

---

# Projected Newton-type Methods in Machine Learning

**Mark Schmidt**

*University of British Columbia  
Vancouver, BC, V6T 1Z4*

`schmidtmarkw@gmail.com`

**Dongmin Kim**

*University of Texas at Austin  
Austin, Texas 78712*

`dmkim@cs.utexas.edu`

**Suvrit Sra**

*Max Planck Institute for Biological Cybernetics  
72076, Tübingen, Germany*

`suvrit@tuebingen.mpg.de`

*We consider projected Newton-type methods for solving large-scale optimization problems arising in machine learning and related fields. We first introduce an algorithmic framework for projected Newton-type methods by reviewing a canonical projected (quasi-)Newton method. This method, while conceptually pleasing, has a high computation cost per iteration. Thus, we discuss two variants that are more scalable, namely, two-metric projection and inexact projection methods. Finally, we show how to apply the Newton-type framework to handle non-smooth objectives. Examples are provided throughout the chapter to illustrate machine learning applications of our framework.*

---

## 11.1 Introduction

We study Newton-type methods for solving the optimization problem

$$\min_{\mathbf{x}} f(\mathbf{x}) + r(\mathbf{x}), \quad \text{subject to } \mathbf{x} \in \Omega, \quad (1)$$

where  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is twice continuously differentiable and convex;  $r : \mathbb{R}^n \rightarrow \mathbb{R}$  is continuous and convex but not necessarily differentiable everywhere; while  $\Omega$  is a “simple” convex constraint-set. This formulation is general and captures numerous problems in machine learning, especially where  $f$  corresponds to a loss, and  $r$  to a regularizer. Let us, however, defer concrete examples of (1) until we have developed some theoretical background.

We propose to solve (1) via Newton-type methods, a certain class of second-order methods that are known to often work well for unconstrained problems. For constrained problems too, we may consider Newton-type methods that, akin to their unconstrained versions, iteratively minimize a quadratic approximation to the objective, this time subject to constraints. This idea dates back to Levitin and Polyak (1966, §7), and it is referred to as a *projected Newton* method.

Projected Newton methods for optimization over convex sets share many of the appealing properties of their unconstrained counterparts. For example, their iterations are guaranteed to improve the objective function for a small enough step size; global convergence can be shown under a variant of the Armijo condition; and rapid local convergence rates can be shown around local minima satisfying strong convexity (Bertsekas, 1999). In a similar vein, we may consider projected quasi-Newton methods, where we interpolate differences in parameter and gradient values to approximate the Hessian matrix. The resulting *Newton-type* methods are the subject of this chapter, and we will particularly focus on the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton approximation. The main appeal of the L-BFGS approximation is its linear-time iteration complexity, and its strong empirical performance on a variety of problems.

**Summary of remainder.** We first restrict ourselves to smooth optimization, where  $r(\mathbf{x}) = 0$ . For this setting, we describe projected Newton-type methods (§11.2), covering basic implementation issues such as Hessian approximation and line-search. Then, we describe two-metric projection methods (§11.3), followed by inexact (or truncated) projected-Newton methods (§11.4). Finally, we discuss the nonsmooth setting, where  $r(\mathbf{x}) \neq 0$ , for which we describe two Newton-type methods (§11.5).

---

## 11.2 Projected Newton-type Methods

Projected Newton-type methods optimize their objective iteratively. At iteration  $k$ , they first approximate the objective function around the current

iterate  $\mathbf{x}^k$  by the quadratic model

$$\mathcal{Q}^k(\mathbf{x}, \alpha) \triangleq f(\mathbf{x}^k) + (\mathbf{x} - \mathbf{x}^k)^T \nabla f(\mathbf{x}^k) + \frac{1}{2\alpha} (\mathbf{x} - \mathbf{x}^k)^T \mathbf{H}^k (\mathbf{x} - \mathbf{x}^k). \quad (2)$$

This model is parametrized by a positive stepsize  $\alpha$ , and it uses a positive-definite matrix  $\mathbf{H}^k$  to approximate the Hessian  $\nabla^2 f(\mathbf{x}^k)$ . To generate the next iterate that decreases the objective while remaining feasible, the methods minimize the quadratic model (2) over the (*convex*) constraint-set  $\Omega$ . Thus, for a fixed  $\alpha > 0$ , they compute the unique element

$$\bar{\mathbf{x}}_\alpha^k = \operatorname{argmin}_{\mathbf{x} \in \Omega} \mathcal{Q}^k(\mathbf{x}, \alpha), \quad (3)$$

which is then used to obtain the new iterate by simply setting

$$\mathbf{x}^{k+1} \leftarrow \mathbf{x}^k + \beta(\bar{\mathbf{x}}_\alpha^k - \mathbf{x}^k), \quad (4)$$

where  $\beta \in (0, 1]$  is another stepsize. To ensure a sufficient decrease in the objective value, one typically begins by setting  $\alpha = \beta = 1$ , and then decreases one of them until  $\mathbf{x}^{k+1}$  satisfies the following Armijo condition<sup>1</sup>

$$f(\mathbf{x}^{k+1}) \leq f(\mathbf{x}^k) + \nu \langle \nabla f(\mathbf{x}^k), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle, \quad \nu \in (0, 1). \quad (5)$$

We collect the above described steps into Algorithm 11.1, which we present as the general framework for projected Newton-type methods.

---

**Algorithm 11.1** A projected Newton-type method.

---

Given  $\mathbf{x}^0 \in \Omega$ ,  $\mathbf{H}^0 \succ \mathbf{0}$   
**for**  $k = 0, \dots$ , until some stopping criteria met **do**  
  **Step I:** Build  $\mathcal{Q}^k(\mathbf{x}, \alpha)$  using (2)  
  **repeat**  
    **Step IIa:** Minimize  $\mathcal{Q}^k(\mathbf{x}, \alpha)$  over  $\Omega$   
    **Step IIb:** Update  $\mathbf{x}^{k+1} \leftarrow \mathbf{x}^k + \beta(\bar{\mathbf{x}}_\alpha^k - \mathbf{x}^k)$   
    **Step III:** Update  $\alpha$  and/or  $\beta$   
  **until** descent condition (5) is satisfied  
**end for**

---

Convergence properties of various forms of this method are discussed, for example, in Bertsekas (1999, §2.3). In particular, convergence to a stationary point can be shown under the assumption that the eigenvalues of  $\mathbf{H}^k$  are bounded between two positive constants. Also, if  $\mathbf{x}^*$  is a minimizer of  $f(\mathbf{x})$  over  $\Omega$  satisfying certain conditions, and once  $\mathbf{x}^k$  is sufficiently close to  $\mathbf{x}^*$ ,

---

1. A typical value for the sufficient decrease parameter  $\nu$  is  $10^{-4}$ .

then  $\alpha = 1$  and  $\beta = 1$  are accepted as stepsizes and the sequence  $\|\mathbf{x}^k - \mathbf{x}^*\|$  converges to zero at a superlinear rate.

Algorithm 11.1 is conceptually simple, and thus appealing. It can, however have numerous variants depending on how each step is implemented. For example, in Step I, which particular quadratic model is used; in Step II, how we minimize the model function; and in Step III, how we compute the stepsizes  $\alpha$  and  $\beta$ . For each of these three steps one has multiple possible choices, and consequently, different combinations lead to methods of differing character. We describe some popular implementation choices below.

### 11.2.1 Building a quadratic model

When a positive-definite Hessian is readily available, we can simply set  $\mathbf{H}^k = \nabla^2 f(\mathbf{x}^k)$ . By doing so, the quadratic model (2) becomes merely the approximation obtained via a second-order Taylor expansion of  $f$ . This model leads to computation of an *exact* Newton-step at each iteration. At the other extreme, if we select  $\mathbf{H}^k = \mathbf{I}$ , the identity matrix of appropriate size, then the search direction of the resulting method reduces to the negative gradient, essentially yielding the projected gradient method. These two strategies often contrast each other in terms of computing a search (descent) direction: the Newton-step is considered one of the most sophisticated, while the gradient-step is regarded as one of the simplest. In cases where we can efficiently compute the Euclidean projection operator, projected gradient steps have a low per-iteration computational cost. However, this benefit comes at the expense of linear convergence speed. The Newton-step is usually more expensive; Step IIb will typically be costly to solve even if we can efficiently compute the Euclidean projection onto the constraint set. However, the more expensive Newton-step generally enjoys a local superlinear convergence rate.

Despite its theoretical advantages, an exact Newton-step often is resource intensive, especially when computing the exact Hessian is expensive. To circumvent some of the associated computational issues, one usually approximates the Hessian: this idea underlies the well-known quasi-Newton approximation. Let us therefore briefly revisit the BFGS update that approximates the exact Hessian.

**BFGS update.** There exist several approximations to the Hessian: e.g., Powell-Symmetric-Broyden (PSB), Davidson-Fletcher-Powell (DFP), or the Broyden-Fletcher-Goldfarb-Shanno (BFGS). We focus on BFGS, as it is believed to be the most effective in general (Gill et al., 1981; Bertsekas, 1999).

First, define the difference vectors  $\mathbf{g}$  and  $\mathbf{s}$  as follows:

$$\mathbf{g} = \nabla f(\mathbf{x}^{k+1}) - \nabla f(\mathbf{x}^k), \quad \text{and} \quad \mathbf{s} = \mathbf{x}^{k+1} - \mathbf{x}^k.$$

Now, assume we already have  $\mathbf{H}^k$ , the current approximation to the Hessian. Then, the BFGS update adds a rank-two correction to  $\mathbf{H}^k$  to obtain

$$\mathbf{H}^{k+1} = \mathbf{H}^k - \frac{\mathbf{H}^k \mathbf{s} \mathbf{s}^T \mathbf{H}^k}{\mathbf{s}^T \mathbf{H}^k \mathbf{s}} + \frac{\mathbf{g} \mathbf{g}^T}{\mathbf{s}^T \mathbf{g}}. \quad (6)$$

We can plug  $\mathbf{H}^{k+1}$  into (2) to obtain an updated model  $\mathcal{Q}^{k+1}$ . But depending on the implementation of subsequent steps (3) and (4), it might be more convenient and computationally efficient to update an estimate to the inverse of  $\mathbf{H}^k$  instead. For this case, we can apply the Sherman-Morrison-Woodbury formula to (6), thus obtaining the update

$$\mathbf{S}^{k+1} = \mathbf{S}^k + \left(1 + \frac{\mathbf{g}^T \mathbf{S}^k \mathbf{g}}{\mathbf{s}^T \mathbf{g}}\right) \frac{\mathbf{s} \mathbf{s}^T}{\mathbf{s}^T \mathbf{g}} - \frac{(\mathbf{S}^k \mathbf{g} \mathbf{s}^T + \mathbf{s} \mathbf{g}^T \mathbf{S}^k)}{\mathbf{s}^T \mathbf{g}}, \quad (7)$$

where  $\mathbf{S}^k$  is the inverse of  $\mathbf{H}^k$ , also known as the *gradient scaling* matrix.

**Limited memory BFGS update.** Though the BFGS update may greatly relieve the burden of Hessian computation, it still require the same storage:  $\mathcal{O}(n^2)$  for dense  $\mathbf{H}^k$  or  $\mathbf{S}^k$ , which is troublesome for large-scale problems. This difficulty is addressed by the limited memory BFGS (L-BFGS) update, where, instead of using full matrices  $\mathbf{H}^k$  and  $\mathbf{S}^k$ , a small number of vectors, say  $m$ , are used to approximate the Hessian or its inverse. The standard L-BFGS approach (Nocedal, 1980) can be implemented using the following formula (Nocedal and Wright, 2000)

$$\begin{aligned} \mathbf{S}^k &= \frac{\mathbf{s}_{k-1}^T \mathbf{g}_{k-1}}{\mathbf{g}_{k-1}^T \mathbf{g}_{k-1}} \bar{\mathbf{V}}_{k-M}^T \bar{\mathbf{V}}_{k-M} + \rho_{k-M} \bar{\mathbf{V}}_{k-M+1}^T \mathbf{s}_{k-M} \mathbf{s}_{k-M}^T \bar{\mathbf{V}}_{k-M+1} \\ &\quad + \rho_{k-M+1} \bar{\mathbf{V}}_{k-M+2}^T \mathbf{s}_{k-M+1} \mathbf{s}_{k-M+1}^T \bar{\mathbf{V}}_{k-M+2} \\ &\quad + \cdots \\ &\quad + \rho_{k-1} \mathbf{s}_{k-1} \mathbf{s}_{k-1}^T, \end{aligned} \quad (8)$$

for  $k \geq 1$ ; the scalars  $\rho_k$ , and matrices  $\bar{\mathbf{V}}_{k-M}$  are defined by

$$\rho_k = 1/(\mathbf{s}_k^T \mathbf{g}_k), \quad \bar{\mathbf{V}}_{k-M} = [\mathbf{V}_{k-M} \cdots \mathbf{V}_{k-1}], \quad \text{and} \quad \mathbf{V}_k = \mathbf{I} - \rho_k \mathbf{s}_k \mathbf{g}_k^T.$$

The L-BFGS approximation only requires  $\mathcal{O}(mn)$  storage; moreover, multiplication of  $\mathbf{S}^k$  by a vector can also be performed at this cost.

For both BFGS and L-BFGS, a choice that can significantly impact performance is the initial approximation  $\mathbf{H}^0$ . A typical strategy to select  $\mathbf{H}^0$  is

to first set it to the negative gradient direction on the first iteration, and then to set  $\mathbf{H}^0 = (\mathbf{g}^T \mathbf{g}) / (\mathbf{g}^T \mathbf{s}) \mathbf{I}$  on the next iteration. This choice was proposed by Shanno and Phua (1978) to optimize the condition number of the approximation. In the L-BFGS method we can reset  $\mathbf{H}^0$  using this formula after *each* iteration (Nocedal and Wright, 2000, §7.2). With this strategy, the unit stepsizes of  $\alpha = 1$  and  $\beta = 1$  are typically accepted, which may remove the need for a line search on most iterations.

Provided that  $\mathbf{H}^0$  is positive-definite, the subsequent (implicit) Hessian approximations  $\mathbf{H}^k$  generated by the L-BFGS update are guaranteed to be positive-definite as long as  $\mathbf{g}^T \mathbf{s}$  is positive (Nocedal and Wright, 2000, §6.1). This positivity is guaranteed if  $f(\mathbf{x})$  is strongly-convex, but when  $f(\mathbf{x})$  is not strongly-convex a more advanced strategy is required, see e.g., (Nocedal and Wright, 2000, §18.3).

### 11.2.2 Solving the subproblem

With our quadratic approximation  $\mathcal{Q}^k(\mathbf{x}, \alpha)$  in hand, the next step is to solve the subproblem (3). For  $\alpha \neq 0$ , simple rearrangement shows that<sup>2</sup>

$$\bar{\mathbf{x}}_\alpha^k = \operatorname{argmin}_{\mathbf{x} \in \Omega} \mathcal{Q}^k(\mathbf{x}, \alpha) = \operatorname{argmin}_{\mathbf{x} \in \Omega} \frac{1}{2} \|\mathbf{x} - \mathbf{y}^k\|_{\mathbf{H}^k}^2, \quad (9)$$

where  $\|\mathbf{x}\|_{\mathbf{H}^k}$  is defined by the norm  $\sqrt{\mathbf{x}^T \mathbf{H}^k \mathbf{x}}$ , and  $\mathbf{y}^k$  is the unconstrained Newton step:  $\mathbf{y}^k = \mathbf{x}^k - \alpha [\mathbf{H}^k]^{-1} \nabla f(\mathbf{x}^k)$ . In words,  $\bar{\mathbf{x}}_\alpha^k$  is obtained by projecting the Newton-step onto the constraint-set  $\Omega$ , where projection is with respect to the metric defined by the Hessian approximation  $\mathbf{H}^k$ .

One major drawback of (9) is that it can be computationally challenging, even when  $\Omega$  has relatively simple structure. To ease the computational burden, instead of using the metric defined by  $\mathbf{H}^k$ , we could compute the projection under the standard Euclidean norm, while slightly modifying the Newton step to ensure convergence. This is the subject of §11.3. Alternately, in §11.4 we consider computing an *approximate* solution to (9) itself.

*Remark:* If we replace  $\mathbf{H}^k$  by  $\mathbf{I}$  both as the projection metric, as well as in the Newton step, we recover gradient projection methods.

### 11.2.3 Computing the stepsizes

Consider the stepsizes  $\alpha$  and  $\beta$  in (3) and (4). Generally speaking, any positive  $\alpha$  and  $\beta$  that generate  $\mathbf{x}^{k+1}$  satisfying the descent-condition (5) are

---

2. If we use  $\mathbf{S}^k$ , the inverse of Hessian, then  $\bar{\mathbf{x}}_\alpha^k$  may be equivalently obtained by solving  $\operatorname{argmin}_{\mathbf{x} \in \Omega} \frac{1}{2} \|\mathbf{x} - (\mathbf{x}^k - \alpha \mathbf{S}^k \nabla f(\mathbf{x}^k))\|_{[\mathbf{S}^k]^{-1}}^2 = \operatorname{argmin}_{\mathbf{x} \in \Omega} \frac{1}{2} \|\mathbf{x} - \mathbf{y}^k\|_{[\mathbf{S}^k]^{-1}}^2$ .

acceptable. Practical choices are discussed below.

**Backtracking.** Suppose we fix  $\alpha = 1$  for all  $k$ , and let  $\mathbf{d}^k = \bar{\mathbf{x}}_1^k - \mathbf{x}^k$ . Then we obtain the following update:

$$\mathbf{x}^{k+1} \leftarrow \mathbf{x}^k + \beta \mathbf{d}^k.$$

To select  $\beta$ , we can simply start with  $\beta = 1$  and iteratively decrease  $\beta$  until the resulting  $\mathbf{x}^{k+1}$  satisfies (5).<sup>3</sup> More formally, we set  $\beta = \tau \cdot \sigma^m$  for some  $\tau > 0$  and  $\sigma \in (0, 1)$ , where  $m \geq 0$  is the first integer that satisfies

$$f(\mathbf{x}^{k+1}) \leq f(\mathbf{x}^k) + \tau \cdot \sigma^m \nabla f(\mathbf{x}^k)(\mathbf{x}^{k+1} - \mathbf{x}^k).$$

Several strategies are available to reduce the number of backtracking iterations. For example, rather than simply dividing the stepsize by a constant we can use information collected about the function during the line search to make a more intelligent choice. For example, if some trial value of  $\beta$  is not accepted then we can set the next  $\beta$  to the minimum of the quadratic polynomial that has a value of  $f(\mathbf{x}^k)$  at zero,  $f(\mathbf{x}^k + \beta \mathbf{d}^k)$  at  $\beta$ , and a slope of  $\nabla f(\mathbf{x}^k)^T \mathbf{d}^k$  at zero (Nocedal and Wright, 2000, §3.5). This choice gives the optimal stepsize if  $f(\mathbf{x})$  is a quadratic function, and often drastically reduces the number of backtracking iterations needed. For some functions, quadratic interpolation can also be used to give a more intelligent choice than  $\beta = 1$  for the *first* trial value of  $\beta$ , while cubic interpolation can be used if we have tested more than one value of  $\beta$  or if we compute  $\nabla f(\mathbf{x}^k + \beta \mathbf{d}^k)$  for the trial values of  $\beta$  (Nocedal and Wright, 2000, §3.5).

**Backtracking (Armijo) along projection arc (Bertsekas, 1999)** Alternatively, we can set  $\beta = 1$  for all  $k$  to obtain

$$\mathbf{x}^{k+1} \leftarrow \bar{\mathbf{x}}_\alpha^k,$$

and then determine an  $\alpha$  satisfying (5). Similar to simple backtracking, we compute  $\alpha = s \cdot \sigma^m$  for some  $s, \tau > 0$ , and  $\sigma \in (0, 1)$ , where  $m \geq 0$  is the first integer that satisfies

$$f(\bar{\mathbf{x}}_{s, \sigma^m}^k) \leq f(\mathbf{x}^k) + \tau \nabla f(\mathbf{x}^k)(\bar{\mathbf{x}}_{s, \sigma^m}^k - \mathbf{x}^k).$$

Unlike simple backtracking that searches along a line segment as  $\beta$  varies, this strategy searches along a potentially non-linear path as  $\alpha$  varies. Because of this, polynomial interpolation to select trial values of  $\alpha$  is more

---

3. Also known as Armijo backtracking along feasible direction.

tenuous than for simple backtracking, but on many problems polynomial interpolation still significantly decreases the number of trial values evaluated.

This stepsize computation might be more involved when computing projections onto  $\Omega$  is expensive, since it requires solving an optimization problem to compute  $\bar{\mathbf{x}}_\alpha^k$  for each trial value of  $\alpha$ . However, it can still be appealing because it is more likely to yield iterates that lie on the boundaries of the constraints. This property is especially useful when the boundaries of the constraints represent a solution of interest, e.g., *sparse* solutions with the constraint set  $\Omega = \mathbb{R}_+^n$ .

In the sequel, we consider a few specific instantiations of the general framework introduced in this section. Specifically, we first consider two-metric projection methods for the specific case of bound-constrained problems (§11.3). Subsequently, we consider inexact projected-Newton methods for optimization over more general simple convex sets (§11.4). Finally, we explore the versatility of the framework by extending it to problems with non-smooth objective functions (§11.5).

---

### 11.3 Two-metric Projection Methods

As mentioned earlier, computing the projection with respect to a quadratic norm defined by  $\mathbf{H}^k$  can be computationally challenging. However, we often encounter problems with simple convex domains, onto which we can efficiently compute Euclidean projections. For optimization over such domains, we might therefore prefer projecting the Newton-step under the Euclidean norm. Indeed, this choice is made by the well-known *two-metric projection* method, named so, because it uses a different matrices (metrics) for scaling the gradient and for computing the projection.

In two-metric projection algorithms, we can benefit from low iteration complexity if we use L-BFGS approximations. However, some problems still persist: the “obvious” procedure with an unmodified Newton-step may not improve on the objective function, even for an arbitrarily small positive stepsize. Nevertheless, there are many cases where one can derive a two-metric projection method that can dodge this drawback without giving up the attractive properties of its unconstrained counterpart. A particular example is Bertsekas’ projected-Newton method (Bertsekas, 1982), and we discuss it below for the case where  $\Omega$  consists of bound constraints.

The projected-Newton method may be viewed in light of Algorithm 11.1. Specifically, it takes the Hessian  $\nabla^2 f(\mathbf{x}^k)$  and modifies its *inverse* so that the gradient scaling matrix  $\mathbf{S}^k$  has a special structure. It subsequently



invokes orthogonal projection in Step II, and then in Step III, it computes its stepsize using backtracking along the projection arc. The key variation from Algorithm 11.1 lies in how to modify the inverse Hessian to obtain a valid gradient scaling; the details follow below.

### 11.3.1 Bound constrained smooth convex problems

Consider the following special case of (1)

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}), \quad \text{subject to } \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \quad (10)$$

where  $\mathbf{l}$  and  $\mathbf{u}$  are fixed vectors, and inequalities are taken componentwise (which can be set to  $\infty$  or  $-\infty$  if the variables are unbounded). The function  $f$  is assumed to be convex and twice continuously differentiable. Such bound-constrained problems arise when we have natural restrictions (e.g., nonnegativity) on the variables, or as Lagrangian duals of problems with convex inequality constraints. For bound-constrained problems the projection under the Euclidean norm is the standard orthogonal projection obtained by taking componentwise medians among  $l_i$ ,  $u_i$  and  $x_i$ ,

$$[\mathcal{P}(\mathbf{x})]_i \triangleq \text{mid}\{l_i, x_i, u_i\}.$$

At each iteration, we partition the variables into two groups: *free* and *restricted*. Restricted variables are defined as a particular subset of the variables close to their bounds, based on the sign of the corresponding components in the gradient. Formally, the set of *restricted* variables is:

$$\mathcal{J}^k \triangleq \{i \mid x_i^k \leq l_i + \varepsilon \wedge \partial_i f(\mathbf{x}^k) > 0, \quad \text{or} \quad x_i^k \geq u_i - \varepsilon \wedge \partial_i f(\mathbf{x}^k) < 0\},$$

for some small positive  $\varepsilon$ . The set  $\mathcal{J}^k$  collects variables that are near their bounds, *and* for which the objective  $f(\mathbf{x})$  can be decreased by moving the variables towards (or past) their bounds. The set of *free* variables, denoted  $\mathcal{F}^k$ , is simply defined as the complement of  $\mathcal{J}^k$  in the set  $\{1, 2, \dots, n\}$ .

Without loss of generality, let us assume that  $\mathcal{F}^k = \{1, 2, \dots, N\}$  and  $\mathcal{J}^k = \{N+1, \dots, n\}$ . Now define a diagonal matrix  $\mathbf{D}^k \in \mathbb{R}^{n-N \times n-N}$  that scales the *restricted* variables, a typical choice being the identity matrix. We denote the scaling with respect to the free variables as  $\bar{\mathbf{S}}^k \in \mathbb{R}^{N \times N}$ , which, for the projected-Newton method, is given by the principal sub-matrix of the inverse of the Hessian  $\nabla^2 f(\mathbf{x}^k)$ , as induced by the free variables. In symbols,

$$\bar{\mathbf{S}}^k \leftarrow [\nabla^2 f(\mathbf{x}^k)]_{\mathcal{F}^k}^{-1}. \quad (11)$$

With these definitions, we are now ready to present the main step of the two-metric projection algorithm. This step can be written as the Euclidean

projection of a Newton-step that uses a gradient scaling  $\mathbf{S}^k$  of the form

$$\mathbf{S}^k \triangleq \begin{bmatrix} \bar{\mathbf{S}}^k & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^k \end{bmatrix}. \quad (12)$$

The associated stepsize  $\alpha$  can be selected by backtracking along the projection arc until the Armijo condition is satisfied. Note that this choice of the stepsize computation does not increase the computational complexity of the method, since computing the orthogonal projection after each backtracking step is trivial. Combining this gradient scaling with orthogonal projection, we obtain the projected-Newton update:

$$\begin{aligned} \mathbf{x}^{k+1} \leftarrow \bar{\mathbf{x}}_\alpha^k &= \operatorname{argmin}_{\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \frac{1}{2} \|\mathbf{x} - (\mathbf{x}^k - \alpha \mathbf{S}^k \nabla f(\mathbf{x}^k))\|_{[\mathbf{S}^k]^{-1}}^2 \\ &\approx \operatorname{argmin}_{\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \frac{1}{2} \|\mathbf{x} - (\mathbf{x}^k - \alpha \mathbf{S}^k \nabla f(\mathbf{x}^k))\|_{\mathbf{I}}^2 \\ &= \mathcal{P}[\mathbf{x}^k - \alpha^k \mathbf{S}^k \nabla f(\mathbf{x}^k)], \end{aligned} \quad (13)$$

where  $\alpha^k$  is computed by backtracking along the projection arc.

This algorithm has been shown to be globally convergent (Bertsekas, 1982; Gafni and Bertsekas, 1984), and under certain conditions achieves local superlinear convergence.

**Theorem 11.1** (Convergence). *Assume that  $\nabla f$  is Lipschitz continuous on  $\Omega$ , and  $\nabla^2 f$  has bounded eigenvalues. Then every limit point of  $\{\mathbf{x}^k\}$  generated by iteration (13) is a stationary point of (10).*

**Theorem 11.2** (Convergence rate). *Let  $f$  be strictly convex and twice continuously differentiable. Let  $\mathbf{x}^*$  be the non-degenerate optimum of Problem (13) and assume that for some  $\delta > 0$ ,  $\nabla^2 f(\mathbf{x})$  has bounded eigenvalues for all  $\mathbf{x}$  that satisfy  $\|\mathbf{x} - \mathbf{x}^*\| < \delta$ . Then the sequence  $\{\mathbf{x}^k\}$  generated by iteration (13) converges to  $\mathbf{x}^*$ , and the rate of convergence in  $\{\|\mathbf{x}^k - \mathbf{x}^*\|\}$  is superlinear.*

We remark that though the convergence rate of the two-metric projection method has been shown for  $\mathbf{S}^k$  derived from the Hessian, the convergence itself merely requires a positive-definite gradient scaling  $\mathbf{S}^k$  with bounded eigenvalues for all  $k$  (Bertsekas, 1982). Thus, the quasi-Newton approximations introduced in §11.2 are viable choices to derive convergent methods,<sup>4</sup> we present such variations of the two-metric method in the following exam-

---

4. In a simpler but still globally-convergent variation of the two-metric projection method, we could simply set  $\mathbf{S}^k$  to be a diagonal matrix with positive diagonal elements.

ple.

**Example 11.3** (Nonnegative least-squares). *A problem of considerable importance in the applied sciences is the nonnegative least-squares (NNLS):*

$$\min_{\mathbf{x}} \quad \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2, \quad \text{subject to } \mathbf{x} \geq \mathbf{0}, \quad (14)$$

where  $\mathbf{A} \in \mathbb{R}^{m \times n}$ . This problem is essentially an instance of (10).

Given Algorithm 11.1, one can simply implement the update (13) and then use BFGS or L-BFGS to obtain  $\mathbf{S}^k$ . However, we can further exploit the simple constraint  $\mathbf{x} \geq \mathbf{0}$  and improve the computational (empirical) efficiency of the algorithm. To see how, consider the restricted variables in the update (13). When variable  $i \in \mathcal{J}^k$  and  $\varepsilon$  becomes sufficiently small, we obtain

$$\mathcal{P}[\mathbf{x}^k - \alpha^k \mathbf{S}^k \nabla f(\mathbf{x}^k)]_i = \mathcal{P}[x_i^k - \alpha^k [\mathbf{D}^k]_{ii} \cdot \partial_i f(\mathbf{x}^k)] = 0.$$

In other words, if  $i \in \mathcal{J}^k$ , then  $x_i^{k+1} = 0$ , whereby we can safely ignore these variables throughout the update. In an implementation, this means that we can confine computations to free variables, which can save a large number of floating point operations especially when  $|\mathcal{F}^k| \ll |\mathcal{J}^k|$ .

**Example 11.4** (Linear SVM). *Consider the standard binary classification task with inputs  $(\mathbf{x}_i, y_i)_{i=1}^m$ , where  $\mathbf{x}_i \in \mathbb{R}^n$  and  $y_i \in \pm 1$ . Assume for simplicity that we wish to learn a bias-free decision function  $f(\mathbf{x}) = \text{sgn}(\mathbf{w}^T \mathbf{x})$  by solving either the SVM primal*

$$\begin{aligned} \underset{\mathbf{w}}{\text{minimize}} \quad & \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{i=1}^m \xi_i \\ \text{subject to} \quad & y_i (\mathbf{w}^T \mathbf{x}_i) \geq 1 - \xi_i, \quad \xi_i \geq 0, \quad 1 \leq i \leq m, \end{aligned} \quad (15)$$

or its (more familiar) dual

$$\begin{aligned} \underset{\boldsymbol{\alpha}}{\text{minimize}} \quad & \frac{1}{2} \boldsymbol{\alpha}^T \mathbf{Y} \mathbf{X}^T \mathbf{X} \mathbf{Y} \boldsymbol{\alpha} - \boldsymbol{\alpha}^T \mathbf{1} \\ \text{subject to} \quad & 0 \leq \alpha_i \leq C, \end{aligned} \quad (16)$$

where  $\mathbf{Y} = \text{Diag}(y_1, \dots, y_m)$  and  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m] \in \mathbb{R}^{n \times m}$ . The dual (16) is a special case of (10), and can be solved by adapting the two-metric projection method in a manner similar to that for NNLS.

**Example 11.5** (Sparse Gaussian Graphical Models). *The dual to a standard formulation for learning sparse Gaussian graphical models takes the form (Banerjee et al., 2006):*

$$\min_{\tilde{\boldsymbol{\Sigma}} + \mathbf{X} \succ \mathbf{0}} \quad -\log \det(\tilde{\boldsymbol{\Sigma}} + \mathbf{X}), \quad \text{subject to } |\mathbf{X}_{ij}| \leq \lambda_{ij}, \quad \forall ij. \quad (17)$$

There has been substantial recent interest in solving this problem (Banerjee

et al., 2006). Here  $\tilde{\Sigma}$  represents the empirical covariance of a data set, and the bound constraints on the elements of the matrix encourage the associated graphical model to be sparse for sufficiently large values of the  $\lambda_{ij}$  variables.

Notice that the constraint  $|\mathbf{X}_{ij}| \leq \lambda_{ij}$  is equivalent to the box-constraints:  $-\lambda_{ij} \leq \mathbf{X}_{ij} \leq \lambda_{ij}$ . Thus, provided  $\tilde{\Sigma} + \mathbf{X}$  is positive-definite for the initial  $\mathbf{X}$ , we can apply a simplified two-metric projection algorithm to this problem where we use projection to address the bound constraints and backtracking to modify the iterates when they leave the positive-definite cone.

## 11.4 Inexact Projection Methods

The previous section focused on examples with bound constraints. For optimizing over more general but still simple convex sets, an attractive choice is *inexact* projected-Newton methods. These methods represent a natural generalization of methods for unconstrained optimization that are alternately referred to as *Hessian-free*, *truncated*, or *inexact* Newton methods. In inexact projected-Newton methods, rather than finding the exact minimizer in Step IIa of Algorithm 11.1, we find an approximate minimizer using an iterative solver. That is, we use a single-metric projection, but solve the projection inexactly. Note that the iterative solver can be a first-order optimization strategy, and thus can take advantage of an efficient Euclidean projection operator. Under only mild conditions on the iterative solver, this approximate projection algorithm still leads to an improvement in the objective function. There are many ways to implement an inexact projected-Newton strategy but in this section we focus on the one described in Schmidt et al. (2009). In this method, we use the L-BFGS Hessian approximation, which we combine with simple Armijo backtracking and a variant of the projected gradient algorithm for iteratively solving subproblems. In the next section we review an effective iterative solver, and in the sequel we discuss using it within an inexact projected-Newton method.

### 11.4.1 Spectral Projected Gradient

The traditional motivation for examining projected Newton methods is that the basic gradient projection method may take a very large number of iterations to reach an acceptably accurate solution. However, there has been substantial recent interest on variants of gradient projection that exhibit much better empirical convergence properties. For example, Birgin et al. (2000) presented several *spectral* projected gradient (SPG) methods. In SPG methods, either  $\alpha$  or  $\beta$  is set to one, and the other step size is set to one of

the step sizes proposed by (Barzilai and Borwein, 1988). For example, we might set  $\beta = 1$  and  $\alpha$  to

$$\alpha^{bb} \triangleq \frac{\mathbf{g}^T \mathbf{s}}{\mathbf{g}^T \mathbf{g}}, \quad \text{where } \mathbf{g} = \nabla f(\mathbf{x}^{k+1}) - \nabla f(\mathbf{x}^k), \text{ and } \mathbf{s} = \mathbf{x}^{k+1} - \mathbf{x}^k. \quad (18)$$

Subsequently, backtracking along one of the two step sizes is used to satisfy the *non-monotonic* Armijo condition (Grippo et al., 1986)

$$f(\mathbf{x}^{k+1}) \leq \max_{i=k-m:k} \{f(\mathbf{x}^i)\} + \tau \nabla f(\mathbf{x}^k)(\mathbf{x}^{k+1} - \mathbf{x}^k), \quad \tau \in (0, 1),$$

which, unlike the ordinary Armijo condition (5), allows some temporary increase in the objective. This non-monotonic Armijo condition typically accepts the initial step length even if it increases the objective function, while still ensuring global convergence of the method.<sup>5</sup> Experimentally, these two simple modifications lead to large improvements in the convergence speed of the method. Indeed, due to its strong empirical performance, SPG has recently also been explored in several other applications (Dai and Fletcher, 2005; Figueiredo et al., 2007; van den Berg and Friedlander, 2008).

An alternative to SPG for accelerating the basic projected gradient method is the method of Nesterov (2004, §2.2.4). In this strategy, an extra extrapolation step is added to the iteration, thereby allowing the method to achieve the optimal worst-case convergence rate among a certain class of algorithms. Besides SPG and this optimal gradient algorithm, there can be numerous alternative iterative solvers. But we restrict our discussion to an SPG-based method and consider for it some implementation and theoretical details.

#### 11.4.2 SPG-based Inexact Projected-Newton

Recall Step IIa in the general framework of Algorithm 11.1.<sup>6</sup>

$$\bar{\mathbf{x}}_1^k = \operatorname{argmin}_{\mathbf{x} \in \Omega} \mathcal{Q}^k(\mathbf{x}, 1), \quad (19)$$

where the quadratic model is

$$\mathcal{Q}^k(\mathbf{x}, 1) = f(\mathbf{x}^k) + (\mathbf{x} - \mathbf{x}^k)^T \nabla f(\mathbf{x}^k) + \frac{1}{2} (\mathbf{x} - \mathbf{x}^k)^T \mathbf{H}^k (\mathbf{x} - \mathbf{x}^k).$$

5. A typical value for the number  $m$  of previous function values to consider is 10.

6. We assume that  $\alpha = 1$  and backtrack along  $\beta$  so that the iterative solver is invoked only once for each iteration; however, the inexact Newton method does not rule out the possibility of fixing  $\beta$  (and invoking the iterative solver for each backtracking step).

For the remainder of this section, we denote  $\mathcal{Q}^k(\mathbf{x}, 1)$  by  $\mathcal{Q}^k$  when there is no confusion. In inexact projected-Newton methods solve the subproblem (19) only approximately; we denote this approximate solution by  $\mathbf{z}^k$  below.

At each iteration of an SPG-based inexact projected-Newton method, we first compute the gradient  $\nabla f(\mathbf{x}^k)$  and (implicitly) compute the quadratic term  $\mathbf{H}^k$  in  $\mathcal{Q}^k$ . Subsequently, we try to minimize this  $\mathcal{Q}^k$  over the feasible set using iterations of an SPG algorithm. Even if  $f$  or  $\nabla f$  are difficult to compute, this SPG subroutine can be efficient if  $\mathcal{Q}^k$  and  $\nabla \mathcal{Q}^k$  can be evaluated rapidly. Given  $f(\mathbf{x}^k)$  and  $\nabla f(\mathbf{x}^k)$ , the dominant cost in evaluating  $\mathcal{Q}^k$  and  $\nabla \mathcal{Q}^k$  is pre-multiplication by  $\mathbf{H}^k$ . By taking the compact representation of Byrd et al. (1994)

$$\mathbf{H}^k = \sigma^k \mathbf{I} - \mathbf{N} \mathbf{M}^{-1} \mathbf{N}^T, \quad \text{where } \mathbf{N} \in \mathbb{R}^{n \times 2m}, \mathbf{M} \in \mathbb{R}^{2m \times 2m}, \quad (20)$$

we can compute  $\mathcal{Q}^k$  and  $\nabla \mathcal{Q}^k$  in  $\mathcal{O}(mn)$  under the L-BFGS Hessian approximation.

In addition to  $\mathcal{Q}^k$  and  $\nabla \mathcal{Q}^k$ , the SPG-subroutine also requires computing the Euclidean projection  $\mathcal{P}_\Omega$  onto the feasible set  $\Omega$ . However, note that the SPG subroutine does not evaluate  $f$  or  $\nabla f$ . Hence, the SPG-based inexact projected Newton method is most effective on problems where computing the projection is much less expensive than evaluating the objective function.<sup>7</sup>

Although in principle we could use SPG to exactly solve problem (19), in practice this is expensive and ultimately unnecessary. Thus, we terminate the SPG-subroutine before the exact solution is found. One might be concerned about terminating the SPG subroutine early, especially because an approximate solution to (3) will in general not be a descent direction. Fortunately, we can guarantee that the SPG subroutine yields a descent direction even under early termination if we initialize it with  $\mathbf{x}^k$  and we perform at least one SPG iteration. To see this, note that positive-definiteness of  $\mathbf{H}^k$  implies that a sufficient condition for  $\mathbf{z}^k - \mathbf{x}^k$  to be a descent direction for some vector  $\mathbf{z}^k$ , is that  $\mathcal{Q}^k(\mathbf{z}^k, \alpha^k) < f(\mathbf{x}^k)$ , since this implies the inequality

$$(\mathbf{z}^k - \mathbf{x}^k)^T \nabla f(\mathbf{x}^k) < 0.$$

By substituting  $\mathcal{Q}^k(\mathbf{x}^k, \alpha^k) = f(\mathbf{x}^k)$ , we see that

$$\mathcal{Q}^k(\mathbf{z}^k, \alpha^k) < \mathcal{Q}^k(\mathbf{x}^k, \alpha^k) = f(\mathbf{x}^k),$$

where  $\mathbf{z}^k$  is the first point satisfying the Armijo condition when we initialize

---

7. This is different than many classical optimization problems like quadratic programming, where evaluating the objective function may be relatively inexpensive but computing the projection may be as difficult as solving the original problem.

SPG with  $\mathbf{x}^k$ .<sup>8</sup> In other words, if we initialize the SPG subroutine with  $\mathbf{x}^k$ , then the SPG iterate gives a descent direction after the first iteration, and every subsequent iteration. Thus, it can be safely terminated early. In an implementation we can parametrize the maximum number of the SPG iterations by  $c$ , which results in an  $\mathcal{O}(mnc)$  iteration cost for the inexact Newton method, assuming that projection requires  $\mathcal{O}(n)$  time.

**Example 11.6** (Blockwise-Sparse Gaussian Graphical Models). *Consider a generalization of Example 11.5 where instead of constraining the absolute values of matrix elements, we constrain the norms of a disjoint set of groups (indexed by  $g$ ) of elements:*

$$\min_{\tilde{\Sigma} + \mathbf{X} \succ 0} \frac{1}{2} \log \det(\tilde{\Sigma} + \mathbf{X}), \quad \text{subject to } \|\mathbf{X}_g\|_2 \leq \lambda_g, \forall g. \quad (21)$$

*This generalization is similar to those examined in (Duchi et al., 2008; Schmidt et al., 2009), and it encourages the Gaussian graphical model to be sparse across groups of variables (i.e. all edges in a group  $g$  will either be included or excluded from the graph). Thus, formulation (21) encourages the precision-matrix to have a ‘blockwise’ sparsity pattern. Unfortunately, this generalization can no longer be written as a problem with bound constraints, nor can we characterize the feasible set with a finite number of linear constraints (though it is possible to write the feasible set using quadratic constraints). Nevertheless, it is easy to compute the projection onto the norm constraints; to project a matrix  $\mathbf{X}$  onto the feasible set with respect to the norm constraints we simply set  $\mathbf{X}_g = \lambda_g / \|\mathbf{X}_g\|_2$  for each group  $g$ . Considering the potentially high cost of evaluating the log-determinant function (and its derivative), this simple projection suggests inexact projected-Newton methods are well-suited for solving (21).*

## 11.5 Toward Nonsmooth Objectives

In this section we reconsider Problem (1), but unlike previous sections, we now allow  $r(\mathbf{x}) \neq 0$ . The resulting *composite optimization* problem occurs frequently in machine learning and statistics, especially with  $r(\mathbf{x})$  being a *sparsity promoting* regularizer—see e.g., Chapter 2.

How should we deal with the non-differentiability of  $r(\mathbf{x})$  in the context of Newton-like methods? While there are many possible answers to this

8. We will be able to satisfy the Armijo condition provided that  $\mathbf{x}^k$  is not already a minimizer.

question, we outline two simple but effective solutions that align well with the framework laid out so far.

### 11.5.1 Two-Metric Subgradient Projection Methods

We first consider the following special case of (1)

$$\min_{\mathbf{x} \in \mathbb{R}^n} \mathcal{F}(\mathbf{x}) = f(\mathbf{x}) + \sum_i r_i(x_i), \quad (22)$$

where  $r(\mathbf{x})$  has the separable form  $r(\mathbf{x}) = \sum_i r_i(x_i)$  and each  $r_i : \mathbb{R} \rightarrow \mathbb{R}$  is continuous and convex but not necessarily differentiable. A widely used instance of this problem is when we have  $r_i(x_i) = \lambda_i |x_i|$  for fixed  $\lambda_i > 0$ , corresponding to  $\ell_1$ -regularization. Note that this problem has a similar structure to the bound-constrained optimization problem (10); problem (10) has *separable constraints*, while problem (22) has a *separable nonsmooth* term. We can use separability of the nonsmooth term to derive a *two-metric subgradient projection* method for (22), analogous to the two-metric gradient projection method discussed in §11.3. The main idea is to choose an appropriately defined *steepest descent* direction and then to take a step resembling a two-metric projection iteration in this direction.

To define an appropriate steepest descent direction we note that even though the objective in (22) is not differentiable, its directional derivatives always exist. Thus, analogous to the differentiable case, we can define the steepest descent direction as the direction that minimizes the directional derivative; among all vectors with unit norm, the steepest descent direction locally decreases the objective most quickly. This direction is closely related to the element of the subdifferential of a function  $\mathcal{F}(\mathbf{x})$  with minimum norm:

**Definition 11.7** (Minimum-norm subgradient). *Let*

$$\mathbf{z}^k = \operatorname{argmin}_{\mathbf{z} \in \partial \mathcal{F}(\mathbf{x})} \|\mathbf{z}\|_2 \quad (23)$$

Following an argument outlined in (Bertsekas et al., 2003, §8.4),<sup>9</sup> the steepest descent direction for a convex function  $\mathcal{F}(\mathbf{x})$  at a point  $\mathbf{x}^k$  is in the direction  $-\mathbf{z}^k$ , where the subdifferential of (22) is given by

$$\partial \mathcal{F}(\mathbf{x}) = \partial \{f(\mathbf{x}^k) + r(\mathbf{x}^k)\} = \nabla f(\mathbf{x}^k) + \partial r(\mathbf{x}^k).$$

Using the separability of  $r(\mathbf{x})$ , we see that the minimum-norm subgradi-

---

9. Replacing maximization with minimization and concavity with convexity.



ent (23) with respect to a variable  $x_i$  is given by

$$z_i^k = \begin{cases} 0, & \text{if } -\nabla_i f(\mathbf{x}^k) \in \{\partial^- r_i(x_i^k), \partial^+ r_i(x_i^k)\}, \\ \min \{|\nabla_i f(\mathbf{x}^k) + \partial^- r_i(x_i^k)|, |\nabla_i f(\mathbf{x}^k) + \partial^+ r_i(x_i^k)|\}, & \text{otherwise,} \end{cases}$$

where the directional derivative  $\partial^+ r_i(x_i^k)$  is given by

$$\partial^+ r_i(x_i^k) = \lim_{\delta \rightarrow 0^+} \frac{r_i(x_i^k + \delta) - r_i(x_i^k)}{\delta},$$

the directional derivative  $\partial^- r_i(x_i^k)$  is defined similarly with  $\delta$  going to zero from below. Thus, it is easy to compute  $\mathbf{z}^k$  given  $\nabla f(\mathbf{x}^k)$ , as well as the left and right partial derivatives ( $\partial^- r_i(x_i^k)$  and  $\partial^+ r_i(x_i^k)$ ) for each  $r_i(x_i^k)$ . Observe that when  $r_i(x_i^k)$  is differentiable,  $\partial^- r_i(x_i^k) = \partial^+ r_i(x_i^k)$ , whereby the minimum norm subgradient is simply  $\nabla_i f(\mathbf{x}^k) + \nabla_i r(\mathbf{x}^k)$ . Further, note that  $\mathbf{z}^k = \mathbf{0}$  at a global optimum; otherwise  $-\mathbf{z}^k$  yields a descent direction and we can use it in place of the negative gradient within a line search method.

Similar to steepest-descent for smooth functions, a generalized steepest-descent for nonsmooth functions may converge slowly, and thus we seek a Newton-like variant. A natural question is whether we can merely use a scaling matrix  $\mathbf{S}^k$  to scale the steepest-descent direction. Similar to the two-metric projection algorithm, the answer is “no” for essentially the same reason: in general a scaled version of the steepest-descent direction may turn out to be an *ascent* direction.

However, we can still use a similar solution to the problem. If we make the positive-definite scaling matrix  $\mathbf{S}^k$  *diagonal* with respect to the variables  $x_i$  that are close to locations where  $r_i(x_i)$  is non-differentiable, then we can still ensure that the method generates descent directions. Thus, we obtain a simple Newton-like method for nonsmooth optimization that uses iterates of the form

$$\mathbf{x}^{k+1} \leftarrow \mathbf{x}^k - \alpha \mathbf{S}^k \mathbf{z}^k. \quad (24)$$

Here, matrix  $\mathbf{S}^k$  has the same structure as (12), but now the variables that receive a diagonal scaling are variables close to non-differentiable values. Formally, the set of *restricted* variables is:

$$\mathcal{J}^k \triangleq \{i \mid \min_{d_i \in \mathcal{D}_i} |d_i - x_i| \leq \varepsilon\}, \quad (25)$$

where  $\mathcal{D}_i$  is the (countable) set containing all locations where  $r_i(x_i)$  is non-differentiable.

In many applications where we seek to solve a problem of the form (22),

we expect the function to be non-differentiable with respect to several of the variables at a solution. Further, it may be desirable that intermediate iterations of the algorithm lie at non-differentiable points. For example, these might represent sparse solutions if a non-differentiability occurs at zero. In these cases, we can add a projection step to the iteration that encourages intermediate iterates to lie at points of non-differentiability. Specifically, if a variable  $x_i$  crosses a point of non-differentiability, we project onto the point of non-differentiability. Since we use a diagonal scaling with respect to the variables that are close to points of non-differentiability, this projection reduces to computing the Euclidean projection onto bound constraints, where the upper and lower bounds are given by the nearest upper and lower points of non-differentiability. Thus, each iteration is effectively a two-metric subgradient projection iteration. To make our description concrete, let us look at a specific example below.

**Example 11.8** ( $\ell_1$ -Regularization). *A prototypical composite minimization problem in machine learning is the  $\ell_1$ -regularized task*

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + \sum_{i=1}^n \lambda_i |x_i|. \quad (26)$$

The scalars  $\lambda_i \geq 0$  control the degree of regularization, and for sufficiently large  $\lambda_i$ , the parameter  $x_i$  is encouraged to be exactly zero.

To apply our framework, we need to efficiently compute the minimum-norm subgradient  $\mathbf{z}^k$  for (26); this gradient may be computed as

$$\mathbf{z}_i^k \triangleq \begin{cases} \nabla_i f(\mathbf{x}) + \lambda_i \operatorname{sgn}(x_i), & |x_i| > 0 \\ \nabla_i f(\mathbf{x}) + \lambda_i, & x_i = 0, \nabla_i f(\mathbf{x}) < -\lambda_i \\ \nabla_i f(\mathbf{x}) - \lambda_i, & x_i = 0, \nabla_i f(\mathbf{x}) > \lambda_i \\ 0, & x_i = 0, -\lambda_i \leq \nabla_i f(\mathbf{x}) \leq \lambda_i \end{cases} \quad (27)$$

For this problem, the restricted variable set (25) corresponds to those variables sufficiently close to zero,  $\{i \mid |x_i| \leq \varepsilon\}$ . Making  $\mathbf{S}^k$  partially diagonal with respect to the restricted variables as before, we define the two-metric projection step for  $\ell_1$ -regularized optimization as

$$\mathbf{x}^{k+1} = \mathcal{P}_{\circ}[\mathbf{x}^k - \alpha \mathbf{S}^k \mathbf{z}^k, \mathbf{x}^k]. \quad (28)$$

Here, the orthant projection (that sets variables to exactly zero) is defined as

$$\mathcal{P}_{\circ}(\mathbf{y}, \mathbf{x})_i \triangleq \begin{cases} 0, & \text{if } x_i y_i < 0, \\ y_i, & \text{otherwise.} \end{cases}$$

Applying this projection is effective at sparsifying the parameter vector since it sets variables that change sign to exactly zero, and it also ensures that the line search does not cross points of non-differentiability. Provided that  $\mathbf{x}^k$  is not stationary, the steps in (28) are guaranteed to improve the objective for sufficiently small  $\alpha$ . The step size  $\alpha$  is selected by a backtracking line search along the projection arc to satisfy a variant of the Armijo condition where the gradient is replaced by the minimum-norm subgradient. If at some iteration the algorithm identifies the correct set of non-zero variables and then maintains the orthant of the optimal solution, then the algorithm essentially reduces to an unconstrained Newton-like method applied to the non-zero variables.

In the two-metric projection algorithm for bound-constrained optimization the choice of the diagonal scaling matrix  $\mathbf{D}^k$  simply controls the rate that very small variables move towards zero, and does not have a significant impact on the performance of the algorithm. However, the choice of  $\mathbf{D}^k$  in the algorithm for  $\ell_1$ -regularization can have a significant effect on the performance of the method, since if  $\mathbf{D}^k$  is too large we may need to perform several backtracking steps before the step length is accepted, while too small of a value will require many iterations to set very small variables to exactly zero. One possibility is to compute the Barzilai-Borwein scaling  $\alpha^{bb}$  of the variables given by (18), and set  $\mathbf{D}^k$  to  $\alpha^{bb}\mathbf{I}$ .

### 11.5.2 Proximal Newton-like methods

The method of the previous section crucially relies on separability of the nonsmooth function  $r(\mathbf{x})$ . For more general nonsmooth  $r(\mathbf{x})$ , an attractive choice is to tackle the non-differentiability of  $r(\mathbf{x})$  via proximity operators (Moreau, 1962; Combettes and Wajs, 2005; Combettes and Pesquet, 2009). These operators are central to forward-backward splitting methods (Combettes and Pesquet, 2009)<sup>10</sup>, as well as to methods based on surrogate optimization (Figueiredo and Nowak, 2003; Daubechies et al., 2004), separable approximation (Wright et al., 2009), gradient-mapping (Nesterov, 2007), or in a proximal trust-region framework (Kim et al., 2010).

The idea of a proximity operator is simple. Let  $r : X \subseteq \mathbb{R}^d \rightarrow (-\infty, \infty]$  be a lower semicontinuous, proper convex function. For a point  $\mathbf{y} \in X$ , the *proximity-operator* for  $r$  applied to  $\mathbf{y}$  is defined as

$$\text{prox}_r(\mathbf{y}) = \underset{\mathbf{x} \in X}{\text{argmin}} \quad \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_2^2 + r(\mathbf{x}). \quad (29)$$

10. Including *iterative shrinkage-thresholding* as a special case

This operator generalizes the projection operator, since when  $r(\mathbf{x})$  is the indicator function for a convex set  $C$ , (29) reduces to projection onto  $C$ . This observation suggests that we might be able to replace projection operators by proximity operators. Indeed, this replacement is done in *forward-backward splitting* methods, where one iterates

$$\mathbf{x}^{k+1} = \text{prox}_{\alpha^k r}(\mathbf{x}^k - \alpha^k \nabla f(\mathbf{x}^k));$$

the iteration “splits” the update into differentiable (forward) and non-differentiable (proximal or backward) steps. This method generalizes first-order projected gradient methods, and under appropriate assumptions it can be shown that the sequence  $\{f(\mathbf{x}^k) + r(\mathbf{x}^k)\}$  converges to  $f(\mathbf{x}^*) + r(\mathbf{x}^*)$ , where  $\mathbf{x}^*$  is a stationary point.

At this point, the reader may already suspect how we might use proximity operators in our Newton-like methods. The key idea is simple: build a quadratic model, but *only* for the differentiable part, and tackle the non-differentiable part via a suitable proximity operator. This simple idea was also previously exploited by (Wright et al., 2009; Kim et al., 2010). Formally, we consider the regularized quadratic model

$$\mathcal{Q}^k(\mathbf{x}, \alpha) \triangleq f(\mathbf{x}^k) + (\mathbf{x} - \mathbf{x}^k)^T \nabla f(\mathbf{x}^k) + \frac{1}{2\alpha} (\mathbf{x} - \mathbf{x}^k)^T \mathbf{H}^k (\mathbf{x} - \mathbf{x}^k) + r(\mathbf{x}), \quad (30)$$

whose minimizer can be recast as the *generalized proximity-operator*:

$$\text{prox}_{\alpha^k r}^{\mathbf{H}^k}(\mathbf{y}^k) = \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \quad \frac{1}{2} \|\mathbf{x} - \mathbf{y}^k\|_{\mathbf{H}^k}^2 + \alpha r(\mathbf{x}), \quad (31)$$

where  $\mathbf{y}^k = \mathbf{x}^k - \alpha[\mathbf{H}^k]^{-1} \nabla f(\mathbf{x}^k)$ ; observe that under the transformation  $\mathbf{x} \rightarrow [\mathbf{H}^k]^{1/2} \mathbf{x}$ , (31) may be viewed as a standard proximity operator.

Using this generalized proximity operator, our Newton-like algorithm becomes

$$\mathbf{x}^{k+1} = \text{prox}_{\alpha^k r}^{\mathbf{H}^k}(\mathbf{x}^k - \alpha[\mathbf{H}^k]^{-1} \nabla f(\mathbf{x}^k)). \quad (32)$$

If instead of the true inverse Hessian, we use  $\mathbf{H}^k = \mathbf{I}$ , iteration (32) degenerates to the traditional forward-backward splitting algorithm. Furthermore, it is equally straightforward to implement a quasi-Newton variant of (32), where for example  $\mathbf{H}^k$  is obtained by an L-BFGS approximation to  $\nabla^2 f(\mathbf{x}^k)$ . Another practical choice might be an inexact quasi-Newton variant, where we use iterations of the SPG-like method of (Wright et al., 2009) to approximately minimize  $\mathcal{Q}^k(\mathbf{x}, \alpha)$  under an L-BFGS approximation of  $f(\mathbf{x})$ ; in other words, the generalized proximity operator (31) is computed inexactly.

Similar to inexact projected Newton methods, under only mild assump-

tions we can guarantee that an inexact solution to the generalized proximity operator yields an improvement on the original objective for a sufficiently small step size  $\alpha$ . For example, assume that  $\nabla f(\mathbf{x})$  is Lipschitz-continuous and that we find a value  $\mathbf{y}$  such that  $\mathcal{Q}^k(\mathbf{y}, \alpha) < \mathcal{Q}^k(\mathbf{x}^k, \alpha)$  in (30). Then we have

$$\begin{aligned} f(\mathbf{x}^k) + r(\mathbf{x}^k) &= \mathcal{Q}^k(\mathbf{x}^k, \alpha) \\ &> \mathcal{Q}^k(\mathbf{y}, \alpha) \\ &= f(\mathbf{x}^k) + (\mathbf{y} - \mathbf{x}^k)^T \nabla f(\mathbf{x}^k) + \frac{1}{2\alpha} (\mathbf{y} - \mathbf{x}^k)^T \mathbf{H}^k (\mathbf{y} - \mathbf{x}^k) + r(\mathbf{y}) \\ &\geq f(\mathbf{x}^k) + (\mathbf{y} - \mathbf{x}^k)^T \nabla f(\mathbf{x}^k) + \frac{m}{2\alpha} \|\mathbf{y} - \mathbf{x}^k\|_2^2 + r(\mathbf{y}) \\ &\geq f(\mathbf{y}) + r(\mathbf{y}) \quad (\text{for } 0 < \alpha \leq m/\mathcal{L}), \end{aligned}$$

where  $m$  is the smallest eigenvalue of  $\mathbf{H}^k$  and the last inequality follows from Lipschitz-continuity of the gradient (Bertsekas, 1999, Proposition A.24), where  $\mathcal{L}$  is the Lipschitz constant of the gradient of  $f(\mathbf{x})$ . A similar property holds if  $\nabla f(\mathbf{x})$  is only locally Lipschitz-continuous.

**Example 11.9** (Group  $\ell_1$ -Regularization). *Consider a generalization of Example 11.8 where instead of penalizing the absolute values of each element of  $\mathbf{x}$ , we penalize the  $\ell_2$  norms of a set of disjoint groups indexed by  $g$ :*

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + \sum_g \lambda_g \|\mathbf{x}_g\|_2. \quad (33)$$

The regularizer in (33) is referred to as a group regularizer ( $\ell_{1,2}$ -regularizer), since it encourages sparsity in terms of groups of variables, and dates back to (Bakin, 1999). The regularization term is nonsmooth when an entire group of variables is set to  $\mathbf{0}$ . However, the proximal operator for this regularizer is easily computed: given a vector  $\mathbf{y}$ , with groups  $\mathbf{y}_g$ , we simply set  $\mathbf{x}_g = (\mathbf{y}_g / \|\mathbf{y}_g\|_2) \max\{0, \|\mathbf{y}_g\|_2 - \alpha \lambda_g\}$ . Thus, inexact proximal-Newton methods are well-suited to solving (33).

**Example 11.10** (Group Nuclear-Norm Regularization). *A related problem is optimizing a smooth function of several matrix inputs with regularization of the nuclear norms of the matrices:*

$$\min_{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n} f(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n) + \sum_{i=1}^n \lambda_i \|\mathbf{X}_i\|_*. \quad (34)$$

Here we use  $\|\mathbf{X}\|_*$  to denote the nuclear norm (or trace norm), the sum of the singular values of  $\mathbf{X}$ . This regularization not only encourages sparsity across individual matrices, but also encourages each matrix to be low-rank. The proximal operator for the nuclear norm can be computed by

soft-thresholding the singular values of each  $\mathbf{X}_i$  (Cai et al., 2010). That is, to compute the proximal operator we replace each singular value  $\sigma_j$  of each  $\mathbf{X}_i$  by  $\sigma_j = \max\{0, \sigma_j - \alpha\lambda_i\}$ , where  $\alpha$  is the parameter of the quadratic approximation (30). Thus, inexact proximal-Newton methods are well-suited to solving (34) too, especially when it is more expensive to evaluate  $f$  and  $\nabla f$  than it is to compute the singular value decomposition of each  $\mathbf{X}_i$ .

---

## 11.6 Summary and Discussion

In this chapter, we have concentrated on minimizing twice-differentiable convex functions, both when their exact Hessian is feasible to use, as well as when quasi-Newton choices are more practical. Note that the quasi-Newton approach can also be applied when the objective function is only once-differentiable. Furthermore, we may relax the assumption of convexity if we concede that the stationary point found by the method may not be a local or global minimum.

As for implementational strategies, while we have focused on L-BFGS methods, an alternative restricted-memory strategy is to use implicit Hessian-vector products. For example, in the two-metric projection strategy we can use Hessian-vector products within a linear conjugate gradient iteration to solve the scaling with respect to the free variables as in (Nocedal and Wright, 2000, §7.1), while we can use Hessian-vector products within the SPG subroutine for inexact projected Newton methods. An alternative means to optimize nonsmooth objectives with an L-BFGS approximation is given by (Yu et al., 2010). Variants of the L-BFGS approximation that apply in stochastic scenarios are examined in (Sunehag et al., 2009).

We close by noting some open issues regarding convergence of the methods discussed in this section. First, we note that global convergence of methods based on the minimum-norm subgradient without a diminishing step size can be tenuous because of the lack of continuity in the derivatives of sequences. For example, see the counter-example in (Bertsekas, 1999, Exercise 6.3.8). Andrew and Gao (2007) give a proof of global convergence of a method related to the two-metric subgradient projection method we discuss in Section 11.5.1, but as pointed out by (Yu et al., 2010) their proof does not account for this lack of continuity. Thus, while the algorithms of Andrew and Gao (2007) and Section 11.5.1 appear to be very effective in practice, it remains to show whether they are globally convergent in general without additional assumptions. A related issue is showing whether the two-metric subgradient projection method identifies the correct set of non-zero variables after a finite number of iterations, and then has a super-linear convergence

rate when exact second-order information is available. Also, in Examples 1.5 and 1.6, we use a projection with respect to a subset of the constraints and do not project with respect to the positive-definite constraint that is known to not be active at the solution. Although this strategy has been used by several authors, and seems to not significantly affect empirical convergence of the method when given a suitable starting point, formally examining convergence under this heuristic deserves some theoretical attention.

Finally, we remark that there are not yet formal proofs of global and local convergence for inexact projected Newton methods, but this appears to be a simpler task than showing convergence of the methods discussed in the previous paragraph. For example, it is likely that global convergence can be proved by showing that a suitable gradient-related condition (Bertsekas, 1999, §1.2) applies to the first iteration in the SPG subroutine that satisfies the Armijo condition, while a local convergence rate can likely be shown by using a forcing sequence (Nocedal and Wright, 2000, §7.1) on the solution accuracy of the SPG subroutine.

---

## References

- G. Andrew and J. Gao. Scalable training of  $L_1$ -regularized log-linear models. *Proc. Intl. Conf. Machine Learning*, 2007.
- S. Bakin. *Adaptive regression and model selection in data mining problems*. PhD thesis, Australian National University, Canberra, 1999.
- O. Banerjee, L. E. Ghaoui, A. d’Aspremont, and G. Natsoulis. Convex optimization techniques for fitting sparse gaussian graphical models. *Proc. Intl. Conf. Machine Learning*, 2006.
- J. Barzilai and J. Borwein. Two-point step size gradient methods. *IMA Journal of Numerical Analysis*, 8(1):141–148, 1988.
- D. P. Bertsekas. Projected Newton Methods for Optimization Problems with Simple Constraints. *SIAM Journal on Control and Optimization*, 20(2):221–246, 1982.
- D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 1999.
- D. P. Bertsekas, A. Nedic, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.
- E. G. Birgin, J. M. Martínez, and M. Raydan. Nonmonotone Spectral Projected Gradient Methods on Convex Sets. *SIAM Journal on Optimization*, 10(4):1196–1211, 2000.
- R. H. Byrd, J. Nocedal, and R. B. Schnabel. Representations of quasi-Newton matrices and their use in limited memory methods. *Mathematical Programming*, 63(2):129–156, 1994.
- J. F. Cai, E. J. Candés, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
- P. L. Combettes and J. Pesquet. Proximal Splitting Methods in Signal Processing. *arXiv:0912.3522v2*, Dec. 2009.

- P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling and Simulation*, 4(4):1168–1200, 2005.
- Y. H. Dai and R. Fletcher. Projected Barzilai-Borwein methods for large-scale box-constrained quadratic programming. *Numerische Mathematik*, 100(1):21–47, 2005.
- I. Daubechies, M. Defrise, and C. D. Mol. An Iterative Thresholding Algorithm for Linear Inverse Problems with a Sparsity Constraint. *Communications on Pure and Applied Mathematics*, 11:1413–1457, 2004.
- J. Duchi, S. Gould, and D. Koller. Projected subgradient methods for learning sparse gaussians. *Conference on Uncertainty in Artificial Intelligence*, 2008.
- M. Figueiredo, R. Nowak, and S. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):586–597, 2007.
- M. A. T. Figueiredo and R. D. Nowak. An EM algorithm for wavelet-based image restoration. *IEEE Transactions on Image Processing*, 12(8):906–916, 2003.
- E. M. Gafni and D. P. Bertsekas. Two-metric projection methods for constrained optimization. *SIAM Journal on Control and Optimization*, 22(6):936–964, 1984.
- P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, 1981.
- L. Grippo, F. Lampariello, and S. Lucidi. A nonmonotone line search technique for Newton’s method. *SIAM Journal on Numerical Analysis*, 23(4):707–716, 1986.
- D. Kim, S. Sra, and I. S. Dhillon. A scalable trust-region algorithm with application to mixed-norm regression. In *Int. Conf. Machine Learning (ICML)*, 2010.
- E. S. Levitin and B. T. Polyak. Constrained minimization methods. *USSR Computational mathematics and mathematical physics*, 6:1–50, 1966. English translation of a paper in *Zh. Vychisl. Mat. i Mat. Fiz.* 6, 5, 787–823, 1966.
- J. J. Moreau. Fonctions convexes duales et points proximaux dans un espace hilbertien. *R. Acad. Sci. Paris Sér. A Math*, 255:2897–2899, 1962.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2004.
- Y. Nesterov. Gradient methods for minimizing composite objective function. Technical report, Université Catholique do Louvain, 2007.
- J. Nocedal. Updating Quasi-Newton Matrices with Limited Storage. *Mathematics of Computation*, 95:339–353, 1980.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, second edition, 2000.
- M. Schmidt, E. van den Berg, M. Friedlander, and K. Murphy. Optimizing costly functions with simple constraints: A limited-memory projected quasi-newton algorithm. *Proceedings of International Workshop on Artificial Intelligence and Statistics*, 2009.
- D. F. Shanno and K. H. Phua. Matrix conditioning and nonlinear optimization. *Mathematical Programming*, 14(1):149–160, 1978.
- P. Sunehag, J. Trunpf, S. Vishwanathan, and N. Schraudolph. Variable metric stochastic approximation theory. *Proceedings of International Workshop on Artificial Intelligence and Statistics*, 2009.
- E. van den Berg and M. P. Friedlander. Probing the Pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.



- S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo. Sparse reconstruction by separable approximation. *IEEE Transactions on Signal Processing*, 57(7):2479–2493, 2009.
- J. Yu, S. Vischwanathan, S. Günter, and N. Schraudolph. A Quasi-Newton Approach to Nonsmooth Convex Optimization. *Journal of Machine Learning Research*, 11:1201–1242, 2010.