

Accelerated gradient methods for total-variation-based CT image reconstruction

Jørgensen, Jakob H.; Jensen, Tobias Lindstrøm; Hansen, Per Christian ; Jensen, Søren Holdt; Sidky, Emil Y.; Pan, Xiaochuan

Publication date:
2011

Document Version
Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Jørgensen, J. H., Jensen, T. L., Hansen, P. C., Jensen, S. H., Sidky, E. Y., & Pan, X. (2011). *Accelerated gradient methods for total-variation-based CT image reconstruction*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Accelerated gradient methods for total-variation-based CT image reconstruction

Jakob H. Jørgensen, Tobias L. Jensen, Per Christian Hansen, Søren H. Jensen, Emil Y. Sidky, and Xiaochuan Pan

Abstract—Total-variation (TV)-based CT image reconstruction has shown experimentally to be capable of producing accurate reconstructions from sparse-view data. In particular TV-based reconstruction is very well suited for images with piecewise nearly constant regions. Computationally, however, TV-based reconstruction is much more demanding, especially for 3D imaging, and the reconstruction from clinical data sets is far from being close to real-time. This is undesirable from a clinical perspective, and thus there is an incentive to accelerate the solution of the underlying optimization problem.

The TV reconstruction can in principle be found by any optimization method, but in practice the large scale of the systems arising in CT image reconstruction preclude the use of memory-demanding methods such as Newton's method. The simple gradient method has much lower memory requirements, but exhibits slow convergence.

In the present work we address the question of how to reduce the number of gradient method iterations needed to achieve a high-accuracy TV reconstruction. We consider the use of two accelerated gradient-based methods, GPBB and UPN, to solve the 3D-TV minimization problem in CT image reconstruction. The former incorporates several heuristics from the optimization literature such as Barzilai-Borwein (BB) step size selection and nonmonotone line search. The latter uses a cleverly chosen sequence of auxiliary points to achieve a better convergence rate.

The methods are memory efficient and equipped with a stopping criterion to ensure that the TV reconstruction has indeed been found. An implementation of the methods (in C with interface to Matlab) is available for download from <http://www2.imm.dtu.dk/~pch/TVReg/>.

We compare the proposed methods with the standard gradient method, applied to a 3D test problem with synthetic few-view data. We find experimentally that for realistic parameters the proposed methods significantly outperform the gradient method.

Index Terms—Total-variation, Gradient-based optimization, Strong convexity, Algorithms

I. INTRODUCTION

Algorithm development for image reconstruction from incomplete data has experienced renewed interest in the past years. Incomplete data arises for instance in case of a small number of views, and the development of algorithms for incomplete data thus has the potential for a reduction in imaging time and the delivered dosage.

Jakob H. Jørgensen and Per Christian Hansen are with the Department of Informatics and Mathematical Modelling, Technical University of Denmark, Richard Petersens Plads, Building 321, 2800 Kgs. Lyngby, Denmark. Corresponding author: Jakob H. Jørgensen, E-mail: jakj@imm.dtu.dk.

Tobias L. Jensen and Søren H. Jensen are with the Department of Electronic Systems, Aalborg University, Niels Jernesvej 12, 9220 Aalborg Ø, Denmark.

Emil Y. Sidky and Xiaochuan Pan are with the Department of Radiology, The University of Chicago, 5841 S. Maryland Avenue, Chicago, IL 60637, USA.

Total-variation (TV)-based image reconstruction is a promising direction, as experiments have documented the potential for accurate image reconstruction under conditions such as few-view and limited-angle data, see e.g. [1].

However, it also known that it is difficult to design fast algorithms for obtaining exact TV reconstructions due to non-linearity and nonsmoothness of the underlying optimization problem. Many different approaches have been developed, such as time marching [2], fixed-point iteration [3], and various minimization-based methods such as sub-gradient methods [4], second-order cone programming (SOCP) [5], and duality-based methods [6], [7] – but for large-scale applications such as CT image reconstruction the computational burden is still unacceptable. As a consequence heuristic and much faster techniques such as the one in [1] for approximating the TV solution have been developed. In such inaccurate, but efficient, TV-minimization solvers the resulting image depends on several algorithm parameters, which introduces an unavoidable variability. In contrast, for the accurate TV algorithms considered here, the resulting image can be considered dependent only on the parameters of the optimization problem.

In this work we present two accelerated gradient-based optimization methods that are capable of computing the TV reconstruction of 3D volumes to within an accuracy specified by the user.

II. THEORY

A. Total-variation-based image reconstruction

In this work we consider total-variation (TV)-based image reconstruction for computed tomography. The 3D reconstruction is represented by the vector x^* which is the solution to the minimization problem

$$x^* = \underset{x \in \mathcal{Q}}{\operatorname{argmin}} \phi(x), \quad \phi(x) = \frac{1}{2} \|Ax - b\|_2^2 + \alpha \|x\|_{\text{TV}}. \quad (1)$$

Here, x is the unknown image, \mathcal{Q} is the set of feasible x , A is the system matrix, b is the projection data stacked into a column vector, and α is the TV regularization parameter specifying the relative weighting between the fidelity term and the TV term. $\|x\|_{\text{TV}}$ is the discrete total-variation of x ,

$$\|x\|_{\text{TV}} = \sum_{j=1}^N \|D_j x\|_2, \quad (2)$$

where N is the number of voxels and D_j is the forward difference approximation to the gradient at voxel j .

B. Smooth and strongly convex functions

We recall that a continuously differentiable function f is convex if

$$f(x) \geq f(y) + \nabla f(y)^T(x - y) \quad (3)$$

for all $x, y \in \mathcal{Q}$. A stronger notion of convexity is *strong convexity*: f is said to be *strongly convex* with *strong convexity parameter* μ if there exists a $\mu \geq 0$ such that

$$f(x) \geq f(y) + \nabla f(y)^T(x - y) + \frac{1}{2}\mu\|x - y\|_2^2 \quad (4)$$

for all $x, y \in \mathcal{Q}$. Furthermore, f has *Lipschitz continuous gradient* with *Lipschitz constant* L , if

$$f(x) \leq f(y) + \nabla f(y)^T(x - y) + \frac{1}{2}L\|x - y\|_2^2 \quad (5)$$

for all $x, y \in \mathcal{Q}$. The ratio μ/L is important for the convergence rate of gradient methods we will consider.

The problem (1) can be shown [8] to be strongly convex and have Lipschitz continuous gradient in the case where A specifies a full-rank overdetermined linear system. In the rank deficient or underdetermined case, which occurs for instance for few-view data, the strong convexity assumption is violated. However, as we shall see, this turns out not to pose a problem for the gradient methods we consider.

III. ALGORITHMS

A. Gradient projection methods

The optimization problem (1) can, in principle, be solved by use of a simple gradient projection (**GP**) method

$$x^{(k+1)} = P_{\mathcal{Q}}\left(x^{(k)} - \theta_k \nabla f(x^{(k)})\right), \quad k = 0, 1, 2, \dots, \quad (6)$$

where $P_{\mathcal{Q}}$ denotes projection onto the set \mathcal{Q} of feasible x , and θ_k is the step size at the k th step. The worst-case convergence rate of **GP** with $\mu > 0$ and constant step size is

$$f(x^{(k)}) - f^* \leq \left(1 - \frac{\mu}{L}\right)^k \cdot C_{\text{GP}}, \quad (7)$$

where C_{GP} is a constant [9, §7.1.4].

For large-scale imaging modalities, such as CT, this slow convergence renders the simple gradient method impractical. On the other hand the simplicity and the low memory requirements of the gradient method remain attractive. Various modifications have been suggested in the optimization literature. For instance, a significant acceleration is often observed empirically if the gradient method is equipped with a Barzilai-Borwein (BB) step size strategy and a nonmonotone line search [10], [11], [12], [13], [14], see Algorithm 1: **GPBB** for a pseudo-code. Empirically we have found $K = 2$ and $\sigma = 0.1$ to be satisfactory parameter choices. However, it remains unproven that **GPBB** achieves a better worst-case convergence rate than (7).

Algorithm 1: GPBB

input : $x^{(0)}, K$
output: $x^{(k+1)}$

```

1  $\theta_0 = 1$  ;
2 for  $k = 0, 1, 2, \dots$  do
3   // BB strategy
4   if  $k > 0$  then
5      $\theta_k \leftarrow \frac{\|x^{(k)} - x^{(k-1)}\|_2^2}{\langle x^{(k)} - x^{(k-1)}, \nabla f(x^{(k)}) - \nabla f(x^{(k-1)}) \rangle}$  ;
6      $\beta \leftarrow 0.95$  ;
7      $\bar{x} \leftarrow P_{\mathcal{Q}}(x^{(k)} - \beta\theta_k \nabla f(x^{(k)}))$  ;
8      $\hat{f} \leftarrow \max\{f(x^{(k)}), f(x^{(k-1)}), \dots, f(x^{(k-K)})\}$  ;
9     while  $f(\bar{x}) \geq \hat{f} - \sigma \nabla f(x^{(k)})^T(x^{(k)} - \bar{x})$  do
10       $\beta \leftarrow \beta^2$  ;
11       $\bar{x} \leftarrow P_{\mathcal{Q}}(x^{(k)} - \beta\theta_k \nabla f(x^{(k)}))$  ;
12     $x^{(k+1)} \leftarrow \bar{x}$  ;
```

B. Nesterov's optimal method

Nesterov [15] proposed a gradient-based method that for given $\mu > 0$ achieves the convergence rate

$$f(x^{(k)}) - f^* \leq \left(1 - \sqrt{\frac{\mu}{L}}\right)^k \cdot C_{\text{N}}, \quad (8)$$

where C_{N} is a constant, and he proved the method to be optimal, i.e., that no gradient-based method can achieve better worst-case convergence rate on the class of strongly convex problems.

Comparing (7) and (8), we see how the ratio μ/L affects the predicted worst-case convergence rates: When μ/L decreases, both convergence rates become slower, but less in (8) due to the square root. We therefore expect Nesterov's method to show better convergence for smaller μ/L . Small μ/L arise for instance when the number of views is small, see [8].

Nesterov's method requires that both μ and L are given by the user, and in order for the method to be convergent μ must be chosen sufficiently small and L sufficiently large. For real world applications such as CT, μ and L are seldom known, which makes the method impractical. Taking overly conservative estimates can depreciate the better convergence rate (8); hence, accurate estimates of μ and L are important.

C. Estimating μ and L

A sufficiently large L can be chosen using *back-tracking line search* [16], [17], see Algorithm 2: **BT** for pseudo-code. Essentially, an estimate \tilde{L} of L is increased by multiplication with a constant $\rho_L > 1$ until (5) is satisfied.

Accurately estimating μ , such that (4) is satisfied globally, is more difficult. Here, we propose a simple and computationally inexpensive heuristic: In the k th iteration choose an estimate μ_k as the largest value of μ that satisfies (4) between $x^{(k)}$ and $y^{(k)}$, and make the μ_k -sequence non-increasing:

$$\mu_k = \min \left\{ \mu_{k-1}, \frac{f(x) - f(y) - \nabla f(y)^T(x - y)}{\frac{1}{2}\|x - y\|_2^2} \right\}. \quad (9)$$

Algorithm 2: BT

input : y, \bar{L}
output: x, \tilde{L}
1 $\tilde{L} \leftarrow \bar{L}$;
2 $x \leftarrow P_{\mathcal{Q}} \left(y - \tilde{L}^{-1} \nabla f(y) \right)$;
3 **while** $f(x) > f(y) + \nabla f(y)^T (x - y) + \frac{1}{2} \tilde{L} \|x - y\|_2^2$ **do**
4 $\tilde{L} \leftarrow \rho_L \tilde{L}$;
5 $x \leftarrow P_{\mathcal{Q}} \left(y - \tilde{L}^{-1} \nabla f(y) \right)$;

Algorithm 3: UPN

input : $x^{(0)}, \bar{\mu}, \bar{L}$
output: $x^{(k+1)}$
1 $[x^{(1)}, L_0] \leftarrow \mathbf{BT}(x^{(0)}, \bar{L})$;
2 $\mu_0 = \bar{\mu}, \quad y^{(1)} \leftarrow x^{(1)}, \quad \theta_1 \leftarrow \sqrt{\mu_0 / L_0}$;
3 **for** $k = 1, 2, \dots$ **do**
4 $[x^{(k+1)}, L_k] \leftarrow \mathbf{BT}(y^{(k)}, L_{k-1})$;
5 $\mu_k \leftarrow \min\{\mu_{k-1}, M(x^{(k)}, y^{(k)})\}$;
6 $\theta_{k+1} \leftarrow$ positive root of $\theta^2 = (1 - \theta)\theta_k^2 + (\mu_k / L_k) \theta$;
7 $\beta_k \leftarrow \theta_k(1 - \theta_k) / (\theta_k^2 + \theta_{k+1})$;
8 $y^{(k+1)} \leftarrow x^{(k+1)} + \beta_k(x^{(k+1)} - x^{(k)})$;

We call the Nesterov method equipped with estimation of μ and L Unknown Parameter Nesterov (UPN) and pseudo-code is given in Algorithm 3: **UPN**.

Unfortunately, convergence of **UPN** is not guaranteed, since the estimate (9) can be too large. However, we have found empirically that an estimate sufficient for convergence is typically effectively determined by (9).

It is possible to ensure convergence by introducing a restarting procedure [8] at the price of lowering the convergence rate bound and thereby losing optimality of the method. However, we have found empirically that the restarting procedure is seldom needed, and for realistic parameters the simple heuristic (9) is sufficient.

D. Stopping criterion

For an unconstrained convex optimization problem such as (1) the norm of the gradient is a measure of closeness to the minimizer through the first-order optimality conditions [18]. For a constrained convex optimization problem it is possible to express a similar optimality condition, namely in terms of the *gradient map* defined by

$$G_\nu(x) = \nu \left(x - P_{\mathcal{Q}} \left(x - \nu^{-1} \nabla f(x) \right) \right), \quad (10)$$

where $\nu > 0$ is a scalar. A point x^* is optimal if and only if $G_\nu(x^*) = 0$ for any $\nu > 0$ [17]. We can use this to design a stopping criterion: Stop the algorithm after iteration k if $\|G_\nu(x^{(k)})\|_2 / N \leq \epsilon$, where ϵ is a user-specified tolerance.

For an under-determined problem, e.g. in the few-view case, the objective function in (1) is nearly flat at the minimizer, which makes it difficult to determine when a sufficiently accurate reconstruction has been found. Here, the gradient map provides a simple, yet sensitive, stopping criterion.

IV. SIMULATION RESULTS AND DISCUSSION**A. Simulation setup**

At this point we emphasize that our objective is to obtain an accurate TV reconstruction while reducing the required number of gradient method iterations. We include two reconstructions merely to demonstrate that the methods indeed are successful in solving (1), thereby reconstructing the desired image. In [8] dependence of the convergence with respect to parameter variation is explored.

To demonstrate and compare the convergence of **GP**, **GPBB** and **UPN** we set up a simple test problem. As test image x_{true} we use the three-dimensional FORBILD head phantom discretized into 64^3 voxels. We simulate a parallel beam geometry with view directions evenly distributed over the unit sphere. Projections are computed as the forward mapping of the discretized image subject to additive Gaussian white noise e of relative magnitude $\|e\|_2 / \|Ax_{\text{true}}\|_2 = 0.01$, i.e.,

$$b = Ax_{\text{true}} + e. \quad (11)$$

We enforce nonnegativity by taking $\mathcal{Q} = \mathbb{R}_+^{64^3}$. We consider two reconstructions: A “many-view” using 55 views and a “few-view” using only 19 views of size 91^2 pixels. In the latter case A has less rows than columns, which can be shown [8] to lead to violation of the assumption on strong convexity. The iterations are continued to the tolerance $\epsilon = 10^{-8}$ is met.

B. Simulation results

Fig. 1 shows the middle (33rd) axial voxel slice through the original 3D volume together with many-view and few-view **UPN** reconstructions using $\alpha = 0.01$. Both reconstructions reproduce the original features accurately, except for two small features are missing in the few-view reconstruction. Fig. 2 shows the convergence of the three methods in terms of objective value $\phi^{(k)}$ relative to the true minimal objective value ϕ^* as function of the iterations k . As ϕ^* is unknown, we have approximated it by computing the **UPN** solution for an ϵ two orders of magnitude lower than the value used in the iterations.

In both cases we see that **UPN** converges to a satisfactory accuracy within 2000 iterations, whereas **GP** does not, and **GPBB** only does in the former case. In the many-view case **UPN** and **GPBB** both produce a significant (and comparable) acceleration over **GP**. In the few-view case, we also observe acceleration for both, but **UPN** stands out with much faster convergence. This is in accordance with the expectation stated in Section III-B.

The adequacy of the stopping criterion is evaluated by a simple visual comparison of the few-view simulation gradient map norm decay (Fig. 2 right) and the objective decay (Fig. 2 center). Apart from the erratic decay for **GPBB** (which is caused by highly irregular step length selection) there is a pronounced correspondence, and we therefore consider the stopping criterion effective.

Although **UPN** was designed for strongly convex problems, we conclude that the method also works in the non-strongly convex case of having few-view data – in fact, from the preliminary results the non-strongly convex case is where




Fig. 1. Central axial slices. Left: Original. Center: Many-view **UPN** reconstruction. Right: Few-view **UPN** reconstruction. The display color range is set to $[1.04, 1.07]$ for improved viewing contrast.




Fig. 2. Convergence histories. Left: Many-view simulation. Center: Few-view simulation. Right: Gradient map norm histories for few-view simulation.

UPN shows its biggest potential by exhibiting a much faster convergence than **GP** and **GPBB**.

V. CONCLUSION

We have described the gradient-based optimization methods **GPBB** and **UPN** and their worst-case convergence rates. Our simulations show that both algorithms are able to significantly accelerate high-accuracy TV-based CT image reconstruction compared to a simple gradient method. In particular **UPN** shows much faster convergence when applied to few-view data. The software (implementation in C with an interface to Matlab) is available from <http://www2.imm.dtu.dk/~pch/TVReg/>.

ACKNOWLEDGMENT

This work is part of the project CSI: Computational Science in Imaging, supported by grant 274-07-0065 from the Danish Research Council for Technology and Production Sciences.

REFERENCES

- [1] E. Y. Sidky, C.-M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam ct," *J. X-Ray Sci. Technol.*, vol. 14, pp. 119–139, 2006.
- [2] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, pp. 259–268, 1992.
- [3] C. R. Vogel and M. E. Oman, "Iterative methods for total variation denoising," *SIAM J. Sci. Comput.*, vol. 17, pp. 227–238, 1996.
- [4] P. L. Combettes and J. Luo, "An adaptive level set method for non-differentiable constrained image recovery," *IEEE Trans. Image Proces.*, vol. 11, pp. 1295–1304, 2002.
- [5] D. Goldfarb and W. Yin, "Second-order cone programming methods for total variation-based image restoration," *SIAM J. Sci. Comput.*, vol. 27, pp. 622–645, 2005.
- [6] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imaging Vis.*, vol. 20, pp. 89–97, 2004.
- [7] M. Hintermüller and G. Stadler, "An infeasible primal-dual algorithm for total bounded variation-based INF-convolution-type image restoration," *SIAM J. Sci. Comput.*, vol. 28, pp. 1–23, 2006.
- [8] T. L. Jensen, J. H. Jørgensen, P. C. Hansen, and S. H. Jensen, "Implementation of an optimal first-order method for strongly convex total variation regularization," *submitted*.
- [9] A. S. Nemirovsky and D. B. Yudin, *Problem Complexity and Method Efficiency in Optimization*. Wiley-Interscience, New York, 1983.
- [10] J. Barzilai and J. M. Borwein, "Two-point step size gradient methods," *IMA J. Numer. Anal.*, vol. 8, pp. 141–148, 1988.
- [11] E. G. Birgin, J. M. Martínez, and M. Raydan, "Nonmonotone spectral projected gradient methods on convex sets," *SIAM J. Optim.*, vol. 10, pp. 1196–1211, 2000.
- [12] M. Zhu, S. J. Wright, and T. F. Chan, "Duality-based algorithms for total-variation-regularized image restoration," *Comput. Optim. Appl.*, 2008, doi: 10.1007/s10589-008-9225-2.
- [13] L. Grippo, F. Lampariello, and S. Lucidi, "A nonmonotone line search technique for Newton's method," *SIAM J. Numer. Anal.*, vol. 23, pp. 707–716, 1986.
- [14] M. Raydan, "The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem," *SIAM J. Optim.*, vol. 7, pp. 26–33, 1997.
- [15] Y. Nesterov, *Introductory Lectures on Convex Optimization*. Kluwer Academic Publishers, Dordrecht, 2004.
- [16] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. on Imaging Sciences*, vol. 2, pp. 183–202, 2009.
- [17] L. Vandenbergh, "Optimization methods for large-scale systems," 2009, lecture Notes. www.ee.ucla.edu/~vandenbe/ee236c.html.
- [18] S. Boyd and L. Vandenbergh, *Convex Optimization*. Cambridge University Press, 2004.