SENSING-AWARE KERNEL SVM

Weicong Ding, Prakash Ishwar, Venkatesh Saligrama, and W. Clem Karl

Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA. {dingwc, pi, srv, wckarl}@bu.edu

ABSTRACT

We propose a novel approach for designing kernels for support vector machines (SVMs) when the class label is linked to the observation through a latent state and the likelihood function of the observation given the state (the sensing model) is available. We show that the Bayes-optimum decision boundary is a hyperplane under a mapping defined by the likelihood function. Combining this with the maximum margin principle yields kernels for SVMs that leverage knowledge of the sensing model in an optimal way. We derive the optimum kernel for the bag-of-words (BoWs) sensing model and demonstrate its superior performance over other kernels in document and image classification tasks. These results indicate that such optimum sensing-aware kernel SVMs can match the performance of rather sophisticated state-of-the-art approaches.

Index Terms— Sensing model, Kernel method, SVM, Bag of Words, Supervised Classification

1. INTRODUCTION

This paper presents a new provably optimum method for designing kernels for SVMs that explicitly incorporates information about the structure of the underlying data generating process. We consider the typical classification task where the observed data point $\mathbf{x} \in \mathcal{X}$ and label y follow some joint distribution $p(\mathbf{x}, y)$. In many real world problems, however, observations x are only indirectly related through some generative process to latent variables $\mathbf{z} \in \mathcal{Z}$ via a sensing model $p(\mathbf{x}|\mathbf{z})$ and the latent variables have different distributions conditioned on the underlying label y. An example of such a situation is medical tomography, where the acquired data x consists of X-ray based projections but the diagnosis is done on reconstructed cross-sections of the body z to determine the presence or absence of a disease y. Another example is the classic "Bag-of-Word" (BoW) modeling paradigm widely-used to model text documents, images, etc. [1][2], where the observed document \mathbf{x} is modeled as a collection of iid drawings of words from a latent document-specific distribution \mathbf{z} over the vocabulary. This \mathbf{z} is, in turn, related to the document category y through an unknown $p(\mathbf{z}|y)$.

For supervised classification, a generative approach would make further assumptions on $p(\mathbf{z}|y)$ which may either not hold or lead to an intractable posterior inference problem [3]. On the other hand, the classical distribution-free discriminative paradigm would ignore knowledge of the sensing model $p(\mathbf{x}|\mathbf{z})$ and build classifiers using labeled training data [4, 5]. Yet another approach is a decoupled two-step process where the latent states are first estimated as $\hat{\mathbf{z}}(\mathbf{x})$ using the sensing model and then a classifier is built based on $(\hat{\mathbf{z}}, y)$ pairs. Computing estimates, however, can be costly and can introduce confounding artifacts when data is limited or noisy, and may be unnecessary if the ultimate goal is a simple decision.

Motivated by such problems, we start by examining the structure of the Bayes-optimum classifier, focusing on binary classification, and then connect it to kernel SVMs via the max-margin principle. Our proposed "sensing-aware" kerneldesign arises as the natural consequence of such a procedure. We illustrate the approach on the BoW model and demonstrate its merits on document and image classification tasks.

2. OPTIMUM SENSING-AWARE KERNEL

Let $y \in \{-1, +1\}$, $z \in Z$, and $x \in X$ denote the class label, latent variable, and observation respectively. x and y are independent conditioned on z and the *sensing model* p(x|z) is known. We make the following technical assumptions on the joint distribution which are satisfied in many applications.

Assumption 1: (Square-integrability) $\forall \mathbf{x}, y, p(\mathbf{x}|\mathbf{z}), p(\mathbf{z}, y) \in \mathcal{L}^2(\mathcal{Z})$ as functions of \mathbf{z} .

Assumption 2: (Uniform boundedness) There exists $R \in [0,\infty)$ such that $\sup_{x,z} |p(\mathbf{x}|\mathbf{z})| \leq R$.

Further, let $\langle f, g \rangle_{\mathcal{Z}} := \int_{\mathcal{Z}} f(\mathbf{z}) g(\mathbf{z}) d\mathbf{z}$ denote the usual inner product in $\mathcal{L}^2(\mathcal{Z})$ which induces the norm $||f||_{\mathcal{Z}}^2 \triangleq \langle f, f \rangle_{\mathcal{Z}}$. Let $||f||_1 \triangleq \int_{\mathcal{Z}} |f(\mathbf{z})| d\mathbf{z}$ denote the ℓ_1 -norm in $\mathcal{L}^2(\mathcal{Z})$.

When full knowledge of the generative model $p(\mathbf{x}, y)$ is available, it is well known that the *Maximum Aposteriori Probability* (MAP) estimator of the class label minimizes the error probability, i.e., the Bayes risk. For binary classification, this reduces to a simple likelihood ratio test (LRT):

$$p(y=1|\mathbf{x}) \stackrel{1}{\underset{-1}{\gtrless}} p(y=-1|\mathbf{x}) \tag{1}$$

This article is based upon work supported by the U.S. AFOSR and the U.S. NSF under award numbers #FA9550-10-1-0458 (subaward #A1795) and #1218992 respectively.

Using $p(\mathbf{x}, \mathbf{z}, y) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z}, y)$ and marginalizing over \mathbf{z} the LRT (1) reduces to,

$$\langle p(\mathbf{x}|\mathbf{z}), w(\mathbf{z}) \rangle_{\mathcal{Z}} \stackrel{1}{\underset{-1}{\gtrless}} 0$$
 (2)

where $w(\mathbf{z}) := p(\mathbf{z}, y = 1) - p(\mathbf{z}, y = -1)$. Note that $w(\mathbf{z})$ satisfies the following constraint

 $||w(\mathbf{z})||_1 \le ||p(\mathbf{z}, y = 1)||_1 + ||p(\mathbf{z}, y = -1)||_1 = 1.$

Theorem: If we interpret $p(\mathbf{x}|\mathbf{z})$ and $w(\mathbf{z})$ as vectors (functions of \mathbf{z} with \mathbf{x} held fixed) in $\mathcal{L}^2(\mathcal{Z})$, then the Bayes-optimal rule (2) is a *linear* classifier in $\mathcal{L}^2(\mathcal{Z})$ defined by the separating hyperplane $w(\mathbf{z})$.

When a set of training samples $\{\mathbf{x}^i, y^i\}_{i=1}^M$ that are generated from $p(\mathbf{x}, y)$ in an iid fashion are given, and $w(\mathbf{z})$ is unknown, the popular **max-margin principle** advocates using the hyperplane $w^*(\mathbf{z})$ that maximizes the separation, i.e., a (soft) margin, between the two classes. This max-margin hyperplane $w^*(\mathbf{z})$ is the optimal solution of the following (infinite-dimensional) constrained optimization problem

$$w^{*}(\mathbf{z}) = \underset{w(\mathbf{z}) \in \mathcal{L}^{2}(\mathcal{Z}), \xi_{i}}{\operatorname{arg\,min}} \quad \frac{1}{2} \|w(\mathbf{z})\|_{\mathcal{Z}}^{2} + C \sum_{i=1}^{N} \xi_{i}$$

such that
$$y^{i} \langle p(\mathbf{x}^{i} | \mathbf{z}), w(\mathbf{z}) \rangle \geq 1 - \xi_{i}, \quad \xi_{i} \geq 0$$

$$\|w(\mathbf{z})\|_{1} \leq 1$$
(3)

And the optimum classifier is then given by

$$f^*(\mathbf{x}) = \operatorname{sign}\left(\langle p(\mathbf{x}|\mathbf{z}), w^*(\mathbf{z}) \rangle\right).$$

If we write $w^*(\mathbf{z}) = \sum_{i=1}^M \beta_i p(\mathbf{x}^i|\mathbf{z}) + w^{\perp}(\mathbf{z})$, where w^{\perp} is orthogonal to the linear span of $\{p(\mathbf{x}_i|\mathbf{z})\}_{i=1}^M$, and ignore the

orthogonal to the linear span of $\{p(\mathbf{x}_i | \mathbf{z})\}_{i=1}^{M}$, and ignore the constraint $||w(\mathbf{z})|| \le 1$, the optimization problem (3) reduces to

$$\min_{\beta_{i},\xi_{i},w^{\perp}} \quad \frac{1}{2} \sum_{i,j=1}^{M} \beta_{i}\beta_{j}K(\mathbf{x}^{i},\mathbf{x}^{j}) + \|w^{\perp}\|_{2}^{2} + C \sum_{i=1}^{N} \xi_{i}$$
such that
$$y^{i} \sum_{j=1}^{M} \beta_{j}K(\mathbf{x}^{i},\mathbf{x}^{j}) \ge 1 - \xi_{i}, \quad \xi_{i} \ge 0$$
(4)

where

$$K(\mathbf{x}, \mathbf{x}') \triangleq \langle p(\mathbf{x}|\mathbf{z}), \ p(\mathbf{x}'|\mathbf{z}) \rangle = \int_{\mathcal{Z}} p(\mathbf{x}|\mathbf{z}) p(\mathbf{x}'|\mathbf{z}) d\mathbf{z}.$$
 (5)

Note that the optimal solution of (4) has $w^{\perp} = 0$ making it a finite-dimensional constrained optimization problem. Equation (4) is, in fact, a kernel SVM problem with the kernel defined by equation (5). We will refer to (5) as the "sensing-aware kernel" in what follows. The optimal classifier $f^*(\mathbf{x})$ corresponding to the solution to 4 has the following form

$$f(\mathbf{x}) = \operatorname{sign}\left(\sum_{i=1}^{n} \beta_i^* K(\mathbf{x}, \mathbf{x}^i)\right).$$
(6)

The sensing-aware kernel thus provides a *principled* way to incorporate the partial information about the generative model $p(\mathbf{x}|\mathbf{z})$. For a wide range of $p(\mathbf{x}|\mathbf{z})$, $K(\mathbf{x}, \mathbf{x}')$ has a

closed form expression and can be calculated efficiently. In this paper, we focus on the BoW generative model.

Remark 1: If $\mathbf{t}(\mathbf{x})$ is a sufficient statistic for \mathbf{z} , i.e., $\mathbf{x} - \mathbf{t}(\mathbf{x}) - \mathbf{z}$ is a Markov chain, then $\mathbf{t}(\mathbf{x})$ can replace \mathbf{x} in the above analysis and the resulting kernel SVM will have the kernel $K(\mathbf{x}, \mathbf{x}') = \int_{\mathcal{Z}} p(\mathbf{t}(\mathbf{x})|\mathbf{z})p(\mathbf{t}(\mathbf{x}')|\mathbf{z})d\mathbf{z}$.

Remark 2: There is a vast literature dedicated to kernel methods for classification[5]. While the standard RBF kernel works well in many problems, various kernels have been designed for specific tasks. Among them, a line of work on model-based kernels is related to our approach in the sense that knowledge of a generative model is utilized. The Probability Product Kernel (PPK) proposed in [6] first estimates latent variable \hat{z}^i using observation x^i and defines the kernel in terms of the estimate as $K(\mathbf{x}^i, \mathbf{x}^j) := \int_{\mathcal{X}} p(\mathbf{x} | \hat{\mathbf{z}}^i) p(\mathbf{x} | \hat{\mathbf{z}}^j) d\mathbf{x}$. Heat or Diffusion kernels (Diff) that exploit the Fisher information metric on the probability manifold were proposed in [7]. Kernels based on the KL divergence were proposed in [8]. Although each approach has its own strength, unlike ours, the aforementioned probabilistic kernels are not designed to directly minimize the classification error. Moreover, they require full model knowledge which is unreasonable for large and complex datasets like text or images.

3. KERNEL FOR BAG OF WORDS MODELS

In the "Bag of words" (BoW), or "Bag of Features" (BoF) model, there is a collection of *documents* composed of words from a vocabulary of size W. Each document is modeled as being generated by N i.i.d drawings of words from a latent $W \times 1$ document word-distribution vector \mathbf{z} . Since the ordering is ignored, a document is represented as a $W \times 1$ empirical word-count vector $\mathbf{x} = (x_1, \ldots, x_W)^T$. Therefore, $p(\mathbf{x}|\mathbf{z})$ is a multinomial distribution

$$p(\mathbf{x}|\mathbf{z}) = \binom{N}{x_1, \dots, x_W} \prod_{w=1}^W z_w^{x_w}$$

From this it can be shown that the sensing kernel for two documents $\mathbf{x}^i, \mathbf{x}^j$ defined in 5 has the following form,

$$K(\mathbf{x}^{i}, \mathbf{x}^{j}) = \prod_{w=1}^{W} \frac{\Gamma(x_{w}^{i} + x_{w}^{j} + 1)}{\Gamma(x_{w}^{i} + 1)\Gamma(x_{w}^{j} + 1)} \frac{\Gamma(N_{i} + 1)\Gamma(N_{j} + 1)}{\Gamma(N_{i} + N_{j} + W)}$$
$$= \prod_{w=1}^{W} \frac{(x_{w}^{i} + x_{w}^{j})!}{(x_{w}^{i})!(x_{w}^{j})!} \frac{N_{i}!N_{j}!}{(N_{i} + N_{j} + W - 1)!}$$
(7)

where N_i is the number of words in the *i*-th document and $\Gamma(t)$ is the ordinary Gamma function. One practical concern is the size of the factorials. Typically, W is a large constant and N_i varies across documents. Therefore, the product of factorials will have a large dynamic range and lead to overflow errors (even for small W). The product is sensitive to the differences in the N_i values across documents. These issues can be addressed by using one of the following approximations to the original kernel 7.

- Sensing 0 kernel: $K^0(\mathbf{x}^i, \mathbf{x}^j) := \log(K(\mathbf{x}^i, \mathbf{x}^j)).$
- Sensing 1 kernel: Following (7),

$$K_n^1(\mathbf{x}^i, \mathbf{x}^j) := \log\left\{\prod_{w=1}^W \frac{\Gamma(n\bar{x}_w^i + n\bar{x}_w^j + 1)}{\Gamma(n\bar{x}_w^i + 1)\Gamma(n\bar{x}_w^j + 1)}\right\}$$

where $\bar{\mathbf{x}}^i := \mathbf{x}^i / N_i$ is the word-frequency (normalized word-count) and n is set to be some constant.

• Sensing 2 kernel: $K_N^2(\mathbf{x}^i, \mathbf{x}^j) = \log(K(\hat{\mathbf{x}}_N^i, \hat{\mathbf{x}}_N^j))$ where $\hat{\mathbf{x}}_N^i$ are the words counts obtained by randomly resampling the *i*-th document, uniformly, N times.

In general, these approximations balance the number of words of documents and avoid numerical issues by a log-mapping that compresses the dynamic range. In typical experimental settings, we have found that these approximations yield similar performance.

4. DOCUMENT CLASSIFICATION

In this section, we consider the problem of text document classification. In this application context, the terms "word", "vocabulary", and "document" have their natural meanings. We selected the 20 *Newsgourps* dataset which contains 18774 news postings from 20 newsgroups (classes) to test performance. The average number of words per document in this dataset is 117. Following standard practice [9], we removed a standard list of stop words from the vocabulary.

4.1. Binary Classification

We chose to distinguish postings in the newsgroup *alt.atheism* from *talk.religion.misc*. This is a difficult task due to the similarity of content in these two groups [3, 10]. The training set contains 856 documents and the test set contains 569 documents. The average number of words per document is 132.

 Table 1. Binary classification accuracy for *alt.atheism* vs.

 talk.religion.misc in the 20 Newsgroups dataset.

Method	CCR %	Method	CCR %
Sensing1 SVM	82.07	RBF SVM	78.91
Sensing2 SVM	83.35	DiscLDA [3]	83.00
PPK SVM[6]	81.02	G-MedLDA [10]	83.57
Diff SVM [7]	79.96		

We used LIBSVM [11] to train our kernel SVMs. We report results only for the Sensing 1 and Sensing 2 approximations of our sensing-aware kernel since in this dataset, the number of words N_i varies significantly across documents. Table 1 compares the Correct Classification Rate (CCR %) of the proposed sensing-aware kernels against two modelbased kernels, PPK and Diff, and the baseline RBF kernel. Table 1 also shows results for a discriminative Latent Dirichlet Allocation method (denoted by DiscLDA) [3] and a recent method that is based on a max-margin supervised topic model (denoted by G-MedLDA) [10]. Both of these methods posit more complex models and represent the current stateof-the-art. The free model parameters such as n and N in sensing kernels 1 and 2, t in Diff, and σ in the RBF kernel, were tuned using a 5 fold cross-validation. The CCR for the Gibbs MedLDA method is for the best number of topics k as in [10]. Words in the test set that were not in the training set were dropped.

We can see that the proposed Sensing 1 and 2 kernels outperform the RBF kernel and the model-based kernels PPK and Diff. Surprisingly, the performance of our very simple method is better than that of DiscLDA and is almost the same as that of G-MedLDA (with only 1 less document correctly classified). These state-of-the-art methods make use of fairly complex models and require sophisticated inference algorithms while our sensing-aware kernel SVM makes minimal modeling assumptions and has much lower time complexity.

4.2. Multi-class Classification

We next studied multi-class classification in the 20 Newsgroups dataset with all 20 classes. We adopted a widely used training/test split where the training set consists of 11,269 documents, and the test set consists of 7,505 documents. We used the "one-versus-all" strategy following [10] to do multiclass classification with binary classifiers. We followed the same settings as in the binary case. For the Sensing 1 and 2 kernels, we used n = N = 150. Table 2 shows the CCRs for the proposed sensing-aware kernels, the RBF baseline, the two model-based kernels PPK and Diff, and the G-MedLDA [10] algorithm. Also shown is the CCR for a recently developed deep learning method based on a Restricted Boltzmann Machine (RBM) [12] which outperforms the standard RBF SVM. Deep learning has recently emerged as a powerful generic approach and has attained state-of-the-art performance in many applications. Since [12] uses the same training settings, we simply quote the results reported in that paper. Observe that the approximate Sensing kernels 1 and 2

 Table 2. Multi-class classification accuracy for the 20 Newsgroups dataset.

Method	CCR %	Method	CCR %
Sensing1 SVM	79.5	RBF SVM	75.9
Sensing2 SVM	80.5	RBM [12]	76.2
PPK SVM[6]	78.2	G-MedLDA [10]	80.9
Diff SVM [7]	78.0		

outperform all other kernel SVMs and the RBM. As in binary classification, their performance is close to the G-MedLDA method that uses complex models.

5. IMAGE CLASSIFICATION

In this section, we consider the problem of recognizing image scene categories. We use the Natural Scene category dataset first introduced in [2, 13]. This dataset consists of 15 scene categories, e.g., office, street view, forests, etc., with 200-400 grayscale images in each category (total 4485 images) and an average image size of 300×250 pixels.

5.1. Modeling Images as Bags of Features

There is extensive literature on BOF models for images. A typical model consists of 1) local feature extraction, 2) visual-vocabulary construction, and 3) image representation as BoFs. In order to highlight the effect of kernels, we follow the baseline approach proposed in [2, 13, 14].

1) Local feature extraction. For each image, the SIFT descriptors of all the 16×16 image patches centered at grid points spaced 8 pixels apart are computed.

2) Visual vocabulary construction. D patches are randomly selected from training images. W-means clustering is performed over the corresponding D SIFT descriptors. The W mean vectors form a visual vocabulary, labeled as $1, \ldots, W$. 3) BoF representation. For each image, the SIFT descriptor of each patch is assigned the label of its nearest neighbor in the visual vocabulary. An image is then a collection of visual words from the vocabulary and the spatial correlation is ignored.

To incorporate spatial structure, we use the pyramid matching scheme proposed in [13]. This uses a sequence of coarse-to-fine grids. At level L, an image is split into $2^L \times 2^L$ non-overlapping cells, and each cell is treated as a separate "document". Hence, at level L, each image \mathbf{x}^i is represented as a collection of 2^{2L} word-frequency vectors $\{\mathbf{x}_{L,c}^i, c = 1, \ldots, 2^{2L}\}$, ordered according to the spatial position of the cell. The kernel at level L for two images $\mathbf{x}^i, \mathbf{x}^j$ is defined as $K^L(\mathbf{x}^i, \mathbf{x}^j) = \sum_{c=1}^{2^{2L}} \kappa(\mathbf{x}_{L,c}^i, \mathbf{x}_{L,c}^j)$, where κ is the kernel function for two BoF documents.

The above kernel is only for a single level. The pyramid kernel K for two images is defined as a weighted sum of single-level kernels: $K(\mathbf{x}^i, \mathbf{x}^j) = \sum_{l=0}^{L} \alpha^l K^l(\mathbf{x}^i, \mathbf{x}^j)$. We use the same weighting factors α^l as in [13].

5.2. Experimental Results

We follow the settings in [13, 2]. We set the vocabulary size as W = 400. We randomly select 100 images per class for training and leave the rest for testing. We repeat experiments 10 times with a randomly selected training/testing split for each run. We use the "one-versus-all" strategy for multi-class classification and report the average CCR over 10 random runs.

Following [13], we consider three setups. "L = 0, single" uses K^0 as the kernel for classification. This simply views the entire image as a document. Similarly, "L = 2, single" uses K^2 as the kernel for classification. In this case, the image is divided into $2^{2L} = 16$ regions. "L = 2, Pyramid" uses $K = \sum_{l=0}^{2} \alpha^l K^l$ as the kernel value.

We compare our Sensing 0 and Sensing 2 kernels with the probabilistic model-based PPK and Diff kernels and the standard RBF kernel. We also compare with the classic Spatial Pyramid (SP) kernel of [13]. This kernel is specially crafted for this problem and has good empirical performance in several Computer Vision tasks. We also consider a recent work called Reconfiguration BoW (RBoW), which makes additional modeling assumptions on $p(\mathbf{z}|y)$. Finally, we consider a recently developed deep learning algorithm [15] (denoted by MRF) which achieves state-of-the-art performance on the same task.

For the Sensing 2 kernel, we simply set N = 500. As in text classification, all the other free parameters are tuned by a 5-fold cross-validation. We measure performance using the overall classification accuracy (CCR %). The results are summarized in Table 3. Since the settings are identical, we simply quote the results for MRF and RBoW as reported in [15] and [16] respectively.

 Table 3. CCR % for the scene category database for the three different settings

	L = 0, single	L=2, single	L=2,Pyramid
Sensing 0	72.3	79.1	80.7
Sensing 2	72.9	79.7	81.3
PPK [6]	73.6	77.7	78.6
Diff [7]	74.7	77.1	77.8
RBF	65.6	72.2	73.5
SP [13]	74.8	79.3	81.4
RBOW [16]	N/A	78.6	N/A
MRF [15]		81.2	

We can see that in the "L = 2, single-level" setting, the proposed Sensing kernel 2 clearly improves over the classic kernel SP and outperforms other kernels (RBF, PPK and Diff) as well as the more complex RBoW method. However in the "L = 0 single-level" setting, the sensing-aware kernels do perform worse than PPK, Diff, and SP. This could be because L = 0 corresponds to the whole image. It has been strongly suggested in the Computer Vision literature that modeling the entire image as a BoF is unrealistic in many datasets [14] [13]. Our proposed sensing-aware approach only makes sense if the data is actually generated from the sensing structure. Our experimental results accord with the fact that BoF is a reasonable model for local regions, but not for the entire image.

We also did experiments for the "L = 2, pyramid" setting. We can see that the Sensing 2 kernel achieves almost the same performance as the state-of-the-art SP and MRF. The improvement is not as significant as in the "L = 2 singlelevel" case. This could be because it involves the L = 0 level kernel values, which does not fit the BoF model well.

6. CONCLUSION

In this paper, we proposed a novel kernel design principle to incorporate partial model information. On two distinct types of data, we showed that with minimal modeling assumptions, the sensing-aware kernel improves upon other standard kernels and handcrafted kernels for specific domains. We also observed that the sensing-aware kernel can match the performance of the state-of-the-art approaches.

7. REFERENCES

- D. Blei, "Probabilistic topic models," *Commun. of the ACM*, vol. 55, no. 4, pp. 77–84, 2012.
- [2] Fei-Fei Li and P. Perona, "A bayesian hierarchical model for learning natural scene categories," in *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (CVPR'05), San Diego, CA, Jun. 2005, pp. 524–531.
- [3] S. Lacoste-Julien, F. Sha, and M. Jordan, "Disclda: Discriminative learning for dimensionality reduction and classification," in *Advances in Neural Information Processing Systems 21(NIPS'09)*, pp. 897–904. 2009.
- [4] V. N. Vapnik, *The nature of statistical learning theory*, Springer-Verlag New York, Inc., New York, NY, USA, 1995.
- [5] T. Hastie, R. Tibshirani, and J. Friedman, *Elements of Statistical Learning*, Springer-Verlag New York, Inc., New York, NY, USA, 2009.
- [6] T. Jebara, R. Kondor, and A. Howard, "Probability product kernels," *J. Mach. Learn. Res.*, vol. 5, pp. 819–844, Dec. 2004.
- [7] J. Lafferty and G. Lebanon, "Diffusion kernels on statistical manifolds," J. Mach. Learn. Res., vol. 6, pp. 129–163, Dec. 2005.
- [8] Pedro J. Moreno, Purdy P. Ho, and Nuno Vasconcelos, "A kullback-leibler divergence based kernel for svm classification in multimedia applications," in Advances in Neural Information Processing Systems 16 (NIPS'03). 2004.
- [9] D. D. Lewis, Y. Yang, T. G. Rose, and F. Li, "Rcv1: A new benchmark collection for text categorization research," *J. Mach. Learn. Res.*, vol. 5, pp. 361–397, Dec. 2004.
- [10] J. Zhu, N. Chen, H. Perkins, and B. Zhang, "Gibbs maxmargin topic models with fast sampling algorithms," in *the 30th Int. Conf. on Machine Learning (ICML'13)*, Atlanta, GA, Jun. 2013.
- [11] C. Chang and C. Lin, "LIBSVM: A library for support vector machines," ACM Transactions on Intelligent Systems and Technology, vol. 2, pp. 27:1–27:27, 2011.
- [12] H. Larochelle, M. Mandel, R. Pascanu, and Y. Bengio, "Learning algorithms for the classification restricted boltzmann machine," *J. Mach. Learn. Res.*, vol. 13, pp. 643–669, Mar. 2012.

- [13] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, New York, NY, Jun. 2006, pp. 2169–2178.
- [14] K. Grauman and T. Darrell, "The pyramid match kernel: Discriminative classification with sets of image features," in *Proc. the 10th IEEE International Conference on Computer Vision(ICCV'05)*, Beijing, China, Oct. 2005, pp. 1458–1465.
- [15] M. Ranzato, V. Mnih, J.M. Susskind, and G.E. Hinton, "Modeling natural images using gated mrfs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 9, pp. 2206–2222, 2013.
- [16] Sobhan Naderi Parizi, John Oberlin, and Pedro F. Felzenszwalb, "Reconfigurable models for scene recognition," in *Proceedings of the 2012 IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 2775–2782.