

Cross-Validation Optimization for Large Scale Structured Classification Kernel Methods

Matthias W. Seeger

*Max Planck Institute for Biological Cybernetics
Spemannstr. 38
Tübingen, Germany*

SEEGER@TUEBINGEN.MPG.DE

Editor: Zoubin Ghahramani

Abstract

We propose a highly efficient framework for penalized likelihood kernel methods applied to multi-class models with a large, structured set of classes. As opposed to many previous approaches which try to decompose the fitting problem into many smaller ones, we focus on a Newton optimization of the complete model, making use of model structure and linear conjugate gradients in order to approximate Newton search directions. Crucially, our learning method is based entirely on matrix-vector multiplication primitives with the kernel matrices and their derivatives, allowing straightforward specialization to new kernels, and focusing code optimization efforts to these primitives only.

Kernel parameters are learned automatically, by maximizing the cross-validation log likelihood in a gradient-based way, and predictive probabilities are estimated. We demonstrate our approach on large scale text classification tasks with hierarchical structure on thousands of classes, achieving state-of-the-art results in an order of magnitude less time than previous work.

Parts of this work appeared in the conference paper Seeger (2007).

Keywords: multi-way classification, kernel logistic regression, hierarchical classification, cross validation optimization, Newton-Raphson optimization

1. Introduction

In recent years, machine learning researchers started to address problems with kernel machines which require models with a large number of dependent variables, and whose fitting demand training samples with very many cases. For example, for multi-way classification models with a hierarchically structured label space (Cai and Hofmann, 2004), modern applications call for predictions on thousands of classes, and very large data sets become available. However, if n and C denote data set size and number of classes respectively, nonparametric kernel methods like *support vector machines* (SVMs) or *Gaussian processes* (GPs) typically scale super-linearly in nC , if dependencies between the latent class functions are represented properly.

Furthermore, most large scale kernel methods proposed so far refrain from solving the problem of learning hyperparameters (kernel or loss function parameters), also known as “learning the kernels”. The user has to run cross-validation schemes essentially “by hand”, which is not suitable for learning more than a few hyperparameters. However, many models for modern applications come with a large number of hyperparameters (for example to represent dependencies through “mixing” as in independent components analysis), and adjusting them through optimization must make use of gradients.

We propose a general framework for learning in probabilistic kernel classification models. While the models treated here are not novel, a major feature of our approach is the high computational efficiency with which the primary fitting (for fixed hyperparameters) is done. For example, our framework applied to hierarchical classification with hundreds of classes and thousands of data points requires a few minutes for fitting. The central idea is to step back from what seems to be the dominating approach in machine learning at the moment, namely to solve a large convex optimization problem by iteratively solving very many small ones. A popular approach for these small steps is to minimize the criterion w.r.t. a few variables only, keeping the other ones fixed, and many variations of this theme have been proposed. In this paper, we focus on the opposite approach of trying to find directions which lead to fast descent, no matter how many of the variables are involved. This is essentially Newton’s method, and one aspect of our work is to find approximate Newton directions very efficiently, making use of model structure and linear conjugate gradients in order to reduce the computation to standard linear algebra primitives on large contiguous chunks of memory. Interestingly, such global approaches are generally favoured in the optimization community for problems (such as kernel methods fitting) which cannot be decomposed naturally into parts. While other gradient-based optimizers such as scaled conjugate gradients could be used as well, they require more fine-tuning (for example, preconditioning) to the specific problem they are applied to, while Newton’s method is closer to a “black box” technique and can be transferred to novel situations without many changes.

For multi-way classification, our primary fitting method scales linearly in C , and depends on n mainly via a fixed number of *matrix-vector multiplications* (MVM) with $n \times n$ kernel matrices. In many situations, these MVM primitives can be computed very efficiently, often without having to store the kernel matrices themselves.

We also show how to choose hyperparameters *automatically* by maximizing the cross-validation log likelihood, making use of our primary fitting technology as inner loop in order to compute the CV criterion and its gradient. It is important to note that our hyperparameter learning method works by gradient-based optimization, where the dominating part of the gradient computation does not scale with the number of hyperparameters at all.¹ The gradient computation also requires a number of MVMs with derivatives of kernel matrices, which can be reduced to kernel MVMs for many frequently used kernels (see Section 7.3). Therefore, our approach can in principle be used to learn a large number of hyperparameters without user interaction.

We apply our framework to hierarchical classification with many classes. The hierarchy is represented through an ANOVA setup. While the C latent class functions are fully dependent *a priori*, the scaling of our method stays close to what unstructured (flat) classification with C classes would require. We test our framework on the same tasks treated by Cai and Hofmann (2004), achieving comparable results in at least an order of magnitude less time.

Our proposal to use approximate Newton methods is not novel as such. The Newton method, or a variant of it called Fisher scoring, is the standard approach for fitting generalized linear models in statistics (Green and Silverman, 1994; McCullach and Nelder, 1983), at least if parametric models are fitted to moderately sized samples. Our primary fitting method for flat multi-way classification (see Section 2) appeared in Williams and Barber (1998). However, we demonstrate the usefulness of this principle on a much larger scale, showing how model structure can (and has to) be exploited

1. Such scaling behaviour is fairly standard in Gaussian process marginal likelihood maximization techniques (Williams and Barber, 1998), but has only recently been brought to attention in the SVM community (Keerthi et al., 2007).

in this context. Furthermore, we demonstrate how the secondary task of hyperparameter learning can be reduced to the same underlying primitives.

The structure of the paper is as follows. Our model and method of parameter fitting is given in Section 2. An extension to hierarchical classification is provided in Section 3, and in Section 4 we give our automatic hyperparameter learning procedure. Essential computational details are discussed in Section 5. Experimental results on a very large hierarchical text classification and several standard machine learning problems are given in Section 6. We close with a discussion in Section 7, relating our global direction approach to popular block coordinate descent techniques in Section 7.2, and pointing out future work in Section 7.4.

Optimized C++ software for our framework is available as part of the LHOTSE toolbox for adaptive statistical models, which is freely available for non-commercial purposes.² The implementation contains the linear kernel case used in Section 6.1 (see Appendix D.3), as well as a generic representation described in Appendix D.1, with which the experiments in Section 6.2, Section 6.3 have been done. It is fairly simple to include new kernels or (approximate) kernel MVM implementations.

2. Penalized Multiple Logistic Regression

In this section, we introduce our framework on a multi-way classification model with C classes, where structure between classes is not modelled. We refer to this setup as *flat classification*, in that the label set is flat (unstructured).

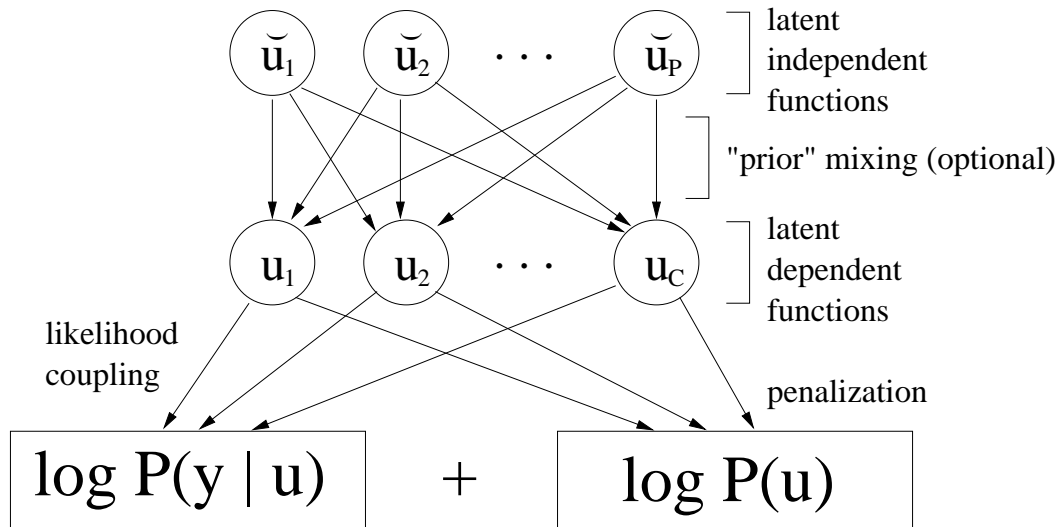


Figure 1: Structure of penalized likelihood optimization.

In general, our framework is applicable to models of the form depicted in Figure 1. A set of latent (unobserved) functions $u_c(\cdot)$ is fitted to observed data by penalized likelihood maximization. For many models, the penalisation term (also called regulariser) corresponds to the logarithm of a prior density over the $u_c(\cdot)$. This primary fitting step corresponds to a convex optimization problem over finitely many variables. Structure in such models is represented either as couplings in the log

². Available at www.kyb.tuebingen.mpg.de/bs/people/seeger/lhotse/.

likelihood function, or in the penalisation (or log prior) term. The latter can be realized through the linear mixing of *a priori* independent functions $\tilde{u}_p(\cdot)$, in other words the penaliser over the latter decouples w.r.t. p (our main example of such mixing is hierarchical classification, developed in Section 3).

We now apply this general framework to flat classification, where $y \in \{1, \dots, C\}$ is to be predicted from $x \in \mathcal{X}$, given some i.i.d. data $D = \{(x_i, y_i) \mid i = 1, \dots, n\}$. Our notation convention for vectors and matrices is detailed in Appendix A, where we also collect all major notational definitions in a table. We code y_i as $y_i \in \{0, 1\}^C$, $1^T y_i = 1$ (zero-one coding).³ We employ the *multiple logistic regression model*, consisting of C latent class functions $u_c(\cdot)$ feeding into the multiple logistic (or softmax) likelihood $P(y_{ic} = 1 \mid x_i, u_i(\cdot)) = e^{u_c(x_i)} / (\sum_{c'} e^{u_{c'}(x_i)})$.

We write $u_c(\cdot) = f_c(\cdot) + b_c$ for intercept (or bias) parameters $b_c \in \mathbb{R}$ and functions $f_c(\cdot)$ living in a reproducing kernel Hilbert space (RKHS) with kernel $K^{(c)} = K^{(c)}(\cdot, \cdot)$ (Schölkopf and Smola, 2002), and consider the *penalized negative log likelihood*

$$\Phi = - \sum_{i=1}^n \log P(y_i \mid u_i) + (1/2) \sum_{c=1}^C \|f_c(\cdot)\|_c^2 + (1/2) \sigma^{-2} \|b\|^2, \quad u_i = (u_c(x_i))_c \in \mathbb{R}^C,$$

which we minimize for primary fitting. Here, $\|\cdot\|_c$ is the RKHS norm for kernel $K^{(c)}$. The idea is that deviations in f_c from desired functional properties encoded in $K^{(c)}$ are penalized by a large $\|f_c(\cdot)\|_c^2$. For example, for the Gaussian kernel (7), non-smooth f_c are penalized, and for the linear kernel (Appendix D.3), $\|f_c(\cdot)\|_c^2$ is the squared norm of the weight vector. Details on penalized likelihood kernel methods and RKHS penalisation can be found in Green and Silverman (1994) and Schölkopf and Smola (2002).

The model can also be understood in a Bayesian context, where the penalisation terms come from zero mean Gaussian process priors on the functions $f_c(\cdot)$, and b has a zero mean Gaussian prior with variance σ^2 . From this viewpoint, we do a maximum a-posteriori (MAP) approximation here, without however taking covariances into account properly (which would be much more expensive to do). Details on Gaussian processes for machine learning can be found in Seeger (2004) and Rasmussen and Williams (2006).

Since the likelihood depends on the $f_c(\cdot)$ only through the values $f_c(x_i)$ at the data points, every minimizer of Φ must be a kernel expansion: $f_c(\cdot) = \sum_i \alpha_{ic} K^{(c)}(\cdot, x_i)$. This fact is known as representer theorem (Green and Silverman, 1994; Wahba, 1990). Plugging this in, the regulariser becomes $(1/2) \alpha^T K \alpha + (1/2) \sigma^{-2} \|b\|^2$, where $K^{(c)} = (K^{(c)}(x_i, x_j))_{i,j} \in \mathbb{R}^{n,n}$, and $K = \text{diag}(K^{(c)})_c$ is block-diagonal. The kernels $K^{(c)}$ can in general be different, although sharing kernels among classes can lead to computational savings, in that some of the blocks $K^{(c)}$ are identical. Our implementation of block sharing is described in Appendix D.1.

We show in Section 5.1.1 that the b_c may be eliminated as $b = \sigma^2 (I \otimes 1^T) \alpha$. Thus, if $\tilde{K} = K + \sigma^2 (I \otimes 1)(I \otimes 1^T)$, then our criterion Φ becomes

$$\Phi = \Phi_{lh} + \frac{1}{2} \alpha^T \tilde{K} \alpha, \quad \Phi_{lh} = -y^T u + 1^T l, \quad l_i = \log 1^T \exp(u_i), \quad u = \tilde{K} \alpha. \quad (1)$$

Φ is strictly convex in α , being a sum of linear, quadratic, and *logsumexp* terms of the form $\log 1^T \exp(u_i)$ (Boyd and Vandenberghe, 2002), so it has a unique minimum point $\hat{\alpha}$. The corre-

3. We switch between the formats y_i, y_i . Note that y_{ic} denotes a component in $y_i = (y_{ic})_c$.

sponding kernel expansions are

$$\hat{u}_c(\cdot) = \sum_i \hat{\alpha}_{ic} (K^{(c)}(\cdot, x_i) + \sigma^2).$$

Estimates of the conditional probability on test points x_* are obtained by plugging $\hat{u}_c(x_*)$ into the likelihood. These estimates are asymptotically consistent, although better finite sample estimates could probably be obtained by a more Bayesian treatment.

We note that this setup is related to the multi-class SVM (Crammer and Singer, 2001), where $-\log P(y_i|u_i)$ is replaced by the margin loss $-u_{y_i}(x_i) + \max_c \{u_c(x_i) + 1 - \delta_{c,y_i}\}$. Here, $\delta_{a,b} = \mathbb{I}_{\{a=b\}}$. The negative log multiple logistic likelihood has similar properties, but is smooth as a function of u , and the primary fitting of α does not require constrained convex optimization. Furthermore, universal consistency for estimates of $P(y_*|x_*)$ can be established for the multiple logistic loss, but fails to hold for the SVM variant (Bartlett and Tewari, 2004).

We will minimize Φ using the *Newton-Raphson* (NR) algorithm. The computation of Newton search directions requires solving a system with the Hessian and the gradient of Φ , which we will do approximately using the *linear conjugate gradients* (LCG) algorithm. This can be done without fully computing, storing, or inverting the Hessian, all of which would not be possible for large nC . In fact, the task is reduced to computing $k_1(k_2 + 2)$ MVMs with K , where k_1 is the number of NR iterations, k_2 the number of LCG steps for computing each Newton direction. Since NR is a second-order convergent method, k_1 is generally small. k_2 determines the quality of each Newton direction, and again, fairly small values seem sufficient (see Section 6.1). Details are provided in Section 5.1.

Finally, some readers may wonder why we favour the NR algorithm here, which in practice can be fairly complicated to implement, while we could do a simpler gradient-based optimization of Φ w.r.t. α , for example by scaled (non-linear) conjugate gradients (SCG). The problem is that on tasks of the size we want to address, non-invariant methods such as SCG tend to fail completely if not properly preconditioned, and we experienced exactly that in preliminary experiments. In contrast to that, NR is invariant to the choice of optimization variables, so does not have to be preconditioned. It is by far the preferred method in the optimization literature (Bertsekas, 1999; Boyd and Vandenberghe, 2002), and many ideas for preconditioning or Quasi-Newton try to approximate the NR directions. We think that a proper SCG implementation can be at least as efficient as NR, but needs fine-tuning to the specific problem, which in the case of hierarchical classification (discussed next) is already quite difficult. More details on this point are given in Section 5.4 and also Section 7.2.

3. Hierarchical Classification

So far we dealt with flat classification, the classes being independent *a priori*, with block-diagonal kernel matrix K . However, if the label set has a known structure,⁴ we can benefit from representing it in the model. Here we focus on hierarchical classification, the label set $\{1, \dots, C\}$ being the leaf nodes of a tree. Classes with lower common ancestor should be more closely related. In this section, we propose a model for this setup and show how it can be dealt with in our framework with minor modifications and reasonable extra cost.

In flat classification, the latent class functions $u_c(\cdot)$ are modelled as *a priori* independent, in that the penaliser (or the log prior in the GP view) is a sum of individual terms for each c , without

4. Learning an unknown label set structure may be achieved by expectation maximization techniques, but this is subject to future work.




Figure 2: Example of a tree-structured target space, where labels correspond to leaf nodes (shaded).

interaction terms. Analysis of variance (ANOVA) models go beyond such independent designs, they have previously been applied to text classification by Cai and Hofmann (2004), see also Shahbaba and Neal (2007). Let $\{0, \dots, P\}$ be the nodes of the tree, 0 being the root, and the numbers are assigned breadth first (1, 2, ... are the root's children). The tree is determined by P and n_p , $p = 0, \dots, P$, the number of children of node p . Let L be the set of leaf nodes, $|L| = C$. Assign a pair of latent functions u_p, \check{u}_p to each node, except the root. The \check{u}_p are assumed *a priori* independent, as in flat classification. u_p is the sum of $\check{u}_{p'}$, where p' is running over the nodes (including p) on the path from the root to p . An example is given in Figure 2. The class functions to be fed into the classification likelihood are the $u_{L(c)}$ of the leaves. This setup represents similarities according to the hierarchy. For example, if leaves $L(c), L(c')$ have the common parent p , then $u_{L(c)} = u_p + \check{u}_{L(c)}$, $u_{L(c')} = u_p + \check{u}_{L(c')}$, so the class functions *share* the effect u_p . Since regularisation forces all independent effects $\check{u}_{p'}$ to be smooth, the classes c, c' are urged to behave similarly *a priori*.

Let $u = (u_p(x_i))_{i,p}$, $\check{u} = (\check{u}_p(x_i))_{i,p} \in \mathbb{R}^{nP}$. The vectors are linearly related as $u = (\Phi \otimes I)\check{u}$, $\Phi \in \{0, 1\}^{P,P}$, a special case of the mixing of Figure 1. Importantly, Φ has a simple structure which allows MVM with Φ or Φ^T to be computed easily in $O(P)$, without having to compute or store Φ explicitly. Let $cs_p = \sum_{p' < p} n_{p'}$, and define $\Phi_p \in \mathbb{R}^{d,d}$, $d = cs_p + n_p$, to be the upper left block of Φ , so that $\Phi = \Phi_p$. If p is a leaf node, then $\Phi_p = \Phi_{p-1}$. Otherwise, Φ_p is obtained from Φ_{p-1} by attaching rows $(\delta_p^T \Phi_{p-1}, \delta_j^T)$, $j = 1, \dots, n_p$, where $\delta_p^T \Phi_{p-1}$ is the p -th row of Φ_{p-1} . This is because $u_{cs_p+j} = u_p + \check{u}_{cs_p+j}$ for the functions of the children of p . Formally,

$$\Phi_p = \begin{pmatrix} \Phi_{p-1} & 0 \\ \delta_p^T \Phi_{p-1} & I \end{pmatrix},$$

where the lower right $I \in \mathbb{R}^{n_p, n_p}$. Note that Φ is lower triangular with $\text{diag } \Phi = I$. This recursive definition directly implies simple methods for computing $v \mapsto \Phi v$ and $v \mapsto \Phi^T v$.

Under the hierarchical model, the class functions $u_{L(c)}$ are strongly dependent *a priori*. Representing this prior coupling in our framework amounts to simply plugging in the implied kernel matrix

$$K = (\Phi_{L,\cdot} \otimes I) \check{K} (\Phi_{L,\cdot}^T \otimes I), \quad (2)$$

into the flat classification model of Section 2. Here, the inner \tilde{K} is block-diagonal, while in the flat model, K itself had this property. In the hierarchical case, K is not sparse and certainly not block-diagonal, but we are still able to compute kernel MVMs efficiently: pre- and post-multiplying by Φ is very cheap, and \tilde{K} is block-diagonal just as in the flat case.

In fact, the step from flat to hierarchical classification requires minor modifications of existing code only. If code for representing a block-diagonal K is available, we can use it to represent the inner \tilde{K} , just replacing C by P . This simplicity carries through to the hyperparameter learning method (see Section 4). The cost of a kernel MVM is increased⁵ by a factor $P/C < 2$, which in most hierarchies in practice is close to 1.

However, it would be wrong to claim that hierarchical classification in general comes as cheap as flat classification. In fact, primary fitting becomes more costly, precisely because there is more coupling between the variables. In the flat case, the Hessian of Φ (1) is close to block-diagonal. The LCG algorithm to compute Newton directions converges quickly, because it nearly decomposes into C independent ones, and fewer NR steps are required. In the hierarchical case, this “near-decomposition” does not hold, and both LCG and NR need more iterations to attain the same accuracy, although each LCG step comes at about the same cost as in the flat case.

In numerical mathematics, much work has been done to approximately decouple linear systems by preconditioning. In some of these strategies, knowledge about the structure of the system matrix (in our case: the hierarchy) can be used to drive preconditioning. An important point for future research is to find a good preconditioning strategy for the system (5). However, in all our experiments so far the fitting of the hierarchical model took less than twice the time required for the flat model on the same task.

4. Hyperparameter Learning

Our framework comes with an automatic method for setting free hyperparameters h , by gradient-based maximization of the cross-validation (CV) log likelihood. Our primary fitting method of Section 2 is used here as principal subroutine. Such a setup is commonplace in Bayesian statistics, where (marginal) inference is typically employed as subroutine in parameter learning.

Recall that primary fitting works by minimizing Φ (1) w.r.t. α . Let $\{J_k\}$ be a partition of the data set range $\{1, \dots, n\}$, with $J_k = \{1, \dots, n\} \setminus I_k$, and let

$$\Phi_{J_k} = u_{[J_k]}^T ((1/2)\alpha_{[J_k]} - y_{J_k}) + 1^T l_{[J_k]}$$

be the negative log likelihood of the subset J_k of the data. Here, $u_{[J_k]} = \tilde{K}_{J_k} \alpha_{[J_k]}$. The $\alpha_{[J_k]}$ are independent variables, *not* part of a common⁶ α . The cross-validation criterion is

$$\Psi = \sum_k \Psi_{I_k}, \quad \Psi_{I_k} = -y_{I_k}^T u_{[I_k]} + 1^T l_{[I_k]}, \quad u_{[I_k]} = \tilde{K}_{I_k, J_k} \alpha_{[J_k]}, \quad (3)$$

where $\alpha_{[J_k]}$ is the minimizer of Φ_{J_k} . Since for each k , we fit and evaluate the likelihood on disjoint parts of y , Ψ is an unbiased estimator of the true negative expected log likelihood.

In order to adjust h , we pick a fixed partition at random, then do gradient-based minimization of Ψ w.r.t. h . To this end, we maintain the set $\{\alpha_{[J_k]}\}$ of primary variables, and iterate between re-fitting

5. Nodes with a single child only can be pruned from the hierarchy. Note that our formalism does not require all leaf nodes to have the same depth.

6. Which is why they are not referred to as α_{J_k} .

those for each fold k , and computing Ψ and $\nabla_h \Psi$. The gradient can be determined analytically, using a computation which is equivalent to the Newton direction computations for $\alpha_{[J_k]}$, meaning that the same code can be used. Details are given in Section 5.2. Note that Ψ is not a convex objective.

As for computational complexity, suppose there are q folds. The update of the $\alpha_{[J_k]}$ requires q primary fitting applications, but since they are initialized with the previous values $\alpha_{[J_k]}$, they do converge very rapidly, especially during later iterations. Computing Ψ based on the $\alpha_{[J_k]}$ comes basically for free. The gradient computation decomposes into two parts: accumulation, and kernel derivative MVMs. The accumulation part requires solving q systems of size $((q-1)/q)nC$, thus qk_3 kernel MVMs on the \tilde{K}_{J_k} if linear conjugate gradients (LCG) is used, k_3 being the number of LCG steps. We also need two buffer matrices E, F of qnC elements each. Note that the accumulation step is *independent* of the number of hyperparameters. The second part consists of q kernel derivative MVMs for each independent component of h . This second part is much simpler than the accumulation one, consisting entirely of large matrix operations, which can be run very efficiently using specialized numerical linear algebra code. The method for computing Ψ and $\nabla_h \Psi$ can be plugged into a custom gradient-based optimizer, such as Quasi-Newton or conjugate gradients, in order to learn h .

As shown in Section 5.3, the extension of hyperparameter learning to the hierarchical case of Section 3 is done by wrapping the accumulation part with Φ MVMs, the coding and additional memory effort being minimal.

We finally note from our findings in practice (see Section 6.3) that on large tasks, our automatic method can require some fine-tuning. This is due to the delicate dependencies between the different approximations used. The accuracy of Ψ and $\nabla_h \Psi$ depends on how accurate the inner NR optimizations for $\alpha_{[J_k]}$ turn out, and the latter depend on how many iterations of LCG are done in order to compute search directions. Fortunately, Φ_{J_k} and its gradient w.r.t. $u_{[J_k]}$ can be computed exactly in order to assess inner optimization convergence, so we do at least know when things go wrong. In our implementation, we deem an evaluation of Ψ and $\nabla_h \Psi$ usable if the average of $\|\nabla_{u_{[J_k]}} \Phi_{J_k}\|$ over folds is below a threshold, which depends on the problem and on time constraints. A failed evaluation leads to a right bracket there for the outer optimization line search, in that step sizes beyond the failed one are not accessed. We can now tune the basic running time parameters k_1, k_2 so that Ψ evaluations do not fail too often. In this context, it is important to regard the $\{\alpha_{[J_k]}\}$ as an *inner state* alongside the hyperparameter vector h . Although inner optimizations are convex, for large problems and reasonable k_1, k_2 , successive minima are attained only when we start from the previous best inner state. This is true especially during later stages, where for certain problems (see Section 6.3) h attends “extreme” values and the inner optimizations become quite hard.⁷ Therefore, the inner state used to initialize a given Ψ evaluation is the final one for the last recent *successful* evaluation.⁸ Inner states attained during failed evaluations are discarded.

7. Although inner optimizations are convex, speed of convergence of NR depends strongly on the value of h . For “extreme” values, the Newton direction computation by LCG is harder, and search directions can become large in early NR iterations. The latter may be because we work in u rather than α space, but only the former is really feasible.

8. Within outer line searches, we use $\{\alpha_{[J_k]}\}$ from the last recent successful evaluation *to the left* (along the search direction).

5. Computational Details

In this section, we provide details for the material above. The techniques given here do characterize our framework, they are novel in this combination, and some of them may be useful in other contexts as well. More specific details of our implementation can be found in Appendix D.

5.1 Details for Flat Classification

In this section, we provide details for the primary fitting optimization in the case of flat multi-way classification, introduced in Section 2. Note that this fitting method appeared in Williams and Barber (1998) in the context of approximate Gaussian process inference, although some fairly essential ideas here are novel to our knowledge (symmetrisation of Newton system, pair optimization line search, numerical stability considerations).

Recall that we want to minimize the strictly convex criterion Φ (1) w.r.t. α , using the Newton-Raphson (NR) method. Modern variants of this algorithm iterate line searches along the *Newton directions* $-H^{-1}g$, where g, H are gradient and Hessian of Φ at the current α . We will start with the Newton direction computation in Section 5.1.1, commenting on the line searches afterwards in Section 5.1.2 (it turns out that it basically comes for free). An overview of the fitting algorithm is given in Section 5.1.3.

5.1.1 COMPUTING THE NEWTON DIRECTION

Recall Φ and related variables from (1). Let $\pi_{ic} = P(y_{ic} = 1|u_i)$, $\pi = \exp(u - 1 \otimes l)$, and recall that Φ_{lh} is the likelihood part in Φ . Now,

$$g := \nabla \Phi_{lh} = \pi - y, \quad W := \nabla \nabla \Phi_{lh} = D - DP_{cls}D, \quad P_{cls} = (1 \otimes I)(1^T \otimes I).$$

Here, $D = \text{diag } \pi$, and gradient and Hessian are taken w.r.t. u (*not* w.r.t. α). Our convention for nC vectors and matrices and the use of \otimes is explained in Appendix A. The form of W can be understood by noting that W is block-diagonal in a *different* ordering, which uses c (classes) as inner and i (data points) as outer index, then switching to our standard ordering.

It is easy to compute gradient and Hessian of Φ w.r.t. α, b . A full (classical) Newton step is given by the system

$$\begin{aligned} (I + WK)\alpha' + W(I \otimes 1)b' &= Wu - g, \\ (I \otimes 1^T)WK\alpha' + (I \otimes 1^T)W(I \otimes 1)b' + \sigma^{-2}b' &= (I \otimes 1^T)(Wu - g), \end{aligned}$$

and the Newton search direction is obtained as the difference $\alpha' - \alpha, b' - b$. Subtracting $(I \otimes 1^T)$ times the first from the second, we obtain $b' = \sigma^2(I \otimes 1^T)\alpha'$, and plugging this into the first equation, we have

$$(I + W(K + \sigma^2 P_{data}))\alpha' = Wu - g, \quad P_{data} = (I \otimes 1)(I \otimes 1^T). \quad (4)$$

Note that $P_{data}a = (\sum_{i'} a_{i'})_i$, which does the same as P_{cls} , but on index i rather than c . We denote

$$\tilde{K} = K + \sigma^2 P_{data},$$

noting that this corresponds to $\tilde{K}^{(c)} = K^{(c)} + \sigma^2 11^T$. The correct way of incorporating intercept parameters is to add the constant σ^2 to the kernels, then to obtain $b_c = \sigma^2 \sum_i \alpha_{ic}$. This is the meaning

of “eliminating b ” in Section 2. While we could optimize σ^2 as a hyperparameter, we consider it fixed and given for simplicity.⁹ In the sequel, we consider b being eliminated from the model by replacing $K \rightarrow \tilde{K}$ everywhere. We have $u = \tilde{K}\alpha$.

We can solve the system (4) exactly if we can tolerate a scaling of $O(n^3 C)$ and $O(n^2 C)$ memory. Note that this scaling is linear rather than cubic in C . The exact solution is derived in Appendix C. It is efficient for moderate n , and generally useful for code debugging, and is supported by our implementation. In the remainder of this section, we focus on approximate computations.

Although we could solve the system using a bi-conjugate gradients solver, we can do much better by transforming it into symmetric positive definite form. First, note that W is positive semidefinite, but singular. This can be seen by noting that the parameterization of our likelihood in terms of u_i is overcomplete, in that $u_i + \kappa 1$ gives the same likelihood values for all κ . We could fix one of the u_i components, which would however lead to subtle dependencies between the remaining $C - 1$ functions $u_c(\cdot)$. In order to justify our *a priori* independent treatment of these functions, we have to retain the overcomplete likelihood. The nullspace $\ker W$ is given by $\{(d)_c \mid d \in \mathbb{R}^n\}$ and has dimension n . This can be seen by noting that $Wa = 0$ iff $a = (\bar{a})_c$, $\bar{a} = \sum_c a^{(c)}$. W has rank $n(C - 1)$. We have $a \in \text{ran } W$ iff $\sum_c a^{(c)} = (1^T \otimes I)a = 0$ (recall that $\ker W$ and $\text{ran } W$ are orthogonal, and their direct sum is \mathbb{R}^{nC}). From (4) we see that $\alpha' + g$ lies in $\text{ran } W$. Note that $\sum_c g^{(c)} = \sum_c (\pi^{(c)} - y^{(c)}) = 1 - 1 = 0$, therefore $g \in \text{ran } W$, thus $\alpha' \in \text{ran } W$. We see that the dual coefficients must fulfill the constraint $\alpha \in \text{ran } W$. Note that $\text{ran } W$ is in fact independent of D . Whatever starting value is used for α , it should be projected onto $\text{ran } W$, which is done by subtracting $C^{-1}P_{cls}\alpha$. The NR updates then make sure that the constraint remains fulfilled.

Next, we need a decomposition $W = VV^T$ of W . Such a V exists (because W is positive semidefinite). In fact,

$$W = ADA^T, \quad A = I - DP_{cls}.$$

This follows easily from $(1^T \otimes I)D(1 \otimes I) = \sum_c D^{(c)} = I$. Thus, $W = VV^T$ with $V = AD^{1/2}$. The matrix A has fixed points $\text{ran } W$, namely if $a \in \text{ran } W$, then $(1^T \otimes I)a = 0$, so that $Aa = a$.

Since there exists some \tilde{v} (not unique) s.t. $\alpha' = W\tilde{v}$, we can rewrite the system (4) as

$$V(I + V^T \tilde{K}V)V^T \tilde{v} = V(V^T u - \tilde{g}),$$

where \tilde{g} is s.t. $g = V\tilde{g}$ (such a vector exists because $g \in \text{ran } W$). This suggests the following procedure for finding α' :

$$(I + V^T \tilde{K}V)\beta = V^T u - \tilde{g}, \quad \alpha' = V\beta. \tag{5}$$

To see the validity of this approach, simply multiply both sides of (5) by V from the left, which shows that $V\beta$ solves the original system. Since the latter has a unique solution (strict convexity!), we must have $V\beta = \alpha'$. Finally, we note that $\tilde{g} = D^{-1/2}g$ does the job, because $VD^{-1/2}g = Ag = g$. The latter follows because $g \in \text{ran } W$.

Thus, in exact arithmetic, the Newton direction computation is implemented in a three-stage procedure. First, compute $\tilde{g} = D^{-1/2}g$. Second, solve the system (5) for β . This is a symmetric positive definite system with the typically well-conditioned matrix $I + V^T \tilde{K}V$, and can be solved efficiently using the linear conjugate gradients (LCG) algorithm (Saad, 1996). The cost of each step is dominated by the MVM $v \mapsto Kv$, which scales linearly in C , due to the block-diagonal structure of K . Third, set $\alpha' = V\beta$. The Newton direction is obtained as $\alpha' - \alpha$.

9. In our experience so far, a good value of σ^2 is fairly robust across different tasks for the same problem, but may differ strongly between different problems. It can be chosen based on some initial experiments.

We can start the LCG run from a good guess as follows. Let α be the current dual vector which solved the last recent system. We would like to initialize β s.t. $\alpha = V\beta = AD^{1/2}\beta$. If we assume that $D^{1/2}\beta \in \text{ran}W$, then $\alpha = D^{1/2}\beta$. Therefore, a good initialization is $\beta = D^{-1/2}\alpha$. Alternatively, we may also retain β from the last recent system.

Issues of numerical stability are addressed in Appendix B. Furthermore, the LCG algorithm is hardly ever run without some sort of preconditioning. Our present implementation uses diagonal preconditioning, as described in Appendix B. We have already noted in Section 3 that a non-diagonal preconditioning strategy could be valuable, but this is subject to future work.

5.1.2 THE LINE SEARCH

The classical NR algorithm proceeds doing full steps $\alpha \rightarrow \alpha'$, but modern variants typically employ a line search along the Newton direction $\alpha' - \alpha$. In the non-convex case, this ensures global convergence, and even for our convex objective Φ , a line search saves time and leads to numerically more stable behaviour. Interestingly, the special structure of our problem leads to the fact that line searches essentially come for free, certainly compared with the effort of obtaining Newton directions. We refer to this simple idea as *pair optimization*, the reader may be reminded of similar tricks in primal-dual schemes for SVM.

Let $s = \alpha' - \alpha$ be the NR direction, computed as shown above, and set α_0 to α . The line search minimizes (or sufficiently decreases) Φ on the line segment $\alpha_0 + \lambda s$, $\lambda \in (0, 1]$, starting with $\lambda = 1$ (which is the classical Newton step). The idea is to treat Φ as a function of the pair (u, α) , where $u = \tilde{K}\alpha$. The corresponding line segment is $u = u_0 + \lambda \tilde{s}$, $\tilde{s} = \tilde{K}s$, requiring a single kernel MVM for computing \tilde{s} . Let $j = \text{argmax} |\tilde{s}_j|$. For an evaluation of Φ at u , we reconstruct $\lambda = (u_j - u_{0,j})/\tilde{s}_j$ and $\alpha = \alpha_0 + \lambda s$, then

$$\Phi = u^T ((1/2)\alpha - y) + 1^T l, \quad \nabla\Phi = \pi - y + \alpha,$$

so that an evaluation comes at the cost $O(nC)$ and does not require additional kernel MVM applications. We now do the line minimization of Φ in the variable u . The driving feature of pair optimization is that we can go back and forth between α and u without significant cost, once the search direction is known w.r.t. both variables.

5.1.3 OVERVIEW OF THE OPTIMIZATION ALGORITHM

In Algorithm 1, we give a schematic overview of the primary fitting algorithm, written in terms of a MVM primitive $v \mapsto Kv$. For simplicity, we do not include the measures discussed in Appendix B to increase numerical stability.

5.2 Details for Hyperparameter Learning

In this section, we provide details for the CV hyperparameter learning scheme, introduced in Section 4. The gradient of the CV criterion Ψ (3) is computed as follows. Ψ is a sum of terms Ψ_{I_k} , one for each fold. We focus on a single term and write $I = I_k$, $J = J_k$. $\alpha_{[J]}$ is determined by the stationary equation $\alpha_{[J]} + g_{[J]} = 0$ (all terms of subscript $[J]$ are as in Section 5.1.1, but for the subset J of the data, and w.r.t. $\alpha_{[J]}$). Taking derivatives gives

$$d\alpha_{[J]} = -W_{[J]} ((dK_J)\alpha_{[J]} + \tilde{K}_J(d\alpha_{[J]})),$$

Algorithm 1 Newton-Raphson optimization to find posterior mode $\hat{\alpha}$.

Require: Starting values for α, b . Targets y .

$\alpha = \alpha - C^{-1}(\sum_c \alpha^{(c)})_c$, so that $\alpha \in \text{ran } W$. $u = \tilde{K}\alpha$.

repeat

 Compute $l, \log(\pi)$ from u . Compute Φ .

if relative improvement in Φ small enough **then**

 Terminate outer loop.

else if maximum number of iterations done **then**

 Terminate outer loop.

end if

 Initialize $\beta = D^{-1/2}\alpha$. Compute r.h.s. $r = V^T u - \tilde{g}$, $\tilde{g} = D^{-1/2}g$.

 Compute preconditioner $\text{diag}(I + V^T \tilde{K}V)$.

 Run preconditioned CG algorithm in order to solve the system (5) approximately. The CG code is configured by a primitive to compute $v \mapsto (I + V^T \tilde{K}V)v$, which in turn calls the primitive for $v \mapsto Kv$.

 Compute $\alpha' = AD^{1/2}\beta'$.

 Do line search along $s = \alpha' - \alpha$. This is done in u , along $\tilde{s} = \tilde{K}s$.

 Assign line minimizer to α, u .

until forever

since $dg_{[j]} = W_{[j]} du_{[j]}$. We obtain a system for $d\alpha_{[j]}$ which is symmetrised as in Section 5.1.1:

$$\left(I + V_{[j]}^T \tilde{K}_J V_{[j]}\right) \beta = -V_{[j]}^T (dK_J) \alpha_{[j]}, \quad d\alpha_{[j]} = V_{[j]} \beta.$$

Also,

$$d\Psi_I = (\pi_{[I]} - y_I)^T \left((dK_{I,J}) \alpha_{[J]} + \tilde{K}_{I,J} (d\alpha_{[J]}) \right).$$

With

$$f = I_{.,I}(\pi_{[I]} - y_I) - I_{.,J} V_{[J]} \left(I + V_{[J]}^T \tilde{K}_J V_{[J]} \right)^{-1} V_{[J]}^T \tilde{K}_{J,I} (\pi_{[I]} - y_I),$$

we have that $d\Psi_I = (I_{.,J} \alpha_{[J]})^T (dK) f$.

If we collect these vectors as columns of $E, F \in \mathbb{R}^{n_C, q}$, q the number of folds, we have that

$$d\Psi = \text{tr } E^T (dK) F \tag{6}$$

for the complete criterion. The computation of E, F was called ‘‘accumulation’’ in Section 4. It involves a loop over folds, in which $\alpha_{[j_k]}$ is determined by NR optimization, starting from its previous value, then f (column of F) is computed by solving one more system of the same form as is required to compute Newton directions. Importantly, this accumulation phase is independent of the number of hyperparameters. The gradient computation then requires to compute (6) for each component, using kernel derivative MVMs. First of all, $\partial K / \partial h_p$ is block-diagonal just as K , and for many standard kernels, it is a simple expression, involving K itself (see Section 7.3), so one may be able to share computations between the different gradient components. Importantly, the computation of (6) is easily broken down into large numerical linear algebra primitives, for which very efficient code may be used (see Section 7.2). This is a significant advantage in the presence of

many hyperparameters. For moderately many hyperparameters, the accumulation clearly dominates the CV criterion and gradient computation.

The dominating part of the accumulation is the re-optimisation of the $\alpha_{[j]}$, which are done by calling the optimized code for primary fitting (Section 5.1) as subroutine. Here, a feature of our implementation becomes important. Instead of representing each K_{J_k} separately, we represent the full K only for all subset kernel MVMs. The representation depends on the covariance function, and in general on how kernel MVMs are actually done. A generic representation is described in Appendix D.1. In order to work on the data subset J_k , we *shuffle* the representation such that in the permuted kernel matrix, K_{J_k} forms the upper left corner. This means that linear algebra primitives with K_{J_k} can be run without mapping matrix coordinates through an index, which would be many times slower. Details on ‘‘covariance shuffling’’ are given in Appendix D.2.

As mentioned in Section 5.1.1 and detailed in Appendix C, we can also compute Newton directions exactly in $O(Cn^3)$ in the flat classification case. This exact treatment can be extended to the computation of Ψ and its gradient, as is shown in Appendix C. Exact computations lead to more robust behaviour, and may actually run faster for small to moderate n . Exact computations are also useful for debugging purposes.

5.3 Details for Hierarchical Classification

In this section, we provide details for hierarchical classification method, introduced in Section 3. Recall that $u = (\Phi \otimes I)\check{u}$ for an indicator matrix Φ of simple structure, and that MVM with Φ or Φ^T can be computed easily in $O(P)$, without having to store Φ . Since the $\check{u}_p(\cdot)$ are given independent priors (or regularisers), the corresponding kernel matrix \check{K} is block-diagonal. The induced covariance matrix K over u_L is given by (2), and hierarchical classification differs from the flat variant only in that this non-block-diagonal matrix is used.

The MVM primitive $v \mapsto Kv$ is computed in three steps. MVM with $(\Phi_{L,\cdot} \otimes I)$ and $(\Phi_{L,\cdot}^T \otimes I)$ works by computing $S \mapsto S\Phi$, $S \mapsto S\Phi^T$ for $S \in \mathbb{R}^{n,P}$. In between, MVM with \check{K} has to be done in the same way as for flat classification, only that \check{K} has P rather than C diagonal blocks.

The diagonal preconditioning of LCG (see Appendix B) requires the computation of $\text{diag} K \in \mathbb{R}^{nC}$. We have

$$K_{ic,ic} = (\delta_p^T \Phi \otimes \delta_i^T) \check{K} (\Phi^T \delta_p \otimes \delta_i) = \delta_p^T \Phi (\text{diag}(\check{K}_i^{(p')})_{p'}) \Phi^T \delta_p, \quad p = L(c),$$

where $\delta_p^T \Phi$ is the p -th row of Φ . From the recursive structure of Φ we know that if $n_p > 0$, then $\delta_{cs_p+j}^T \Phi = \delta_p^T \Phi + \delta_{cs_p+j}^T$, $j = 1, \dots, n_p$, so if

$$d_{i(cs_p+j)} = d_{ip} + \check{K}_i^{(cs_p+j)}, \quad j = 1, \dots, n_p,$$

then $\text{diag} K = d_L$.

Hyperparameter learning (see Section 5.2) is easily extended to the hierarchical case, recalling (2) and the fact that Φ does not depend on hyperparameters. Define $\tilde{E} = (\Phi_{L,\cdot}^T \otimes I)E \in \mathbb{R}^{n^p,q}$, \tilde{F} accordingly, with E, F given in Section 5.2. The gradient components (6) translate to $\text{tr} \tilde{E}^T (d\check{K}) \tilde{F}$, where \check{K} is block-diagonal as before. In our implementation, we reserve buffer space for \tilde{E}, \tilde{F} , yet build E, F there during accumulation. We then transform them to \tilde{E}, \tilde{F} using in-place computations.

The step from flat to hierarchical classification requires only minor modifications of existing code. Wrappers for MVM and the other primitives essentially pre- and post-multiply their input

with Φ and Φ^T respectively, calling the existing “flat” primitives for \tilde{K} in between (block-diagonal with P rather than C blocks).

5.4 Why Newton Raphson?

Why do we propose to use the second-order NR method for minimizing Φ , instead of using a simpler gradient-based technique such as scaled conjugate gradients (SCG)? We already motivated our choice at the end of Section 2, but give more details concerning this important point here.

The convex problems we are interested in here live in very high-dimensional spaces and come with complicated couplings between the components which cannot be characterized simply. Certainly, there is no simple decomposition into parts. It is well known in the optimization literature (Bertsekas, 1999) that simple gradient-based techniques such as SCG require well-chosen preconditioning in order to work effectively in such cases.

For example, we could optimize $\Phi(1)$ w.r.t. α directly, the gradient requires a single MVM with K rather than solving a system. However, this problem is very ill-conditioned, the Hessian being $\tilde{K}W\tilde{K} + \tilde{K}$ (large kernel matrices are typically very ill-conditioned, and here we deal with K^2), and SCG runs exceedingly slowly to the point of being essentially useless (as we determined in experiments). It can be saved (to our knowledge) only by preconditioning, which in our case requires to solve a system again. Another idea is to optimize Φ w.r.t. u by SCG, which works better. The Hessian is $W + \tilde{K}^{-1}$, whose condition number is similar to K . In preliminary direct comparisons, the NR method still works more efficiently, meaning that SCG would require additional preconditioning specific to the problem at hand, which would likely be different for flat and hierarchical classification. From our experience, and also from the predominance of NR in the optimization literature, we opted for this method which comes with self-tuning capabilities, making it easier to transfer the framework to novel problems.

6. Experiments

In this section, we provide experimental results for our method on a range of flat and hierarchical classification tasks.

6.1 Hierarchical Classification: Patent Texts

We use the WIPO-alpha collection,¹⁰ many thanks to L. Cai, T. Hofmann for providing us with the count data and dictionary. We did Porter stemming, stop word removal, and removal of empty categories. The attributes are bag-of-words over the dictionary. All input vectors x_i were scaled to unit norm. Many thanks to Peter Gehler for helping us with the preprocessing.

These tasks have previously been studied by Cai and Hofmann (2004), where patents (title and claim text) are to be classified w.r.t. the standard taxonomy *IPC*, a tree with 4 levels and 5229 nodes. Sections A, B, . . . , H form the first level. As in Cai and Hofmann (2004), we concentrate on the 8 subtasks rooted at the sections, ranging in size from D ($n = 1140, C = 160, P = 187$) to B ($n = 9794, C = 1172, P = 1319$). We use linear kernels (see Appendix D.3) with variance parameters v_c .

All experiments are averaged over three training/test splits, different methods using the same ones. The CV criterion Ψ is used with a different (randomly drawn) 5-partition per section and

10. Available at www.wipo.int/tools/en/dbindex.html, or google for “Data Collections hosted by WIPO”.

split, the same across all methods. Our method outputs a predictive distribution $p_j \in \mathbb{R}^C$ for each test case x_j . The standard prediction $y(x_j) = \operatorname{argmax}_c p_{jc}$ maximizes expected accuracy, classes are ranked as $r_j(c) \leq r_j(c')$ iff $p_{jc} \geq p_{jc'}$, where $r_j(c) \in \{1, \dots, C\}$ is the rank of class c for case x_j . Let y_j denote the true label for x_j . The test scores we use here are the same as in Cai and Hofmann (2004): *accuracy* (acc) $m^{-1} \sum_j \mathbf{1}_{\{y(x_j)=y_j\}}$, *precision* (prec) $m^{-1} \sum_j r_j(y_j)^{-1}$, *parent accuracy* (pacc) $m^{-1} \sum_j \mathbf{1}_{\{\operatorname{par}(y(x_j))=\operatorname{par}(y_j)\}}$, $\operatorname{par}(c)$ being the parent of leaf node $L(c)$ (recall that $L(c)$ corresponds to class c). Here, m is the test set size. Let $\Delta(c, c')$ be half the length of the shortest path between leafs $L(c), L(c')$. The *taxo-loss* (taxo) is $m^{-1} \sum_j \Delta(y(x_j), y_j)$. These scores are motivated in Cai and Hofmann (2004). For taxo-loss and parent accuracy, we better choose $y(x_j)$ to minimize expected loss,¹¹ which is different in general than the standard prediction (the latter maximizes expected accuracy and precision).

We compare methods F1, F2, H1, H2 (F: flat, not using IPC; H: hierarchical). F1: all v_c shared (1); H1: v_c shared across each level of the tree (3). F2, H2: v_c shared across each subtree rooted at root’s children (A: 15, B: 34, C: 17, D: 7, E: 7, F: 17, G: 12, H: 5). The numbers in parentheses are the total number of hyperparameters. Recall that there are three parameters determining the running time (see Section 2, Section 4). For hyperparameter learning: $k_1 = 8, k_2 = 4, k_3 = 15$ (F1, F2); $k_1 = 10, k_2 = 4, k_3 = 25$ (H1, H2).¹² For the final fitting (after hyperpars have been learned): $k_1 = 25, k_2 = 12$ (F1, F2); $k_1 = 30, k_2 = 17$ (H1, H2). The optimization is started from $v_c = 5$ for all methods. We set $\sigma^2 = 0.01$ throughout. Results are given in Table 1.

The hierarchical model outperforms the flat one consistently, especially w.r.t. taxo-loss and parent accuracy. Also, minimizing expected loss is consistently better than using the standard rule for the latter, although the differences are not significant. H1 and H2 do not perform differently: choosing many different v_c in the linear kernel seems no advantage here (but see Section 6.2). The results are quite similar to the ones of Cai and Hofmann (2004), obtained with a support vector machine variant. However, for our method, the recommendation in Cai and Hofmann (2004) to use $v_c = 1$ (not further motivated there) leads to significantly worse results in all scores. The v_c chosen by our method are generally larger. Note that their code has not been made publicly available, so a direct comparison with “all other things equal” could not be done.

In Table 2, we present running times¹³ for the final fitting and for a single fold during hyperparameter optimization (5 of these are required for $\Psi, \nabla_h \Psi$). In comparison, a final fitting time of 2200s on the D section is quoted in Cai and Hofmann (2004), using a SVM variant, while we require 119s (more than six times faster).¹⁴ It is precisely this high efficiency of primary fitting, which allows us to use it as inner loop for automatic hyperparameter learning (Cai and Hofmann, 2004, do not adjust hyperparameters to the data). Possible reasons for the performance difference are given in Section 7.2.

11. For parent accuracy, let $p(j)$ be the node with maximal mass (under p_j) of its children which are leafs, then $y(x_j)$ must be a child of $p(j)$.

12. Except for section C, where $k_1 = 14, k_2 = 6, k_3 = 35$.

13. Processor time on 64bit 2.33GHz AMD machines.

14. Cai and Hofmann average over three training/test splits. The timing figure 2200s in their paper is for three splits (thanks to one of the reviewers to point this out).

	acc (%)				prec (%)				taxo			
	F1	H1	F2	H2	F1	H1	F2	H2	F1	H1	F2	H2
A	40.6	41.9	40.5	41.9	51.6	53.4	51.4	53.4	1.27	1.19	1.29	1.19
B	32.0	32.9	31.7	32.7	41.8	43.8	41.6	43.7	1.52	1.44	1.55	1.44
C	33.7	34.7	34.1	34.5	45.2	46.6	45.4	46.4	1.34	1.26	1.35	1.27
D	40.0	40.6	39.7	40.8	52.4	54.1	52.2	54.3	1.19	1.11	1.18	1.11
E	33.0	34.2	32.8	34.1	45.1	47.1	45.0	47.1	1.39	1.31	1.38	1.31
F	31.4	32.4	31.4	32.5	42.8	44.9	42.8	45.0	1.43	1.34	1.43	1.34
G	40.1	40.7	40.2	40.7	51.2	52.5	51.3	52.5	1.32	1.26	1.32	1.26
H	39.3	39.6	39.4	39.7	52.4	53.3	52.5	53.4	1.17	1.15	1.17	1.14

	taxo[0-1]				pacc (%)				pacc[0-1] (%)			
	F1	H1	F2	H2	F1	H1	F2	H2	F1	H1	F2	H2
A	1.28	1.19	1.29	1.18	58.9	61.6	58.2	61.5	57.2	61.3	56.9	61.4
B	1.54	1.44	1.56	1.44	53.6	56.4	52.7	56.6	51.9	55.9	51.4	55.9
C	1.33	1.26	1.32	1.26	58.9	62.6	58.5	62.0	58.6	61.8	58.9	61.6
D	1.20	1.12	1.22	1.12	64.6	67.0	64.4	67.1	63.5	67.1	62.6	67.0
E	1.43	1.33	1.44	1.34	56.0	59.1	56.2	59.2	54.0	58.2	53.5	57.9
F	1.43	1.34	1.44	1.34	56.8	59.7	56.8	59.8	54.9	58.7	54.6	58.9
G	1.32	1.26	1.32	1.26	58.0	59.7	57.6	59.6	56.8	59.2	56.6	58.9
H	1.19	1.16	1.19	1.15	61.6	62.5	61.8	62.5	59.9	61.6	60.0	61.8

Table 1: Results on patent text classification tasks A-H. Methods F1, F2 flat, H1, H2 hierarchical. taxo[0-1], pacc[0-1] for $\operatorname{argmax}_c p_{jc}$ standard prediction rule, rather than minimization of expected loss.

	Final NR (s)		CV Fold (s)			Final NR (s)		CV Fold (s)	
	F1	H1	F1	H1		F1	H1	F1	H1
A	2030	3873	573	598	E	131.5	203.4	32.2	49.6
B	3751	8657	873	1720	F	1202	2871	426	568
C	4237	7422	719	1326	G	1342	2947	232	579
D	56.3	118.5	9.32	20.2	H	971.7	1052	146	230

Table 2: Running times for tasks A-H. Method F1 flat, H1 hierarchical. Final NR: Final fitting with Newton-Raphson. CV Fold: Re-optimization of $\alpha_{[J]}$ and gradient accumulation for single fold J .

6.2 Flat Classification: Remote Sensing

We use the *satimage* remote sensing task from the *statlog* repository.¹⁵ This task has been used in the extensive SVM multi-class study of Hsu and Lin (2002), where it is among the data sets on which the different methods show the most variance. It has $n = 4435$ training, $m = 2000$ test cases, and $C = 6$ classes. Covariates x have 36 attributes with values in $\{0, \dots, 255\}$. No preprocessing was done.

We use the isotropic *Gaussian* (RBF) covariance function

$$K^{(c)}(x, x') = v_c \exp\left(-\frac{w_c}{2} \|x - x'\|^2\right), \quad v_c, w_c > 0. \quad (7)$$

We compare the methods *mc-sep* (ours with separate kernels for each class; 12 hyperparameters), *mc-tied* (ours with a single shared kernel; 2 hyperparameters), *mc-semi* (ours with single kernel $M^{(1)}$, but different v_c ; 7 hyperparameters), *Irest* (one-against-rest; 12 hyperparameters). For *Irest*, C binary classifiers are fitted on the tasks of separating class c from all others. They are combined afterwards by the rule $x_* \mapsto \operatorname{argmax}_c \hat{P}_c(+1|x_*)$, where $\hat{P}_c(+1|x_*)$ is the predictive probability estimate of the c -classifier.¹⁶ Note that *Irest* is arguably the most efficient method, in that its binary classifiers can be fitted separately and in parallel. Even if run sequentially, *Irest* typically requires less memory by a factor of C than a joint multi-class method, although this is not true if the kernel matrices are dominating the memory requirements and they are shared between classes in a multi-class method (as in *mc-tied* and *mc-semi* here).

We use our 5-fold CV criterion Ψ for each method. Results here are averaged over ten randomly drawn 5-partitions of the training set (the same partitions are used for the different methods). All optimizations are started from $v_c = 10$, $w_c = (\sum_j \operatorname{Var}[x_j])^{-1} = 0.017$, $\operatorname{Var}[x_j]$ being the empirical variance of attribute j . We set $\sigma^2 = 16$ throughout. The parameters determining the running time (see Section 2, Section 4) are set to $k_1 = 13$, $k_2 = 25$, $k_3 = 40$ during hyperparameter learning, and $k_1 = 30$, $k_2 = 50$ for final fitting (these are very conservative settings). Error-reject curves are shown in Figure 3.

Test errors are 7.95% ($\pm 0.15\%$) for *mc-sep*, 8.00% ($\pm 0.10\%$) for *Irest*, 8.10% ($\pm 0.13\%$) for *mc-semi*, and 8.35% ($\pm 0.20\%$) for *mc-tied*. Therefore, using a single fixed kernel for all $K^{(c)}$ does significantly worse than allowing for an individual $K^{(c)}$ per class. The test error difference between *mc-sep* and *Irest* is not significant here, but the error-reject curve is significantly better for our method *mc-sep* than for one-against-rest, especially in the domain $\alpha \in [0.025, 0.25]$, arguably most important in practice (where the rejection of a small fraction of test cases may often be an option). This indicates that the predictive probability estimates from our method are better than from one-against-rest, at least w.r.t. their ranking property. The curves for *mc-semi*, *mc-tied* are closer to *Irest*, underlining that different kernels $K^{(c)}$ should be used for each class. The result for *mc-sep* is state-of-the-art. The best SVM technique tested in Hsu and Lin (2002) attained 7.65% (no error-reject curves were given there), and SVM one-against-rest attained 8.3% in this study. To put this into perspective, note that extensive hyperparameter selection by cross-validation is done in Hsu and Lin (2002), in what seems to be a quite user-intensive process, while our method is completely automatic.

15. Available at <http://www.niaad.liacc.up.pt/old/statlog/>.

16. Asymptotically, $\hat{P}_c(+1|x_*)$ converges to the true $P(y_* = c|x_*)$, and this combination rule is optimal. We use our method with $C = 2$ in order to implement the binary classifiers.




Figure 3: Error-reject curves (averaged over 10 runs) for different methods on the *satimage* task. Curve obtained by allowing the method to abstain from prediction on fraction α of test set, counting errors for predictions only. Depends on ranking of test points. Ranking score (over test points x_*): $\max_c \hat{P}(y_* = c | x_*)$ (*mc-sep*, *mc-semi*, *mc-tied*), $\max_c \hat{P}_c(+1 | x_*)$ (*1rest*).

6.3 Flat Classification: Handwritten Digits

We use the *USPS* handwritten digits recognition task (LeCun et al., 1989). The covariates x are 16×16 gray-scale images with values in $\{16k + 15 \mid k = 0, \dots, 30\}$. The task has $n = 7291$ training, $m = 2007$ test cases, and $C = 10$ classes. No preprocessing was done.

We use Gaussian kernels (7) once more, different ones for each class. We do not optimize the 5-fold CV criterion Ψ using the full training set, but subsets of size $n' = 2000$. Our results are averaged over five runs with different randomly drawn training subsets for hyperparameter learning, while we use the full training set for final fitting. All optimizations are started from $v_c = 10$, $w_c = (\sum_j \text{Var}[x_j])^{-1} = 0.0166$, and we set $\sigma^2 = 4$ throughout. The parameters determining the running time are $k_1 = 25$, $k_2 = 35$, $k_3 = 40$ during hyperparameter learning (on $n' = 2000$ points), and $k_1 = 45$, $k_2 = 80$ for final fitting (on $n = 7291$ points). The settings for hyperparameter learning are quite conservative, and the final fitting ones were sufficient for convergence on three of the five runs, whereas on two we had to add another $k_1 = 25$ iterations with $k_2 = 90$. An error-reject curve is shown in Figure 4.




Figure 4: Error-reject curves (averaged over 5 runs) for different methods on the *USPS* task. Curve obtained by allowing the method to abstain from prediction on fraction α of test set, counting errors for predictions only. Depends on ranking of test points.

Test errors are $4.77\% (\pm 0.18\%)$. These results are state-of-the-art for kernel classification. Seeger (2003) reports 4.98% for the IVM (Sect. 4.8.4), where hyperparameters are learned automatically. Csató (2002) states 5.15% for his sparse online method with multiple sweeps over the data (Sect. 5.2). Results for the support vector machine are given in Schölkopf and Smola (2002), Table 18.1, method SV-254, where a combination heuristic based on kernel PCA was used to attain a test error of 4.4% . Crammer and Singer (2001) quote a test error of 4.38% , kernel parameters having been selected by 5-fold cross-validation. All these used the Gaussian kernel as well. The latter studies do not quote fluctuations w.r.t. choices such as the fold partition in CV, which is not negligible in our case here. The SVM-based methods do not attempt test set rankings or predictive probability estimation, and the corresponding studies do not show error-reject (or ROC) curves. Seeger (2003) gives an error-reject curve, which is very similar to ours here.

Note that the harder settings of k_1, k_2 for the final fitting are necessary due to the problem size, and are motivated in Section 4. There are 72910 parameters, and the hyperparameters found through optimizing Ψ spread by 3 orders of magnitude, so that the corresponding final fitting problems are computationally hard to solve without a good initialization of α (in the absence of such, we start with $\alpha = 0$). If we solve for Newton directions using too few LCG steps, the approximations often do not lead to much (or any) descent. Such “stalling” of NR line searches does happen now and then even

after $k_2 = 80$ LCG steps.¹⁷ Lessons learned from these large scale experiments are commented on in Section 4. There are delicate dependencies between k_1, k_2 and the running time to convergence, which need to be explored in large scale settings, but this was not done thoroughly here.

7. Discussion

We have presented a general framework for learning kernel-based penalized likelihood classification methods from data. A central feature of the framework is its high computational efficiency, even though all classes are treated jointly. This is achieved by employing approximate Newton-Raphson optimization for the parameter fitting, which requires few large steps only for convergence. These steps are reduced to matrix-vector multiplication (MVM) primitives with kernel matrices. For general kernels, these MVM primitives can be reduced to large numerical linear algebra primitives, which can be computed very efficiently on modern computer architectures. This is very much in contrast to many chunking algorithms for kernel method fitting, which have been proposed in machine learning, and the advantages of our approach are detailed in Section 7.2. Dependencies between classes can be encoded *a priori* with minor additional efforts, as has been demonstrated for the case of hierarchical classification. Our method provides estimates of predictive probabilities which are asymptotically correct. Hyperparameters can be adjusted automatically, by optimizing a cross-validation log likelihood score in a gradient-based manner, and these computations are once more reduced to the same MVM primitives. This means that within our framework, all code optimization efforts can be concentrated on these essential primitives (see also Section 7.3), rather than having to tune a set of further heuristics.

7.1 Related Work

Our primary fitting optimization for flat multi-way classification appeared in Williams and Barber (1998), although some fairly essential features are novel here. They also did not consider large scale problems or class structures. Empirical Bayesian criteria such as the marginal likelihood are routinely used for hyperparameter learning in Gaussian process models (Williams and Barber, 1998; Williams and Rasmussen, 1996). However, in cases other than regression estimation with Gaussian noise, the marginal likelihood for a GP model cannot be computed analytically, and approximations differ strongly in terms of accuracy and computational complexity. For the multi-class model, Williams and Barber (1998) use an MAP approximation for fixed hyperparameters, just as we do, but their second-order approximation to the marginal likelihood is quite different from our criterion, conceptually as well as computationally (see below). Approximately solving large linear system by linear conjugate gradients (LCG) is standard in numerical mathematics, and has been used in machine learning as well (Gibbs, 1997; Williams and Barber, 1998; Keerthi and DeCoste, 2005).

The idea of optimizing approximations to a cross-validation score for hyperparameter learning is not novel (Craven and Wahba, 1979; Qi et al., 2004). Our approach is different to these, in that the CV score and gradient computations are reduced to elementary steps of the primary fitting method,

17. We cannot obtain a good initial α value from the final $\alpha_{[k]}$ of hyperparameter learning, because this is done on training subsets only. Moreover, in our implementation, the “stalling” (no improvement) of a NR step means that LCG is restarted from its last recent β , so that eventually an improvement in Φ is still obtained. Of course, the stalled NR iterations counts as such, and we do k_1 iterations in total.

so both can be done with the same code.¹⁸ In contrast, scores like GCV (Craven and Wahba, 1979) or second order marginal likelihood (Williams and Barber, 1998) come in terms of the form $\text{tr}H^{-1}$ or $\log|H|$ for the Hessian H of size nC . There are approximate reductions of computing these terms to solving linear systems (randomized trace estimator, Lanczos), but they rely on additional sampling of Gaussian noise, which introduces significant inaccuracies. In practice, optimizing such “noisy” criteria is quite difficult, whereas our criterion can be optimized using standard optimization code. Qi et al. (2004) propose an interesting approach of approximating leave-one-out CV using expectation propagation, see also Opper and Winther (2000). They use a sparse approximation for efficiency, but they deal with a single-process model only ($C = 1$), and it is not clear how to implement EP efficiently (scaling linearly in C) for the multi-class model. Interestingly, they observe that optimizing their approximate CV score is more robust to overfitting than the marginal likelihood. Finally, none of these papers propose (or achieve) a complete reduction to kernel MVM primitives only, nor do they deal with representing class structures or work on problems of the scale considered here.

Many different multi-class SVM techniques have been proposed, see Crammer and Singer (2001) and Hsu and Lin (2002) for references. These can be split into joint (“all-together”) and decomposition methods. The latter reduce the multi-class problem to a set of binary ones (“one-against-rest” of Section 6.2 is a prominent example), with the advantage that good code is available for the binary case. The problem with these methods is that the binary discriminants are fitted separately without knowledge of each other, or of their role in the final multi-way classifier, so information from data is wasted. Also, their post-hoc combination into a multi-way discriminant is heuristic. Joint methods are like ours here, in that all classes are jointly represented. Fitting is a constrained convex problem, and often fairly sparse solutions (many zeros in α) are found. However, in difficult structured label tasks, the degree of sparsity is usually not high, and in these cases, commonly used chunking algorithms for multi-class SVM can be very inefficient (see Section 7.2). We should note that our approach here cannot be applied directly to multi-class SVMs, since they require the solution of a constrained convex problem, but the principles used here should hold there as well. Some novel suggestions here appear independently in Keerthi et al. (2007). SVM methods typically do not come with efficient automatic kernel parameter learning schemes, and they do not provide estimates of predictive probabilities which are asymptotically correct.

On the other hand, in a direct comparison our implementation would still be slower than the highly optimized multi-class SVM code of Crammer and Singer (2001), at least on standard non-structured tasks such as USPS (Section 6.3) or MNIST. Especially on the latter, sparsity in α is clearly present, and years of experience with the SVM problem led to very effective ways of exploiting it. In contrast, α in our approach is not sparse, and it is not our goal here to find a sparse approximation. Hyperparameters are selected “by hand” in their method, not via gradient-based optimization. For a small number of hyperparameters, this traditional approach is often faster than our optimization-based one here, and importantly, it can be fully parallelized. However, our approach is still workable in situations with many dependent hyperparameters (for example, Section 7.4.1), where CV by hand simply cannot be done.

Our ANOVA setup for hierarchical classification is proposed by Cai and Hofmann (2004), whose use it within a SVM “all-together” method. We compare our method against theirs in Sec-

18. A small drawback of our approach is that our CV score Ψ depends on a partitioning of the training set. In our experiments here, we chose this at random. Leave-one-out (LOO) CV does not depend on a partitioning, but it is not clear how to reduce LOO CV to solving a small number of linear systems.

tion 6.1, achieving quite similar results in an order of magnitude less time. They also do not address the problem of hyperparameter learning.

7.2 Global versus Decomposition Methods

In most kernel methods proposed so far in machine learning, the primary fitting to data (for fixed hyperparameters) translates to a convex minimization problem, where the penalisation terms correspond to quadratic expressions with kernel matrices. While kernel matrices may show a rapidly decaying eigenvalue spectrum, they certainly do couple the optimization variables strongly.¹⁹ While a convex function can be optimized by any method which just guarantees descent in each step, there are huge differences in how fast the minimum is attained to a desired accuracy. In fact, in the absence of local minima, the speed of convergence becomes the most important characteristic of a method, besides robustness and ease of implementation.

In machine learning, the most dominant technique for large scale (structured label) kernel classification is what optimization researchers call *block coordinate descent methods* (BCD), see Bertsekas (1999, Sect. 2.7). The idea is to minimize the objective w.r.t. a few variables at a time, keeping all others fixed, and to iterate this process using some scheduling over the variables. Each step is convex again,²⁰ yet much smaller than the whole, and often the steps can be solved analytically. Ignoring the aspect of scheduling, such methods are simple to implement.

A complementary approach is to find search directions which lead to as fast a descent as possible, these directions typically involve all degrees of freedom of the optimization variables. If local first and second order information can be computed, the optimal search direction is Newton's, which has to be corrected if constraints are present (conditional gradient or gradient projection methods). If the Newton direction cannot be computed feasibly, approximations may be used. Such Newton-like methods are certainly vastly preferred in the optimization community, due to superior convergence rates, but also because features of modern computer architectures are used more efficiently, as is detailed below. In this paper, we advocate to follow this preference for kernel machine fitting in machine learning. We are encouraged not only by our own experiences, but can refer to the fact that (approximate) Newton methods are standard for fitting generalized linear models in statistics, and that such methods are also routinely used for Gaussian process models (Williams and Barber, 1998; Rasmussen and Williams, 2006), albeit typically on problems of smaller scale than treated here.

The dominance of BCD methods for kernel machine fitting, while somewhat surprising, can be attributed to early success stories with SVM training, culminating in the SMO algorithm (Platt, 1998), where only two variables are changed at a time. If an SVM is fitted to a task with low noise, the solution can be highly sparse, and if the active set of "support vectors" is detected early in the optimization process, methods like SMO can be very efficient. Importantly, SMO or other BCD methods are easily implemented. On the other hand, as SVMs are increasingly applied to hard structured label problems which usually do not have very sparse solutions, or whose active sets are hard to find, weaknesses of BCD methods become apparent.

Block coordinate descent methods are often referred to as using the "divide-and-conquer" principle, but this is not the case for kernel method fitting. BCD methods "are often useful in contexts where the cost function and the constraints have a partially decomposable structure with respect to

19. An almost low-rank kernel matrix translates into a coupling of a simple structure, but the dominant couplings are typically strong and not sparse.

20. If $f(x)$ is convex, so is $f(Ax)$ for any matrix A . The same is true for linear constraints.

the problem’s optimization variables” (Bertsekas, 1999, p. 269). In kernel methods, such a decomposable structure is not present, because the penalisation terms couple all variables strongly via the kernel matrices. In such cases, chunking techniques *divide without conquering*, often running for very many steps, because improvements w.r.t. some block of variables tend to erase earlier improvements. This central problem of block coordinate descent methods is well known as “zig-zagging” in optimization. It occurs whenever variables not in the same block are significantly coupled, a situation which is to be expected for kernel machines in general. The situation is similar to a number of cases in machine learning and statistics. Iterative proportional fitting (Della Pietra et al., 1997) is a BCD method for learning the potentials of an undirected graphical model (Markov random field), which is all but superseded now by modern global direction methods such as limited memory Quasi-Newton, running orders of magnitude faster. Gibbs sampling is a basic Markov chain Monte Carlo technique for approximate Bayesian inference, which typically is very simple to implement, but is exceedingly slow in the presence of many coupled variables. Modern techniques such as Hybrid Monte Carlo, or Swendsen-Wang can be seen as “global direction” variants of “block coordinate” Gibbs sampling, and while they are harder to implement, they typically run orders of magnitude faster.

Another significant problem with BCD methods may come more as a surprise, namely because it concerns a characteristic which is often sold as advantage of these methods: they make each step as small as possible. Such small steps can often be dealt with analytically, or by using simple methods. This characteristic certainly makes BCD methods easy to implement. However, in light of modern computer architectures, the advice must be to make each step as *large as possible*, with the aim of requiring fewer steps. Modern systems use many internal parallelisms and hierarchies of caching, with the aim of processing vector operations many times faster than an equivalent loop over scalar operations, and large global steps do make use of these features. In contrast, a method which calls very many small steps in a non-linear ordering, runs contrary to these mechanisms. For example, data transfer between cache levels is done in blocks of significant sizes, and a method which accesses memory element-wise from all over the place, leads to inefficiencies up to cache thrashing, where the majority of cache accesses are misses (see Appendix D.3 for an example).

In well-designed global direction methods, the bottleneck operations (where almost all real-world running time is spent) are large vectorised mappings which access memory contiguously. Even better, these operations should lie in a standard class, for which highly optimized implementations are readily available. In our case here, the bottleneck operations are numerical linear algebra primitives from the *basic linear algebra subroutines* (BLAS), a standardized interface for high-performance dense linear algebra code. Very efficient implementations of the BLAS are available for all computer architectures.²¹

In this paper, we advocate to take a step back and to use global direction methods as approximation to Newton’s method for kernel machine fitting. The prospect seems daunting, since there are many thousands of variables with complicated couplings, and the reader may be reminded of early disastrous trials of applying off-the-shelf QP packages to SVM fitting, or of ongoing efforts to formulate otherwise tractable machine learning problems as semidefinite programs and “solving” them using $O(n^7)$ SDP packages. This association is wrong. Our advice is to *approximate* the global Newton direction, making use of all structure in the model in order to gain efficiency, which is exactly the opposite of running a black box solver or implementing Newton’s method straight

21. ATLAS, a self-tuning BLAS implementation, is available as free software, see <http://math-atlas.sourceforge.net/>.

out of a textbook. In the context of kernel machines, where couplings through large unstructured matrices are present, the large steps of approximating the Newton direction should be reduced to standard linear algebra primitives on dense or sparse matrices, operating on contiguous chunks of memory *as large as possible*, since highly optimized code for such primitives is readily available.

7.3 Matrix-Vector Multiplication

The computational load in our framework is determined by applications of the MVM primitives $v \mapsto K^{(c)}v$ and $v \mapsto (\partial K^{(c)}/\partial h_p)v$. A user only needs to provide those for a kernel of choice. The generic representation of Appendix D.1 applies to general covariance functions, but much more efficient alternatives may be used in special cases (see Appendix D.3).

If the cost for a direct evaluation of these primitives is prohibitive, several well-known approximations may be applied. It has been suggested to use specialized data structures to approximate MVM with matrices from isotropic kernels (Yang et al., 2005; Shen et al., 2006) such as the Gaussian one (7). For such kernels, the derivative MVM can often be addressed using the same techniques. For the Gaussian kernel, we have $K = v \exp(wA)$, $A = (-(1/2)\|x_i - x_j\|^2)_{i,j}$, in which case $(\partial K/\partial \log v) = K$ and $(\partial K/\partial \log w) = wK \circ A$, where \circ denotes component-wise product. Since the specialized data structures concentrate on approximating A , they apply to all required MVM variants. Our public code could fairly easily be extended, given specialized approximate kernel MVM code is available.

7.4 Extensions and Future Work

Some concrete extensions are mentioned just below. In general, we think that many structured label kernel methods proposed in the SVM context can be addressed in our framework as well. For example, the kernel conditional random field (CRF) (Lafferty et al., 2004) for label sequence learning can be treated by recognizing that MVM with the Hessian of the CRF log likelihood can be implemented efficiently using the method described in Pearlmutter (1994). We also plan to address hierarchical multi-label problems, which differ from hierarchical multi-class in that each instance can have multiple associated labels.

7.4.1 MODELLING DEPENDENCIES BETWEEN CLASSES

In the flat classification application of Section 2, we do not explicitly represent dependencies between classes. This is done in the hierarchical classification application of Section 3, but the dependency structure is fixed *a priori*. In this section, we motivate how dependencies between classes can be learned from data.

Let $B \in \mathbb{R}^{C,C}$ be a nonsingular coupling matrix, which will be a part of the model. In fact, B should be regarded as hyperparameters. In the flat model, we have $u_i = f_i + b$, which is now replaced by $u_i = Bf_i + b$, or

$$u = (B \otimes I)K\alpha + b \otimes 1.$$

This is the same modification which led to the hierarchical extension, only that the fixed coupling matrix Φ is replaced by the variable B . Therefore, the same modification of our method can be done, replacing K by $K^{(B)} = (B \otimes I)K(B^T \otimes I)$. Again, MVM with $K^{(B)}$ is of the same complexity as with K , because MVM with $(B \otimes I)$ can be done in $O(C^2n)$. Note that B represents *conditional* depen-

dencies between the $u_c(\cdot)$, its role is comparable to the “mixing matrix” in independent components analysis.

We are also interested in learning B , whose elements are taken as hyperparameters. The corresponding gradient of Ψ is obtained in the same way as described in Section 5.2, only that $dK_{[B]}$ now further decomposes into parts for dB and for dK . Note that learning B by non-automatic cross-validation would not be possible, due to the large number of components.

Consider the case where we have many classes C , but not much data for most of them. We can postulate the assumptions that the behaviour of the C class functions $u_c(\cdot)$ is represented by $p \ll C$ underlying latent factors, which are then modelled as independent. This is achieved in our framework by having a non-square mixing matrix $B \in \mathbb{R}^{C \times p}$ (the “factor loadings”). It is not hard to adapt our framework to this case, in which it is obviously necessary to learn B as hyperparameters from data. A related model in a Bayesian context was considered in Teh et al. (2005).

7.4.2 UNCERTAIN TARGETS IN HIERARCHICAL CLASSIFICATION

Recall the hierarchical classification setup of Section 3. Suppose that for some patterns x_i , the target is unknown, but we know that the path from its class to the root goes through an inner node p . Denote by L_p the set of leaf nodes of the subtree rooted at p , so that $L_p = \{p\}$ for a leaf node $p \in L$, and $L_0 = L$.

We can allow for such uncertain target information by using pseudo-targets $\tilde{y}_i \in \{1, \dots, P\}$. If $\tilde{y}_i \notin L$, it is the lowest inner node we are certain about. The corresponding likelihood factor is

$$\sum_{c \in L_{\tilde{y}_i}} P(y_i = c | u_i).$$

Note that the log likelihood is not a concave function anymore, whenever $|L_{\tilde{y}_i}| > 1$, and in the presence of such factors, primary fitting is not a convex problem. However, an expectation-maximization (EM) (Dempster et al., 1977) approach can be used to deal with uncertain targets. Namely, in “E steps” we compute

$$q_{ic} \propto \mathbf{I}_{\{c \in L_{\tilde{y}_i}\}} P(y_i = c | u_i)$$

for the current $u_c(\cdot)$, where $q_i = (q_{ic})_c$ are distributions. “M steps” consists of Newton-Raphson iterations as before, but using $\sum_c q_{ic} \log P(y_i = c | u_i)$ as log likelihood factors. To this end, we just have to replace the vector $y \in \mathbb{R}^{nC}$ by q . Importantly, we only used the properties $1^T y_i = 1$, $y_i \geq 0$ above (but not that $y_{ic} \in \{0, 1\}$), which are true for q just as well.

7.4.3 LOW RANK APPROXIMATIONS

Our generic kernel matrix representation is described in Appendix D.1. If the data set size n is large, we may not be able to keep the correlation matrices $M^{(l)}$ in memory, and MVM with them becomes prohibitively expensive. We can use standard low rank matrix approximations to deal with this problem (see also Section 7.3). Namely, suppose that $M^{(l)}$ is approximated by $P^{(l)} L^{(l)} L^{(l)T} P^{(l)T}$, where $P^{(l)}$ is a permutation matrix, and $L^{(l)} \in \mathbb{R}^{n, d_l}$ for $d_l \ll n$. Denote $I_l \subset \{1, \dots, n\}$ the active set of size d_l . The approximation may be obtained by an incomplete Cholesky factorization²² (ICF), which has the special property that only a small set of d_l columns of $M^{(l)}$ (along with its diagonal)

22. Matlab code for ICF (in the form required here) can be downloaded from <http://www.kyb.tuebingen.mpg.de/bs/people/seeger/software.html>.

ever have to be evaluated (Bach and Jordan, 2002). In this case, $L_{1\dots d_l}^{(l)}$ is the lower triangular Cholesky factor of $M_{I_l}^{(l)} \in \mathbb{R}^{d_l \times d_l}$, so that $P^{(l)T} M_{I_l}^{(l)} = L^{(l)} L_{1\dots d_l}^{(l)T}$. Note that in the ICF case, we have that

$$\text{diag} \left(P^{(l)T} M^{(l)} P^{(l)} - L^{(l)} L^{(l)T} \right) \geq 0$$

point-wise, because the elements are simply the squared pivots for a potential continuation of the factorization (which has been stopped after d_l steps). Therefore, we can correct the approximation by replacing the diagonal of $L^{(l)} L^{(l)T}$ by the true one, ending up with the approximation

$$M^{(l)} \approx \tilde{M}^{(l)} := (\text{diag}^2 M^{(l)}) + P^{(l)} \left(L^{(l)} L^{(l)T} - (\text{diag}^2 L^{(l)} L^{(l)T}) \right) P^{(l)T}.$$

Snelson and Ghahramani (2006) motivate this diagonal correction in another context. It is clear that MVM with $\tilde{M}^{(l)}$ can be done in $O(n d_l)$.

If $*$ indexes test points different from the training points, then the test-training correlation matrix is

$$M_{*,\cdot}^{(l)} = M_{*,I}^{(l)} (L_{1\dots d_l}^{(l)})^{-T} L^{(l)T} P^{(l)T}.$$

We can also learn parameters of the $M^{(l)}$ functions in this low rank setup by gradient-based optimization, assuming that the choice of I_l does not depend on these kernel parameters, but this is not discussed here.

Acknowledgments

We would like to thank Olivier Chapelle for many useful discussions on approximations of cross-validation optimization, and Peter Gehler for help with preprocessing the WIPO-alpha data, furthermore L. Cai and T. Hofmann for providing us with counts and dictionary data. Dilan Görür gave helpful comments on the NIPS paper related to this work. Supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778.

Appendix A. Notation

In this section, we describe the notation used in this paper. We denote vectors and matrices by bold-face lower-case and upper-case letters, scalars and scalar functions are set normally. Subscripts select parts of objects, they can be single indexes or index sets. For example, $a = (a_i)_i$ is a vector with components a_i , $A = (a_{i,j})_{i,j}$ a matrix with entries $a_{i,j}$. We also write $a = (a_i)$, $A = (a_{i,j})$ if the indexes are clear from context. $A_{\cdot,i}$ is the i -th column of A (“ \cdot ” is short for the full index set). \otimes denotes the Kronecker product, $A \otimes B = (a_{i,j} B)_{i,j}$, $\mathbf{1}$ ($\mathbf{0}$) the vector of all ones (vector/matrix of all zeros), I the identity matrix, and $\delta_j = (I_{\{i=j\}})_i$ (columns of I). For a matrix A , $\text{diag} A = (a_{i,i})_i$ extracts the diagonal. For a vector v , $\text{diag} v$ is the corresponding diagonal matrix. We also use this for matrix-valued vectors, an example is the diagonal kernel matrix $K = \text{diag}(K^{(c)})_c$ in flat classification.

Many vectors and matrices are indexed by data points (i) and classes (c) at the same time, for example $u = (u_{ic}) \in \mathbb{R}^{nC}$. We use double indexes ic for these, which are flattened as $i + (c - 1)n$, so the component ordering²³ is $u = (u_{11}, u_{21}, \dots, u_{n1}, u_{12}, \dots)$. In this context, selection index sets

23. In Matlab, `reshape(u, n, C)` would give a matrix in $\mathbb{R}^{n \times C}$.

I are applied to the i (data point) index only: $u_I = (u_{ic})_{i \in I, c} \in \mathbb{R}^{|I|C}$. Kronecker product notation works nicely with this double index convention. If $A \otimes B$ is applied to u , A has C , B n columns. We frequently use $(1^T \otimes I)u = \sum_c u^{(c)}$, where $u^{(c)} = (u_{ic})_i \in \mathbb{R}^n$, or $(1 \otimes I)v$ for $v \in \mathbb{R}^n$, which stacks v on top of each other C times. The matrix $P_{cls} = (1 \otimes I)(1^T \otimes I)$ (introduced in Section 5.1) combines these operations:

$$P_{cls}x = \left(\begin{array}{c} \bar{x} \\ \bar{x} \\ \vdots \end{array} \right) \Bigg\} C, \quad \bar{x} = \sum_c x^{(c)},$$

and P_{data} does the same, but operating on the i rather than the c index.

All major notational definitions are listed in Table 3 for reference. For kernel matrices (for example, $K^{(c)}$), we do not list the kernel functions (here: $K^{(c)}$), and for evaluation vectors (for example, u), we do not list the underlying functions (here: $u^{(c)}(\cdot)$).

n	Number data points	2	LCG	Linear Conjugate Gradients	2
C	Number classes	2	P	Number nodes (hierarchy)	3
y	Targets (zero-one)	2	L	Leaf nodes (hierarchy)	3
x_i	Input points	2	\tilde{u}	Latent output (before mixing)	3
u	Latent output (after mixing)	2	Φ	Hierarchy mixing matrix	3
b	Intercepts	2	\tilde{K}	Kernel matrix (before mixing)	(2)
σ^2	Penalizing constant for b	2	I_k, J_k	Partitions for CV criterion	4
Φ	Criterion for primary fitting	2	Ψ	CV criterion	4
α	Dual variables	2	q	Number of folds	4
K	Kernel matrix (after mixing)	2	h	Hyperparameters	4
$K^{(c)}$	Kernel matrix block	2	k_3	Complexity parameter	4
\tilde{K}	Kernel matrix (b eliminated)	2	g, W	Gradient, Hessian Φ_{lh}	5.1
l	Logsumexp vector	(1)	P_{cls}	Sum-distribute matrix	5.1
k_1, k_2	Complexity parameters	2	P_{data}	Sum-distribute matrix	5.1
NR	Newton-Raphson	2	E, F	Accumulation matrices	5.2

Table 3: Reference for notational definitions. k : Section of definition; (k) : Equation of definition.

Appendix B. Details for Primary Fitting Algorithm

In this section, we discuss further details of the primary fitting algorithm of Section 2, in addition to Section 5.1.

We need to counter the problem that roundoff errors may lead to numerical instabilities. The criterion we minimize is strictly convex, even if the kernel matrix K is singular (or nearly so). However, problems could arise from components in π becoming very small. Recall that $\log \pi_{ic} = u_{ic} - l_i$. We make use of a threshold $\kappa < 0$ and define

$$I = \{(i, c) \mid \log \pi_{ic} < \kappa, y_{ic} > 0\}, \quad I_0 = \{(i, c) \mid \log \pi_{ic} < \kappa, y_{ic} = 0\}.$$

The indices in I can be problematic due to the corresponding component $\tilde{g}_{ic} \approx y_{ic}/\pi_{ic}^{1/2}$ becoming large. Note that this happens only if (x_i, y_i) is a strong outlier w.r.t. the current predictor. Now, from

the system (5) we see that $D^{1/2}\beta = DA^T(u - \tilde{K}\alpha') - g$. Therefore, if $(i, c) \in I$, then $(D^{1/2}\beta)_{ic} \approx -g_{ic} \approx y_{ic}$. The idea is to solve the reduced system on the components in $\setminus I$ for $(D^{1/2}\beta)_{\setminus I}$ and to plug in $(D^{1/2}\beta)_I = y_I$. Finally, within $\setminus I$, the components in I_0 may be problematic when computing the starting value $\beta = D^{-1/2}\alpha$ for the CG run. However, in this case $\tilde{g}_{ic} \approx 0$, leading to $\beta_{ic} \approx 0$ from (5). The corresponding components in the starting value β can therefore be set to zero.

Next, the LCG algorithm for solving systems of the form (5) needs to be preconditioned. Suppose we want to solve $Ax = b$. If we have an approximation \tilde{A} to A so that $v \mapsto \tilde{A}^{-1}v$ can be computed efficiently (essentially in linear time in the size of v), the preconditioned CG algorithm solves the system $\tilde{A}^{-1}Ax = \tilde{A}^{-1}b$ instead. The idea is that $\tilde{A}^{-1}A$ typically has a lower condition number than A , and LCG converges faster and less erratically. Our implementation does preconditioning with the diagonal of the system matrix $I + V^T\tilde{K}V$. Note that $V\delta_{ic} = \pi_{ic}^{1/2}(\delta_c - \pi_i) \otimes \delta_i$, so that

$$(I + V^T\tilde{K}V)_{ic,ic} = 1 + \pi_{ic} \left((1 - 2\pi_{ic})(K_i^{(c)} + \sigma^2) + \sum_{c'} \pi_i^{(c')^2} (K_i^{(c')} + \sigma^2) \right).$$

Therefore, the diagonal can be computed based on the $\text{diag}K^{(c)}$ vectors. If the joint kernel matrix K is not block-diagonal (as in hierarchical classification, see Section 3), $\text{diag}K$ is not sufficient for computing the system matrix diagonal. Let $v \in \mathbb{R}^{nC}$ be defined via $v_i = K_i\pi_i$, where $K_i = (I \otimes \delta_i^T)K(I \otimes \delta_i) \in \mathbb{R}^{C,C}$. Then, the system matrix diagonal has elements

$$1 + \pi_{ic} (K_{ic} + \sigma^2 - 2w_{ic} + \pi_i^T w_i), \quad w = v + \sigma^2\pi.$$

Appendix C. Solving Systems Exactly

In this section, we show how to implement our flat multi-class scheme using exact rather than approximate solutions of linear systems, yet still scaling linearly in C (at present, we do not know how to implement hierarchical classification exactly with such scaling).

For a Newton step, we need to solve $(I + W\tilde{K})\alpha' = r$ with $W = D - DP_{cls}D$. This can be written as

$$(A - UV^T)D^{-1/2}\alpha' = D^{-1/2}r, \quad A = I + D^{1/2}\tilde{K}D^{1/2}, \\ U = D^{1/2}(1 \otimes I), \quad V = (A - I)U.$$

We now use the Sherman-Morrison-Woodbury formula together with the fact that $U^TU = \sum_c D^{(c)} = I$ to obtain

$$\alpha' = D^{1/2} (A^{-1} + A^{-1}U(U^TA^{-1}U)^{-1}U^T(I - A^{-1}))D^{-1/2}r.$$

We used that $V^TA^{-1} = U^T(I - A^{-1})$. Note that A^{-1} is block-diagonal, and that

$$U^TA^{-1}U = \sum_c D^{(c)1/2}A^{(c)-1}D^{(c)1/2}.$$

We maintain Cholesky factors of all $A^{(c)}$, as well as the Cholesky decomposition $U^TA^{-1}U = RR^T$ (where $A^{(c)-1}$ are obtained from the Cholesky factors).

For hyperparameter learning, we consider the partitions (I, J) sequentially. Since the $\alpha_{[J]}$ are different across folds, we cannot obtain the $A_{[J]}, H_{[J]}$ as parts of underlying common matrices. Recall Section 5.2. $\alpha_{[J]} + g_{[J]} = 0$ gives $H_{[J]}(d\alpha_{[J]}) = -W_{[J]}(dK_J)\alpha_{[J]}$. With

$$f = I_{\cdot,I}(\pi_{[I]} - y_I) - I_{\cdot,J}W_{[J]}H_{[J]}^{-T}\tilde{K}_{J,I}(\pi_{[I]} - y_I),$$

we have that $d\Psi_I = (I_{\cdot,j}\alpha_{[j]})^T(dK)f$. Again, these vectors are accumulated in matrices E, F . Solving a system with $H_{[j]}^T$ is an obvious variant of the procedure discussed above.

Appendix D. Further Details of the Implementation

Our implementation is designed to be as efficient as possible, while still being general and easy to extend to novel situations. This is achieved mainly by breaking down the problems to calling sequences of MVM primitives. These are then reduced to large numerical linear algebra primitives, where matrices are organized contiguously in memory, in order to exploit modern caching architectures (see Section 7.2).

D.1 A Generic Kernel Matrix Representation

A kernel matrix representation is some data structure which allows to compute kernel matrix MVMs $v \mapsto K^{(c)}v$ efficiently, being the principal primitives of our primary fitting method. Further requirements arise if additional features of our framework are used. For example, if hyperparameters are to be learned as well, derivative MVMs $v \mapsto (\partial K^{(c)}/\partial h_p)v$ are required as well, and ‘‘covariance shuffling’’ should be possible (see Section 5.2).

An efficient representation depends strongly on the covariance function used, and also on whether kernel matrix MVMs are approximated rather than computed exactly. For example, for linear kernels a special representation is used (see Appendix D.3). In this section, we describe a generic representation, which is part of our implementation.

The generic representation can be used with any covariance function, in that no special structure is assumed. It requires kernel matrices to be stored explicitly, which may not be possible for very large n . In general, we allow for different covariance functions $K^{(c)}$ for each class c , although sharing of kernels is supported, in that $K^{(c)}(\cdot, \cdot) = v_c M^{(l_c)}(\cdot, \cdot)$ and $v_c > 0$. Here, $l_c = l_{c'}$ is allowed for $c \neq c'$. The matrices $M^{(l)}$ are stored explicitly. Note that the flexibility of using different variance parameters v_c with the same $M^{(l)}$ does come at no extra cost, except for the fact that these have to be adjusted individually.

Since the $M^{(l)}$ are symmetric, two can be stored each in a $n \times n$ block, say the odd-numbered ones in the lower triangles. Here, the $\text{diag} M^{(l)}$ are stored separately, and whenever a specific $M^{(l)}$ is required explicitly, the diagonal is copied into the block. It is important to note that the BLAS (see Section 7.2) directly supports symmetric matrices which are stored in the lower or upper triangle of a rectangular block.

The reader may wonder whether space could be saved by storing intermediates of the $M^{(l)}$ instead. For example, if the $M^{(l)}$ are isotropic kernels of the form $f^{(l)}(\|x - x'\|)$, we could store the inner product matrix $(x_i^T x_j)_{i,j}$ only. In practice, this turns out to be significantly slower (by a factor), the reason being that the optimized BLAS primitives are many times more efficient than applying a non-linear function $f^{(l)}$ point-wise to a matrix, even if the matrix is stored contiguously. For the same reason, computing MVMs on the fly without storing matrices is even more costly.

D.2 Shuffling the Kernel Matrix Representation

Covariance matrix shuffling has been motivated in Section 5.2. It is required during hyperparameter optimization, because the MVM primitives for sub-matrices K_{J_k} have to be driven by a single representation of the complete K (note that each K_{J_k} is of size $n(q-1)/q$, thus almost as large as K).

A simple approach would be to use sub-indexed matrix-vector multiplication code, but this is very inefficient (usually more than one order of magnitude slower than the flat BLAS functions).

Instead, when dealing with fold k , we shuffle the representation so that K_{J_k} moves to the upper left corner of the matrix. How this is done, depends on the representation. In this context, it is important to note that the underlying BLAS explicitly allows working on sub-matrices within upper left corners of larger frames, with virtually no loss in efficiency.²⁴ In the generic representation of Appendix D.1, we simply permute the kernel matrices $K^{(c)}$ using the index (J_k, I_k) . A corresponding de-shuffling operation has to restore the old representation for K .

D.3 The Linear Kernel

Our application described in Section 6.1 uses the linear kernel $K^{(c)}(x, x') = v_c x^T x'$, where x is very high-dimensional (word counts over a dictionary), but also extremely sparse (by far the most entries are zero). The linear kernel fits the setup of Appendix D.1 with a single $M^{(1)} = XX^T$, where $X \in \mathbb{R}^{n,d}$ is the design matrix. X is very sparse, and in our implementation is represented using a standard sparse matrix format.

An MVM is done as $v \mapsto v_c(XX^T v)$, where X is sparse. More generally, we do $S \mapsto XX^T S$ with large matrices S . Kernel matrix shuffling (Appendix D.2) is implemented by simply reordering the non-zero positions for X . In this context, it is interesting to remark a finding which underlines the arguments in Section 7.2. The sparse matrix format is such that $XX^T S$ is reduced to so-called *daxpy* operations ($a = a + \alpha b$) on the *rows* of S . By Fortran (and BLAS) convention, S is stored in column-order, so that rows can only be accessed directly by using a striding value > 1 (the distance between consecutive vector elements in memory). We added a simple trick (called *dimension flipping*) to the implementation, which in essence switches our default ordering of Cn vectors $v = (v_{11}, v_{21}, v_{31}, \dots)^T$ to $(v_{11}, v_{12}, v_{13}, \dots)^T$ before major kernel MVM computations are done. This simple modification led to a direct five-times speedup, which underlines the importance of contiguous memory access in the bottleneck computations of a method (which allows optimal usage of cache hierarchies).

References

- F. Bach and M. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- P. Bartlett and A. Tewari. Sparseness vs estimating conditional probabilities: Some asymptotic results. In *Conference on Computational Learning Theory 17*, pages 564–578. Springer, 2004.
- D. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 1999.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2002.
- L. Cai and T. Hofmann. Hierarchical document categorization with support vector machines. In *CIKM 13*, pages 78–87, 2004.
- K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2:265–292, 2001.

24. Matrices in BLAS are stored column-wise, each column has to be contiguous in memory, but the striding value (offset in memory required to jump to the next column) can be larger than the number of rows.

- P. Craven and G. Wahba. Smoothing noisy data with spline functions: Estimating the correct degree of smoothing by the method of generalized cross-validation. *Numerische Mathematik*, 31:377–403, 1979.
- L. Csató. *Gaussian Processes — Iterative Sparse Approximations*. PhD thesis, Aston University, Birmingham, UK, March 2002.
- S. Della Pietra, V. Della Pietra, and J. Lafferty. Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(4):380–393, 1997.
- A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of Roy. Stat. Soc. B*, 39:1–38, 1977.
- M. Gibbs. *Bayesian Gaussian Processes for Regression and Classification*. PhD thesis, University of Cambridge, 1997.
- P.J. Green and B. Silverman. *Nonparametric Regression and Generalized Linear Models*. Monographs on Statistics and Probability. Chapman & Hall, 1994.
- C.-W. Hsu and C.-J. Lin. A comparison of methods for multi-class support vector machines. *IEEE Transactions on Neural Networks*, 13:415–425, 2002.
- S. Keerthi and D. DeCoste. A modified finite Newton method for fast solution of large scale linear SVMs. *Journal of Machine Learning Research*, 6:341–361, 2005.
- S. Keerthi, V. Sindhvani, and O. Chapelle. An efficient method for gradient-based adaptation of hyperparameters in SVM models. In B. Schölkopf, J. Platt, and T. Hofmann, editors, *Advances in Neural Information Processing Systems 19*. MIT Press, 2007.
- J. Lafferty, Y. Liu, and X. Zhu. Kernel conditional random fields: Representation, clique selection, and semi-supervised learning. Technical Report CMU-CS-04-115, Carnegie Mellon University, 2004.
- Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1:541–551, 1989.
- P. McCullach and J.A. Nelder. *Generalized Linear Models*. Number 37 in Monographs on Statistics and Applied Probability. Chapman & Hall, 1st edition, 1983.
- M. Opper and O. Winther. Gaussian processes for classification: Mean field algorithms. *Neural Computation*, 12(11):2655–2684, 2000.
- B. Pearlmutter. Fast exact multiplication by the Hessian. *Neural Computation*, 6(1):147–160, 1994.
- J. Platt. Fast training of support vector machines using sequential minimal optimization. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods: Support Vector Learning*, pages 185–208. MIT Press, 1998.
- Y. Qi, T. Minka, R. Picard, and Z. Ghahramani. Predictive automatic relevance determination by expectation propagation. In C. Brodley, editor, *International Conference on Machine Learning 21*. Morgan Kaufmann, 2004.

- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Y. Saad. *Iterative Methods for Sparse Linear Systems*. International Thomson Publishing, 1st edition, 1996.
- B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, 1st edition, 2002.
- M. Seeger. *Bayesian Gaussian Process Models: PAC-Bayesian Generalisation Error Bounds and Sparse Approximations*. PhD thesis, University of Edinburgh, July 2003. See www.kyb.tuebingen.mpg.de/bs/people/seeger.
- M. Seeger. Gaussian processes for machine learning. *International Journal of Neural Systems*, 14(2):69–106, 2004.
- M. Seeger. Cross-validation optimization for large scale hierarchical classification kernel methods. In B. Schölkopf, J. Platt, and T. Hofmann, editors, *Advances in Neural Information Processing Systems 19*, pages 1233–1240. MIT Press, 2007.
- B. Shahbaba and R. Neal. Improving classification when a class hierarchy is available using a hierarchy-based prior. *Bayesian Analysis*, 2(1):221–238, 2007.
- Y. Shen, A. Ng, and M. Seeger. Fast Gaussian process regression using KD-trees. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*. MIT Press, 2006.
- E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*. MIT Press, 2006.
- Y.-W. Teh, M. Seeger, and M. I. Jordan. Semiparametric latent factor models. In Z. Ghahramani and R. Cowell, editors, *Workshop on Artificial Intelligence and Statistics 10*, 2005.
- G. Wahba. *Spline Models for Observational Data*. CBMS-NSF Regional Conference Series. Society for Industrial and Applied Mathematics, 1990.
- C. Williams and C. Rasmussen. Gaussian processes for regression. In D. Touretzky, M. Mozer, and M. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*. MIT Press, 1996.
- C. K. I. Williams and D. Barber. Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1342–1351, 1998.
- C. Yang, R. Duraiswami, and L. Davis. Efficient kernel machines using the improved fast Gauss transform. In L. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1561–1568. MIT Press, 2005.