

Automatic Optimization of Alignment Parameters for Tomography Datasets

Folkert Bleichrodt and K. Joost Batenburg

Centrum Wiskunde & Informatica, Science Park 123,
1098 XG Amsterdam, The Netherlands
{f.bleichrodt,k.j.batenburg}@cwi.nl

Abstract. As tomographic imaging is being performed at increasingly smaller scales, the stability of the scanning hardware is of great importance to the quality of the reconstructed image. Instabilities lead to perturbations in the geometrical parameters used in the acquisition of the projections. In particular for electron tomography and high-resolution X-ray tomography, small instabilities in the imaging setup can lead to severe artifacts. We present a novel alignment algorithm for recovering the true geometrical parameters *after* the object has been scanned, based on measured data. Our algorithm employs an optimization algorithm that combines alignment with reconstruction. We demonstrate that problem-specific design choices made in the implementation are vital to the success of the method. The algorithm is tested in a set of simulation experiments. Our experimental results indicate that the method is capable of aligning tomography datasets with considerably higher accuracy compared to standard cross-correlation methods.

Keywords: alignment, parameter estimation, tomography.

1 Introduction

Tomography deals with the problem of reconstructing an object from projections [1]. Projections are measured by a scanning device at varying orientations with respect to the object. Each projection consists of measurements for a set of parallel lines, where the line integral of the object density is measured. By rotating around the object and creating a large number of projections, enough information is gathered to reconstruct an object that resembles the ground truth. The result is an ill conditioned inverse problem that can be solved using numerical methods.

Tomography has a wide range of applications, ranging from industrial quality control of large objects using X-rays down to imaging of nanomaterials by electron microscopy. In particular at the smallest scales, problems with the *alignment* of the projection data form a key bottleneck for the quality of the reconstructed image. The geometry that is presumed in the reconstruction process is different from the actual geometry of the projection acquisition.

In electron tomography, the specimen has to be recentered for each recorded image as the sample stage is not eucentric, causing lateral shifts in the projection

images [2]. In high-resolution X-ray tomography, the rotation axis may not be perfectly centered at the detector, leading to structured shifts in the projections. In addition, limited accuracy of the rotation stage leads to uncertainties about the exact projection angles. Such instabilities result in small deviations in the parameters describing the geometry. As a result, inconsistencies are present in the system of equations governing the reconstruction problem. These inconsistencies must be removed to obtain accurate reconstructions.

We remark that the alignment problem for tomography is fundamentally different from some other problems also named “alignment” in the image processing literature [3–5]. Compared to, for example, the alignment of photographs in a stitching problem [3], the key difference is that for tomographic alignment, the 3D object itself is related to the (unaligned) projections by a complex inverse problem. Therefore, projections from different angles can often not be directly compared and can only be related to each other by solving this inverse problem. This also makes it impossible to use image registration methods [5] for the type of alignment we consider.

A range of tomographic alignment algorithms have been proposed, which can generally be divided in two classes: methods using fiducial markers and methods based on automatic, markerless alignment. Marker based alignment is often applied in electron tomography [6]. Small, dense particles are distributed among the sample, which can be tracked accurately in consecutive projections. A system of equations, relating the marker positions in the projection domain and their position in the sample, can be solved to compute the alignment parameters with a high degree of accuracy. The method requires a long preparation time and the use of markers can result in artifacts in the reconstructed image. Instead of fiducial markers, features in the projection data can also act as markers, [7].

For algorithms that do not use markers, a well known approach is cross-correlation [8]. Here it is assumed that consecutive projections are similar and differ in a smooth way, thereby making strong assumptions about the unknown object. By finding the maximum cross-correlation between successive projections, it is possible to make a rough estimation of the alignment parameters that can be described as an affine transformation of the projections. The main problem of this method is its low accuracy.

Other markerless methods are based on minimizing the inconsistencies between the forward projections of the reconstructed image and the measured projections. These methods, called *projection distance minimization* methods henceforth, are a more general approach to the alignment problem. See for example [9–11]. Other methods focus on an error measure based on the reconstruction [12], or use passive auto-focus [13].

In this paper, a new markerless alignment method based on projection distance minimization is presented. We propose the Levenberg-Marquardt Projection Distance Minimization algorithm (LMPDM). Similar to the algorithm proposed in [11], the alignment and reconstruction problem are solved simultaneously. The objective of combined alignment and reconstruction is posed as

an nonlinear least squares optimization problem and a numerical method is employed for solving it.

Instead of the Quasi-Newton BFGS method used in [11], we choose Levenberg-Marquardt (LM), which has been shown to yield better convergence for certain least squares problems, as discussed in chapter 10 of [14]. When implementing a numerical scheme for this nonlinear least-squares problem, several design choices must be made, with respect to computation of numerical derivatives and image resolution. We demonstrate that these design choices are crucial to the success of the LM algorithm in recovering the alignment parameters.

Our experimental results, based on simulated projection data, show that if a multi-resolution scheme is combined with local smoothing of the Jacobian, our LMPDM algorithm is capable of recovering the alignment parameters with high accuracy. Also, the underlying tomography software library is implemented on the GPU, which makes the algorithm scalable.

This paper is structured as follows. Mathematical background and implementation details are discussed in Section 2. In Section 3 and 4, a series of experiments is described and the results are presented. Section 5 contains a discussion of the results. Section 6 concludes this paper.

2 Methods and Implementation

This section will formulate the alignment problem in a mathematical context and introduce the notation. Subsequently, the LMPDM method and its implementation details will be discussed. Furthermore, design choices are explained that improve the accuracy of the alignment algorithm.

2.1 Model and Notation

The object from which the projections are acquired can be modeled by a gray value image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. A projection at angle θ is the collection of line integrals over the lines $l_{\theta,t} = \{(x, y) : x \cos \theta + y \sin \theta = t\}$ for detector positions $t \in \mathcal{T} \subset \mathbb{R}$, where \mathcal{T} denotes the discrete set of detector positions, see Figure 1. The relation between the object and its projections $P(\theta, t)$ is given by the Radon transform

$$P(\theta, t) = \mathcal{R}(f)(\theta, t) = \int \int_{-\infty}^{\infty} f(x, y) \bar{\delta}(x \cos \theta + y \sin \theta - t) dx dy, \quad (1)$$

with $\bar{\delta}$ the Dirac delta function. By discretizing the image f , the set of angles, and the set of detector positions, and numerically approximating the Radon transform we arrive at the algebraic representation of the tomography problem. In this form the object and its projections are related by a linear operator

$$W\mathbf{x} = \mathbf{p}, \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^N$ represents the unknown object, $W \in \mathbb{R}^{M \times N}$ is the projection operator and $\mathbf{p} \in \mathbb{R}^M$ is the measured set of projections [1]. From this point




Fig. 1. Geometry for the two dimensional case: (a) A tomographic scan: the dark gray region represents the object along with its projection below. The detector-source pair rotates around the object; (b) Projection acquisition at angle θ_i with angular offset ϕ_i . The object has a shift of δ_i in the detector plane with respect to its assumed position.

on, we focus on the reconstruction of a single slice of the object, i.e. a 2D image from a set of 1D projections. The object is represented as a two dimensional pixel grid with N pixels. Let K be the number of projections of the object that have been acquired by a detector having D discrete elements. The total number of line projections is then given by $M \equiv KD$. The projection operator is a sparse matrix with $w_{i,j}$ modeling the contribution of pixel j to the projection value measured by detector i . So the inner product of row i of W and the object \mathbf{x} gives a discrete approximation of the line integral over a line perpendicular to detector i .

Projections of the object are recorded at a discrete set of angles

$$\begin{aligned} \boldsymbol{\theta} &= \theta_1, \dots, \theta_K, \\ 0 &\leq \theta_1 < \dots < \theta_K \leq \pi. \end{aligned}$$

Up until now, we have assumed that the measurements correspond perfectly with the Radon transform. In practice, each of the projections have a perturbation in the angles as well as the object position. These are represented in the alignment parameters

$$\begin{aligned} \boldsymbol{\delta} &= \delta_1, \dots, \delta_K, \\ \boldsymbol{\phi} &= \phi_1, \dots, \phi_K. \end{aligned} \tag{3}$$

as illustrated in Figure 1(b). Accordingly, in the continuous case, a single line projection at angle θ_i and detector offset t is represented by the Radon transform including the alignment parameters:

$$\mathcal{R}(f)(\theta_i + \phi_i, t + \delta_i) = \int \int_{-\infty}^{\infty} f(x, y) \bar{\delta}(x \cos(\theta_i + \phi_i) + y \sin(\theta_i + \phi_i) - (t + \delta_i)) dx dy.$$

In the discrete model, the coefficients in our projection operator depend on the geometry,

$$W(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\delta})\mathbf{x} = \mathbf{p}. \quad (4)$$

This expression for W is not easily available in closed form, but it can be evaluated numerically.

In an experimental setup, the projections contain noise and the perturbations of the geometrical parameters are not known. Therefore, the system in Eq. (4) is inconsistent. Alignment involves estimating the unknown alignment parameters in Eq. (3). Minimizing the least-squares residual of Eq. (4) seems to be a good approach, because in the absence of noise and when the alignment parameters are known exactly, then Eq. (4) is consistent.

Now we can define the objective of combined alignment and reconstruction as a minimization problem of the *projection distance*, defined by the following cost function

$$\min_{\mathbf{x}, \boldsymbol{\phi}, \boldsymbol{\delta}} \frac{1}{2} \|\mathbf{r}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\delta}, \mathbf{x})\|_2^2 := \min_{\mathbf{x}, \boldsymbol{\phi}, \boldsymbol{\delta}} \frac{1}{2} \|W(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\delta})\mathbf{x} - \mathbf{p}\|_2^2, \quad (5)$$

with \mathbf{r} the residual. The 2-norm is chosen because it allows us to use least square solvers and it has some nice properties, due to its simplicity. Alternative distance or similarity measures such as mutual information can be employed here and might give satisfying results as well.

In Eq. (5), the minimization with respect to \mathbf{x} is a linear inverse problem that yields a reconstructed image. The minimization with respect to $\boldsymbol{\delta}$ and $\boldsymbol{\phi}$ can be seen as a non-linear model fitting problem. The combination in the full cost function is, hence, a non-linear least squares problem.

Projection matching algorithms such as [10], consider the same cost function as in Eq. (5), however, an alternating approach is employed. In those methods, the alignment estimation of $\boldsymbol{\phi}$ and $\boldsymbol{\delta}$ is separated from the calculation of the gray values \mathbf{x} . Such methods are heuristic in nature and it is not guaranteed that this approach converges to a local minimum. This is why we chose to minimize over the full set of variables at the same time.

The cost function seems suitable to solve by using one of the standard algorithms from numerical optimization. A method specifically aimed at these kinds of problems is the Newton-type algorithm Levenberg-Marquardt. However, due to numerical problems, a straightforward implementation often does not yield an accurate alignment. In the following sections, we will demonstrate that problem-specific design choices in the implementation are essential for accurate parameter estimation.

2.2 Levenberg-Marquardt

Levenberg-Marquardt, see chapter 10 of [14], is an iterative method that generates a sequence of input vectors $\{\mathbf{y}_k\} = \{(\mathbf{x}^{(k)}, \boldsymbol{\phi}^{(k)}, \boldsymbol{\delta}^{(k)})\}$ that have monotonically decreasing cost function values. Each iteration has the basic form

$$\mathbf{y}_{k+1} = \mathbf{y}_k + \boldsymbol{\eta}_k, \quad (6)$$

where the descent direction $\boldsymbol{\eta}_k$ is found by minimizing a quadratic model of the objective function using gradient information:

$$\min_{\boldsymbol{\eta}_k} \left\| \begin{bmatrix} J_k \\ \sqrt{\lambda_k} \end{bmatrix} \boldsymbol{\eta}_k + \begin{bmatrix} \mathbf{r}_k \\ 0 \end{bmatrix} \right\|^2 \quad (7)$$

with J_k the Jacobian of the residual \mathbf{r}_k and λ_k a regularization parameter. This parameter limits the norm of the search direction and acts as a trust-region. It is adjusted based on the accuracy of the quadratic model.

The linear least squares problem in Eq. (7) can be solved using one of the many available least squares solvers.

2.3 Computing the Jacobian

For computing the Jacobian of the residual we use a combination of an analytical expression and a numerical approximation. With respect to the image \mathbf{x} the Jacobian is given by $J_{\mathbf{x}} = W$, but for the derivative with respect to the parameters $\boldsymbol{\delta}$ and $\boldsymbol{\phi}$ we do not have such an expression. Therefore we approximate the gradients in the Jacobian by a central finite differences scheme:

$$\nabla_{\delta_i} \mathbf{r}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\delta}, \mathbf{x}) = \frac{W(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\delta} + h\hat{\mathbf{e}}_i)\mathbf{x} - W(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\delta} - h\hat{\mathbf{e}}_i)\mathbf{x}}{2h}, \quad (8)$$

where $\hat{\mathbf{e}}_i$ is the i th basis vector. A similar expression is used for $\boldsymbol{\phi}$.

As illustrated in Figure 2, our GPU-implementation of the cost function in Eq. (5) shows irregularities at small scales. These are introduced by the discretization of the problem domain, by floating-point errors involved in computing the cost function, and by noise in the projection data. This behavior makes the accuracy of the numerical Jacobian in Eq. (8) highly dependent on the step size h . Therefore, a robust method for choosing a good step size h is needed.

Methods proposed in literature for computing numerical derivatives on discrete, noisy data are not feasible in our implementation, due to their computational intensity [15, 16].

As an alternative, we propose the following method. We sample the cost function in the direction of $\boldsymbol{\phi}$ (and similarly for $\boldsymbol{\delta}$):

$$s_{\phi}(\alpha) = \frac{1}{2} \|W(\boldsymbol{\theta}, \boldsymbol{\phi} + \alpha\mathbf{1}, \boldsymbol{\delta})\mathbf{x} - \mathbf{p}\|^2, \quad (9)$$

at the equidistant points

$$\alpha = -8h, -7h, \dots, 7h, 8h.$$

Here $\mathbf{1}$ is a vector of which each element is 1. The sample points with odd indices are used to generate a spline. If the cost function is smooth at the current scale h , we can assume that the spline is a close approximation to the cost function. As an error measure for this we compute the difference between the sample points with even indices and the generated spline and normalize to yield a relative error.




Fig. 2. Comparison of the cost function at different scales. The plots show the cost function in the range $\frac{1}{2}\|W(\theta, \phi, \delta \pm h\mathbf{1})\mathbf{x} - \mathbf{p}\|^2$, together with its approximating spline. The irregular behavior starts to disappear for $h = 10^{-3}$.

By computing this error for several scales h , we can select the scale for which the error is minimal. The cost function at this scale does not show irregularities due to the discretization. This h is then used in Eq. (8) as step size. Figure 2 illustrates this method.

Sampling these values to compute a step size h is costly, hence the step sizes are computed once at the beginning of the algorithm. It is recomputed only after a transition between resolutions, because our algorithm employs a multi-resolution technique as discussed in the next section.

2.4 Multi-resolution

One of the main difficulties in applying the alignment algorithm in practice is the computational scale. It is not uncommon to have datasets containing billions of detector values. A conventional approach to reduce the computation time is to apply multi-resolution techniques. We utilize this technique by running the algorithm repeatedly, going from a coarse to a fine representation of the data. The output of one run serves as the input of the next. Low frequency components of the error are removed first at coarse grids. This approach refines the solution by gradually removing higher frequency components of the alignment error.

In our case the domain of the multi-resolution technique is the reconstructed image and a sinogram (set of projections). We have chosen to match the pixel size of the image with the size of a detector element. This makes the implementation easier, since the sinogram and reconstructed image can simply be resized when going from coarse to fine representation.

Lowering the resolution makes the images smoother, hence multi-resolution acts as a regularization of the optimization problem in Eq. (5). For example, the detector shift is measured in the number of detector elements. So on a coarse grid, the detector shift is reduced by the same factor by which the grid has been resized. Essentially, the initial values become closer to the optimal values. This makes it more likely to find the global minimum and possibly skip local minima. The effect of applying multi-resolution is shown in Section 4.

3 Experiments

A series of simulation experiments was carried out to evaluate the capabilities of the LMPDM algorithm. In the simulations we used the following hardware: a workstation with an Intel(R) Core(TM) i7-2600K CPU @ 3.40GHz combined with a Geforce GTX 570 GPU. For the forward and back projection operations, a GPU implementation was used.

We have applied LMPDM to three datasets based on the phantom images, shown in the left column of Figure 3. The datasets consist of projections at 100 angles, which were generated from the phantom images. The equidistant angles are in the range $[0, \pi)$ and random, uniformly distributed offsets $\phi_i \in [-0.9, 0.9]$ were added. The error in the angles is at most $\pm 0.9^\circ$, such that the ordering of the angles is preserved. Also, for each angle a uniform random shift $\delta_i \in [-10, 10]$ was applied. The maximum shift of 10 detector pixels is approximately 5 percent of the image size, which is 256×256 . The detector has 256 detector elements per projection. Poisson noise was applied to the projections using a photon count of 10^5 , to simulate moderate experimental noise. The projection matrix W is computed by the method of Joseph [17], using a GPU implementation.

The method we employ for solving the quadratic model in Eq. (7) is LSMR [18]. As a stop criterion for LMPDM, the change in parameters relative to their norm is monitored. If this falls below a certain threshold, the algorithm stops. The same holds for the norm of the gradient $\|J^T \mathbf{r}\|$ of the cost function in Eq. (5). In any case the algorithm transitions to a higher resolution, or is terminated, when a total of 100 iterations is reached.

For comparison, we have also employed a cross-correlation algorithm. This method estimates object shifts by correlating consecutive projections. Cross-correlation on two discrete real signals f and g is defined as:

$$(f \star g)(i) := \sum_{j=1}^L f(j)g(i+j) \quad (10)$$

where L is the length of the reference signal f . Usually zero-padding of g is needed. The cross-correlation attains its maximum value when the two signals align, or match as closely as possible. The corresponding i gives us the shift between the signals. To allow sub-pixel precision in the alignment, prior to the cross-correlation, the projections were upsampled by a factor of ten.

A region of the first projection, that is in view for all projections, acts as reference. To this, the second projection is correlated, estimating the relative shift. Then the second projection acts as reference to which the third projection is aligned and so on. Note that we assume here, that the first projection is perfectly aligned. If this was not the case, the projections are shifted away from the center of rotation, which still produces alignment artifacts.

4 Results

The qualitative results are given in Figure 3. Column 2 shows the unaligned reconstructions, 300 iterations of the algebraic reconstruction method SIRT were performed [1]. These show the impact of small perturbations in the geometry. Details are blurred and the background is filled with stripes.

The third column shows SIRT reconstructions using the alignment parameters found by the cross-correlation method. Since this method cannot retrieve angular offsets, the resulting artifacts are still visible. For the mandible bone dataset, cross-correlation clearly fails. Due to the fact that the sample is flat, projections from different angles have very different width. Therefore, without stretching of the projections, their correlation is rather limited. Many streaks inside the objects remain. This is because the shift parameters are not found accurately. If we look at the difference between the found alignment parameters and their true values for the Shepp-Logan phantom in Figure 5, it is clear that cross-correlation does not yield sub-pixel accuracy. The LMPDM method, however, achieves an accuracy of approximately one tenth of the pixel size.

The alignment results of our method LMPDM are given in the last column of Figure 3. Here, the details are much clearer and the streaks are almost gone. Overall, the reconstructions are lacking some sharpness. Note that in the LMPDM aligned Shepp-Logan image, a shift has occurred with respect to the phantom image. This is because the alignment parameters are invariant to a global shift or rotation of the object. In our error measurements, this global shift and rotation have been removed first.

In Figure 4, the convergence history is shown. The curves show step-wise convergence behavior. This is the result of the multi-resolution approach. At some point, the algorithm cannot improve the parameters at the current resolution. Therefore, a transition to a higher resolution occurs. At the higher resolution, finer details can be resolved and the errors can be reduced further. Note that jumps occur in the residual at these resolution transitions. The residual is not invariant with respect to the image size. Therefore, this behavior is expected and does not indicate a convergence problem. For the Shepp-Logan and particle dataset, we see that the error in ϕ starts to drop at higher resolutions (64×64), while the shifts are refined at all resolution. The reason for this is that the alignment of the projection angles requires details to be present in the reconstruction. The shifts however can align quite well to a low quality image.

The importance of the multi-resolution approach combined with an automatically selected step size in Eq. (8) becomes apparent when the Levenberg-Marquardt routine is used on a single resolution, with fixed step sizes of $h = 10^{-6}$ in Eq. (8). These step sizes have an order of magnitude that is generally considered to give accurate finite differences. The step sizes produced by our spline method are in the order of $h_\delta = 1$ and $h_\phi = 0.1$. The results in Figure 5(c) point out that the alignment parameters are not found, the error increases even. This shows that the proposed methods for multi-resolution and local smoothing of the Jacobian are essential to achieve high accuracy.




Fig. 3. Overview of the results: (a) phantom images of size 256×256 , the Shepp-Logan head phantom, a mandible bone and particles phantom respectively; (b) the unaligned SIRT reconstructions; (c) SIRT reconstructions aligned by cross-correlation; (d) SIRT reconstructions using alignment parameters found by LMPDM

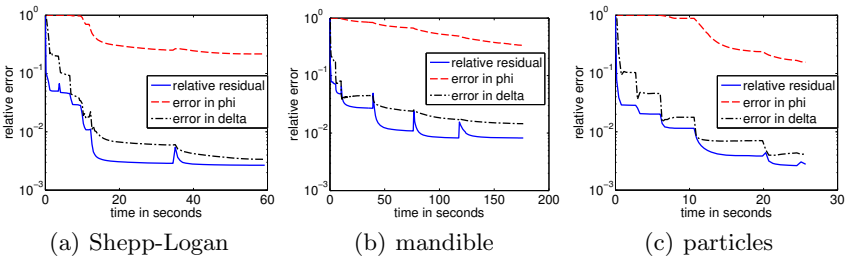


Fig. 4. Convergence history of the simulations using multi-resolution and automatically selected step size. The horizontal axis shows the wall clock time. On the vertical axis, the relative residuals are shown $\|W\mathbf{x} - \mathbf{p}\|_2 / \|\mathbf{p}\|_2$, $\|\phi - \phi_{\text{true}}\|_2 / \|\phi_{\text{true}}\|_2$ and $\|\delta - \delta_{\text{true}}\|_2 / \|\delta_{\text{true}}\|_2$.




Fig. 5. These plots show the difference between the found alignment parameters and their true values for the Shepp-Logan dataset. Since cross-correlation methods cannot find perturbations in angles, the stars in the scatter plot in (b) are the initial values of ϕ . In (c), convergence is shown for LMPDM without using multi-resolution and with step sizes of 10^{-6} for computing the Jacobian with respect to both δ and ϕ .

5 Discussion

From the results we can see that our proposed method, LMPDM, performs well on the selected phantom data. However, a straightforward, naive implementation of Levenberg-Marquardt is bound to fail. The reason for this is the irregular behavior of the objective function (5) due to the single precision code. The methods we have introduced, an automatically selected step size combined with a multi-resolution technique, are sufficient to solve this problem. On the one hand, the improved accuracy of the Jacobian yields more accurate descent directions for LM, which improves convergence. Moreover, multi-resolution helps to find a minimum of (5) even if the perturbations in the parameters are large. Most projection matching alignment algorithms like [10] require an initial coarse alignment if this is the case.

The results from the multi-resolution LMPDM in Figure 4 show, that at low resolutions (16×16 , 32×32), the errors in the shifts decrease rapidly and only a few iterations are needed for convergence. This suggests that there might exist a more optimal selection of the resolutions and their order. Perhaps a multigrid method with an efficient intergrid transfer operator could improve performance.

The run times we have measured, in the order of a minute, show that LMPDM is an efficient method, suitable for experimental datasets for the reconstruction of 2D slices.

6 Conclusions

A new markerless alignment algorithm based on projection matching has been proposed. Using a robust technique to compute the Jacobian combined with a multi-resolution scheme, the accuracy of the LM optimization algorithm can be improved substantially. The resulting LMPDM algorithm performs well even if the perturbations in the alignment parameters are large. The timing results show that the method is efficient enough to be used on 2D experimental datasets.

For future research, it is interesting to generalize the algorithm to 3D, which adds a challenge in computational scale, as well as the added complexity of the geometrical parameters. Also, one can experiment with extra terms in the cost function in Eq. (5), such as prior knowledge, or use other distance measures.

Acknowledgments. This research was supported by the Netherlands Organisation for Scientific Research (NWO), programme 639.072.005.

References

1. Slaney, M., Kak, A.: Principles of computerized tomographic imaging. Society for Industrial and Applied Mathematics (1988)
2. Jing, Z., Sachs, F.: Alignment of tomographic projections using an incomplete set of fiducial markers. *Ultramicroscopy* 35(1), 37–43 (1991)
3. Szeliski, R.: Image alignment and stitching: A tutorial. *Foundations and Trends® in Computer Graphics and Vision* 2(1), 1–104 (2006)
4. Viola, P., Wells III, W.: Alignment by maximization of mutual information. *International Journal of Computer Vision* 24(2), 137–154 (1997)
5. Zitova, B., Flusser, J.: Image registration methods: a survey. *Image and Vision Computing* 21(11), 977–1000 (2003)
6. Frank, J.: *Electron tomography: Three-dimensional imaging with the transmission electron microscope*. Plenum Pub. Corp. (1992)
7. Brandt, S., Heikkonen, J., Engelhardt, P.: Automatic alignment of transmission electron microscope tilt series without fiducial markers. *Journal of Structural Biology* 136(3), 201–213 (2001)
8. Fitchard, E., Aldridge, J., Reckwerdt, P., Mackie, T.: Registration of synthetic tomographic projection data sets using cross-correlation. *Physics in Medicine and Biology* 43(6), 1645 (1999)
9. Houben, L., Bar Sadan, M.: Refinement procedure for the image alignment in high-resolution electron tomography. *Ultramicroscopy* (2011)
10. Parkinson, D., Knoechel, C., Yang, C., Larabell, C., Le Gros, M.: Automatic alignment and reconstruction of images for soft X-ray tomography. *Journal of Structural Biology* (2011)
11. Yang, C., Ng, E., Penczek, P.: Unified 3-D structure and projection orientation refinement using quasi-Newton algorithm. *Journal of Structural Biology* 149(1), 53–64 (2005)
12. Katsevich, A., Silver, M., Zamyatin, A.: Local tomography and the motion estimation problem. *SIAM Journal on Imaging Sciences* 4(1), 200–219 (2011)
13. Kingston, A., Sakellariou, A., Varslot, T., Myers, G., Sheppard, A.: Reliable automatic alignment of tomographic projection data by passive auto-focus. *Medical Physics* 38(9), 4934 (2011)
14. Nocedal, J., Wright, S.: *Numerical optimization*, 2nd edn. Springer (2006)
15. Hyman, J., Larroutourol, B.: The numerical differentiation of discrete functions using polynomial interpolation methods. *Applied Mathematics and Computation* 10, 487–506 (1982)
16. Chartrand, R.: *Numerical differentiation of noisy, nonsmooth data*. Los Alamos National Laboratory 13 (2005)
17. Joseph, P.: An improved algorithm for reprojecting rays through pixel images. *IEEE Transactions on Medical Imaging* 1(3), 192–196 (1982)
18. Fong, D., Saunders, M.: LSMR: An iterative algorithm for sparse least-squares problems. Arxiv preprint arXiv:1006.0758 (2010)