Journal of Computational and Applied Mathematics 313 (2017) 243–258

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Tomographic image reconstruction using training images[☆]

Sara Soltani [a], Martin S. Andersen [b], Per Christian Hansen [b,*]

[a] Fingerprint Cards AB, Lyskær 3CD, DK-2730 Herlev, Denmark

[b] Department of Applied Mathematics and Computer Science, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

ARTICLE INFO

Article history:
Received 15 August 2015
Received in revised form 5 August 2016

MSC:
primary 65F22
secondary 65K10

Keywords:
Tomography
Dictionary learning
Inverse problems
Regularization
Sparse representation
Image reconstruction

ABSTRACT

We describe and examine an algorithm for tomographic image reconstruction where prior knowledge about the solution is available in the form of training images. We first construct a non-negative dictionary based on prototype elements from the training images; this problem is formulated within the framework of sparse learning as a regularized non-negative matrix factorization. Incorporating the dictionary as a prior in a convex reconstruction problem, we then find an approximate solution with a sparse representation in the dictionary. The dictionary is applied to non-overlapping patches of the image, which reduces the computational complexity compared to previous formulations. Computational experiments clarify the choice and interplay of the model parameters and the regularization parameters, and we show that in few-projection low-dose settings our algorithm is competitive with total variation regularization and tends to include more texture and more correct edges.

1. Introduction

Computed tomography (CT) is a technique to compute an image of the interior of an object from measurements obtained by sending X-rays through the object and recording the damping of each ray. CT is used routinely in medical imaging, materials science, nondestructive testing and many other applications.

CT is an inverse problem [1] and it is challenging to obtain sharp and reliable reconstructions in low-dose measurements where we face underdetermined systems of equations, because we must limit the accumulated amount of X-rays for health reasons or because measurement time is limited. In these circumstances the classic methods of CT, such as filtered back projection [2] and algebraic reconstruction techniques [3], are often incapable of producing satisfactory reconstructions because they fail to incorporate adequate prior information [4]. To overcome these difficulties it is necessary to incorporate prior information about the solution that can compensate for the lack of data.

A popular prior is that the image is piecewise constant, leading to total variation (TV) regularization schemes [5,6]. These methods can be very powerful when the solution is approximately composed of homogeneous regions separated by sharp boundaries.

A completely different approach is to use prior information in the form of "training images" that characterize the geometrical or visual features of interest, e.g., from high-accuracy reconstructions or from pictures of specimen slices. The goal of this work is to elaborate on this approach. In particular we consider the two-stage framework where the most important features of the training data are first extracted and then integrated in the reconstruction problem.

☆ This work is part of the project HD-Tomo funded by Advanced Grant No. 291405 from the European Research Council.

* Corresponding author.
E-mail addresses: sara.soltani@fingerprints.com (S. Soltani), mskan@dtu.dk (M.S. Andersen), pcha@dtu.dk (P.C. Hansen).

http://dx.doi.org/10.1016/j.cam.2016.09.019

A natural way to extract and represent prior information from training images is to form a *dictionary* that sparsely encodes the information [7]. Learning the dictionary from given training data appears to be very suited for incorporating priors that are otherwise difficult to formulate in a closed form, such as image texture. Dictionary learning – combined with sparse representation [8–10] – is now used in many image processing areas including denoising [11,12], inpainting [13], and deblurring [14]. Elad and Ahron [15] address the image denoising problem using a process that combines dictionary learning and reconstruction. They use a dictionary trained from a noise-free image using the K-SVD algorithm [16] combined with an adaptive dictionary trained on patches of the noisy image.

The use of dictionary learning in tomographic imaging has also emerged recently, e.g., in X-ray CT [17–19], magnetic resonance imaging [20,21], electron tomography [22], positron emission tomography [23], and phase-contrast tomography [24]. Two different approaches have emerged—either one constructs the dictionary from the given data in a joint learning-reconstruction algorithm [23,20,22,21], or one constructs the dictionary from training images in a separate step before the reconstruction [17,24,18,19]. Most of these works use K-SVD to learn the dictionary (except [17] that uses an "online dictionary learning method" [25]), and all the methods regularize the reconstruction by means of a penalty that is applied to a patch around every pixel in the image. In other words, all patches in the reconstruction are required be close to the subspace spanned by the dictionary images. While all these methods perform better than classical reconstruction methods, they show no significant improvement over the TV-regularized approach.

In simultaneous learning and reconstruction, where the dictionary is learned from the given data, the prior is purely data-driven. Hence, one can argue that it violates a fundamental principle of inverse problems where a data-independent prior is incorporated to eliminate unreasonable models that fit the data. For this reason we prefer to separate the two steps (which requires that reliable training images are available). We describe and examine a two-stage framework where we first construct a dictionary that contains prototype elements from these images, and then we use the dictionary as a prior to regularize the reconstruction problem via computing a solution that has a sparse representation in the dictionary.

Our two-stage algorithm is inspired by the work in [17] and, to some extent, [18]. The algorithm in [17] is tested on a simple and ideal tomography setup with no noise in the data, and in [18] the dictionary is trained from an image reconstructed by a high-dose X-ray exposure and then used to reconstruct the same image with fewer X-ray projections.

The focus of this paper is a new formulation of tomographic reconstruction where training images are used as a strong prior. Our algorithm utilizes the dictionary in a different way than previous formulations, by using *non-overlapping* blocks of the image which reduces the number of unknowns and thus the computational work. We use state-of-the-art numerical optimization methods to handle the large-scale optimization problems, and we apply this algorithm applied to underdetermined problems where a strong prior is necessary. We perform comprehensive studies of the influence of the learned dictionary structure and the dictionary parameters on the CT reconstruction, and we compare our algorithm with both classical methods and with TV regularization.

Our paper is organized as follows. In Section 2 we briefly discuss dictionary learning methods and present a framework for solving the image reconstruction problem using dictionaries, and in Section 3 we describe the implementation details of algorithm. Section 4 presents careful numerical experiments where we study the influence of the algorithm and design parameters. Section 5 summarizes our work. We use the following notation, where A is an arbitrary matrix:

$$\|A\|_{\mathrm{F}} = \left(\sum_{ij} A_{ij}^2\right)^{1/2}, \qquad \|A\|_{\mathrm{sum}} = \sum_{ij} |A_{ij}|, \qquad \|A\|_{\max} = \max_{ij} |A_{ij}|.$$

2. The reconstruction framework

X-ray CT is based on the principle that if we send X-rays through an object and measure the damping of each ray then, with infinitely many rays, we can perfectly reconstruct the object. The attenuation of an X-ray is proportional to the object's attenuation coefficient, as described by Lambert–Beer's law [26, Section 2.3.1]. We divide the domain onto pixels whose unknown non-negative attenuation coefficients are organized in the vector $x \in \mathbb{R}^n$. Similarly we organize the measured damping of the rays into the vector $b \in \mathbb{R}^m$. Then we obtain a linear system of equations $Ax = b$ with a large sparse *system matrix* A governed solely by the geometry of the measurements: element a_{ij} is the length of the ith ray passing through pixel j, and the matrix is sparse because each ray only hits a small number of pixels [1].

The matrix A is ill-conditioned, and often rank deficient, due to the ill-posedness of the underlying inverse problem and therefore the solution is very sensitive to noise in the data b. For this reason, a simple least squares approach with non-negativity constraints fails to produce a meaningful solution, and we must use regularization to incorporate prior information about the solution [27].

This work is concerned with underdetermined problems where $m < n$, and the need for regularization is even more pronounced. Classical reconstruction methods such as filtered back projection and algebraic iterative methods are not suited for these problems because they fail to incorporate enough prior information. TV regularization, which is suited for edge-preserving reconstructions, takes the form

$$\min_x \frac{1}{2}\|Ax - b\|_2^2 + \lambda_{\mathrm{TV}} \sum_{1\le i\le n} \left\|D_i^{\mathrm{fd}} x\right\|_2 \quad \text{subject to } x \ge 0, \tag{1}$$