
Multi-class SVMs: From Tighter Data-Dependent Generalization Bounds to Novel Algorithms

Yunwen Lei

Department of Mathematics
City University of Hong Kong
yunwelei@cityu.edu.hk

Ürün Dogan

Microsoft Research
Cambridge CB1 2FB, UK
udogan@microsoft.com

Alexander Binder

ISTD Pillar
Singapore University of Technology and Design
Machine Learning Group, TU Berlin
alexander.binder@sutd.edu.sg

Marius Kloft

Department of Computer Science
Humboldt University of Berlin
kloft@hu-berlin.de

Abstract

This paper studies the generalization performance of multi-class classification algorithms, for which we obtain—for the first time—a data-dependent generalization error bound with a *logarithmic* dependence on the class size, substantially improving the state-of-the-art linear dependence in the existing data-dependent generalization analysis. The theoretical analysis motivates us to introduce a new multi-class classification machine based on ℓ_p -norm regularization, where the parameter p controls the complexity of the corresponding bounds. We derive an efficient optimization algorithm based on Fenchel duality theory. Benchmarks on several real-world datasets show that the proposed algorithm can achieve significant accuracy gains over the state of the art.

1 Introduction

Typical multi-class application domains such as natural language processing [1], information retrieval [2], image annotation [3] and web advertising [4] involve tens or hundreds of thousands of classes, and yet these datasets are still growing [5]. To handle such learning tasks, it is essential to build algorithms that scale favorably with respect to the number of classes. Over the past years, much progress in this respect has been achieved on the algorithmic side [4–7], including efficient stochastic gradient optimization strategies [8].

Although also theoretical properties such as consistency [9–11] and finite-sample behavior [1, 12–15] have been studied, there still is a discrepancy between algorithms and theory in the sense that the corresponding theoretical bounds do often not scale well with respect to the number of classes. This discrepancy occurs the most strongly in research on *data-dependent* generalization bounds, that is, bounds that can measure generalization performance of prediction models purely from the training samples, and which thus are very appealing in model selection [16]. A crucial advantage of these bounds is that they can better capture the properties of the distribution that has generated the data, which can lead to tighter estimates [17] than conservative data-independent bounds.

To our best knowledge, for multi-class classification, the first data-dependent error bounds were given by [14]. These bounds exhibit a quadratic dependence on the class size and were used by [12] and [18] to derive bounds for kernel-based multi-class classification and multiple kernel learning (MKL) problems, respectively. More recently, [13] improve the quadratic dependence to a linear dependence by introducing a novel surrogate for the multi-class margin that is independent on the true realization of the class label.

However, a heavy dependence on the class size, such as linear or quadratic, implies a poor generalization guarantee for large-scale multi-class classification problems with a massive number of classes. In this paper, we show data-dependent generalization bounds for multi-class classification problems that—for the first time—exhibit a *sublinear* dependence on the number of classes. Choosing appropriate regularization, this dependence can be as mild as logarithmic. We achieve these improved bounds via the use of Gaussian complexities, while previous bounds are based on a well-known structural result on Rademacher complexities for classes induced by the maximum operator. The proposed proof technique based on Gaussian complexities exploits potential coupling among different components of the multi-class classifier, while this fact is ignored by previous analyses.

The result shows that the generalization ability is strongly impacted by the employed regularization. Which motivates us to propose a new learning machine performing block-norm regularization over the multi-class components. As a natural choice we investigate here the application of the proven ℓ_p norm [19]. This results in a novel ℓ_p -norm multi-class support vector machine (MC-SVM), which contains the classical model by Crammer & Singer [20] as a special case for $p = 2$. The bounds indicate that the parameter p crucially controls the complexity of the resulting prediction models.

We develop an efficient optimization algorithm for the proposed method based on its Fenchel dual representation. We empirically evaluate its effectiveness on several standard benchmarks for multi-class classification taken from various domains, where the proposed approach significantly outperforms the state-of-the-art method of [20].

The remainder of this paper is structured as follows. Section 2 introduces the problem setting and presents the main theoretical results. Motivated by which we propose a new multi-class classification model in Section 3 and give an efficient optimization algorithm based on Fenchel duality theory. In Section 4 we evaluate the approach for the application of visual image recognition and on several standard benchmark datasets taken from various application domains. Section 5 concludes.

2 Theory

2.1 Problem Setting and Notations

This paper considers multi-class classification problems with $c \geq 2$ classes. Let \mathcal{X} denote the input space and $\mathcal{Y} = \{1, 2, \dots, c\}$ denote the output space. Assume that we are given a sequence of examples $S = \{(x_1, y_1), \dots, (x_n, y_n)\} \in (\mathcal{X} \times \mathcal{Y})^n$, independently drawn according to a probability measure P defined on the sample space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. Based on the training examples S , we wish to learn a prediction rule h from a space H of hypotheses mapping from \mathcal{Z} to \mathbb{R} and use the mapping $x \rightarrow \arg \max_{y \in \mathcal{Y}} h(x, y)$ to predict (ties are broken by favoring classes with a lower index, for which our loss function defined below always counts an error). For any hypothesis $h \in H$, the margin $\rho_h(x, y)$ of the function h at a labeled example (x, y) is $\rho_h(x, y) := h(x, y) - \max_{y' \neq y} h(x, y')$. The prediction rule h makes an error at (x, y) if $\rho_h(x, y) \leq 0$ and thus the expected risk incurred from using h for prediction is $R(h) := \mathbb{E}[1_{\rho_h(x, y) \leq 0}]$.

Any function $h : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ can be equivalently represented by the vector-valued function (h_1, \dots, h_c) with $h_j(x) = h(x, j), \forall j = 1, \dots, c$. We denote by $\tilde{H} := \{\rho_h : h \in H\}$ the class of margin functions associated to H . Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a Mercer kernel with ϕ being the associated feature map, i.e., $k(x, \tilde{x}) = \langle \phi(x), \phi(\tilde{x}) \rangle$ for all $x, \tilde{x} \in \mathcal{X}$. We denote by $\|\cdot\|_*$ the dual norm of $\|\cdot\|$, i.e., $\|w\|_* := \sup_{\|\tilde{w}\| \leq 1} \langle w, \tilde{w} \rangle$. For a convex function f , we denote by f^* its Fenchel conjugate, i.e., $f^*(v) := \sup_w [\langle w, v \rangle - f(w)]$. For any $\mathbf{w} = (w_1, \dots, w_c)$ we define the $\ell_{2,p}$ -norm by $\|\mathbf{w}\|_{2,p} := [\sum_{j=1}^c \|w_j\|_2^p]^{1/p}$. For any $p \geq 1$, we denote by p^* the dual exponent of p satisfying $1/p + 1/p^* = 1$ and $\bar{p} := p(2-p)^{-1}$. We require the following definitions.

Definition 1 (Strong Convexity). A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is said to be β -strongly convex w.r.t. a norm $\|\cdot\|$ iff $\forall x, y \in \mathcal{X}$ and $\forall \alpha \in (0, 1)$, we have

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y) - \frac{\beta}{2} \alpha(1 - \alpha) \|x - y\|^2.$$

Definition 2 (Regular Loss). We call ℓ a L -regular loss if it satisfies the following properties:

- (i) $\ell(t)$ bounds the 0-1 loss from above: $\ell(t) \geq 1_{t \leq 0}$;
- (ii) ℓ is L -Lipschitz in the sense $|\ell(t_1) - \ell(t_2)| \leq L|t_1 - t_2|$;

(iii) $\ell(t)$ is decreasing and it has a zero point c_ℓ , i.e., $\ell(c_\ell) = 0$.

Some examples of L -regular loss functions include the hinge $\ell_h(t) = (1 - t)_+$ and the margin loss

$$\ell_\rho(t) = 1_{t \leq 0} + (1 - t\rho^{-1})1_{0 < t \leq \rho}, \quad \rho > 0. \quad (1)$$

2.2 Main results

Our discussion on data-dependent generalization error bounds is based on the established methodology of Rademacher and Gaussian complexities [21].

Definition 3 (Rademacher and Gaussian Complexity). *Let H be a family of real-valued functions defined on \mathcal{Z} and $S = (z_1, \dots, z_n)$ a fixed sample of size n with elements in \mathcal{Z} . Then, the empirical Rademacher and Gaussian complexities of H with respect to the sample S are defined by*

$$\mathfrak{R}_S(H) = \mathbb{E}_\sigma \left[\sup_{h \in H} \frac{1}{n} \sum_{i=1}^n \sigma_i h(z_i) \right], \quad \mathfrak{G}_S(H) = \mathbb{E}_g \left[\sup_{h \in H} \frac{1}{n} \sum_{i=1}^n g_i h(z_i) \right],$$

where $\sigma_1, \dots, \sigma_n$ are independent random variables with equal probability taking values $+1$ or -1 , and g_1, \dots, g_n are independent $N(0, 1)$ random variables.

Note that we have the following comparison inequality relating Rademacher and Gaussian complexities (Cf. Section 4.2 in [22]):

$$\mathfrak{R}_S(H) \leq \sqrt{\frac{\pi}{2}} \mathfrak{G}_S(H) \leq 3\sqrt{\frac{\pi}{2}} \sqrt{\log n} \mathfrak{R}_S(H). \quad (2)$$

Existing work on data-dependent generalization bounds for multi-class classifiers [12–14, 18] builds on the following structural result on Rademacher complexities (e.g., [12], Lemma 8.1):

$$\mathfrak{R}_S(\max\{h_1, \dots, h_c\} : h_j \in H_j, j = 1, \dots, c) \leq \sum_{j=1}^c \mathfrak{R}_S(H_j), \quad (3)$$

where H_1, \dots, H_c are c hypothesis sets. This result is crucial for the standard generalization analysis of multi-class classification since the margin ρ_h involves the maximum operator, which is removed by (3), but at the expense of a linear dependency on the class size. In the following we show that this linear dependency is suboptimal because (3) does not take into account the coupling among different classes. For example, a common regularizer used in multi-class learning algorithms is $r(h) = \sum_{j=1}^c \|h_j\|_2^2$ [20], for which the components h_1, \dots, h_c are correlated via a $\|\cdot\|_{2,2}$ regularizer, and the bound (3) ignoring this correlation would not be effective in this case [12–14, 18].

As a remedy, we here introduce a new structural complexity result on function classes induced by general classes via the maximum operator, while allowing to preserve the correlations among different components meanwhile. Instead of considering the Rademacher complexity, Lemma 4 concerns the structural relationship of Gaussian complexities since it is based on a comparison result among different Gaussian processes.

Lemma 4 (Structural result on Gaussian complexity). *Let H be a class of functions defined on $\mathcal{X} \times \mathcal{Y}$ with $\mathcal{Y} = \{1, \dots, c\}$. Let g_1, \dots, g_{nc} be independent $N(0, 1)$ distributed random variables. Then, for any sample $S = \{x_1, \dots, x_n\}$ of size n , we have*

$$\mathfrak{G}_S(\{\max\{h_1, \dots, h_c\} : h = (h_1, \dots, h_c) \in H\}) \leq \frac{1}{n} \mathbb{E}_g \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i), \quad (4)$$

where \mathbb{E}_g denotes the expectation w.r.t. to the Gaussian variables g_1, \dots, g_{nc} .

The proof of Lemma 4 is given in Supplementary Material A. Equipped with Lemma 4, we are now able to present a general data-dependent margin-based generalization bound. The proof of the following results (Theorem 5, Theorem 7 and Corollary 8) is given in Supplementary Material B.

Theorem 5 (Data-dependent generalization bound for multi-class classification). *Let $H \subset \mathbb{R}^{\mathcal{X} \times \mathcal{Y}}$ be a hypothesis class with $\mathcal{Y} = \{1, \dots, c\}$. Let ℓ be a L -regular loss function and denote $B_\ell := \sup_{(x,y), h} \ell(\rho_h(x, y))$. Suppose that the examples $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ are independently*

drawn from a probability measure defined on $\mathcal{X} \times \mathcal{Y}$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, the following multi-class classification generalization bound holds for any $h \in H$:

$$R(h) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_h(x_i, y_i)) + \frac{2L\sqrt{2\pi}}{n} \mathbb{E}_{\mathbf{g}} \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i) + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}},$$

where g_1, \dots, g_{nc} are independent $N(0, 1)$ distributed random variables.

Remark 6. Under the same condition of Theorem 5, [12] derive the following data-dependent generalization bound (Cf. Corollary 8.1 in [12]):

$$R(h) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_h(x_i, y_i)) + \frac{4Lc}{n} \mathfrak{R}_S(\{x \rightarrow h(x, y) : y \in \mathcal{Y}, h \in H\}) + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

This linear dependence on c is due to the use of (3). For comparison, Theorem 5 implies that the dependence on c is governed by the term $\sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i)$, an advantage of which is that the components h_1, \dots, h_c are jointly coupled. As we will see, this allows us to derive an improved result with a favorable dependence on c , when a constraint is imposed on (h_1, \dots, h_c) . \square

The following Theorem 7 applies the general result in Theorem 5 to kernel-based methods. The hypothesis space is defined by imposing a constraint with a general strongly convex function.

Theorem 7 (Data-dependent generalization bound for kernel-based multi-class learning algorithms and MC-SVMs). *Suppose that the hypothesis space is defined by*

$$H := H_{f, \Lambda} = \{h^{\mathbf{w}} = (\langle \mathbf{w}_1, \phi(x) \rangle, \dots, \langle \mathbf{w}_c, \phi(x) \rangle) : f(\mathbf{w}) \leq \Lambda\},$$

where f is a β -strongly convex function w.r.t. a norm $\|\cdot\|$ defined on H satisfying $f^*(0) = 0$. Let ℓ be a L -regular loss function and denote $B_\ell := \sup_{(x, y), h} \ell(\rho_h(x, y))$. Let g_1, \dots, g_{nc} be independent $N(0, 1)$ distributed random variables. Then, for any $\delta > 0$, with probability at least $1 - \delta$ we have

$$R(h^{\mathbf{w}}) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_{h^{\mathbf{w}}}(x_i, y_i)) + \frac{4L}{n} \sqrt{\frac{\pi\Lambda}{\beta} \mathbb{E}_{\mathbf{g}} \sum_{i=1}^n \left\| (g_{(j-1)n+i} \phi(x_i))_{j=1, \dots, c} \right\|_*^2} + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

We now consider the following specific hypothesis spaces using a $\|\cdot\|_{2,p}$ constraint:

$$H_{p, \Lambda} := \{h^{\mathbf{w}} = (\langle \mathbf{w}_1, \phi(x) \rangle, \dots, \langle \mathbf{w}_c, \phi(x) \rangle) : \|\mathbf{w}\|_{2,p} \leq \Lambda\}, \quad 1 \leq p \leq 2. \quad (5)$$

Corollary 8 (ℓ_p -norm MC-SVM generalization bound). *Let ℓ be a L -regular loss function and denote $B_\ell := \sup_{(x, y), h} \ell(\rho_h(x, y))$. Then, with probability at least $1 - \delta$, for any $h^{\mathbf{w}} \in H_{p, \Lambda}$ the generalization error $R(h^{\mathbf{w}})$ can be upper bounded by:*

$$\frac{1}{n} \sum_{i=1}^n \ell(\rho_{h^{\mathbf{w}}}(x_i, y_i)) + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}} + \frac{2L\Lambda}{n} \sqrt{\sum_{i=1}^n k(x_i, x_i)} \times \begin{cases} \sqrt{e}(4 \log c)^{1 + \frac{1}{2 \log c}}, & \text{if } p^* \geq 2 \log c, \\ (2p^*)^{1 + \frac{1}{p^*}} c^{\frac{1}{p^*}}, & \text{otherwise.} \end{cases}$$

Remark 9. The bounds in Corollary 8 enjoy a mild dependence on the number of classes. The dependence is polynomial with exponent $1/p^*$ for $2 < p^* < 2 \log c$ and becomes logarithmic if $p^* \geq 2 \log c$. Even in the theoretically unfavorable case of $p = 2$ [20], the bounds still exhibit a radical dependence on the number of classes, which is substantially milder than the quadratic dependence established in [12, 14, 18] and the linear dependence established in [13]. Our generalization bound is data-dependent and shows clearly how the margin would affect the generalization performance (when ℓ is the margin loss ℓ_ρ): a large margin ρ would increase the empirical error while decrease the model's complexity, and vice versa. \square

2.3 Comparison of the Achieved Bounds to the State of the Art

Related work on data-independent bounds. The large body of theoretical work on multi-class learning considers data-independent bounds. Based on the ℓ_∞ -norm covering number bound of linear operators, [15] obtain a generalization bound exhibiting a linear dependence on the class size, which is improved by [9] to a radical dependence of the form $O(n^{-\frac{1}{2}} (\log^{\frac{3}{2}} n)^{\frac{\sqrt{c}}{\rho}})$. Under conditions

analogous to Corollary 8, [23] derive a class-size independent generalization guarantee. However, their bound is based on a delicate definition of margin, which is why it is commonly not used in the mainstream multi-class literature. [1] derive the following generalization bound

$$\mathbb{E} \left[\frac{1}{p} \log \left(1 + \sum_{\tilde{y} \neq y} e^{p(\rho - \langle \tilde{\mathbf{w}}_y - \tilde{\mathbf{w}}_{\tilde{y}}, \phi(x) \rangle)} \right) \right] \leq \inf_{\mathbf{w} \in H} \mathbb{E} \left[\frac{1}{p} \log \left(1 + \sum_{\tilde{y} \neq y} e^{p(\rho - \langle \mathbf{w}_y - \mathbf{w}_{\tilde{y}}, \phi(x) \rangle)} \right) \right] + \frac{\lambda n}{2(n+1)} \|\mathbf{w}\|_{2,2}^2 + \frac{2 \sup_{x \in \mathcal{X}} k(x, x)}{\lambda n}, \quad (6)$$

where ρ is a margin condition, $p > 0$ a scaling factor, and λ a regularization parameter. Eq. (6) is class-size independent, yet Corollary 8 shows superiority in the following aspects: first, for SVMs (i.e., margin loss ℓ_ρ), our bound consists of an empirical error ($\frac{1}{n} \sum_{i=1}^n \ell_\rho(\rho_{h, \mathbf{w}}(x_i, y_i))$) and a complexity term divided by the margin value (note that $L = 1/\rho$ in Corollary 8). When the margin is large (which is often desirable) [14], the last term in the bound given by Corollary 8 becomes small, while—on the contrary—the bound (6) is an increasing function of ρ , which is undesirable. Secondly, Theorem 7 applies to general loss functions, expressed through a strongly convex function over a general hypothesis space, while the bound (6) only applies to a specific regularization algorithm. Lastly, all the above mentioned results are conservative data-independent estimates.

Related work on data-dependent bounds. The techniques used in above mentioned papers do not straightforwardly translate to data-dependent bounds, which is the type of bounds in the focus of the present work. The investigation of these was initiated, to our best knowledge, by [14]: with the structural complexity bound (3) for function classes induced via the maximal operator, [14] derive a margin bound admitting a quadratic dependency on the number of classes. [12] use these results in [14] to study the generalization performance of MC-SVMs, where the components h_1, \dots, h_c are coupled with an $\|\cdot\|_{2,p}, p \geq 1$ constraint. Due to the usage of the suboptimal Eq. (3), [12] obtain a margin bound growing quadratically w.r.t. the number of classes. [18] develop a new multi-class classification algorithm based on a natural notion called the multi-class margin of a kernel. [18] also present a novel multi-class Rademacher complexity margin bound based on Eq. (3), and the bound also depends quadratically on the class size. More recently, [13] give a refined Rademacher complexity bound with a linear dependence on the class size. The key reason for this improvement is the introduction of $\rho_{\theta, h} := \min_{y' \in \mathcal{Y}} [h(x, y) - h(x, y') + \theta 1_{y'=y}]$ bounding margin ρ_h from below, and since the maximum operation in $\rho_{\theta, h}$ is applied to the set \mathcal{Y} rather than the subset $\mathcal{Y} - \{y_i\}$ for ρ_h , one needs not to consider the random realization of y_i . We also use this trick in our proof of Theorem 5. However, [13] fail to improve this linear dependence to a logarithmic dependence, as we achieved in Corollary 8, due to the use of the suboptimal structural result (3).

3 Algorithms

Motivated by the generalization analysis given in Section 2, we now present a new multi-class learning algorithm, based on performing empirical risk minimization in the hypothesis space (5). This corresponds to the following ℓ_p -norm MC-SVM ($1 \leq p \leq 2$):

Problem 10 (Primal problem: ℓ_p -norm MC-SVM).

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \left[\sum_{j=1}^c \|\mathbf{w}_j\|_2^p \right]^{\frac{2}{p}} + C \sum_{i=1}^n \ell(t_i), \\ \text{s.t. } \quad & t_i = \langle \mathbf{w}_{y_i}, \phi(x_i) \rangle - \max_{y \neq y_i} \langle \mathbf{w}_y, \phi(x_i) \rangle, \end{aligned} \quad (\text{P})$$

For $p = 2$ we recover the seminal multi-class algorithm by Crammer & Singer [20] (CS), which is thus a special case of the proposed formulation. An advantage of the proposed approach over [20] can be that, as shown in Corollary 8, the dependence of the generalization performance on the class size becomes milder as p decreases to 1.

3.1 Dual problems

Since the optimization problem (P) is convex, we can derive the associated dual problem for the construction of efficient optimization algorithms. The derivation of the following dual problem is deferred to Supplementary Material C. For a matrix $\alpha \in \mathbb{R}^{n \times c}$, we denote by α_i the i -th row. Denote by e_j the j -th unit vector in \mathbb{R}^c and $\mathbf{1}$ the vector in \mathbb{R}^c with all components being zero.

Problem 11 (Completely dualized problem for general loss). *The Lagrangian dual of (10) is:*

$$\begin{aligned} \sup_{\alpha \in \mathbb{R}^{n \times c}} & -\frac{1}{2} \left[\sum_{j=1}^c \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^{\frac{p}{p-1}} \right]^{\frac{2(p-1)}{p}} - C \sum_{i=1}^n \ell^* \left(-\frac{\alpha_{iy_i}}{C} \right) \\ \text{s.t. } & \alpha_{ij} \leq 0 \wedge \alpha_i \cdot \mathbf{1} = 0, \quad \forall j \neq y_i, i = 1, \dots, n. \end{aligned} \quad (\text{D})$$

Theorem 12 (REPRESENTER THEOREM). *For any dual variable $\alpha \in \mathbb{R}^{n \times c}$, the associated primal variable $\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_c)$ minimizing the Lagrangian saddle problem can be represented by:*

$$\mathbf{w}_j = \left[\sum_{\tilde{j}=1}^c \left\| \sum_{i=1}^n \alpha_{i\tilde{j}} \phi(x_i) \right\|_2^{p^*} \right]^{\frac{2}{p^*}-1} \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^{p^*-2} \left[\sum_{i=1}^n \alpha_{ij} \phi(x_i) \right].$$

For the hinge loss $\ell_h(t) = (1-t)_+$, we know its Fenchel-Legendre conjugate is $\ell_h^*(t) = t$ if $-1 \leq t \leq 0$ and ∞ otherwise. Hence $\ell_h^* \left(-\frac{\alpha_{iy_i}}{C} \right) = -\frac{\alpha_{iy_i}}{C}$ if $-1 \leq -\frac{\alpha_{iy_i}}{C} \leq 0$ and ∞ otherwise. Now we have the following dual problem for the hinge loss function:

Problem 13 (Completely dualized problem for the hinge loss (ℓ_p -norm MC-SVM)).

$$\begin{aligned} \sup_{\alpha \in \mathbb{R}^{n \times c}} & -\frac{1}{2} \left[\sum_{j=1}^c \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^{\frac{p}{p-1}} \right]^{\frac{2(p-1)}{p}} + \sum_{i=1}^n \alpha_{iy_i} \\ \text{s.t. } & \alpha_i \leq \mathbf{e}_{y_i} \cdot C \wedge \alpha_i \cdot \mathbf{1} = 0, \quad \forall i = 1, \dots, n. \end{aligned} \quad (7)$$

3.2 Optimization Algorithms

The dual problems (D) and (7) are not quadratic programs for $p \neq 2$, and thus generally not easy to solve. To circumvent this difficulty, we rewrite Problem 10 as the following equivalent problem:

$$\begin{aligned} \min_{\mathbf{w}, \beta} & \sum_{j=1}^c \frac{\|\mathbf{w}_j\|_2^2}{2\beta_j} + C \sum_{i=1}^n \ell(t_i) \\ \text{s.t. } & t_i \leq \langle \mathbf{w}_{y_i}, \phi(x_i) \rangle - \langle \mathbf{w}_y, \phi(x_i) \rangle, \quad y \neq y_i, i = 1, \dots, n, \\ & \|\beta\|_{\bar{p}} \leq 1, \bar{p} = p(2-p)^{-1}, \beta_j \geq 0. \end{aligned} \quad (8)$$

The class weights β_1, \dots, β_c in Eq. (8) play a similar role as the kernel weights in ℓ_p -norm MKL algorithms [19]. The equivalence between problem (P) and Eq. (8) follows directly from Lemma 26 in [24], which shows that the optimal $\beta = (\beta_1, \dots, \beta_c)$ in Eq. (8) can be explicitly represented in closed form. Motivated by the recent work on ℓ_p -norm MKL, we propose to solve the problem (8) via alternately optimizing \mathbf{w} and β . As we will show, given temporarily fixed β , the optimization of \mathbf{w} reduces to a standard multi-class classification problem. Furthermore, the update of β , given fixed \mathbf{w} , can be achieved via an analytic formula.

Problem 14 (Partially dualized problem for a general loss). *For fixed β , the partial dual problem for the sub-optimization problem (8) w.r.t. \mathbf{w} is*

$$\begin{aligned} \sup_{\alpha \in \mathbb{R}^{n \times c}} & -\frac{1}{2} \sum_{j=1}^c \beta_j \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^2 - C \sum_{i=1}^n \ell^* \left(-\frac{\alpha_{iy_i}}{C} \right) \\ \text{s.t. } & \alpha_{ij} \leq 0 \wedge \alpha_i \cdot \mathbf{1} = 0, \quad \forall j \neq y_i, i = 1, \dots, n. \end{aligned} \quad (9)$$

The primal variable \mathbf{w} minimizing the associated Lagrangian saddle problem is

$$\mathbf{w}_j = \beta_j \sum_{i=1}^n \alpha_{ij} \phi(x_i). \quad (10)$$

We defer the proof to Supplementary Material C. Analogous to Problem 13, we have the following partial dual problem for the hinge loss.

Problem 15 (Partially dualized problem for the hinge loss (ℓ_p -norm MC-SVM)).

$$\begin{aligned} \sup_{\alpha \in \mathbb{R}^{n \times c}} & f(\alpha) := -\frac{1}{2} \sum_{j=1}^c \beta_j \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^2 + \sum_{i=1}^n \alpha_{iy_i} \\ \text{s.t. } & \alpha_i \leq \mathbf{e}_{y_i} \cdot C \wedge \alpha_i \cdot \mathbf{1} = 0, \quad \forall i = 1, \dots, n. \end{aligned} \quad (11)$$

The Problems 14 and 15 are quadratic, so we can use the dual coordinate ascent algorithm [25] to very efficiently solve them for the case of linear kernels. To this end, we need to compute the gradient and solve the restricted problem of optimizing only one $\alpha_i, \forall i$, keeping all other dual variables fixed [25]. The gradient of f can be exactly represented by \mathbf{w} :

$$\frac{\partial f}{\partial \alpha_{ij}} = -\beta_j \sum_{\tilde{i}=1}^n \alpha_{\tilde{i}j} k(x_i, x_{\tilde{i}}) + 1_{y_i=j} = 1_{y_i=j} - \langle \mathbf{w}_j, \phi(x_i) \rangle. \quad (12)$$

Suppose the additive change to be applied to the current α_i is $\delta \alpha_i$, then from (12) we have

$$f(\alpha_1, \dots, \alpha_{i-1}, \alpha_i + \delta \alpha_i, \alpha_{i+1}, \dots, \alpha_n) = \sum_{j=1}^c \frac{\partial f}{\partial \alpha_{ij}} \delta \alpha_{ij} - \frac{1}{2} \sum_{j=1}^c \beta_j k(x_i, x_i) [\delta \alpha_{ij}]^2 + \text{const.}$$

Therefore, the sub-problem of optimizing $\delta \alpha_i$ is given by

$$\begin{aligned} \max_{\delta \alpha_i} \quad & -\frac{1}{2} \sum_{j=1}^c \beta_j k(x_i, x_i) [\delta \alpha_{ij}]^2 + \sum_{j=1}^c \frac{\partial f}{\partial \alpha_{ij}} \delta \alpha_{ij} \\ \text{s.t.} \quad & \delta \alpha_i \leq \mathbf{e}_{y_i} \cdot C - \alpha_i \wedge \delta \alpha_i \cdot \mathbf{1} = 0. \end{aligned} \quad (13)$$

We now consider the subproblem of updating class weights β with temporarily fixed \mathbf{w} , for which we have the following analytic solution. The proof is deferred to the Supplementary Material C.1.

Proposition 16. (Solving the subproblem with respect to the class weights) Given fixed \mathbf{w}_j , the minimal β_j optimizing the problem (8) is attained at

$$\beta_j = \|\mathbf{w}_j\|_2^{2-p} \left(\sum_{\tilde{j}=1}^c \|\mathbf{w}_{\tilde{j}}\|_2^p \right)^{\frac{p-2}{p}}. \quad (14)$$

The update of β_j based on Eq. (14) requires calculating $\|\mathbf{w}_j\|_2^2$, which can be easily fulfilled by recalling the representation established in Eq. (10).

The resulting training algorithm for the proposed ℓ_p -norm MC-SVM is given in Algorithm 1. The algorithm alternates between solving a MC-SVM problem for fixed class weights (Line 3) and updating the class weights in a closed-form manner (Line 5). Recall that Problem 11 establishes a completely dualized problem, which can be used as a sound stopping criterion for Algorithm 1.

Algorithm 1: Training algorithm for ℓ_p -norm MC-SVM.

input: examples $\{(x_i, y_i)_{i=1}^n\}$ and the kernel k .

initialize $\beta_j = \sqrt[p]{1/c}, \mathbf{w}_j = 0$ for all $j = 1, \dots, c$

while Optimality conditions are not satisfied **do**

optimize the multi-class classification problem (9)
compute $\|\mathbf{w}_j\|_2^2$ for all $j = 1, \dots, c$, according to Eq. (10)
update β_j for all $j = 1, \dots, c$, according to Eq. (14)

end

4 Empirical Analysis

We implemented the proposed ℓ_p -norm MC-SVM algorithm (Algorithm 1) in C++ and solved the involved MC-SVM problem using dual coordinate ascent [25]. We experiment on six benchmark datasets: the Sector dataset studied in [26], the News 20 dataset collected by [27], the Rcv1 dataset collected by [28], the Birds 15, Birds 50 as a part from [29] and the Caltech 256 collected by griffin2007caltech. We used fc6 features from the BVLC reference caffe-net from [30]. Table 1 gives an information on these datasets.

We compare with the classical CS in [20], which constitutes a strong baseline for these datasets [25]. We employ a 5-fold cross validation on the training set to tune the regularization parameter C by grid search over the set $\{2^{-12}, 2^{-11}, \dots, 2^{12}\}$ and p from 1.1 to 2 with 10 equidistant points. We repeat the experiments 10 times, and report in Table 2 on the average accuracy and standard deviations attained on the test set.

Dataset	No. of Classes	No. of Training Examples	No. of Test Examples	No. of Attributes
Sector	105	6,412	3,207	55,197
News 20	20	15,935	3,993	62,060
Rcv1	53	15,564	518,571	47,236
Birds 15	200	3,000	8,788	4,096
Birds 50	200	9,958	1,830	4,096
Caltech 256	256	12,800	16,980	4,096

Table 1: Description of datasets used in the experiments.

Method / Dataset	Sector	News 20	Rcv1	Birds 15	Birds 50	Caltech 256
ℓ_p -norm MC-SVM	94.20\pm0.34	86.19\pm0.12	85.74\pm0.71	13.73\pm1.4	27.86\pm0.2	56.00\pm1.2
Crammer & Singer	93.89 \pm 0.27	85.12 \pm 0.29	85.21 \pm 0.32	12.53 \pm 1.6	26.28 \pm 0.3	54.96 \pm 1.1

Table 2: Accuracies achieved by CS and the proposed ℓ_p -norm MC-SVM on the benchmark datasets.

We observe that the proposed ℓ_p -norm MC-SVM consistently outperforms CS [20] on all considered datasets. Specifically, our method attains 0.31% accuracy gain on Sector, 1.07% accuracy gain on News 20, 0.53% accuracy gain on Rcv1, 1.2% accuracy gain on Birds 15, 1.58% accuracy gain on Birds 50, and 1.04% accuracy gain on Birds 15. We perform a Wilcoxon signed rank test between the accuracies of CS and our method on the benchmark datasets, and the p-value is 0.03, which means our method is significantly better than CS at the significance level of 0.05. These promising results indicate that the proposed ℓ_p -norm MC-SVM could further lift the state of the art in multi-class classification, even in real-world applications beyond the ones studied in this paper.

5 Conclusion

Motivated by the ever growing size of multi-class datasets in real-world applications such as image annotation and web advertising, which involve tens or hundreds of thousands of classes, we studied the influence of the class size on the generalization behavior of multi-class classifiers. We focus here on data-dependent generalization bounds enjoying the ability to capture the properties of the distribution that has generated the data. Of independent interest, for hypothesis classes that are given as a maximum over base classes, we developed a new structural result on Gaussian complexities that is able to preserve the coupling among different components, while the existing structural results ignore this coupling and may yield suboptimal generalization bounds. We applied the new structural result to study learning rates for multi-class classifiers, and derived, for the first time, a data-dependent bound with a logarithmic dependence on the class size, which substantially outperforms the linear dependence in the state-of-the-art data-dependent generalization bounds.

Motivated by the theoretical analysis, we proposed a novel ℓ_p -norm MC-SVM, where the parameter p controls the complexity of the corresponding bounds. This class of algorithms contains the classical CS [20] as a special case for $p = 2$. We developed an effective optimization algorithm based on the Fenchel dual representation. For several standard benchmarks taken from various domains, the proposed approach surpassed the state-of-the-art method of CS [20] by up to 1.5%.

A future direction will be to derive a data-dependent bound that is completely independent of the class size (even overcoming the mild logarithmic dependence here). To this end, we will study more powerful structural results than Lemma 4 for controlling complexities of function classes induced via the maximum operator. As a good starting point, we will consider ℓ_∞ -norm covering numbers.

Acknowledgments

We thank Mehryar Mohri for helpful discussions. This work was partly funded by the German Research Foundation (DFG) award KL 2698/2-1.

References

- [1] T. Zhang, “Class-size independent generalization analysis of some discriminative multi-category classification,” in *Advances in Neural Information Processing Systems*, pp. 1625–1632, 2004.
- [2] T. Hofmann, L. Cai, and M. Ciaramita, “Learning with taxonomies: Classifying documents and words,” in *NIPS workshop on syntax, semantics, and statistics*, 2003.

- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255, IEEE, 2009.
- [4] A. Beygelzimer, J. Langford, Y. Lifshits, G. Sorkin, and A. Strehl, "Conditional probability tree estimation analysis and algorithms," in *Proceedings of UAI*, pp. 51–58, AUAI Press, 2009.
- [5] S. Bengio, J. Weston, and D. Grangier, "Label embedding trees for large multi-class tasks," in *Advances in Neural Information Processing Systems*, pp. 163–171, 2010.
- [6] P. Jain and A. Kapoor, "Active learning for large multi-class problems," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 762–769, IEEE, 2009.
- [7] O. Dekel and O. Shamir, "Multiclass-multilabel classification with more classes than examples," in *International Conference on Artificial Intelligence and Statistics*, pp. 137–144, 2010.
- [8] M. R. Gupta, S. Bengio, and J. Weston, "Training highly multiclass classifiers," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1461–1492, 2014.
- [9] T. Zhang, "Statistical analysis of some multi-category large margin classification methods," *The Journal of Machine Learning Research*, vol. 5, pp. 1225–1251, 2004.
- [10] A. Tewari and P. L. Bartlett, "On the consistency of multiclass classification methods," *The Journal of Machine Learning Research*, vol. 8, pp. 1007–1025, 2007.
- [11] T. Glasmachers, "Universal consistency of multi-class support vector classification," in *Advances in Neural Information Processing Systems*, pp. 739–747, 2010.
- [12] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2012.
- [13] V. Kuznetsov, M. Mohri, and U. Syed, "Multi-class deep boosting," in *Advances in Neural Information Processing Systems*, pp. 2501–2509, 2014.
- [14] V. Koltchinskii and D. Panchenko, "Empirical margin distributions and bounding the generalization error of combined classifiers," *Annals of Statistics*, pp. 1–50, 2002.
- [15] Y. Guermeur, "Combining discriminant models with new multi-class svms," *Pattern Analysis & Applications*, vol. 5, no. 2, pp. 168–179, 2002.
- [16] L. Oneto, D. Anguita, A. Ghio, and S. Ridella, "The impact of unlabeled patterns in rademacher complexity theory for kernel classifiers," in *Advances in Neural Information Processing Systems*, pp. 585–593, 2011.
- [17] V. Koltchinskii and D. Panchenko, "Rademacher processes and bounding the risk of function learning," in *High Dimensional Probability II*, pp. 443–457, Springer, 2000.
- [18] C. Cortes, M. Mohri, and A. Rostamizadeh, "Multi-class classification with maximum margin multiple kernel," in *ICML-13*, pp. 46–54, 2013.
- [19] M. Kloft, U. Brefeld, S. Sonnenburg, and A. Zien, "Lp-norm multiple kernel learning," *The Journal of Machine Learning Research*, vol. 12, pp. 953–997, 2011.
- [20] K. Crammer and Y. Singer, "On the algorithmic implementation of multiclass kernel-based vector machines," *The Journal of Machine Learning Research*, vol. 2, pp. 265–292, 2002.
- [21] P. L. Bartlett and S. Mendelson, "Rademacher and gaussian complexities: Risk bounds and structural results," *J. Mach. Learn. Res.*, vol. 3, pp. 463–482, 2002.
- [22] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: isoperimetry and processes*, vol. 23. Berlin: Springer, 1991.
- [23] S. I. Hill and A. Doucet, "A framework for kernel-based multi-category classification," *J. Artif. Intell. Res.(JAIR)*, vol. 30, pp. 525–564, 2007.
- [24] C. A. Micchelli and M. Pontil, "Learning the kernel function via regularization," *Journal of Machine Learning Research*, pp. 1099–1125, 2005.
- [25] S. S. Keerthi, S. Sundararajan, K.-W. Chang, C.-J. Hsieh, and C.-J. Lin, "A sequential dual method for large scale multi-class linear svms," in *14th ACM SIGKDD*, pp. 408–416, ACM, 2008.
- [26] J. D. Rennie and R. Rifkin, "Improving multiclass text classification with the support vector machine," tech. rep., AIM-2001-026, MIT, 2001.
- [27] K. Lang, "Newsweeder: Learning to filter netnews," in *Proceedings of the 12th international conference on machine learning*, pp. 331–339, 1995.
- [28] D. D. Lewis, Y. Yang, T. G. Rose, and F. Li, "Rev1: A new benchmark collection for text categorization research," *The Journal of Machine Learning Research*, vol. 5, pp. 361–397, 2004.
- [29] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona, "Caltech-UCSD Birds 200," Tech. Rep. CNS-TR-2010-001, California Institute of Technology, 2010.
- [30] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.

Supplementary Material

A Proofs of Structural Results on Gaussian Complexities

Our discussion on complexity bound is based on the following comparison result of Gaussian processes due to [1].

Lemma A.1 (E.g., Theorem 1 in [2]). *Let $\{\mathfrak{X}_\theta : \theta \in \Theta\}$ and $\{\mathfrak{Y}_\theta : \theta \in \Theta\}$ be two mean-zero separable Gaussian processes indexed by the same set Θ and suppose that*

$$\mathbb{E}[(\mathfrak{X}_\theta - \mathfrak{X}_{\bar{\theta}})^2] \leq \mathbb{E}[(\mathfrak{Y}_\theta - \mathfrak{Y}_{\bar{\theta}})^2], \quad \forall \theta, \bar{\theta} \in \Theta. \quad (\text{A.1})$$

Then,

$$\mathbb{E}[\sup_{\theta \in \Theta} \mathfrak{X}_\theta] \leq \mathbb{E}[\sup_{\theta \in \Theta} \mathfrak{Y}_\theta].$$

Proof of Lemma 4. Define two mean-zero separable Gaussian processes indexed by the finite dimensional Euclidean space $\{(h(x_1), \dots, h(x_n)) : h = (h_1, \dots, h_c) \in H\}$ (for simplicity, we use here the index h to denote $(h(x_1), \dots, h(x_n))$)

$$\begin{aligned} \mathfrak{X}_h &:= \sum_{i=1}^n g_i \max\{h_1(x_i), h_2(x_i), \dots, h_c(x_i)\}, \\ \mathfrak{Y}_h &:= \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i), \quad \forall h \in H. \end{aligned}$$

For any $h = (h_1, \dots, h_c), \bar{h} = (\bar{h}_1, \dots, \bar{h}_c) \in H$, the independence of the g_i and the equalities $\mathbb{E}g_i^2 = 1$ imply that

$$\begin{aligned} \mathbb{E}[(\mathfrak{X}_h - \mathfrak{X}_{\bar{h}})^2] &= \sum_{i=1}^n [\max\{h_1(x_i), \dots, h_c(x_i)\} - \max\{\bar{h}_1(x_i), \dots, \bar{h}_c(x_i)\}]^2 \\ \mathbb{E}[(\mathfrak{Y}_h - \mathfrak{Y}_{\bar{h}})^2] &= \sum_{i=1}^n \sum_{j=1}^c |h_j(x_i) - \bar{h}_j(x_i)|^2. \end{aligned} \quad (\text{A.2})$$

For any $\mathbf{a} = (a_1, \dots, a_c), \mathbf{b} = (b_1, \dots, b_c) \in \mathbb{R}^c$, it can be directly checked that

$$|\max\{a_1, \dots, a_c\} - \max\{b_1, \dots, b_c\}| \leq \max\{|a_1 - b_1|, \dots, |a_c - b_c|\} \leq \sum_{i=1}^c |a_i - b_i|. \quad (\text{A.3})$$

Applying the above inequality with $\mathbf{a} = (h_1(x_i), \dots, h_c(x_i)), \mathbf{b} = (\bar{h}_1(x_i), \dots, \bar{h}_c(x_i)), i = 1, \dots, n$, yields directly the following bounds relating the increments of the two Gaussian processes $\mathfrak{X}_h, \mathfrak{Y}_h$:

$$\begin{aligned} \mathbb{E}[(\mathfrak{X}_h - \mathfrak{X}_{\bar{h}})^2] &\stackrel{(\text{A.2})}{=} \sum_{i=1}^n [\max\{h_1(x_i), \dots, h_c(x_i)\} - \max\{\bar{h}_1(x_i), \dots, \bar{h}_c(x_i)\}]^2 \\ &\stackrel{(\text{A.3})}{\leq} \sum_{i=1}^n \max\{|h_1(x_i) - \bar{h}_1(x_i)|, \dots, |h_c(x_i) - \bar{h}_c(x_i)|\}^2 \\ &= \sum_{i=1}^n \max\{|h_1(x_i) - \bar{h}_1(x_i)|^2, \dots, |h_c(x_i) - \bar{h}_c(x_i)|^2\} \\ &\stackrel{(\text{A.3})}{\leq} \sum_{i=1}^n \sum_{j=1}^c |h_j(x_i) - \bar{h}_j(x_i)|^2 \stackrel{(\text{A.2})}{=} \mathbb{E}[(\mathfrak{Y}_h - \mathfrak{Y}_{\bar{h}})^2], \quad \forall h, \bar{h} \in H. \end{aligned}$$

That is, the condition (A.1) holds and therefore Lemma A.1 can be applied here to yield the stated result. \square

The following structural lemma regarding the Gaussian complexity of simplistic multi-class hypothesis spaces (not involving any argmax operator) will be used further below in the proof of Theorem 5.

Lemma A.2. *Let H be a class of functions defined on $\mathcal{X} \times \mathcal{Y}$ with $\mathcal{Y} = \{1, \dots, c\}$. Let $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ be a sequence of examples. Let g_1, \dots, g_{nc} be independent $N(0, 1)$ distributed random variables. Then the empirical Gaussian complexity of H can be controlled by:*

$$\mathfrak{G}_S(H) \leq \frac{1}{n} \mathbb{E}_{\mathbf{g}} \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i).$$

Proof. Define two Gaussian processes indexed by H :

$$\mathfrak{X}_h := \sum_{i=1}^n g_i h_{y_i}(x_i), \quad \mathfrak{Y}_h := \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i), \quad \forall h \in H.$$

For any $h, \bar{h} \in H$, it is obvious that

$$\begin{aligned} \mathbb{E}[(\mathfrak{X}_h - \mathfrak{X}_{\bar{h}})^2] &= \sum_{i=1}^n [h_{y_i}(x_i) - \bar{h}_{y_i}(x_i)]^2 \\ &\leq \sum_{i=1}^n [(h_1(x_i) - \bar{h}_1(x_i))^2 + \dots + (h_c(x_i) - \bar{h}_c(x_i))^2] \\ &= \mathbb{E}[(\mathfrak{Y}_h - \mathfrak{Y}_{\bar{h}})^2]. \end{aligned}$$

Now the stated inequality follows directly from Lemma A.1. \square

B Proof of Generalization Bounds for Multi-class Classification

B.1 Proof of Generalization Bound for General Multi-Class Classification (Theorem 5)

One of the main results of this paper is proved in this section. We first give a concentration inequality attributed to [3].

Lemma B.1 (McDiarmid inequality [3]). *Let Z_1, \dots, Z_n be independent random variables taking values in a set \mathcal{Z} , and assume that $f : \mathcal{Z}^n \rightarrow \mathbb{R}$ satisfies*

$$\sup_{\substack{z_1, \dots, z_n \\ \bar{z}_k \in \mathcal{Z}}} |f(z_1, \dots, z_n) - f(z_1, \dots, z_{k-1}, \bar{z}_k, z_{k+1}, \dots, z_n)| \leq c_i \quad (\text{B.1})$$

for $1 \leq k \leq n$. Then, for any $0 < \delta < 1$, with probability at least $1 - \delta$ we have

$$f(Z_1, \dots, Z_n) \leq \mathbb{E}f(Z_1, \dots, Z_n) + \sqrt{\frac{\sum_{k=1}^n c_k^2 \log(1/\delta)}{2}}.$$

Proof of Theorem 5. For any $\theta > 0$, introduce the following function bounding $\rho_h(x, y)$ from below:

$$\rho_{\theta, h}(x, y) = h(x, y) - \max_{y' \in \mathcal{Y}} [h(x, y') - \theta 1_{y'=y}] = \min_{y' \in \mathcal{Y}} [h(x, y) - h(x, y') + \theta 1_{y'=y}].$$

It can be checked that $\rho_{\theta, h}(x, y) = \min(\rho_h(x, y), \theta)$. Introduce two function classes derived from $\rho_{\theta, h}$:

$$\widetilde{H}_{\theta} = \{\rho_{\theta, h}(x, y) : h \in H\}, \quad \ell \circ \widetilde{H}_{\theta} = \{\ell(\rho_{\theta, h}(x, y)) : h \in H\}.$$

According to the definition of L -regular loss function and the relationship $\rho_{\theta, h} \leq \rho_h$, we have

$$R(h) = \mathbb{E}[1_{\rho_h(X, Y) \leq 0}] \leq \mathbb{E}[1_{\rho_{\theta, h}(X, Y) \leq 0}] \leq \mathbb{E}[\ell(\rho_{\theta, h}(X, Y))],$$

which, together with McDiarmid inequality [3] and the symmetrization technique (e.g., Theorem 4.4 in [4]), yields the following inequality

$$R(h) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_{\theta, h}(x_i, y_i)) + 2\mathfrak{R}_S(\ell \circ \widetilde{H}_{\theta}) + 3B_{\ell} \sqrt{\frac{\log \frac{2}{\delta}}{2n}}, \quad \forall h \in H \quad (\text{B.2})$$

with probability at least $1 - \delta$.

For the fixed parameter $\theta = c_\ell$, we observe that $\rho_{\theta,h}(x, y) = \min(\rho_h(x, y), c_\ell)$. If $\rho_h(x, y) > c_\ell$, the definition of L -regular loss implies that

$$\ell(\rho_{\theta,h}(x, y)) = \ell(c_\ell) = 0 = \ell(\rho_h(x, y)).$$

Otherwise, we have $\rho_{\theta,h}(x, y) = \rho_h(x, y)$. Therefore, for any (x, y) we have $\ell(\rho_{\theta,h}(x, y)) = \ell(\rho_h(x, y))$, which, coupled with the Lipschitz property of ℓ and Eq. (B.2), yields the following inequality with probability at least $1 - \delta$:

$$R(h) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_h(x_i, y_i)) + 2L\mathfrak{R}_S(\widetilde{H}_\theta) + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}}, \quad \forall h \in H. \quad (\text{B.3})$$

The Rademacher complexity of \widetilde{H}_θ satisfies the following inequality:

$$\begin{aligned} \mathfrak{R}_S(\widetilde{H}_\theta) &= \frac{1}{n} \mathbb{E}_\sigma \left[\sup_{h \in H} \sum_{i=1}^n \sigma_i (h(x_i, y_i) - \max_{y \in \mathcal{Y}} (h(x_i, y) - \theta 1_{y=y_i})) \right] \\ &\leq \frac{1}{n} \mathbb{E}_\sigma \left[\sup_{h \in H} \sum_{i=1}^n \sigma_i h(x_i, y_i) \right] + \frac{1}{n} \mathbb{E}_\sigma \left[\sup_{h \in H} \sum_{i=1}^n \sigma_i \max_{y \in \mathcal{Y}} (h(x_i, y) - \theta 1_{y=y_i}) \right] \\ &\leq \sqrt{\frac{\pi}{2}} \mathfrak{G}_S(H) + \frac{1}{n} \sqrt{\frac{\pi}{2}} \mathbb{E}_g \left[\sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n g_i \max(h_1(x_i) - \theta 1_{y_i=1}, \dots, h_c(x_i) - \theta 1_{y_i=c}) \right], \end{aligned} \quad (\text{B.4})$$

where the last step follows from the relationship between Gaussian and Rademacher processes expressed in Eq. (2). Furthermore, according to Lemma 4, the last term of the above inequality can be addressed by

$$\begin{aligned} &\mathbb{E}_g \left[\sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n g_i \max\{h_1(x_i) - \theta 1_{y_i=1}, \dots, h_c(x_i) - \theta 1_{y_i=c}\} \right] \\ &\stackrel{\text{Lemma 4}}{\leq} \mathbb{E}_g \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} (h_j(x_i) - \theta 1_{y_i=j}) \\ &= \mathbb{E}_g \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i) - \underbrace{\mathbb{E}_g \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} \theta 1_{y_i=j}}_{=0} \\ &= \mathbb{E}_g \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i). \end{aligned}$$

With this inequality and using Lemma A.2 to tackle $\mathfrak{G}_S(H)$, we immediately derive the following bound on $\mathfrak{R}_S(\widetilde{H}_\theta)$:

$$\mathfrak{R}_S(\widetilde{H}_\theta) \leq \frac{\sqrt{2\pi}}{n} \mathbb{E}_g \sup_{h=(h_1, \dots, h_c) \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i).$$

Plugging this Rademacher complexity bound back into Eq. (B.3), we obtain the stated result. \square

B.2 Proof of Generalization Bound for Kernel-Based Multi-Class Classification and MC-SVMs (Theorem 7)

To apply Theorem 5, we need to control the term $\sup_{h \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j(x_i)$, which we tackle by the following lemma due to [5].

Lemma B.2 (Corollary 4 in [5]). *If f is β -strongly convex w.r.t. $\|\cdot\|$ and $f^*(\mathbf{0}) = 0$, then, for any sequence v_1, \dots, v_n and for any μ we have*

$$\sum_{i=1}^n \langle v_i, \mu \rangle - f(\mu) \leq \sum_{i=1}^n \langle \nabla f^*(v_{1:i-1}), v_i \rangle + \frac{1}{2\beta} \sum_{i=1}^n \|v_i\|_*^2,$$

where $v_{1:i}$ denotes the sum $\sum_{j=1}^i v_j$.

Proof of Theorem 7. For the hypothesis space H and any $\lambda > 0$, applying Lemma B.2 with $\mu = (\mathbf{w}_1, \dots, \mathbf{w}_c)$ and $v_i = \lambda(g_i\phi(x_i), g_{n+i}\phi(x_i), \dots, g_{(c-1)n+i}\phi(x_i))$, we have

$$\begin{aligned} \lambda \sup_{h^{\mathbf{w}} \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j^{\mathbf{w}}(x_i) &= \sup_{h^{\mathbf{w}} \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} \langle \mathbf{w}_j, \lambda \phi(x_i) \rangle \\ &= \sup_{h^{\mathbf{w}} \in H} \sum_{i=1}^n \langle (\mathbf{w}_1, \dots, \mathbf{w}_c), (\lambda g_i \phi(x_i), \lambda g_{n+i} \phi(x_i), \dots, \lambda g_{(c-1)n+i} \phi(x_i)) \rangle \\ &\leq \sup_{h^{\mathbf{w}} \in H} f(\mathbf{w}_1, \dots, \mathbf{w}_c) + \sum_{i=1}^n \langle \nabla f^*(v_{1:i-1}), v_i \rangle + \frac{\lambda^2}{2\beta} \sum_{i=1}^n \|(g_i \phi(x_i), g_{n+i} \phi(x_i), \dots, g_{(c-1)n+i} \phi(x_i))\|_*^2. \end{aligned}$$

Taking expectation on both sides w.r.t. the Gaussian variables g_1, \dots, g_{nc} , the term $\sum_{i=1}^n \langle \nabla f^*(v_{1:i-1}), v_i \rangle$ vanishes, and therefore we obtain

$$\mathbb{E}_{\mathbf{g}} \sup_{h^{\mathbf{w}} \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j^{\mathbf{w}}(x_i) \leq \frac{\Lambda}{\lambda} + \frac{\lambda}{2\beta} \sum_{i=1}^n \mathbb{E}_{\mathbf{g}} \|(g_i \phi(x_i), g_{n+i} \phi(x_i), \dots, g_{(c-1)n+i} \phi(x_i))\|_*^2.$$

Choosing $\lambda = \sqrt{\frac{2\beta\Lambda}{\sum_{i=1}^n \mathbb{E}_{\mathbf{g}} \|(g_i \phi(x_i), g_{n+i} \phi(x_i), \dots, g_{(c-1)n+i} \phi(x_i))\|_*^2}}$, the above inequality translates to

$$\mathbb{E}_{\mathbf{g}} \sup_{h^{\mathbf{w}} \in H} \sum_{i=1}^n \sum_{j=1}^c g_{(j-1)n+i} h_j^{\mathbf{w}}(x_i) \leq \sqrt{\frac{2\Lambda}{\beta} \sum_{i=1}^n \mathbb{E}_{\mathbf{g}} \|(g_i \phi(x_i), g_{n+i} \phi(x_i), \dots, g_{(c-1)n+i} \phi(x_i))\|_*^2}.$$

Putting the above complexity bound into Theorem 5, we obtain the stated result. \square

B.3 Proof of Generalization Bound for ℓ_p -norm Multi-class SVMs (Corollary 8)

The following simple lemma controls the p -th moment of a $N(0, 1)$ distributed random variable. We give the proof here for completeness.

Lemma B.3. Let g be $N(0, 1)$ distributed. For any $p > 0$, the p -th moment of g can be bounded by

$$[\mathbb{E}|g|^p]^{\frac{1}{p}} \leq (2p)^{\frac{1}{2} + \frac{1}{p}}.$$

Proof. Let $\forall n \in \mathbb{N}_+ : \Gamma(n) = (n-1)!$ be the Gamma function. The p -th moment of a $N(0, 1)$ distributed random variable can be exactly expressed via Gamma function [6]:

$$\begin{aligned} \mathbb{E}|g|^p &= \frac{2^{\frac{p}{2}}}{\sqrt{\pi}} \Gamma\left(\frac{p+1}{2}\right) \leq \frac{2^{\frac{p}{2}}}{\sqrt{\pi}} \Gamma\left(\lceil \frac{p+1}{2} \rceil\right) \\ &= \frac{2^{\frac{p}{2}}}{\sqrt{\pi}} \lceil \frac{p-1}{2} \rceil! \leq \frac{2^{\frac{p}{2}}}{\sqrt{\pi}} \sqrt{2\pi} \lceil \frac{p-1}{2} \rceil^{\lceil \frac{p-1}{2} \rceil + \frac{1}{2}} \\ &\leq (2p)^{\frac{p}{2} + 1}, \end{aligned}$$

where in the above deduction we have used Stirling's approximation [7]:

$$n! \leq \sqrt{2\pi n} n^{n+\frac{1}{2}} e^{-n+1/(12n)}.$$

\square

Proof of Corollary 8. Let g_1, \dots, g_{nc} be independent $N(0, 1)$ distributed random variables. Denote by $\tau_s = [\mathbb{E}|g_1|^s]^{\frac{1}{s}}$ the s th moment of a $N(0, 1)$ distributed random variable. Let q be any number satisfying $p \leq q \leq 2$. Introduce the function $f_q(\mathbf{w}) := \frac{1}{2} \|\mathbf{w}\|_{2,q}^2$. Any $h^{\mathbf{w}} \in H_{q,\Lambda}$ satisfies the inequality

$$f_q(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|_{2,q}^2 \leq \frac{1}{2} \Lambda^2.$$

Since $f_q(\mathbf{w})$ is $1/q^*$ -strongly convex w.r.t. the norm $\|\cdot\|_{2,q}$, and the dual norm of $\|\cdot\|_{2,q}$ is $\|\cdot\|_{2,q^*}$ (Cf. section 4.2 in [8]), the summation of the squared dual norm in Theorem 7 can be rewritten as

follows:

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E}_{\mathbf{g}} \| (g_i \phi(x_i), \dots, g_{(c-1)n+i} \phi(x_i)) \|_{2,q^*}^2 &= \sum_{i=1}^n \mathbb{E}_{\mathbf{g}} \left[\sum_{j=1}^c \| g_{(j-1)n+i} \phi(x_i) \|_2^{q^*} \right]^{\frac{2}{q^*}} \\
&= \sum_{i=1}^n \mathbb{E}_{\mathbf{g}} \left[\sum_{j=1}^c |g_{(j-1)n+i}|^{q^*} \right]^{\frac{2}{q^*}} k(x_i, x_i) \\
&\stackrel{\text{symmetry}}{=} \mathbb{E}_{\mathbf{g}} \left[\sum_{j=1}^c |g_j|^{q^*} \right]^{\frac{2}{q^*}} \sum_{i=1}^n k(x_i, x_i) \\
&\stackrel{\text{Jensen}}{\leq} c^{\frac{2}{q^*}} \tau_{q^*}^2 \sum_{i=1}^n k(x_i, x_i).
\end{aligned}$$

From which Theorem 7 immediately implies the following bounds, with probability at least $1 - \delta$ and for any $h^{\mathbf{w}} \in H_{q,\Lambda}$:

$$R(h^{\mathbf{w}}) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_{h^{\mathbf{w}}}(x_i, y_i)) + \frac{4L\Lambda c^{1/q^*} \tau_{q^*}}{n} \sqrt{\frac{\pi q^*}{2} \sum_{i=1}^n k(x_i, x_i)} + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

From the trivial inequality $\|\mathbf{w}\|_{2,p} \geq \|\mathbf{w}\|_{2,q}$, we immediately conclude $H_{p,\Lambda} \subset H_{q,\Lambda}$. Therefore, for any $h^{\mathbf{w}} \in H_{p,\Lambda}$, we have

$$R(h^{\mathbf{w}}) \leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_{h^{\mathbf{w}}}(x_i, y_i)) + \inf_{p \leq q \leq 2} \frac{4L\Lambda c^{1/q^*} \tau_{q^*}}{n} \sqrt{\frac{\pi q^*}{2} \sum_{i=1}^n k(x_i, x_i)} + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

It can be directly checked that the function $t \rightarrow \sqrt{t} c^{1/t}$ is decreasing along the interval $(0, 2 \log c)$ and increasing along the interval $(2 \log c, \infty)$. Therefore, the above generalization bound satisfies the inequality

$$\begin{aligned}
R(h^{\mathbf{w}}) &\leq \frac{1}{n} \sum_{i=1}^n \ell(\rho_{h^{\mathbf{w}}}(x_i, y_i)) + 3B_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}} + \\
&\quad \frac{L\Lambda}{n} \sqrt{8 \sum_{i=1}^n k(x_i, x_i)} \times \begin{cases} \sqrt{2e \log c} \tau_{2 \log c}, & \text{if } p^* \geq 2 \log c, \\ c^{\frac{1}{p^*}} \tau_{p^*} \sqrt{p^*}, & \text{otherwise.} \end{cases}
\end{aligned}$$

Applying Lemma B.3 to bound the moments of Gaussian variables, the stated result follows immediately. \square

C Proofs on the Dual Problems

C.1 Equivalent Representation of ℓ_p -norm Multi-class Classification

The equivalence between Problem (P) and Eq. (8) follows directly from the following lemma due to [9].

Lemma C.1 ([9]). *Let $a_i \geq 0, i \in \mathbb{N}_d$ and $1 \leq r < \infty$. Then*

$$\min_{\eta: \eta_i \geq 0, \sum_{i \in \mathbb{N}_d} \eta_i^r \leq 1} \sum_{i \in \mathbb{N}_d} \frac{a_i}{\eta_i} = \left(\sum_{i \in \mathbb{N}_d} a_i^{\frac{r}{r+1}} \right)^{1+\frac{1}{r}}$$

and the minimum is attained at

$$\eta_i = \frac{a_i^{\frac{1}{r+1}}}{\left(\sum_{k \in \mathbb{N}_d} a_k^{\frac{r}{r+1}} \right)^{\frac{1}{r}}}.$$

Proof of Proposition 16. Fixing \mathbf{w} , the sub-optimization of Eq. (8) w.r.t. β is

$$\begin{aligned} \min_{\beta} \quad & \sum_{j=1}^c \frac{\|\mathbf{w}_j\|_2^2}{2\beta_j} \\ \text{s.t.} \quad & \|\beta\|_{\bar{p}} \leq 1, \bar{p} = p(2-p)^{-1}, \beta_j \geq 0. \end{aligned}$$

The stated result now follows directly by applying Lemma C.1 with $r = \bar{p}$ and $\alpha_j = \|\mathbf{w}_j\|_2^2$. \square

C.2 Derivation of the Completely Dualized Problem (Problem 11)

Derivation of Problem 11. Problem (P) translates to the following equivalent problem

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \left[\sum_{j=1}^c \|\mathbf{w}_j\|_2^p \right]^{\frac{2}{p}} + C \sum_{i=1}^n \ell(t_i) \\ \text{s.t.} \quad & t_i \leq \langle \mathbf{w}_{y_i}, \phi(x_i) \rangle - \langle \mathbf{w}_y, \phi(x_i) \rangle, \quad y \neq y_i, i = 1, \dots, n. \end{aligned} \quad (\text{C.1})$$

The Lagrangian of the above convex optimization problem is

$$\mathcal{L} = \frac{1}{2} \left[\sum_{j=1}^c \|\mathbf{w}_j\|_2^p \right]^{\frac{2}{p}} + C \sum_{i=1}^n \ell(t_i) + \sum_{i=1}^n \sum_{j \neq y_i} \tilde{\alpha}_{ij} (t_i + \langle \mathbf{w}_j, \phi(x_i) \rangle - \langle \mathbf{w}_{y_i}, \phi(x_i) \rangle),$$

with Lagrangian variables $0 \leq \tilde{\alpha} \in \mathbb{R}^{n \times (c-1)}$. For the last term of the Lagrangian, we have the following identity:

$$\begin{aligned} \sum_{i=1}^n \sum_{j \neq y_i} \tilde{\alpha}_{ij} \langle \mathbf{w}_j - \mathbf{w}_{y_i}, \phi(x_i) \rangle &= \sum_{i=1}^n \sum_{j \neq y_i} \tilde{\alpha}_{ij} \langle \mathbf{w}_j, \phi(x_i) \rangle - \sum_{i=1}^n \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} \langle \mathbf{w}_{y_i}, \phi(x_i) \rangle \\ &= \sum_{j=1}^c \langle \mathbf{w}_j, \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) \rangle - \sum_{j=1}^c \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \langle \mathbf{w}_j, \phi(x_i) \rangle \quad (\text{C.2}) \\ &= \sum_{j=1}^c \langle \mathbf{w}_j, \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \rangle. \end{aligned}$$

With this identity, the Lagrangian translates to

$$\begin{aligned} \mathcal{L} &= \frac{1}{2} \left[\sum_{j=1}^c \|\mathbf{w}_j\|_2^p \right]^{\frac{2}{p}} + \sum_{j=1}^c \langle \mathbf{w}_j, \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \rangle + \\ &\quad C \sum_{i=1}^n [\ell(t_i) + \frac{1}{C} \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} t_i]. \quad (\text{C.3}) \end{aligned}$$

According to the definition of Fenchel conjugate function, it holds that

$$\begin{aligned} \inf_{\mathbf{w}, \mathbf{t}} \mathcal{L} &= - \sup_{\mathbf{w}} \left[- \frac{1}{2} \left[\sum_{j=1}^c \|\mathbf{w}_j\|_2^p \right]^{\frac{2}{p}} - \sum_{j=1}^c \langle \mathbf{w}_j, \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \rangle \right] \\ &\quad - C \sum_{i=1}^n \sup_{t_i} [-\ell(t_i) - \sum_{j \neq y_i} \frac{1}{C} \tilde{\alpha}_{ij} t_i] \\ &= - \left[\frac{1}{2} \left\| \left(- \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) + \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right)_{j=1}^c \right\|_{2, \frac{p}{p-1}}^2 \right]^* \\ &\quad - C \sum_{i=1}^n \ell^* \left(- \frac{1}{C} \sum_{j \neq y_i} \tilde{\alpha}_{ij} \right) \\ &= - \frac{1}{2} \left\| \left(\sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right)_{j=1}^c \right\|_{2, \frac{p}{p-1}}^2 - C \sum_{i=1}^n \ell^* \left(- \frac{1}{C} \sum_{j \neq y_i} \tilde{\alpha}_{ij} \right), \quad (\text{C.4}) \end{aligned}$$

where in the last step of the above deduction we have used the identity: $(\frac{1}{2}\|\cdot\|^2)^* = \frac{1}{2}\|\cdot\|_*^2$ and the fact that the dual norm of $\|\cdot\|_{2,p}$ is $\|\cdot\|_{2,\frac{p}{p-1}}$. Consequently, the dual problem becomes

$$\begin{aligned} \sup_{\tilde{\alpha} \in \mathbb{R}^{n \times (c-1)}} & -\frac{1}{2} \left[\sum_{j=1}^c \left\| \sum_{i:y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i:y_i=j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right\|_2^{\frac{p}{p-1}} \right]^{\frac{2(p-1)}{p}} - C \sum_{i=1}^n \ell^* \left(-\frac{1}{C} \sum_{j \neq y_i} \tilde{\alpha}_{ij} \right), \\ \text{s.t. } & \tilde{\alpha} \geq 0. \end{aligned}$$

Introducing $\alpha \in \mathbb{R}^{n \times c}$ via the substitution:

$$\alpha_{ij} = \begin{cases} -\tilde{\alpha}_{ij} & \text{if } j \neq y_i \\ \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} & \text{if } j = y_i, \end{cases} \quad (\text{C.5})$$

we have

$$\sum_{i:y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i:y_i=j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) = - \sum_{i:y_i \neq j} \alpha_{ij} \phi(x_i) - \sum_{i:y_i=j} \alpha_{ij} \phi(x_i), \quad (\text{C.6})$$

from which the stated dual problem follows directly. \square

C.3 Proof of the Representer Theorem (Theorem 12)

Let H_1, \dots, H_c be c Hilbert spaces and $p \geq 1$. Define the function $g_p(v_1, \dots, v_c) : H_1 \times \dots \times H_c \rightarrow \mathbb{R}$ by

$$g_p(v_1, \dots, v_c) = \frac{1}{2} \|(v_1, \dots, v_c)\|_{2,p}^2, \quad p \geq 1.$$

Lemma C.2. *The gradient of g_p is*

$$\frac{\partial g_p(v_1, \dots, v_c)}{\partial v_j} = \left[\sum_{\tilde{j}=1}^c \|v_{\tilde{j}}\|_2^p \right]^{\frac{2}{p}-1} \|v_j\|_2^{p-2} v_j.$$

Proof. By the chain rule, we have

$$\begin{aligned} \frac{\partial g_p(v_1, \dots, v_c)}{\partial v_j} &= \frac{1}{p} \left[\sum_{\tilde{j}=1}^c \|v_{\tilde{j}}\|_2^p \right]^{\frac{2}{p}-1} \frac{\partial \langle v_j, v_j \rangle^{\frac{p}{2}}}{\partial v_j} \\ &= \frac{1}{2} \left[\sum_{\tilde{j}=1}^c \|v_{\tilde{j}}\|_2^p \right]^{\frac{2}{p}-1} \frac{\partial \langle v_j, v_j \rangle}{\partial v_j} \langle v_j, v_j \rangle^{\frac{p}{2}-1} \\ &= \left[\sum_{\tilde{j}=1}^c \|v_{\tilde{j}}\|_2^p \right]^{\frac{2}{p}-1} \|v_j\|_2^{p-2} v_j. \end{aligned}$$

\square

Proof of Representer Theorem (Theorem 12). In our derivation of the dual problem (see Eq. (C.4)), the variable \mathbf{w} should meet the optimality in the sense that

$$\mathbf{w} = \arg \max_{\mathbf{v}} -\frac{1}{2} \left[\sum_{j=1}^c \|v_j\|_2^p \right]^{\frac{2}{p}} + \sum_{j=1}^c \langle v_j, \sum_{i=1}^n \alpha_{ij} \phi(x_i) \rangle.$$

Since $(\nabla f)^{-1} = \nabla f^*$ for any convex function f , and the Fenchel-conjugate of g_p is g_{p^*} , we obtain the following representation of \mathbf{w} :

$$\begin{aligned} \mathbf{w} &= \nabla g_p^{-1} \left(\sum_{i=1}^n \alpha_{i1} \phi(x_i), \dots, \sum_{i=1}^n \alpha_{ic} \phi(x_i) \right) \\ &= \nabla g_{p^*} \left(\sum_{i=1}^n \alpha_{i1} \phi(x_i), \dots, \sum_{i=1}^n \alpha_{ic} \phi(x_i) \right) \\ &= \left[\sum_{j=1}^c \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^{p^*} \right]^{\frac{2}{p^*}-1} \left(\left\| \sum_{i=1}^n \alpha_{i1} \phi(x_i) \right\|_2^{p^*-2} \left[\sum_{i=1}^n \alpha_{i1} \phi(x_i) \right], \dots, \left\| \sum_{i=1}^n \alpha_{ic} \phi(x_i) \right\|_2^{p^*-2} \left[\sum_{i=1}^n \alpha_{ic} \phi(x_i) \right] \right). \end{aligned}$$

That is,

$$\mathbf{w}_j = \left[\sum_{\tilde{j}=1}^c \left\| \sum_{i=1}^n \alpha_{i\tilde{j}} \phi(x_i) \right\|_2^{p^*} \right]^{\frac{2}{p^*}-1} \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^{p^*-2} \left[\sum_{i=1}^n \alpha_{ij} \phi(x_i) \right].$$

□

C.4 Derivation of Partially Dualized Problem (Problem 14)

Derivation of Problem 14. The Lagrangian of the problem (8) w.r.t. \mathbf{w} is

$$\mathcal{L} = \sum_{j=1}^c \frac{\|\mathbf{w}_j\|_2^2}{2\beta_j} + C \sum_{i=1}^n \ell(t_i) + \sum_{i=1}^n \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{ij} (t_i + \langle \mathbf{w}_j, \phi(x_i) \rangle - \langle \mathbf{w}_{y_i}, \phi(x_i) \rangle),$$

with Lagrangian variables $0 \leq \tilde{\alpha} \in \mathbb{R}^{n \times (c-1)}$.

According to the identity (C.2), the Lagrangian translates to

$$\begin{aligned} \mathcal{L} = & \sum_{j=1}^c \frac{\|\mathbf{w}_j\|_2^2}{2\beta_j} + \sum_{j=1}^c \langle \mathbf{w}_j, \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \rangle + \\ & C \sum_{i=1}^n [\ell(t_i) + \frac{1}{C} \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} t_i]. \quad (\text{C.7}) \end{aligned}$$

According to the definition of Fenchel conjugate function, it holds that

$$\begin{aligned} \inf_{\mathbf{w}, \mathbf{t}} \mathcal{L} = & - \sum_{j=1}^c \left[\frac{1}{\beta_j} \sup_{\mathbf{w}_j} \left[-\frac{1}{2} \|\mathbf{w}_j\|_2^2 - \langle \mathbf{w}_j, \beta_j \left(\sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right) \rangle \right] \right] \\ & - C \sum_{i=1}^n \sup_{t_i} [-\ell(t_i) - \sum_{\tilde{j} \neq y_i} \frac{1}{C} \tilde{\alpha}_{i\tilde{j}} t_i] \\ = & - \sum_{j=1}^c \left[\frac{1}{\beta_j} \left[\frac{1}{2} \left\| \beta_j \left(\sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right) \right\|_2^2 \right]^* \right] - C \sum_{i=1}^n \ell^* \left(-\frac{1}{C} \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} \right) \\ = & - \frac{1}{2} \sum_{j=1}^c \beta_j \left\| \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right\|_2^2 - C \sum_{i=1}^n \ell^* \left(-\frac{1}{C} \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} \right), \end{aligned}$$

where in the last step of the above deduction we have used the identity: $(\frac{1}{2} \|\cdot\|^2)^* = \frac{1}{2} \|\cdot\|_*^2$ and the fact that the dual norm of $\|\cdot\|_{2,2}$ is itself. Consequently, the dual problem becomes

$$\begin{aligned} \sup_{\tilde{\alpha} \in \mathbb{R}^{n \times (c-1)}} & - \frac{1}{2} \sum_{j=1}^c \beta_j \left\| \sum_{i: y_i \neq j} \tilde{\alpha}_{ij} \phi(x_i) - \sum_{i: y_i = j} \sum_{\tilde{j} \neq j} \tilde{\alpha}_{i\tilde{j}} \phi(x_i) \right\|_2^2 - C \sum_{i=1}^n \ell^* \left(-\frac{1}{C} \sum_{\tilde{j} \neq y_i} \tilde{\alpha}_{i\tilde{j}} \right), \\ \text{s.t. } & \tilde{\alpha} \geq 0. \end{aligned}$$

Introducing $\alpha \in \mathbb{R}^{n \times c}$ as in Eq. (C.5) and noticing the identity (C.6), the above *dual problem* becomes

$$\begin{aligned} \sup_{\alpha \in \mathbb{R}^{n \times c}} & - \frac{1}{2} \sum_{j=1}^c \beta_j \left\| \sum_{i=1}^n \alpha_{ij} \phi(x_i) \right\|_2^2 - C \sum_{i=1}^n \ell^* \left(-\frac{\alpha_{iy_i}}{C} \right) \\ \text{s.t. } & \sum_{j=1}^c \alpha_{ij} = 0, \quad \forall i = 1, 2, \dots, n, \\ & \alpha_{ij} \leq 0, \quad j \neq y_i, \forall i = 1, \dots, n. \end{aligned} \quad (\text{C.8})$$

Note that in the above derivation of the dual problem, the variable \mathbf{w} should meet the optimality in the sense that

$$\mathbf{w} = \arg \max_{\mathbf{v}} - \frac{1}{2} \sum_{j=1}^c \|\mathbf{v}_j\|_2^2 + \sum_{j=1}^c \beta_j \langle \mathbf{v}_j, \sum_{i=1}^n \alpha_{ij} \phi(x_i) \rangle.$$

The representer theorem stated in Problem 14 follows directly from this optimization condition. □

D Note on Used Features for Caltech256 and UCSD birds

For the purpose of having features, we took the features from a fc6 layer of the BVLC reference caffe-net [10] computed for all images from the UCSD birds dataset [11] and Caltech256 [12]. Note that we neither used fc7 or fc8 layers, nor did we perform finetuning. Images were warped [13] so that they fitted into the quadratic reception field. As the goal was not on maximizing performance but comparing learning machines we resorted to computing one feature per image at training and test time without using the large number of region proposals which yield state of the art in fine-grained classification tasks [14], or mirroring and detection-like approaches like the 500 windows per image as in [15].

E Execution Time Experiments

This section report the training time of the classical CS [16] and the proposed ℓ_p -norm MC-SVM on the benchmark datasets. We repeat the experiments 10 times and report the average as well as standard deviation (in seconds) in Table E.1. For our method, the result is for the single p selected via cross-validation.

Method / Dataset	Sector	News 20	Birds 15	Birds 50
ℓ_p -norm MC-SVM	4914 \pm 64.7	3894 \pm 71.1	912.6 \pm 22.1	518.4 \pm 34.8
Crammer & Singer	3442 \pm 91.8	2227 \pm 43.7	701.7 \pm 50.6	314.1 \pm 17.1

Table E.1: Training time for the classical CS and the proposed ℓ_p -norm MC-SVM on the benchmark datasets.

From Table E.1, we see that our method needs longer training time than CS, but the increase is not that large and is well compensated by its improvement on accuracies.

References

- [1] D. Slepian, "The one-sided barrier problem for gaussian noise," *Bell System Technical Journal*, vol. 41, no. 2, pp. 463–501, 1962.
- [2] R. A. Vitale, "Some comparisons for gaussian processes," *Proceedings of the American Mathematical Society*, pp. 3043–3046, 2000.
- [3] C. McDiarmid, "On the method of bounded differences," in *Surveys in combinatorics* (J. Siemous, ed.), pp. 148–188, Cambridge: Cambridge Univ. Press, 1989.
- [4] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2012.
- [5] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari, "Regularization techniques for learning with matrices," *J. Mach. Learn. Res.*, vol. 13, pp. 1865–1890, 2012.
- [6] A. Winkelbauer, "Moments and absolute moments of the normal distribution," *arXiv preprint arXiv:1209.4340*, 2012.
- [7] H. Robbins, "A remark on stirling's formula," *American Mathematical Monthly*, pp. 26–29, 1955.
- [8] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari, "Applications of strong convexity–strong smoothness duality to learning with matrices," *CoRR*, vol. abs/0910.0610, 2009.
- [9] C. A. Micchelli and M. Pontil, "Learning the kernel function via regularization," *Journal of Machine Learning Research*, pp. 1099–1125, 2005.
- [10] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.
- [11] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona, "Caltech-UCSD Birds 200," Tech. Rep. CNS-TR-2010-001, California Institute of Technology, 2010.
- [12] G. Griffin, A. Holub, and P. Perona, "Caltech-256 object category dataset," 2007.
- [13] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pp. 580–587, 2014.
- [14] N. Zhang, J. Donahue, R. B. Girshick, and T. Darrell, "Part-based r-cnns for fine-grained category detection," in *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I*, pp. 834–849, 2014.
- [15] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pp. 1717–1724, 2014.

486 [16] K. Crammer and Y. Singer, “On the algorithmic implementation of multiclass kernel-based vector ma-
487 chines,” *The Journal of Machine Learning Research*, vol. 2, pp. 265–292, 2002.
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539