Proximal and Projected Newton Methods

Ryan Tibshirani
Convex Optimization 10-725/36-725

Last time: conditional gradient method

For the problem

$$\min_x \ f(x) \ \text{ subject to } \ x \in C$$

where f is convex, smooth and C is a convex set, the conditional gradient (Frank-Wolfe) method chooses an initial $x^{(0)}$ and repeats for $k = 1, 2, 3, \ldots$

$$s^{(k-1)} \in \operatorname*{argmin}_{s \in C} \ \nabla f(x^{(k-1)})^T s$$
$$x^{(k)} = (1 - \gamma_k) x^{(k-1)} + \gamma_k s^{(k-1)}$$

Here γ_k is a step size, either prespecified (as in $\gamma_k = 2/(k+1)$) or chosen by line search

For many problems, linear minimization over C is simpler or more efficient than projection onto C, hence the appeal of Frank-Wolfe

Proximal gradient descent

Recall that proximal gradient descent operates on a problem

$$\min_x \; g(x) + h(x)$$

where g is convex, smooth and h is convex, "simple". We choose initial $x^{(0)}$ and repeat for $k = 1, 2, 3, \ldots$

$$x^{(k)} = \mathrm{prox}_{t_k}\left(x^{(k-1)} - t_k \nabla g(x^{(k-1)})\right)$$

where $\mathrm{prox}_t(\cdot)$ is the proximal operator associated with h,

$$\mathrm{prox}_t(x) = \underset{z}{\mathrm{argmin}} \; \frac{1}{2t}\|x - z\|_2^2 + h(z)$$

- Difficulty of iterations is in applying prox, which only depends on h (assuming that ∇g is computable)
- Proximal gradient descent enjoys same convergence rate as its fully smooth version, hence useful when prox is efficient

Recall the motivation for proximal gradient: iteratively minimize a quadratic expansion in g, plus original h

$$\begin{aligned} x^+ &= \operatorname*{argmin}_z \ \frac{1}{2t}\|x - t\nabla g(x) - z\|_2^2 + h(z) \\ &= \operatorname*{argmin}_z \ \nabla g(x)^T(z-x) + \frac{1}{2t}\|z-x\|_2^2 + h(z) \end{aligned}$$

The quadratic approximation here uses Hessian equal to (a scaled version of) the identity $\frac{1}{t}I$

A fundamental difference between gradient descent and Newton's method was that the latter also iteratively minimized quadratic approximations, but these used the local Hessian of the function in question

So what happens if we replace $\frac{1}{t}I$ in the above with $\nabla^2 g(x)$?

Proximal Newton method

This leads us to the proximal Newton method. Now we must define

$$\text{prox}_H(x) = \underset{z}{\text{argmin}} \ \frac{1}{2}\|x - z\|_H^2 + h(z)$$

where $\|x\|_H^2 = x^T H z$ defines a norm, given a matrix $H \succ 0$. This is a scaled proximal mapping. With $H = \frac{1}{t} I$, we get back previous definition

Starting with $x^{(0)}$, we repeat for $k = 1, 2, 3, \ldots$

$$y^{(k)} = \text{prox}_{H_{k-1}}\left(x^{(k-1)} - H_{k-1}^{-1} \nabla g(x^{(k-1)})\right)$$
$$x^{(k)} = x^{(k-1)} + t_k(y^{(k-1)} - x^{(k-1)})$$

Here $H_{k-1} = \nabla^2 g(x^{(k-1)})$, and t_k is a step size, which we choose by backtracking line search (as in usual Newton)

Let's check this is indeed minimizing a quadratic approximation of g, plus h:

$$\begin{aligned} y &= \operatorname*{argmin}_z \ \frac{1}{2}\|x - H^{-1}\nabla g(x) - z\|_H^2 + h(z) \\ &= \operatorname*{argmin}_z \ \nabla g(x)^T(z-x) + \frac{1}{2}(z-x)^T H(z-x) + h(z) \end{aligned}$$

Notes:

- When $h(z) = 0$, we get back the usual Newton update
- For $H \succ 0$, can check that $\mathrm{prox}_H(\cdot)$ retains many of the nice properties of (unscaled) proximal mappings (Lee et al. 2014). E.g., it is well-defined, since the minimizer is unique
- Difficulty of prox has mostly to do with h, however, now the Hessian of g also plays a role—the structure of this Hessian H can make a difference

Backtracking line search

As with Newton's method in fully smooth problems, pure step sizes $t_k = 1$, $k = 1, 2, 3, \ldots$ need not converge. We need to apply, say, backtracking line search. Set parameters $0 < \alpha \leq 1/2$, $0 < \beta < 1$, and let

$$v = \mathrm{prox}_H\big(x - H^{-1}\nabla g(x)\big) - x$$

be the proximal Newton direction at a given iteration. Start with $t = 1$, and while

$$f(x + tv) > f(x) + \alpha t \nabla g(x)^T v + \alpha\big(h(x + td) - h(x)\big)$$

we shrink $t = \beta t$. (Here $f = g + h$)

Note: this scheme is actually of a different spirit than the one we studied for proximal gradient descent, as it avoids recomputing the prox at each inner backtracking iteration

Wait ... does this even make sense?

Let's back up. One of the main drivers behind proximal gradient descent is that we can transform the problem

$$\min_x \; g(x) + h(x)$$

into a sequence of problems where $g(x)$ is essentially replaced by $\|b - x\|_2^2$. This can be easy, but it depends on h

Now we have transformed into a sequence of problems where $g(x)$ is essentially replaced by $b^T x + x^T A x$. For dense A, this seems like it would rarely be easy, regardless of h ... That is, evaluating the scaled prox

$$\operatorname*{argmin}_z \; \nabla g(x)^T (z - x) + \frac{1}{2}(z - x)^T H (z - x) + h(z)$$

seems to be not an easy subproblem, for a generic Hessian H ...

All this is true, and the prox operator in proximal Newton is usually extremely expensive, and one that we solve with an optimization subroutine (e.g., for $h(x) = \|x\|_1$, prox is standard lasso problem)

What we should hope for: the convergence rate of prox Newton, in terms of the number of iterations (prox evaluations) needed, is like the usual Newton method. This ends up being true

Therefore, if we have a decent inner solver for the prox step, it can be quite efficient to use proximal Newton (e.g., this is true with ℓ_1 regularized generalized linear models). But in general, prox Newton is not to be applied without care

(Well-known implementations using prox Newton: glmnet, QUIC; more on this later)

Convergence analysis

Following Lee et al. (2014), assume that $f = g + h$, where g, h are convex and g is twice smooth. Assume further:

- $mI \preceq \nabla g \preceq LI$, and $\nabla^2 g$ Lipschitz with parameter M
- $\mathrm{prox}_H(\cdot)$ is exactly evaluable

Theorem: Proximal Newton method with backtracking line search satisfies converges globally. Furthermore, for all $k \geq k_0$:

$$\|x^{(k)} - x^\star\|_2 \leq \frac{M}{2m}\|x^{(k-1)} - x^\star\|_2^2$$

Recall that this is called local quadratic convergence. After some point, to get within $f(x^{(k)}) - f^\star \leq \epsilon$, we require $O(\log\log(1/\epsilon))$ iterations. Note: each iteration uses scaled prox evaluation!

Proof sketch

- To prove global convergence, they show that at any step, the backtracking exit condition will be satisfied by

$$t \leq \min\left\{1, \frac{2m}{L}(1-\alpha)\right\}$$

  Use this to show that the update direction converges to zero, which can only happen at the global minimum

- To prove local quadratic convergence, they show that for large enough k, the pure step $t = 1$ eventually satisfies backtracking exit condition. Therefore

$$\begin{aligned}
\|x^+ - x^\star\|_2 &\leq \frac{1}{\sqrt{m}}\|x^+ - x^\star\|_H \\
&\leq \left\|\mathrm{prox}_H\left(x - H^{-1}\nabla g(x)\right) - \mathrm{prox}_H\left(x^\star - H^{-1}\nabla g(x^\star)\right)\right\|_H \\
&\leq \frac{M}{2m}\|x - x^\star\|_2^2
\end{aligned}$$

Glmnet and QUIC

Two notable examples of proximal Newton methods:

- glmnet (Friedman et al. 2009): applies proximal Newton to ℓ_1 regularized generalized linear models, inner probs solved using coordinate descent
- QUIC (Hsiesh et al. 2011): applies proximal Newton to solve graphical lasso problem, uses factorization tricks, inner probs use coordinate descent

Both of these implementations are very widely used for their own purposes. At the proper scale, these are close to state-of-the-art

General note: proximal Newton method will use far less evaluations of (gradient of) g than proximal gradient. When these evaluations are expensive, proximal Newton can win

Example: lasso logistic regression

Example from Lee et al. (2014): ℓ_1 regularized logistic regression, FISTA (accelerated prox grad) versus spaRSA (spectral projected gradient method) versus PN (prox Newton)

Problem with $n = 5000$, $p = 6000$, and a dense feature matrix X

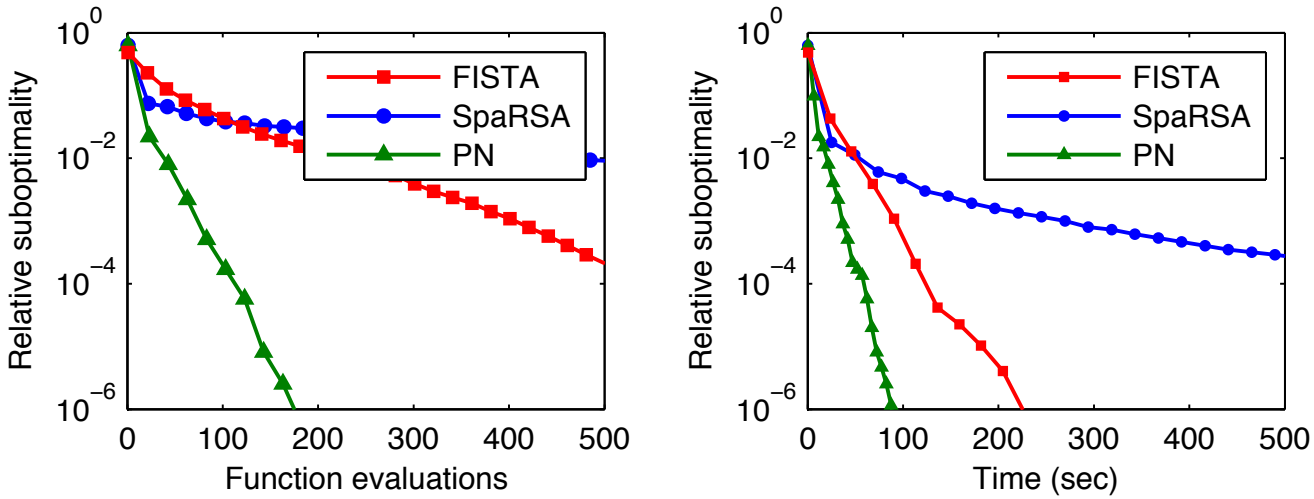

Here g and ∇g require expensive $\exp$ or $\log$ evaluations; dominates computational cost

Now problem with $n = 542,000$, $p = 47,000$, and sparse matrix X

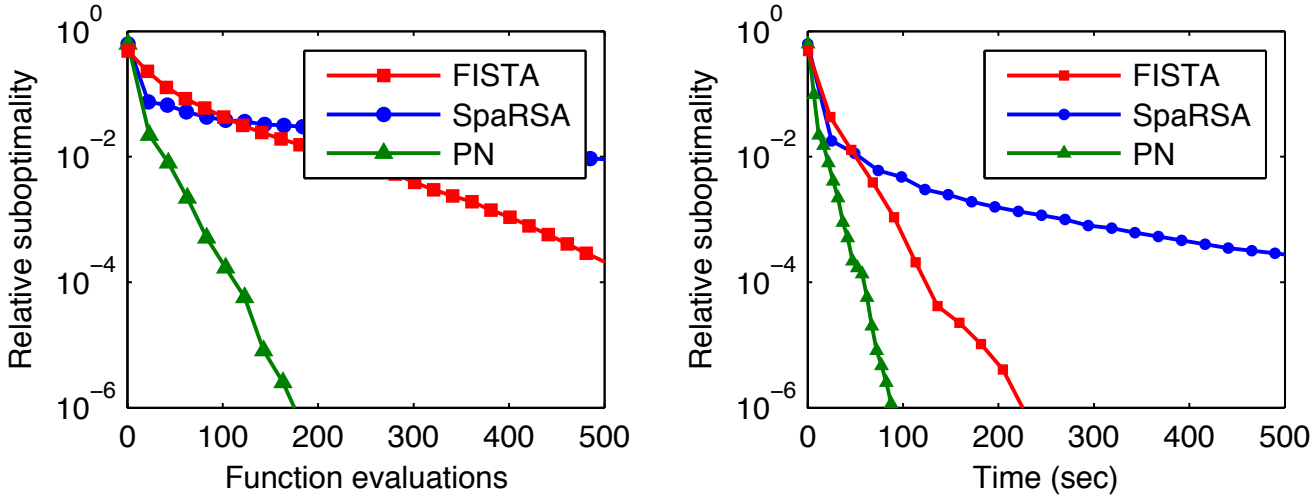

Here g and ∇g require expensive $\exp$ or $\log$ evaluations, but these make up less of total cost, since X is sparse

Inexact prox evaluations

An important note: with proximal Newton, we essentially always perform inexact prox evaluations (not so with proximal gradient)

Example from Lee et al. (2014): graphical lasso estimation, three stopping rules for inner optimizations. Here $n = 72$ and $p = 1255$

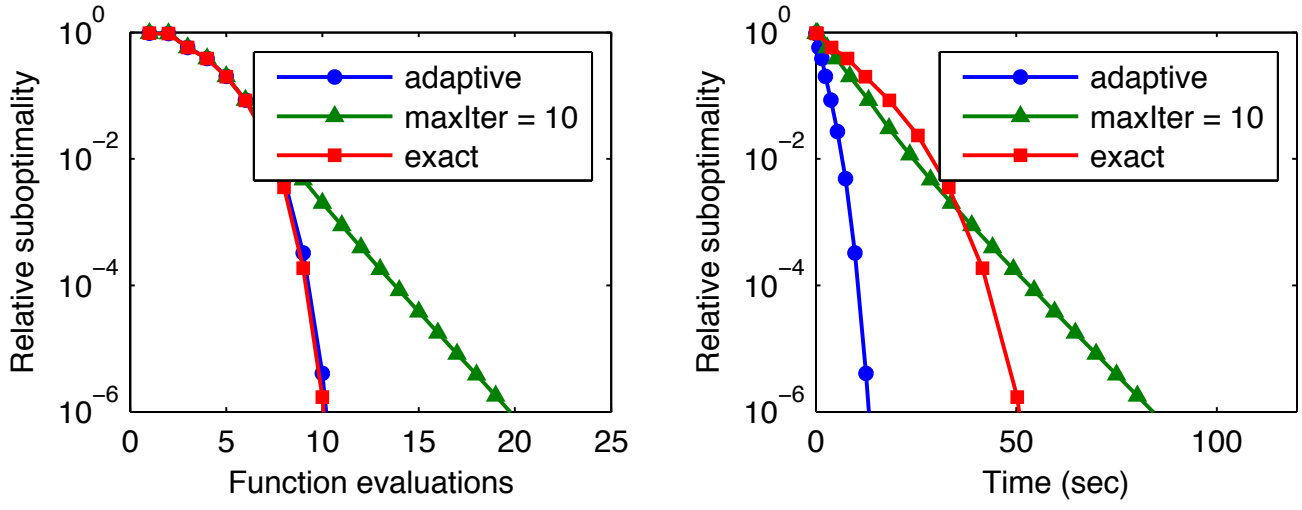

Conclusion is that 10 inner iterations is not enough to ensure fast (quadratic convergence), but their adaptive stopping rule is

For usual (smooth) Newton method, inner problem is to minimize $\tilde{g}_{k-1}(z)$ quadratic approximation to g about $x^{(k-1)}$. Stopping rules based on

$$\|\nabla \tilde{g}_{k-1}(z)\|_2 \leq \eta_k \|\nabla g(x^{(k-1)})\|_2$$

for a specifically chosen "forcing" sequence η_k, $k = 1, 2, 3, \ldots$

For proximal Newton, Lee et al. (2014) advocate the analogy that uses generalized gradients in place of gradients

$$\|G_{\tilde{f}_{k-1}/M}(z)\|_2 \leq \eta_k \|G_{f/M}(x^{(k-1)})\|_2$$

where $\tilde{f}_{k-1} = \tilde{g}_{k-1} + h$, and recall that $m \preceq \nabla^2 g \preceq MI$. Setting

$$\eta_k = \left\{ \frac{m}{2}, \frac{\|G_{\tilde{f}_{k-2}/M}(x^{(k-1)}) - G_{f/M}(x^{(k-1)})\|_2}{\|G_{f/M}(x^{(k-2)})\|_2} \right\}$$

they prove that inexact proximal Newton has local superlinear rate

Proximal quasi-Newton methods

For large problems, computing the Hessian is prohibitive. Proximal quasi-Newton methods avoid exactly forming $H_{k-1} = \nabla g(x^{(k-1)})$ at each step

- Lee et al. (2014) propose iteratively updating H_{k-1} at each step using BFGS-type rules. They show very strong empirical performance, and prove local superlinear convergence
- Tseng and Yun (2010) consider smooth plus block separable problems, and recommend approximating the Hessian in a blockwise fashion, combined with block coordinate descent. This can be very helpful because only small Hessians are ever needed. They prove linear convergence

Note that quasi-Newton methods can not only be helpful when the Hessian is expensive, but also when it is ill-conditioned: singular or close to singular

What's wrong with projected Newton?

Suppose that $h = 1_C(x)$, the indicator function of a convex set C. I.e., consider the problem

$$\min_x \; g(x) \;\; \text{subject to} \;\; C$$

Recall that proximal gradient here reduces to projected gradient. What about proximal Newton? Updates are

$$\begin{aligned} y &= \operatorname*{argmin}_{z \in C} \; \frac{1}{2}\|x - H^{-1}\nabla g(x) - z\|_H^2 \\ &= \operatorname*{argmin}_{z \in C} \; \nabla g(x)^T(z-x) + \frac{1}{2}(z-x)^T H(z-x) \end{aligned}$$

Note when $H = I$ this a projection of $x - \nabla g(x)$ onto C, but this is not a projection in general! In fact, it is much more complicated. Hence, projected Newton does not generally follow from proximal Newton

Projected Newton for box constraints

For some particular constraint sets, projected Newton can be made to work. Box constraints are one such example (Kim et al. 2010, Schmidt et al. 2011). Given a problem

$$\min_x \; g(x) \;\; \text{subject to} \;\; l \leq x \leq u$$

the projected Newton method specifies an initial point $x^{(0)}$, small constant $\epsilon > 0$, and repeats the following steps for $k = 1, 2, 3, \ldots$:

- Define the binding set

$$B_{k-1} = \{i : x_i^{(k-1)} \leq l_i + \epsilon \text{ and } \nabla_i g(x^{(k-1)}) > 0\} \cup \{i : x_i^{(k-1)} \geq u_i - \epsilon \text{ and } \nabla_i g(x^{(k-1)}) < 0\}$$

  These are the variables that are at (close to) boundary, and moving them any further would decrease the criterion. Also define the free set $F_{k-1} = \{1, \ldots n\} \setminus B_{k-1}$

- Define the principal submatrix of the inverse Hessian along the free variables

$$S_{k-1} = \left[\left(\nabla^2 g(x^{(k-1)})\right)^{-1}\right]_{F_{k-1}}$$

- Take a projected Newton step along the free variables only

$$x^{(k)} = P_{[l,u]}\left(x^{(k-1)} - t_k \begin{bmatrix} S_{k-1} & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \nabla_{F_{k-1}} g(x^{(k-1)}) \\ \nabla_{B_{k-1}} g(x^{(k-1)}) \end{bmatrix}\right)$$

  where $P_{[l,u]}$ is the projection onto $[l,u] = [l_1,u_1] \times \ldots [l_n,u_n]$. (Note that the binding set is not really touched)

Projected quasi-Newton is also an option, replacing $\nabla^2 g(x^{(k-1)})$ in the definition of S_{k-1} with an iterated approximation, e.g., using BFGS-style updates

Common box constraints

What kinds of problems have box constraints? Lots!

- Nonnegative least squares
- Nonnegative Kullback-Leibler (KL) divergence minimization:

$$\begin{array}{ll} \min\limits_{u} & \displaystyle\sum_{i=1}^{n} \Big(- y_i \log \frac{(Ku)_i}{y_i} + (Ku)_i \Big) + \lambda R(u) \\ \text{subject to} & u \geq 0 \end{array}$$

- Support vector machine dual
- Dual of ℓ_1 penalized problems, e.g., generalized lasso dual, graphical lasso dual

Example from Kim et al. (2010): image deblurring performed with nonnegative KL divergence minimization

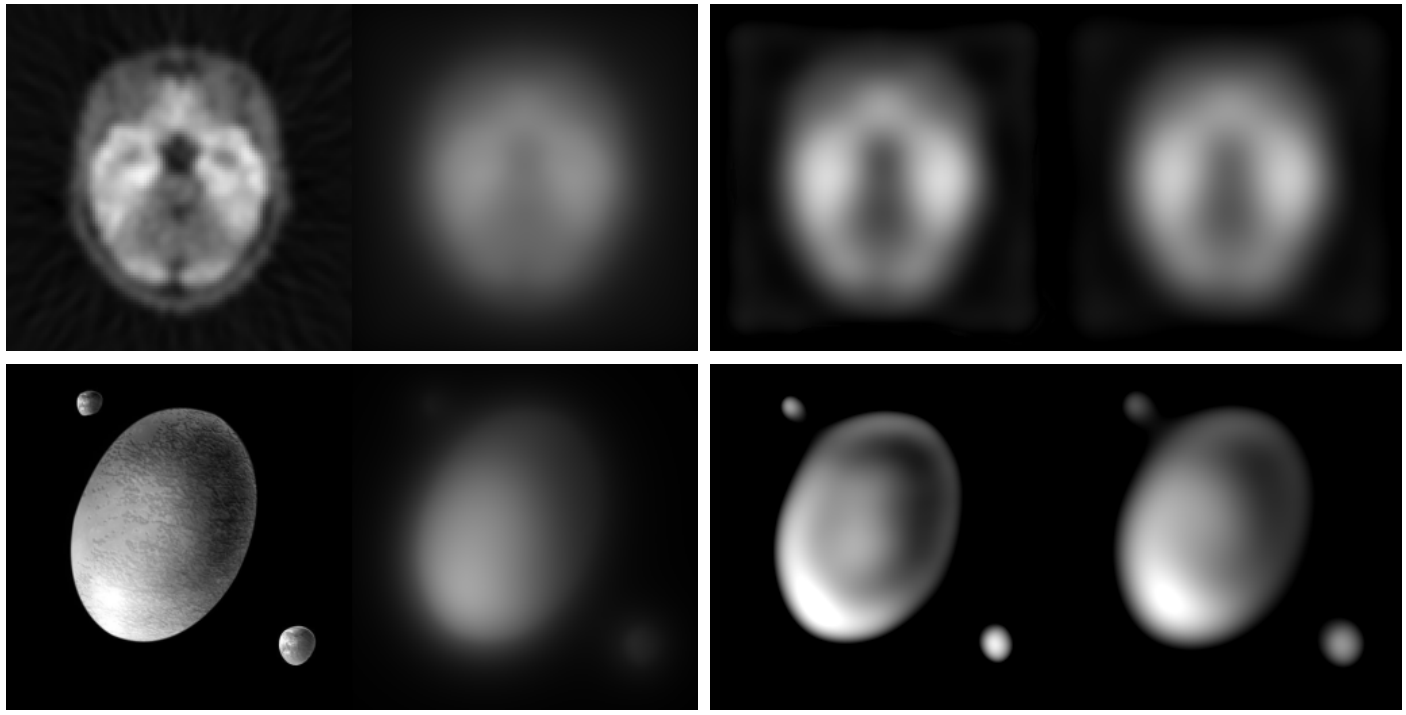

Original image Blurred image PQN-LBFGS LBFGS-B

References

Proximal Newton method:

- J. Friedman and T. Hastie and R. Tibshirani (2009), "Regularization paths for generalized linear models via coordinate descent"
- C.J. Hsiesh and M.A. Sustik and I. Dhillon and P. Ravikumar (2011), "Sparse inverse covariance matrix estimation using quadratic approximation"
- M. Patriksson (1998), "Cost approximation: a unified framework of descent algorithms for nonlinear programs"
- J. Lee and Y. Sun and M. Saunders (2014), "Proximal Newton-type methods for minimizing composite functions"
- P. Tseng and S. Yun (2009), "A coordinate gradient descent method for nonsmooth separable minimization"

Projected Newton method:

- A. Barbero and S. Sra (2011), "Fast Newton-type methods for total variation regularization"
- D. Bertsekas (1982), "Projected Newton methods for optimization problems with simple constraints"
- D. Kim and S. Sra. and I. Dhillon (2010), "Tackling box-constrained optimization via a new projected quasi-Newton approach"
- M. Schmidt and D. Kim and S. Sra (2011), "Projected Newton-type methods in machine learning"