(\mathbf{z})$. For the Hilbert space $L^2(\mathcal{X}, \xi\chi)$ we can thus approximate:

$$\begin{aligned} \langle g, f \rangle_{\xi\chi} &= \int \xi(d\mathbf{z}) \int \chi(d\mathbf{x}|\mathbf{z}) g(\mathbf{x}) f(\mathbf{x}) \\ &\approx \int \xi(d\mathbf{z}) \left(g(\mathbf{z}) f(\mathbf{z}) \int \overbrace{\xi(d\mathbf{x}|\mathbf{z})}^1 + g(\mathbf{z}) \int \overbrace{\chi(d\mathbf{x}|\mathbf{z}) (\mathbf{x} - \mathbf{z})^{\top} \nabla f(\mathbf{z})}^{\mathbf{0} \text{ due to (eq.5)}} \right. \\ &\quad \left. + \int \underbrace{\chi(d\mathbf{x}|\mathbf{z}) (\mathbf{x} - \mathbf{z})^{\top} \nabla g(\mathbf{z}) f(\mathbf{z})}_{\mathbf{0} \text{ due to (eq.5)}} + \nabla g(\mathbf{z})^{\top} \int \underbrace{\chi(d\mathbf{x}|\mathbf{z}) (\mathbf{x} - \mathbf{z}) (\mathbf{x} - \mathbf{z})^{\top} \nabla f(\mathbf{z})}_{\frac{\vartheta}{\xi}(\mathbf{z}) \cdot \Sigma \text{ due to (eq.5)}} \right) \\ &= \langle g, f \rangle_{\xi} + \sum_{k=1}^d \sigma_k^2 \langle \nabla_k g, \nabla_k f \rangle_{\vartheta}. \end{aligned}$$

Using this twice and the zero mean assumption (Eq. 5), we can derive:

$$\begin{aligned} \inf_{g \in \mathcal{F}} \mathcal{C}[f + g|\chi, \mu] &\equiv \inf_{g \in \mathcal{F}} \iint \xi(d\mathbf{z}) \chi(d\mathbf{x}|\mathbf{z}) \left(g^2(\mathbf{x}) - 2g(\mathbf{x}) (\mu(\mathbf{z}) - f(\mathbf{x})) \right) \\ &= \inf_{g \in \mathcal{F}} \langle g, g \rangle_{\xi\chi} + 2\langle g, f \rangle_{\xi\chi} - 2 \int \xi(d\mathbf{z}) \mu(\mathbf{z}) \int \chi(d\mathbf{x}|\mathbf{z}) g(\mathbf{x}) \\ &\approx \inf_{g \in \mathcal{F}} \langle g, g \rangle_{\xi} - 2\langle g, \mu - f \rangle_{\xi} + \sum_{k=1}^d \sigma_k^2 \left(\langle \nabla_k g, \nabla_k g \rangle_{\vartheta} + 2\langle \nabla_k g, \nabla_k f \rangle_{\vartheta} \right) \\ &\equiv \inf_{g \in \mathcal{F}} \|g - (\mu - f)\|_{\xi}^2 + \sum_{k=1}^d \sigma_k^2 \|\nabla_k g + \nabla_k f\|_{\vartheta}^2 = \tilde{\mathcal{C}}[g]. \quad \square \end{aligned}$$

¹² Anything simpler does not converge, as the parameter vectors often evolve along chaotic orbits in \mathbb{R}^{m_k} .

Proof of Proposition 3: The analytical solution to the optimization problem in Equation 6 is derived in Appendix B and has a unique solution if the matrix $\bar{\mathbf{C}}^k$, defined in Equation 11, is of full rank:

$$\bar{\mathbf{C}}^k := \langle \phi^k \cdot \prod_{l \neq k} g^l, \prod_{l \neq k} g^l \cdot \phi^{k\top} \rangle_{\xi} + \sigma_k^2 \dot{\mathbf{C}}^k.$$

For Fourier bases the matrix $\dot{\mathbf{C}}^k$ is diagonal, with \dot{C}_{11}^k being the only zero entry. $\bar{\mathbf{C}}^k$ is therefore full rank if $\sigma_k^2 > 0$ and $\bar{C}_{11}^k > 0$. Because ϑ is *absolutely continuous* w.r.t. ξ , the constraint $\|g^l\|_{\vartheta} = 1, \forall l$, implies that there exist no g^l that is zero on *all* training samples.

As the first Fourier base is a constant, $\langle \phi_1^k \cdot \prod_{l \neq k} g^l, \prod_{l \neq k} g^l \cdot \phi_1^k \rangle_{\xi} > 0$ and the matrix $\bar{\mathbf{C}}^k$ is therefore of full rank. \square

Bibliography

- R. E. Bellman. *Dynamic programming*. Princeton University Press, 1957.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006. ISBN 0387310738.
- W. Böhmer and K. Obermayer. Towards structural generalization: Factored approximate planning. ICRA Workshop on Autonomous Learning, 2013. URL http://autonomous-learning.org/wp-content/uploads/13-ALW/paper_1.pdf.
- W. Böhmer, S. Grünewälder, H. Nickisch, and K. Obermayer. Generating feature spaces for linear algorithms with regularized sparse kernel slow feature analysis. *Machine Learning*, 89(1-2):67–86, 2012.
- W. Böhmer, S. Grünewälder, Y. Shen, M. Musial, and K. Obermayer. Construction of approximation spaces for reinforcement learning. *Journal of Machine Learning Research*, 14:2067–2118, July 2013.
- B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pages 144–152, 1992.
- P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4):547–553, 2009.
- L. Csató and M. Opper. Sparse on-line Gaussian processes. *Neural Computation*, 14(3):641–668, 2002.
- J. H. Friedman. Multivariate adaptive regression splines. *The Annals of Statistics*, 19(1):1–67, 1991.
- J. Gerritsma, R. Onnink, and A. Versluis. Geometry, resistance and stability of the delft systematic yacht hull series. *Int. Shipbuilding Progress*, 28:276–297, 1981.
- M. Gönen and E. Alpaydm. Multiple kernel learning algorithms. *Journal of Machine Learning Research*, 12:2211–2268, July 2011. ISSN 1532-4435.
- S. Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall, 2nd edition, 1998. ISBN 978-0132733502.
- L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: a survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- J. A. Nelder and R. W. M. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society, Series A, General*, 135:370–384, 1972.
- J. Pearl. *Probabilistic reasoning in intelligent systems*. Morgan Kaufmann, 1988.

- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- M. E. Tipping. Sparse bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, September 2001. ISSN 1532-4435.
- P. Tüfekci. Prediction of full load electrical power output of a base load operated combined cycle power plant using machine learning methods. *International Journal of Electrical Power & Energy Systems*, 60:126–140, 2014.
- V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.
- Z. Wang, K. Crammer, and S. Vucetic. Breaking the curse of kernelization: Budgeted stochastic gradient descent for large-scale svm training. *Journal of Machine Learning Research*, 13(1):3103–3131, October 2012. ISSN 1532-4435.
- I-C. Yeh. Modeling of strength of high performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998.

Algorithm 1 (detailed) – LFF-Regression

Input: $\mathbf{X} \in \mathbb{R}^{d \times n}$, $\mathbf{y} \in \mathbb{R}^n$, $\boldsymbol{\sigma}^2 \in \mathbb{R}^d$, $\epsilon, \varepsilon \in \mathbb{R}$
 $\mathbf{C}^k := \langle \phi^k, \phi^k \rangle_{\vartheta^k}$, $\dot{\mathbf{C}}^k := \langle \nabla \phi^k, \nabla \phi^k \rangle_{\vartheta^k}$, $\forall k$ // analytical covariance matrices
 $\phi_{j_t}^k := \phi_j^k(X_{kt})$, $\forall k, \forall j, \forall t$ // optional cache of sample-expansion
 $\mathbf{f} := \mathbf{0} \in \mathbb{R}^n$; $\mathbf{a} := \emptyset$; $\mathbf{B}^k := \emptyset$, $\forall k$; $\Psi := \infty$ // initialize empty $f \in \mathcal{F}^0$
while $\det\left(\frac{1}{n}\Psi\Psi^\top\right) > \varepsilon$ **do**
 $\mathbf{b}^k := \mathbf{1}^k \in \mathbb{R}^{m_k}$, $\forall k$; $\mathbf{g}^k := \mathbf{1} \in \mathbb{R}^n$, $\forall k$ // initialize all g^k as constant
 $\mathbf{h} := \infty \in \mathbb{R}^d$ // initialize estimated improvement
while $\max(\mathbf{h}) > \epsilon$ **do**
for k **in** $\text{randperm}(1, \dots, d)$ **do**
 $\mathbf{R}_k := \dot{\mathbf{C}}^k \mathbf{B}^k [\mathbf{a} \cdot \prod_{s \neq k} \mathbf{B}^{s\top} \mathbf{C}^s \mathbf{b}^s]$ // $\mathbf{R}_k = \frac{\partial}{\partial \mathbf{b}^k} \langle \nabla_k g, \nabla_k f \rangle_{\vartheta}$
 $\mathbf{R}_l := \mathbf{C}^k \mathbf{B}^k [\mathbf{a} \cdot \mathbf{B}^{l\top} \dot{\mathbf{C}}^l \mathbf{b}^l \cdot \prod_{s \neq k, l} \mathbf{B}^{s\top} \mathbf{C}^s \mathbf{b}^s]$, $\forall l \neq k$ // $\mathbf{R}_l = \frac{\partial}{\partial \mathbf{b}^k} \langle \nabla_l g, \nabla_l f \rangle_{\vartheta}$
 $\bar{\mathbf{C}} := \Phi^k \left[\Phi^{k\top} \cdot \prod_{l \neq k} (\mathbf{g}^l)^2 \mathbf{1}^\top \right] + \sigma_k^2 \dot{\mathbf{C}}^k$ // regularized cov. matrix (eq. 11)
 $\mathbf{b}' := \bar{\mathbf{C}}^{-1} \left(\Phi^k \left[(\mathbf{y} - \mathbf{f}) \cdot \prod_{l \neq k} \mathbf{g}^l \right] - \mathbf{R} \boldsymbol{\sigma}^2 \right)$ // unconstrained g_{uc}^k (eq. 11)
 $\mathbf{b}' := \mathbf{b}' / \sqrt{\mathbf{b}'^\top \mathbf{C}^k \mathbf{b}'}$ // enforce constraints (eq. 12)
 $h_k := \frac{2}{n} (\mathbf{b}^k - \mathbf{b}')^\top \left(\Phi^k \left[(\mathbf{y} - \mathbf{f}) \cdot \prod_{l \neq k} \mathbf{g}^l \right] \right)$ // approximate $2 \langle g - g', \mu - f \rangle_{\xi}$
 $h_k := h_k + \mathbf{b}^k \bar{\mathbf{C}} \mathbf{b}^k - \mathbf{b}' \bar{\mathbf{C}} \mathbf{b}' - 2(\mathbf{b}^k - \mathbf{b}')^\top \mathbf{R} \boldsymbol{\sigma}^2$ // cost improvement (eq. 13)
 $\mathbf{b}^k := \mathbf{b}'$; $\mathbf{g}^k := \Phi^{k\top} \mathbf{b}^k$ // update factor function g^k
end for // end function g^k update
end while // end inner loop: cost function converged and thus g optimized
 $\mathbf{B}^k := [\mathbf{B}^k, \mathbf{b}^k]$, $\forall k$; $\Psi := \left[\prod_{k=1}^d \mathbf{B}^{k\top} \Phi^k \right]$ // adding g to the bases functions of f
 $\mathbf{a} := (\Psi\Psi^\top)^{-1} \Psi \mathbf{y}$; $\mathbf{f} := \Psi^\top \mathbf{a}$ // project μ onto new bases
end while // end outer loop: new g no longer linear independent, thus $f \approx \mu$
Output: $\mathbf{a} \in \mathbb{R}^m$, $\{\mathbf{B}^k \in \mathbb{R}^{m_k \times m}\}_{k=1}^d$ // return parameters of $f \in \mathcal{F}^m$
