
Primal-Dual Rates and Certificates

Celestine Düner

IBM Research, Zürich, Switzerland

CDU@ZURICH.IBM.COM

Simone Forte

ETH Zürich, Switzerland

FORTESIMONE90@GMAIL.COM

Martin Takáč

Lehigh University, USA

TAKAC.MT@GMAIL.COM

Martin Jaggi

ETH Zürich, Switzerland

JAGGIM@INF.ETHZ.CH

Abstract

We propose an algorithm-independent framework to equip existing optimization methods with primal-dual certificates. Such certificates and corresponding rate of convergence guarantees are important for practitioners to diagnose progress, in particular in machine learning applications.

We obtain new primal-dual convergence rates, e.g., for the Lasso as well as many L_1 , Elastic Net, group Lasso and TV-regularized problems. The theory applies to any norm-regularized generalized linear model. Our approach provides efficiently computable duality gaps which are globally defined, without modifying the original problems in the region of interest.

1. Introduction

The massive growth of available data has moved data analysis and machine learning to center stage in many industrial as well as scientific fields, ranging from web and sensor data to astronomy, health science, and countless other applications. With the increasing size of datasets, machine learning methods are limited by the scalability of the underlying optimization algorithms to train these models, which has spurred significant research interest in recent years.

However, practitioners face a significant problem arising with the larger model complexity in large-scale machine learning and in particular deep-learning methods - it is increasingly hard to diagnose if the optimization algorithm

used for training works well or not. With the optimization algorithms also becoming more complex (e.g., in a distributed setting), it can often be very hard to pin down if *bad performance of a predictive model either comes from slow optimization, or from poor modeling choices*. In this light, easily verifiable guarantees for the quality of an optimization algorithm are very useful — note that the optimum solution of the problem is unknown in most cases. For convex optimization problems, a primal-dual gap can serve as such a certificate. If available, the gap also serves as a useful stopping criterion for the optimizer.

So far, the majority of popular optimization algorithms for learning applications comes without a notion of primal-dual gap. In this paper, we aim to change this for a relevant class of machine learning problems. We propose a primal-dual framework which is algorithm-independent, and allows to equip existing algorithms with additional primal-dual certificates as an add-on.

Our approach is motivated by the recent analysis of SDCA (Shalev-Shwartz & Zhang, 2013). We extend their setting to the significantly larger class of convex optimization problems of the form

$$\min_{\alpha \in \mathbb{R}^n} f(A\alpha) + g(\alpha)$$

for a given matrix $A \in \mathbb{R}^{d \times n}$, f being smooth, and g being a general convex function. This problem class includes the most prominent regression and classification methods as well as generalized linear models. We will formalize the setting in more details in Section 3, and highlight the associated dual problem, which has the same structure. An overview over some popular examples that can be formulated in this setting either as primal or dual problem is given in Table 1.

Contributions. The main contributions in this work can be summarized as follows:

- Our new primal-dual framework is algorithm-independent, that is it allows users to equip existing algorithms with primal-dual certificates and convergence rates.
- We introduce a new Lipschitzing trick allowing duality gaps (and thus accuracy certificates) which are globally defined, for a large class of convex problems which did not have this property so far. Our approach does not modify the original problems in the region of interest. In contrast to existing methods adding a small strongly-convex (L_2) term as, e.g., in (Shalev-Shwartz & Zhang, 2014; Zhang & Lin, 2015), our approach leaves both the algorithms and the optima unaffected.
- Compared with the well-known duality setting of SDCA (Shalev-Shwartz & Zhang, 2013; 2014) which is restricted to strongly convex regularizers and finite sum optimization problems, our framework encompasses a significantly larger class of problems. We obtain new primal-dual convergence rates, e.g., for the Lasso as well as many L_1 , Elastic Net, group Lasso and TV-regularized problems. The theory applies to any norm-regularized generalized linear model.
- Existing primal-dual guarantees for the class of ERM problems with Lipschitz loss from (Shalev-Shwartz & Zhang, 2013) (e.g., SVM) are valid only for an “average” iteration. We show that the same rate of convergence can be achieved, e.g., for accelerated SDCA, but without the necessity of computing an average iterate.
- Our primal-dual theory captures a more precise notion of data-dependency compared with existing results (which relied on per-coordinate information only). To be precise, our shown convergence rate for the general algorithms is dependent on the *spectral norm of the data*, see also (Takáč et al., 2013; 2015).

2. Related Work

Linearized ADMM solvers. For the problem structure of our interest here, one of the most natural algorithms is the splitting method known as the Chambolle-Pock algorithm (also known as Primal-Dual Hybrid Gradient, Arrow-Hurwicz method, or linearized ADMM) (Chambolle & Pock, 2010). While this algorithm can give rise to a duality gap, it is significantly less general compared to our framework. In each iteration, it requires a complete solution of the proximal operators of both f and g , which can be computationally expensive. Its convergence rate is sensitive to the used step-size (Goldstein et al., 2015). Our framework is not algorithm-specific, and holds for arbitrary iterate sequences. More recently, the SPDC method (Zhang & Lin, 2015) was proposed as a coordinate-wise variant of (Chambolle & Pock, 2010), but only for strongly convex g .

Stochastic coordinate solvers. Coordinate descent/ascent methods have become state-of-the-art for many machine learning problems (Hsieh et al., 2008; Friedman et al., 2010). In recent years, theoretical convergence rate guarantees have been developed for the primal-only setting, e.g., by (Nesterov, 2012; Richtárik & Takáč, 2014), as well as more recently also for primal-dual guarantees, see, e.g., (Lacoste-Julien et al., 2013; Shalev-Shwartz & Zhang, 2013; 2014). The influential Stochastic¹ Dual Coordinate Ascent (SDCA) framework (Shalev-Shwartz & Zhang, 2013) was motivated by the L_2 -regularized SVM, where coordinate descent is very efficient on the dual SVM formulation, with every iteration only requiring access to a single datapoint (i.e. a column of the matrix A in our setup). In contrast to primal stochastic gradient descent (SGD) methods, the SDCA algorithm family is often preferred as it is free of learning-rate parameters, and has a fast linear (geometric) convergence rate. SDCA and recent extensions (Takáč et al., 2013; Shalev-Shwartz & Zhang, 2014; Qu et al., 2014; Shalev-Shwartz, 2015; Zhang & Lin, 2015) require g to be strongly convex.

Under the weaker assumption of weak strong convexity (Necoara, 2015), a linear rate for the primal-dual convergence of SDCA was recently shown by (Ma et al., 2015b).

In the Lasso literature, a similar trend in terms of solvers has been observed recently, but with the roles of primal and dual problems reversed. For those problems, coordinate descent algorithms on the primal formulation have become the state-of-the-art, as in GLMNET (Friedman et al., 2010) and extensions (Shalev-Shwartz & Tewari, 2011; Yuan et al., 2012; 2010). Despite the prototypes of both problem types—SVM for the L_2 -regularized case and Lasso for the L_1 -regularized case—being closely related (Jaggi, 2014), we are not aware of existing primal-dual guarantees for coordinate descent on unmodified L_1 -regularized problems.

Comparison to smoothing techniques. Existing techniques for bringing L_1 -regularized and related problems into a primal-dual setting compatible with SDCA do rely on the classical Nesterov-smoothing approach - By adding a small amount of L_2 to the part g , the objective becomes strongly convex; see, e.g., (Nesterov, 2005; 2012; Tran-Dinh & Cevher, 2014; Richtárik & Takáč, 2014; Shalev-Shwartz & Zhang, 2014). However, the appropriate strength of smoothing is difficult to tune. It will depend on the accuracy level, and will influence the chosen algorithm, as it will change the iterates, the resulting convergence rate as well as the tightness of the resulting duality gaps. The line of work of Tran-Dinh & Cevher (2014; 2015) provides duality gaps for smoothed problems and rates for proximally tractable objectives (Tran-Dinh & Cevher, 2015) and

¹Here ‘stochastic’ refers to randomized selection of the active coordinate.

Table 1. Primal-dual convergence rates of general algorithms applied to (A), for some machine learning and signal processing problems which are examples of our optimization problem formulations (A) and (B). Note that several can be mapped to either (A) or (B). $\lambda > 0$ is a regularization parameter specified by the user. We will discuss the most prominent examples in more detail in Sections 5.2 and 6.

Problem		f/f^*	g/g^*	Convergence
Regularizer	L_1	(A) $f = \ell(A\alpha)$	$g = \lambda \ \alpha\ _1$	Thm 6, Cor 10
	Elastic Net	(A) $f = \ell(A\alpha)$	$g = \lambda \left(\frac{\eta}{2}\ \alpha\ _2^2 + (1-\eta)\ \alpha\ _1\right)$	Thm 2,8, Cor 11
		(B) $f^* = \lambda \left(\frac{\eta}{2}\ \mathbf{w}\ _2^2 + (1-\eta)\ \mathbf{w}\ _1\right)$	$g^* = \ell(-A^T \mathbf{w})$	Thm 2,8, Cor 11
	L_2	(A) $f = \ell(A\alpha)$	$g = \frac{\lambda}{2}\ \alpha\ _2^2$	Thm 2,8
		(B) $f^* = \frac{\lambda}{2}\ \mathbf{w}\ _2^2$	$g^* = \ell(-A^T \mathbf{w})$	Thm 2,8
	Fused Lasso	(A) $f = \ell(A\alpha)$	$g = \lambda \ M\alpha\ _1$	Thm 6
Group Lasso	(A) $f = \ell(A\alpha)$	$g = \lambda \sum_k \ \alpha_{\mathcal{G}_k}\ _2, \{\mathcal{G}_k\}$ part. of $\{1..n\}$	Thm 6	
SVM (hinge loss)	(A) $f = \frac{\lambda}{2}\ A\alpha\ _2^2$	$g = \sum_i -y_i \alpha_i$ s.t. $y_i \alpha_i \in [0, 1]$	Thm 4,9	

Loss ℓ	Logistic Regression	$\ell_{\text{LR}}(\mathbf{z}) := \frac{1}{m} \sum_{j=1}^m \log(1 + \exp(-y_j \mathbf{z}_j))$	for $\mathbf{z} = A\alpha$ or $\mathbf{z} = -A^T \mathbf{w}$
	Least Squares	$\ell_{\text{LS}}(\mathbf{z}) := \frac{1}{2}\ \mathbf{z} - \mathbf{b}\ _2^2$	for $\mathbf{z} = A\alpha$ or $\mathbf{z} = -A^T \mathbf{w}$

also objectives with efficient Fenchel-type operator (Yurtsever et al., 2015). In contrast, our approach preserves all solutions of the original L_1 -optimization — it leaves the iterate sequences of existing algorithms unchanged, which is desirable in practice, and allows the reusability of existing solvers. We do not assume proximal or Fenchel tractability.

Distributed algorithms. For L_1 -problems exceeding the memory capacity of a single computer, a communication-efficient distributed scheme leveraging this Lipschitzing trick is presented in (Smith et al., 2015; Forte, 2015).

3. Setup and Primal-Dual Structure

In this paper, we consider optimization problems of the following primal-dual structure. As we will see, the relationship between primal and dual objectives has many benefits, including computation of the duality gap, which allows us to have certificates for the approximation quality.

We consider the following pair of optimization problems, which are dual² to each other:

$$\min_{\alpha \in \mathbb{R}^n} \left[\mathcal{D}(\alpha) := f(A\alpha) + g(\alpha) \right], \quad (\text{A})$$

$$\min_{\mathbf{w} \in \mathbb{R}^d} \left[\mathcal{P}(\mathbf{w}) := f^*(\mathbf{w}) + g^*(-A^T \mathbf{w}) \right]. \quad (\text{B})$$

The two problems are associated to a given data matrix $A \in \mathbb{R}^{d \times n}$, and the functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}$ are allowed to be arbitrary closed convex functions. Here $\alpha \in \mathbb{R}^n$ and $\mathbf{w} \in \mathbb{R}^d$ are the respective variable vectors. The relation of (A) and (B) is called *Fenchel-Rockafellar Duality* where the functions f^*, g^* in formulation (B) are defined as the *convex conjugates*³ of their corresponding counterparts f, g in (A). The two main powerful features of this general duality structure are first that it includes many more machine learning methods than more traditional du-

ality notions, and secondly that the two problems are fully symmetric, when changing respective roles of f and g . In typical machine learning problems, the two parts typically play the roles of a data-fit (or loss) term as well as a regularization term. As we will see later, those two roles can be swapped, depending on the application.

Optimality Conditions. The first-order optimality conditions for our pair of vectors $\mathbf{w} \in \mathbb{R}^d, \alpha \in \mathbb{R}^n$ in problems (A) and (B) are given as

$$\mathbf{w} \in \partial f(A\alpha), \quad (1a) \quad -A^T \mathbf{w} \in \partial g(\alpha), \quad (2a)$$

$$A\alpha \in \partial f^*(\mathbf{w}), \quad (1b) \quad \alpha \in \partial g^*(-A^T \mathbf{w}) \quad (2b)$$

see, e.g. (Bauschke & Combettes, 2011, Proposition 19.18). The stated optimality conditions are equivalent to α, \mathbf{w} being a saddle-point of the Lagrangian, which is given as $\mathcal{L}(\alpha, \mathbf{w}) = f^*(\mathbf{w}) - \langle A\alpha, \mathbf{w} \rangle - g(\alpha)$ if $\alpha \in \text{dom}(g)$ and $\mathbf{w} \in \text{dom}(f^*)$.

Duality Gap. From the definition of the dual problems in terms of the convex conjugates, we always have $\mathcal{P}(\mathbf{w}) \geq \mathcal{P}(\mathbf{w}^*) \geq -\mathcal{D}(\alpha^*) \geq -\mathcal{D}(\alpha)$, giving rise to the definition of the general duality gap $G(\mathbf{w}, \alpha) := \mathcal{P}(\mathbf{w}) - (-\mathcal{D}(\alpha))$.

For differentiable f , the duality gap can be used more conveniently: Given $\alpha \in \mathbb{R}^n$ s.t. $A\alpha \in \text{dom}(f)$ in the context of (A), a corresponding variable vector $\mathbf{w} \in \mathbb{R}^d$ for problem (B) is given by the first-order optimality condition (1a) as

$$\mathbf{w} = \mathbf{w}(\alpha) := \nabla f(A\alpha). \quad (3)$$

Under strong duality, we have $\mathcal{P}(\mathbf{w}^*) = -\mathcal{D}(\alpha^*)$ and $\mathbf{w}(\alpha^*) = \mathbf{w}^*$, where α^* is an optimal solution of (A). This implies that the suboptimality $\mathcal{P}(\mathbf{w}(\alpha)) - \mathcal{P}(\mathbf{w}^*)$ is always bounded above by the simpler duality gap function

$$G(\alpha) := \mathcal{P}(\mathbf{w}(\alpha)) - (-\mathcal{D}(\alpha)) \geq \mathcal{P}(\mathbf{w}(\alpha)) - \mathcal{P}(\mathbf{w}^*) \quad (4)$$

which hence acts as a certificate of the approximation quality of the variable vector α .

²For a self-contained derivation see Appendix C.

³The conjugate is defined as $h^*(\mathbf{v}) := \sup_{\mathbf{u} \in \mathbb{R}^d} \mathbf{v}^T \mathbf{u} - h(\mathbf{u})$.

4. Primal-Dual Guarantees for Any

Algorithm Solving (A)

In this section we state an important lemma, which will later allow us to transform a suboptimality guarantee of any algorithm into a duality gap guarantee, for optimization problems of the form specified in the previous section.

Lemma 1. *Consider an optimization problem of the form (A). Let f be $1/\beta$ -smooth w.r.t. a norm $\|\cdot\|_f$ and let g be μ -strongly convex with convexity parameter $\mu \geq 0$ w.r.t. a norm $\|\cdot\|_g$. The general convex case $\mu = 0$ is explicitly allowed, but only if g has bounded support.*

Then, for any $\alpha \in \text{dom}(\mathcal{D})$ and any $s \in [0, 1]$, it holds that

$$\begin{aligned} \mathcal{D}(\alpha) - \mathcal{D}(\alpha^*) &\geq sG(\alpha) \\ &+ \frac{s^2}{2} \left(\frac{\mu(1-s)}{s} \|\mathbf{u} - \alpha\|_g^2 - \frac{1}{\beta} \|A(\mathbf{u} - \alpha)\|_f^2 \right), \end{aligned} \quad (5)$$

where $G(\alpha)$ is the gap function defined in (4) and

$$\mathbf{u} \in \partial g^*(-A^\top \mathbf{w}(\alpha)). \quad (6)$$

We note that the improvement bound here bears similarity to the proof of (Bach, 2015, Prop 4.2) for the case of an extended Frank-Wolfe algorithm. In contrast, our result here is algorithm-independent, and leads to tighter results due to the more careful choice of \mathbf{u} , as we'll see in the following.

4.1. Linear Convergence Rates

In this section we assume that we are using an arbitrary optimization algorithm applied to problem (A). It is assumed that the algorithm produces a sequence of (possibly random) iterates $\{\alpha^{(t)}\}_{t=0}^\infty$ such that there exists $C \in (0, 1]$, $D \geq 0$ such that

$$\mathbb{E}[\mathcal{D}(\alpha^{(t)}) - \mathcal{D}(\alpha^*)] \leq (1 - C)^t D. \quad (7)$$

In the next two theorems, we define $\sigma := (\max_{\alpha \neq 0} \|A\alpha\|_f / \|\alpha\|_g)^2$, i.e., the squared spectral norm of the matrix A in the Euclidean norm case.

4.1.1. CASE I. STRONGLY CONVEX g

Let us assume g is μ -strongly convex ($\mu > 0$) (equivalently, its conjugate g^* has Lipschitz continuous gradient with a constant $1/\mu$). The following theorem provides a linear convergence guarantee for any algorithm with given linear convergence rate for the suboptimality $\mathcal{D}(\alpha) - \mathcal{D}(\alpha^*)$.

Theorem 2. *Assume the function f is $1/\beta$ -smooth w.r.t. a norm $\|\cdot\|_f$ and g is μ -strongly convex w.r.t. a norm $\|\cdot\|_g$. Suppose we are using a linearly convergent algorithm as specified in (7). Then, for any*

$$t \geq T := \frac{1}{C} \log \frac{D(\frac{\sigma}{\beta} + \mu)}{\mu\epsilon} \quad (8)$$

it holds that $\mathbb{E}[G(\alpha^{(t)})] \leq \epsilon$.

From (7) we can obtain that after $\frac{1}{C} \log \frac{D}{\epsilon}$ iterations, we

would have a point $\alpha^{(t)}$ such that $\mathbb{E}[D(\alpha^{(t)}) - \mathcal{D}(\alpha^*)] \leq \epsilon$. Hence, comparing with (8) only few more iterations are needed to get the guarantees for the duality gap. The rate (7) is achieved by most of the first order algorithms, including proximal gradient descent (Nesterov, 2013) or SDCA (Richtárik & Takáč, 2014) with $C \sim \mu$ or accelerated SDCA (Lin et al., 2014) with $C \sim \sqrt{\mu}$.

4.1.2. CASE II. GENERAL CONVEX g (OF BOUNDED SUPPORT)

In this section we will assume that g^* is Lipschitz (in contrast to smooth as in Theorem 2) and show that the linear convergence rate is preserved.

Theorem 3. *Assume that the function f is $1/\beta$ -smooth w.r.t. a norm $\|\cdot\|_f$, g^* is L -Lipschitz continuous w.r.t the dual norm $\|\cdot\|_*$, and we are using a linearly convergent algorithm (7). Then, for any*

$$t \geq T := \frac{1}{C} \log \frac{2D \max\{1, 2\sigma L^2/\epsilon\beta\}}{\epsilon} \quad (9)$$

it holds that $\mathbb{E}[G(\alpha^{(t)})] \leq \epsilon$.

In (Wang & Lin, 2014), it was proven that feasible descent methods when applied to the dual of an SVM do improve the objective geometrically as in (7). Later, (Ma et al., 2015b) extended this to stochastic coordinate feasible descent algorithms (including SDCA). Using our new Theorem 3, we can therefore extend their results to linear convergence for the duality gap for the SVM application.

4.2. Sub-Linear Convergence Rates

In this case we will focus only on general L -Lipschitz continuous functions g^* (if g is strongly convex, then many existing algorithms are available and converge with a linear rate).

We will assume that we are applying some (possibly randomized) algorithm on optimization problem (A) which produces a sequence (of possibly random) iterates $\{\alpha^{(t)}\}_{t=0}^\infty$ such that

$$\mathbb{E}[\mathcal{D}(\alpha^{(t)}) - \mathcal{D}(\alpha^*)] \leq \frac{C}{D(t)}, \quad (10)$$

where $D(t)$ is a function which has usually a linear or quadratic growth (i.e. $D(t) \sim \mathcal{O}(t)$ or $D(t) \sim \mathcal{O}(t^2)$).

The following theorem will allow to equip existing algorithms with sub-linear convergence in suboptimality, as specified in (10), with duality gap convergence guarantees.

Theorem 4. *Assume the function f is $1/\beta$ -smooth w.r.t. the norm $\|\cdot\|_f$, g^* is L -Lipschitz continuous, w.r.t. the dual norm $\|\cdot\|_*$, and we are using a sub-linearly convergent algorithm as quantified by (10). Then, for any $t \geq 0$ such that*

$$D(t) \geq \max\left\{\frac{2C\beta}{\sigma L^2}, \frac{2C\sigma L^2}{\beta\epsilon^2}\right\}, \quad (11)$$

it holds that $\mathbb{E}[G(\alpha^{(t)})] \leq \epsilon$.

Let us comment on Theorem 4 stated above. If $D(t) \sim \mathcal{O}(t)$ then this shows a rate of $\mathcal{O}(\epsilon^{-2})$. We note two important facts:

1. The guarantee holds for the duality gap of the iterate $\alpha^{(t)}$ and not for some averaged solution.
2. For the SVM case, this rate is consistent with the result of (Hush et al., 2006). Our result is much more general as it holds for any L -Lipschitz continuous convex function g^* and any β -strongly convex f^* .

Let us make one more important remark. In (Shalev-Shwartz & Zhang, 2013) the authors showed that SDCA (when applied on L -Lipschitz continuous g^*) has $D(t) \sim \mathcal{O}(t)$ and they also showed that an averaged solution (over few last iterates) needs only $\mathcal{O}(\epsilon^{-1})$ iterations to have duality gap $\leq \epsilon$. However, as a direct consequence of our Theorem 4 we can show, e.g., that FISTA (Beck & Teboulle, 2009) (aka. accelerated gradient descent algorithm) or APPROX (Fercoq & Richtárik, 2015)⁴ (a.k.a. accelerated coordinate descent algorithm) will need $\mathcal{O}(\epsilon^{-1})$ iterations to produce an iterate α such that $G(\alpha) \leq \epsilon$. Indeed, e.g., for APPROX, the function $D(t) = ((t-1)\tau + 2n)^2$ (where τ is the size of a mini-batch) and C is a constant which depends on $\alpha^{(0)}$ and α^* and τ . Hence, to obtain an iterate $\alpha^{(t)}$ such that $\mathbb{E}[G(\alpha^{(t)})] \leq \epsilon$ it is sufficient to choose $\tau \geq 1$ such that $t \stackrel{(11)}{\geq} 1 - \frac{2n}{\tau} + \max\{\frac{1}{\tau} \sqrt{\frac{2C\beta}{\sigma L^2}}, \frac{L}{\epsilon\tau} \sqrt{\frac{2C\sigma}{\beta}}\}$ is satisfied.

5. Extending Duality to Non-Lipschitz Functions

5.1. Lipschitzing Trick

In this section we present a trick that allows to generalize our results of the previous section from Lipschitz functions g^* to non-Lipschitz functions. The approach we propose, which we call the *Lipschitzing trick*, will make a convex function Lipschitz on the entire domain \mathbb{R}^d , by modifying its conjugate dual to have bounded support. Formally, the modification is as follows: Let $\mathcal{B} \subset \mathbb{R}^d$ be some closed, bounded convex set. We modify $g : \mathbb{R}^n \rightarrow \mathbb{R}$ by restricting its support to the set \mathcal{B} , i.e.

$$\bar{g}(\alpha) := \begin{cases} g(\alpha) & \text{if } \alpha \in \mathcal{B} \\ +\infty & \text{otherwise} \end{cases}. \quad (12)$$

By definition, this function has bounded support, and hence, by Lemma 5, its conjugate function \bar{g}^* is Lipschitz continuous.

Motivation. We will apply this trick to the part g of optimization problems of the form (A) (such that g^* will become Lipschitz). We want that this modification to have no impact on the outcome of any optimization algorithm run-

⁴APPROX requires g being separable.




Figure 1. Illustration Lipschitzing trick for the scalar function $g(\alpha) = |\cdot|$ with $\mathcal{B} = \{x : |x| \leq B\}$.

ning on (A). Instead, this trick should only affect convergence theory in that it allows us to present a strong primal-dual rate. In the following, we will discuss how this can indeed be achieved by imposing only very weak assumptions on the original problem (A).

The modification is based on the following duality of Lipschitzness and bounded support, as given in Lemma 5 below, which is a generalization of (Rockafellar, 1997, Corollary 13.3.3). We need the following definition:

Definition 1 (B -Bounded Support). A function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ has B -bounded support if its effective domain is bounded by B w.r.t. a norm $\|\cdot\|$, i.e.,

$$g(\mathbf{u}) < +\infty \Rightarrow \|\mathbf{u}\| \leq B. \quad (13)$$

Lemma 5 (Duality between Lipschitzness and L -Bounded Support). Given a proper convex function g , it holds that g has L -bounded support w.r.t. the norm $\|\cdot\|$ if and only if g^* is L -Lipschitz w.r.t. the dual norm $\|\cdot\|_*$.

The following Theorem 6 generalizes our previous convergence results, which were restricted to Lipschitz g^* .

Theorem 6. For an arbitrary optimization algorithm running on problem (A), let $\alpha^{(t)}$ denote its iterate sequence. Assume there is some closed convex set \mathcal{B} containing all these iterates. Then, the same optimization algorithm run on the modified problem — given by Lipschitzing of g^* using \mathcal{B} — would produce exactly the same iterate sequence. Furthermore, Theorem 3 as well as Theorem 4 give primal-dual convergence guarantees for this algorithm (for L such that \mathcal{B} is L -bounded).

Corollary 7. Assume the objective of optimization problem (A) has bounded level sets. For $\alpha^{(t)}$ being the iterate sequence of a monotone optimization algorithm on problem (A) we denote $\delta_t := \mathcal{D}(\alpha^{(t)})$ and let \mathcal{B}_t be the δ_t -level set of \mathcal{D} . Write $B_t > 0$ for a value such that \mathcal{B}_t is B_t -bounded w.r.t. the norm $\|\cdot\|$. Then, at any state t of the algorithm, the set \mathcal{B}_t contains all future iterates and Theorem 6 applies for $L := B_t$.

5.2. Norm-Regularized Problems

We now focus on some applications. First, we demonstrate how the Lipschitzing trick can be applied to find primal-dual convergence rates for problems regularized by an arbitrary norm. We discuss in particular the Lasso problem

and show how the suboptimality bound can be evaluated in practice. In a second part, we discuss the Elastic Net regularizer and show how it fits into our framework.

We consider a special structure of problem (A), namely

$$\min_{\alpha \in \mathbb{R}^n} f(A\alpha) + \lambda \|\alpha\|. \quad (14)$$

where f is some convex non-negative smooth loss function regularized by any norm $\|\cdot\|$. We choose \mathcal{B} to be the $\|\cdot\|$ -norm ball of radius B , such that $\|\alpha\| < B$ for every iterate. The size, B , of this ball can be controlled by the amount of regularization λ . Using the Lipschitzing trick with this \mathcal{B} , the convergence guarantees of Theorem 6 apply to the general class of norm regularized problems.

Note that for any monotone algorithm (initialized at $\mathbf{0}$) applied to (14) we have $\|\alpha\| \leq \frac{1}{\lambda}f(\mathbf{0})$ for every iterate. Furthermore, at every iterate α we can bound $\|\alpha^+\| \leq \frac{1}{\lambda}(f(A\alpha) + \lambda \|\alpha\|)$ for every future iterate α^+ . Hence, $B := \frac{1}{\lambda}f(\mathbf{0})$ is a safe choice, e.g. $B := \frac{1}{2\lambda}\|\mathbf{b}\|_2^2$ for least squares loss and $B := \frac{m}{\lambda} \log(2)$ for logistic regression loss.

Duality Gap. For any problem of the form (14) we can now determine the duality gap. We apply the Lipschitzing trick to $g(\alpha) := \lambda \|\alpha\|$ as in (12), then the convex conjugate of \bar{g} is

$$\bar{g}^*(\mathbf{u}) = \begin{cases} 0 & \|\mathbf{u}\|_* \leq \lambda \\ \max_{\alpha \in \mathcal{B}} \mathbf{u}^\top \alpha - \lambda \|\alpha\| & \text{else} \end{cases}. \quad (15)$$

where $\|\cdot\|_*$ denotes the dual norm of $\|\cdot\|$. Hence, using the primal-dual mapping $\mathbf{w}(\alpha) := \nabla f(A\alpha)$ we can write the duality gap of the modified problem as

$$\bar{G}(\alpha) = \langle \mathbf{w}(\alpha), A\alpha \rangle + \lambda \|\alpha\| + \bar{g}^*(-A^\top \mathbf{w}(\alpha)) \quad (16)$$

Note that the computation of the modified duality gap $\bar{G}(\alpha)$ is not harder than the computation of the original gap $G(\alpha)$ - it requires one pass over the data, assuming the choice of a suitable set \mathcal{B} in (15). Furthermore, note that in contrast to the unmodified duality gap $G(\alpha)$, which is only defined on the set $\{\alpha : \|A^\top \mathbf{w}(\alpha)\|_* \leq \lambda\}$, our new gap $\bar{G}(\alpha)$ is defined on the entire space \mathbb{R}^d .

As an alternative, (Mairal, 2010, Sec. 1.4.1 and App. D.2) has defined a different duality gap by shrinking the dual \mathbf{w} variable until it becomes feasible in $\|A^\top \mathbf{w}\|_* \leq \lambda$.

5.2.1. L_1 -REGULARIZED PROBLEMS

The results from the previous section can be ported to the important special case of L_1 -regularization:

$$\min_{\alpha \in \mathbb{R}^n} f(A\alpha) + \lambda \|\alpha\|_1. \quad (17)$$

We choose \mathcal{B} to be the L_1 -norm ball of radius B and then apply the Lipschitzing trick with \mathcal{B} to the regularization term in (17). An illustration of this modification as well as the impact on its dual are illustrated in Figure 1. Hence, our

theory (Theorem 6) gives primal-dual convergence guarantees for any algorithm applied to the Lasso problem (17). Furthermore, if the algorithm is monotone (initialized at $\alpha := \mathbf{0}$) we know that \mathcal{B} is $\frac{1}{\lambda}f(\mathbf{0})$ -bounded.

Duality Gap. The duality gap (16) for the modified Lasso problem can be computed at every iterate α as

$$\bar{G}(\alpha) = \langle \mathbf{w}, A\alpha \rangle + B \left[\|A^\top \mathbf{w}\|_\infty - \lambda \right]_+ + \lambda \|\alpha\|_1 \quad (18)$$

for $\mathbf{w} = \mathbf{w}(\alpha)$. For the derivation, see Appendix I.1.

5.2.2. GROUP LASSO, FUSED LASSO AND TV-NORM

Our algorithm-independent, primal-dual convergence guarantees and certificates presented so far are not restricted to L_p -regularized problems, but do in fact directly apply to many more general structured regularizers, some of them shown in see Table 1. This includes group Lasso (L_1/L_2) and other norms inducing structured sparsity (Bach et al., 2012), as well as other penalties such as e.g. the fused Lasso $g(\alpha) := \lambda \|M\alpha\|_1$. The total variation denoising problem is obtained for suitable choice of the matrix M .

5.2.3. ELASTIC NET REGULARIZED PROBLEMS

The second application we will discuss is Elastic Net regularization

$$\min_{\alpha \in \mathbb{R}^n} \ell(A\alpha) + \lambda \left(\frac{\eta}{2} \|\alpha\|_2^2 + (1 - \eta) \|\alpha\|_1 \right), \quad (19)$$

for fixed trade-off parameter $\eta \in (0, 1]$.

Our framework allows two different ways to solve problem (19): Either mapping it to formulation (A) (for $\ell=f$ smooth) as in row 2 of Table 1, or to (B) (for general ℓ) as in row 3 of Table 1. In both scenarios, Theorem 2 gives us a fast linear convergence guarantee for the duality gap, if ℓ is smooth. The other theorems apply accordingly for general ℓ when the problem is mapped to (B).

Whether the choice of a dual or primal optimizer will be more beneficial in practice depends on the case, and will be discussed in more detail for coordinate descent methods in Section 6.

Duality Gap. For the Elastic Net problem (19) mapped to (A), we can compute the duality gap (4) as follows:

$$\begin{aligned} G(\alpha) &= \langle \mathbf{w}, A\alpha \rangle + \frac{1}{2\eta\lambda} \sum_{i=1}^n \left[|A_{:,i}^\top \mathbf{w}| - (1 - \eta)\lambda \right]_+^2 \\ &\quad + \lambda \left(\frac{\eta}{2} \|\alpha\|_2^2 + (1 - \eta) \|\alpha\|_1 \right) \end{aligned}$$

with $\mathbf{w} = \mathbf{w}(\alpha) = \nabla \ell(A\alpha)$, see Appendix I.2.

Remark 1. As $\eta \rightarrow 0$ we approach the pure L_1 -case and this gap blows up as $G(\alpha) \rightarrow \infty$. Comparing this to (18), we see that the Lipschitzing trick allows to get certificates even in cases where the duality gap of the unmodified problem is infinity.

6. Coordinate Descent Algorithms

We now focus on a very important class of algorithms, that is coordinate descent methods. In this section, we show how our theory implies much more general primal-dual convergence guarantees for coordinate descent algorithms.

Partially Separable Problems. A widely used subclass of optimization problems arises when one part of the objective becomes separable. Formally, this is expressed as $g(\boldsymbol{\alpha}) = \sum_{i=1}^n g_i(\alpha_i)$ for univariate functions $g_i : \mathbb{R} \rightarrow \mathbb{R}$ for $i \in [n]$. Nicely in this case, the conjugate of g also separates as $g^*(\mathbf{y}) = \sum_i g_i^*(y_i)$. Therefore, the two optimization problems (A) and (B) write as

$$\mathcal{D}(\boldsymbol{\alpha}) := f(A\boldsymbol{\alpha}) + \sum_i g_i(\alpha_i) \quad (\text{SA})$$

$$\mathcal{P}(\mathbf{w}) := f^*(\mathbf{w}) + \sum_i g_i^*(-A_{:,i}^\top \mathbf{w}), \quad (\text{SB})$$

where $A_{:,i} \in \mathbb{R}^d$ denotes the i -th column of A .

The Algorithm. We consider the coordinate descent algorithm described in Algorithm 1. Initialize $\boldsymbol{\alpha}^{(0)} = \mathbf{0}$ and then, at each iteration, sample and update a random coordinate $i \in [n]$ of the parameter vector $\boldsymbol{\alpha}$ to iteratively minimize (SA). Finally, after T iterations output $\bar{\boldsymbol{\alpha}}$, the average vector over the latest $T - T_0$ iterates. The parameter T_0 is some positive number smaller than T .

Algorithm 1 Coordinate Descent on $\mathcal{D}(\boldsymbol{\alpha})$

- 1: **Input:** Data matrix A .
Starting point $\boldsymbol{\alpha}^{(0)} := \mathbf{0} \in \mathbb{R}^n$, $\mathbf{w}^{(0)} = \mathbf{w}(\boldsymbol{\alpha}^{(0)})$.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: Pick $i \in [n]$ randomly
 - 4: Find $\Delta\alpha_i$ minimizing $\mathcal{D}(\boldsymbol{\alpha}^{(t-1)} + \mathbf{e}_i \Delta\alpha_i)$
 - 5: $\boldsymbol{\alpha}^{(t)} \leftarrow \boldsymbol{\alpha}^{(t-1)} + \Delta\alpha_i \mathbf{e}_i$
 - 6: $\mathbf{w}^{(t)} \leftarrow \mathbf{w}(\boldsymbol{\alpha}^{(t)})$
 - 7: **end for**
 - 8: Let $\bar{\boldsymbol{\alpha}} = \frac{1}{T-T_0} \sum_{t=T_0}^{T-1} \boldsymbol{\alpha}^{(t)}$
-

As we will show in the following section, coordinate descent on $\mathcal{D}(\boldsymbol{\alpha})$ is not only an efficient optimizer of the objective $\mathcal{D}(\boldsymbol{\alpha})$, but also provably reduces the duality gap. Therefore, the same algorithm will simultaneously optimize the dual objective $\mathcal{P}(\mathbf{w})$.

6.1. Primal-Dual Analysis for Coordinate Descent

We first show linear primal-dual convergence rate of Algorithm 1 applied to (SA) for strongly convex g_i . Later, we will generalize this result to also apply to the setting of general Lipschitz g_i . This generalization together with the Lipschitzing trick will allow us to derive primal-dual convergence guarantees of coordinate descent for a much broader class of problems, including the Lasso problem.

For the following theorems we assume that the columns of

the data matrix A are scaled such that $\|A_{:,i}\| \leq R$ for all $i \in [n]$ and $\|A_{:,j}\| \leq P$ for all $j \in [d]$, for some norm $\|\cdot\|$.

Theorem 8. Consider Algorithm 1 applied to (SA). Assume f is a $1/\beta$ -smooth function w.r.t. the norm $\|\cdot\|$. Then, if g_i is μ -strongly convex for all i , it suffices to have a total number of iterations of

$$T \geq \left(n + \frac{nR^2}{\mu\beta}\right) \log \left(\left[n + \frac{nR^2}{\mu\beta} \right] \frac{\epsilon_D^{(0)}}{\epsilon} \right)$$

to get $\mathbb{E}[G(\boldsymbol{\alpha}^{(T)})] \leq \epsilon$. Moreover, to obtain an expected duality gap of $\mathbb{E}[G(\bar{\boldsymbol{\alpha}})] \leq \epsilon$ it suffices to have $T > T_0$ with

$$T_0 \geq \left(n + \frac{nR^2}{\mu\beta}\right) \log \left(\left[n + \frac{nR^2}{\mu\beta} \right] \frac{\epsilon_D^{(0)}}{(T-T_0)\epsilon} \right)$$

where $\epsilon_D^{(0)}$ is the initial suboptimality in $\mathcal{D}(\boldsymbol{\alpha})$.

Theorem 8 allows us to upper bound the duality gap, and hence the suboptimality, for every iterate $\boldsymbol{\alpha}^{(T)}$, as well as the average $\bar{\boldsymbol{\alpha}}$ returned by Algorithm 1. In the following we generalize this result to apply to L -Lipschitz functions g_i .

Theorem 9. Consider Algorithm 1 applied to (SA). Assume f is a $1/\beta$ -smooth function w.r.t. the norm $\|\cdot\|$. Then, if g_i^* is L -Lipschitz for all i , it suffices to have a total number of iterations of

$$T \geq \max \left\{ 0, n \log \frac{\epsilon_D^{(0)} \beta}{2L^2 R^2 n} \right\} + n + \frac{20n^2 L^2 R^2}{\beta \epsilon}$$

to get $\mathbb{E}[G(\bar{\boldsymbol{\alpha}})] \leq \epsilon$. Moreover, when $t \geq T_0$ with

$$T_0 = \max \left\{ 0, n \log \frac{\epsilon_D^{(0)} \beta}{2L^2 R^2 n} \right\} + \frac{16n^2 L^2 R^2}{\beta \epsilon}$$

we have the suboptimality bound of $\mathbb{E}[\mathcal{D}(\boldsymbol{\alpha}^{(t)}) - \mathcal{D}(\boldsymbol{\alpha}^*)] \leq \epsilon/2$, where $\epsilon_D^{(0)}$ is the initial suboptimality.

Remark 2. Theorem 9 shows that for Lipschitz g_i^* , Algorithm 1 has $\mathcal{O}(\epsilon^{-1})$ convergence in the suboptimality and $\mathcal{O}(\epsilon^{-1})$ convergence in $G(\bar{\boldsymbol{\alpha}})$. Comparing this result to Theorem 4 which suggests $\mathcal{O}(\epsilon^{-2})$ convergence in $G(\boldsymbol{\alpha})$ for $\mathcal{O}(\epsilon^{-1})$ convergent algorithms, we see that averaging the parameter vector crucially improves convergence in the case of non-smooth f .

Remark 3. Note that our Algorithm 1 recovers the widely used SDCA setting (Shalev-Shwartz & Zhang, 2013) as a special case, when we choose $f^* := \frac{\lambda}{2} \|\cdot\|_2^2$ in (SB). Furthermore, their convergence results for SDCA are consistent with our results and can be recovered as a special case of our analysis. See Corollaries 16, 18, 19 in Appendix J.

6.2. Application to L_1 and Elastic Net Regularized Problems

We now apply Algorithm 1 to the L_1 -regularized problems, as well as Elastic Net regularized problems. We state improved primal-dual convergence rates which are more tailored to the coordinate-wise setting.

Coordinate Descent on L_1 -Regularized Problems. In contrast to the general analysis of L_1 -regularized problems in Section 5.2.1, we can now exploit separability of $g(\alpha) := \lambda \|\alpha\|_1$ and apply the Lipschitzing trick coordinate-wise, choosing $\mathcal{B} := \{\alpha : |\alpha| \leq B\} \subset \mathbb{R}$. This results in the following stronger convergence results:

Corollary 10. *We can use the Lipschitzing trick together with Theorem 9 to derive a primal-dual convergence result for the Lasso problem (17). We find that the g_i^* are B -Lipschitz after applying the Lipschitzing trick to every g_i , and hence the total number of iterations needed on the Lasso problem to get a duality gap of $\mathbb{E}[G(\bar{\alpha})] \leq \epsilon$ is*

$$T \geq \max \left\{ 0, n \log \frac{\beta \epsilon_D^{(0)}}{2B^2 R^2 n} \right\} + n + \frac{20n^2 B^2 R^2}{\beta \epsilon}$$

Remark 4. *We specify the different parameters of Corollary 10 for least squares loss as well as the logistic regression loss (defined in Table 1). Both are 1-smooth (f^* is 1-strongly convex) and we have $\beta := 1$. The initial suboptimality $\epsilon_D^{(0)}$ can be upper bounded by $\frac{1}{2} \|\mathbf{b}\|_2^2$ for the former and by $m \log(2)$ for the latter. For B we choose $\frac{1}{\lambda} f(\mathbf{0})$.*

Coordinate Descent on Elastic Net Regularized Problems. In Section 5.2.3 we discussed how the Elastic Net problem in (19) can be mapped to our setup. In the first scenario (row 2, Table 1) we note that the resulting problem is partially separable and an instance of (SA). In the second scenario we map (19) to (B) (row 3, Table 1). Assuming that the loss function ℓ is separable, this problem is an instance of (SB). The convergence guarantees when applying Algorithm 1 on the primal or on the dual are summarized in Corollary 11.

Corollary 11. *Consider Algorithm 1 for an Elastic Net regularized problem (19), running on either the primal or the dual formulation. Then, to obtain a duality gap of $\mathbb{E}[G(\alpha^{(T)})] \leq \epsilon$, it suffices to have a total of*

$$T \geq (n + \frac{nR^2}{\lambda\eta\zeta}) \log\left(\left[n + \frac{nR^2}{\lambda\eta\zeta}\right] \frac{\epsilon_D^{(0)}}{\epsilon}\right)$$

iterations for coordinate descent on the primal (19) and

$$T \geq (d + \frac{dP^2}{\lambda\eta\zeta}) \log\left(\left[d + \frac{dP^2}{\lambda\eta\zeta}\right] \frac{\epsilon_D^{(0)}}{\epsilon}\right)$$

for coordinate descent on the dual of (19).

According to Corollary 11, the convergence rate is comparable for both scenarios. The constants however depend on the data matrix A – for $d \gg n$ the primal version is beneficial, whereas for $n \gg d$ the dual version is leading.

7. Numerical Experiments

Here we illustrate the usefulness of our framework by showcasing it for two important applications, each one showing two algorithm examples for optimizing (A).

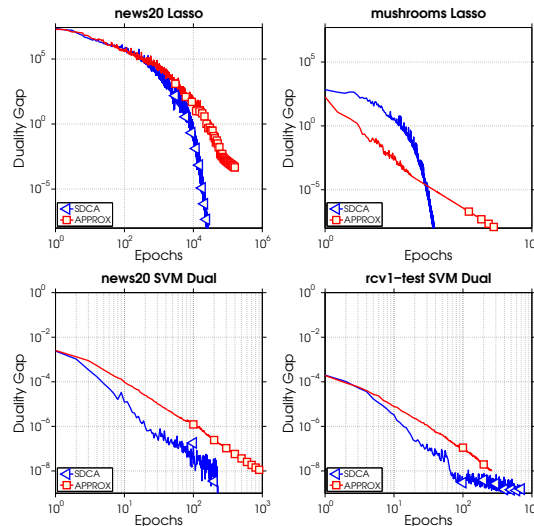


Figure 2. Comparison of CD with its accelerated variant – APPROX on Lasso and SVM problems.

Lasso. The top row of Figure 2 shows the primal-dual convergence of Algorithm 1 (CD) as well as the accelerated variant of CD (APPROX, Fercoq & Richtárik (2015)), both applied to the Lasso problem (A). We have applied the Lipschitzing trick as described in Section 5.1. This makes sure that $\mathbf{w}(\alpha)$ will be always feasible for the modified dual (B), and hence the duality gap can be evaluated.

SVM. It was shown in (Shalev-Shwartz & Zhang, 2013) that if CD (SDCA) is run on the dual SVM formulation, and we consider an “average” solution (over last few iterates), then the duality gap evaluated at averaged iterates has a sub-linear convergence rate $\mathcal{O}(1/t)$. As a consequence of Theorem 4, we have that the APPROX algorithm (Fercoq & Richtárik, 2015) will provide the same sub-linear convergence in duality gap, but holding for the iterates themselves, not only for an average. On the bottom row of Figure 2 we compare CD with its accelerated variant on two benchmark datasets.⁵ We have chosen $\lambda = 1/n$.

8. Conclusions

We have presented a general framework allowing to equip existing optimization algorithms with primal-dual certificates. For future research, it will be interesting to study more applications and algorithms fitting into the studied problem structure, including more cases of structured sparsity, and generalizations to matrix problems.

Acknowledgments. We thank Francis Bach, Michael P. Friedlander, Ching-pei Lee, Dmytro Perekrestenko, Aaditya Ramdas, Virginia Smith and an anonymous reviewer for fruitful discussions.

⁵Available from csie.ntu.edu.tw/~cjlin/libsvmtools/datasets.

References

- Bach, Francis. Duality Between Subgradient and Conditional Gradient Methods. *SIAM Journal on Optimization*, 25(1):115–129, 2015.
- Bach, Francis, Jenatton, Rodolphe, Mairal, Julien, and Obozinski, Guillaume. Optimization with Sparsity-Inducing Penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- Bauschke, Heinz H and Combettes, Patrick L. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer, 2011.
- Beck, Amir and Teboulle, Marc. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- Borwein, J M and Zhu, Q. *Techniques of Variational Analysis and Nonlinear Optimization*. Canadian Mathematical Society Books in Math, Springer, 2005.
- Boyd, Stephen P and Vandenberghe, Lieven. *Convex optimization*. Cambridge University Press, 2004.
- Chambolle, Antonin and Pock, Thomas. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2010.
- Fercoq, Olivier and Richtárik, Peter. Accelerated, Parallel, and Proximal Coordinate Descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- Forte, Simone. Distributed Optimization for Non-Strongly Convex Regularizers. Master’s thesis, ETH Zürich, 2015.
- Friedman, Jerome, Hastie, Trevor, and Tibshirani, Robert. Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, 33(1):1–22, 2010.
- Goldstein, Tom, Li, Min, and Yuan, Xiaoming. Adaptive Primal-Dual Splitting Methods for Statistical Learning and Image Processing. In *NIPS 2015 - Advances in Neural Information Processing Systems 28*, pp. 2080–2088, 2015.
- Hsieh, Cho-Jui, Chang, Kai-Wei, Lin, Chih-Jen, Keerthi, S Sathiya, and Sundararajan, Sellamanickam. A dual coordinate descent method for large-scale linear SVM. In *ICML 2008 - Proceedings of the 25th International Conference on Machine Learning*, pp. 408–415, 2008.
- Hush, Don, Kelly, Patrick, Scovel, Clint, and Steinwart, Ingo. Qp algorithms with guaranteed accuracy and run time for support vector machines. *The Journal of Machine Learning Research*, 7:733–769, 2006.
- Jaggi, Martin. An Equivalence between the Lasso and Support Vector Machines. In *Regularization, Optimization, Kernels, and Support Vector Machines*, pp. 1–26. Chapman and Hall/CRC, 2014.
- Kakade, Sham M, Shalev-Shwartz, Shai, and Tewari, Ambuj. On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. Technical report, Toyota Technological Institute - Chicago, USA, 2009.
- Lacoste-Julien, Simon and Jaggi, Martin. On the Global Linear Convergence of Frank-Wolfe Optimization Variants. In *NIPS 2015 - Advances in Neural Information Processing Systems 28*, 2015.
- Lacoste-Julien, Simon, Jaggi, Martin, Schmidt, Mark, and Pletscher, Patrick. Block-Coordinate Frank-Wolfe Optimization for Structural SVMs. In *ICML 2013 - Proceedings of the 30th International Conference on Machine Learning*, 2013.
- Lin, Qihang, Lu, Zhaosong, and Xiao, Lin. An accelerated proximal coordinate gradient method and its application to regularized empirical risk minimization. *arXiv:1407.1296*, 2014.
- Ma, Chenxin, Smith, Virginia, Jaggi, Martin, Jordan, Michael I, Richtárik, Peter, and Takáč, Martin. Adding vs. averaging in distributed primal-dual optimization. In *ICML 2015 - 32th International Conference on Machine Learning*, 2015a.
- Ma, Chenxin, Tappenden, Rachael, and Takáč, Martin. Linear convergence of the randomized feasible descent method under the weak strong convexity assumption. *arXiv:1506.02530*, 2015b.
- Mairal, Julien. *Sparse coding for machine learning, image processing and computer vision*. PhD thesis, ENS Cachan, 2010.
- Necoara, Ion. Linear convergence of first order methods under weak nondegeneracy assumptions for convex programming. *arXiv:1504.06298*, 2015.
- Nesterov, Yurii. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Nesterov, Yurii. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- Nesterov, Yurii. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- Qu, Zheng, Richtárik, Peter, and Zhang, Tong. Randomized Dual Coordinate Ascent with Arbitrary Sampling. *arXiv*, 2014.
- Richtárik, Peter and Takáč, Martin. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144(1-2):1–38, 2014.
- Rockafellar, R Tyrrell. *Convex analysis*. Princeton University Press, 1997.
- Shalev-Shwartz, Shai. Online Learning and Online Convex Optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- Shalev-Shwartz, Shai. SDCA without duality. *arXiv:1502.06177*, 2015.
- Shalev-Shwartz, Shai and Tewari, Ambuj. Stochastic Methods for l_1 -regularized Loss Minimization. *JMLR*, 12:1865–1892, 2011.
- Shalev-Shwartz, Shai and Zhang, Tong. Stochastic Dual Coordinate Ascent Methods for Regularized Loss Minimization. *JMLR*, 14:567–599, 2013.

- Shalev-Shwartz, Shai and Zhang, Tong. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, Series A:1–41, 2014.
- Smith, Virginia, Forte, Simone, Jordan, Michael I, and Jaggi, Martin. L1-Regularized Distributed Optimization: A Communication-Efficient Primal-Dual Framework. *arXiv*, 2015.
- Takáč, Martin, Bijral, Avleen, Richtárik, Peter, and Srebro, Nathan. Mini-batch primal and dual methods for SVMs. In *In ICML 2013 - 30th International Conference on Machine Learning*, 2013.
- Takáč, Martin, Richtárik, Peter, and Srebro, Nathan. Distributed mini-batch SDCA. *arXiv:1507.08322*, 2015.
- Tran-Dinh, Quoc and Cevher, Volkan. Constrained convex minimization via model-based excessive gap. In *NIPS 2014 - Advances in Neural Information Processing Systems 27*, 2014.
- Tran-Dinh, Quoc and Cevher, Volkan. Splitting the Smoothed Primal-Dual Gap: Optimal Alternating Direction Methods. *arXiv*, 2015.
- Wang, Po-Wei and Lin, Chih-Jen. Iteration complexity of feasible descent methods for convex optimization. *The Journal of Machine Learning Research*, 15(1):1523–1548, 2014.
- Yuan, Guo-Xun, Chang, Kai-Wei, Hsieh, Cho-Jui, and Lin, Chih-Jen. A Comparison of Optimization Methods and Software for Large-scale L1-regularized Linear Classification. *Journal of Machine Learning Research*, 11:3183–3234, 2010.
- Yuan, Guo-Xun, Ho, Chia-Hua, and Lin, Chih-Jen. An Improved GLMNET for L1-regularized Logistic Regression. *JMLR*, 13: 1999–2030, 2012.
- Yurtsever, Alp, Dinh, Quoc Tran, and Cevher, Volkan. A Universal Primal-Dual Convex Optimization Framework. In *NIPS 2015 - Advances in Neural Information Processing Systems 28*, pp. 3132–3140, 2015.
- Zhang, Yuchen and Lin, Xiao. Stochastic Primal-Dual Coordinate Method for Regularized Empirical Risk Minimization. In *ICML 2015 - Proceedings of the 32th International Conference on Machine Learning*, pp. 353–361, 2015.