Adaptive Grid Refinement
for Discrete Tomography

Tristan van Leeuwen[1] and K. Joost Batenburg[1,2,3]

[1] Centrum Wiskunde & Informatica, Amsterdam, The Netherlands
[2] Vision-Lab., University of Antwerp, Antwerp, Belgium
[3] Dept. of Mathematics, University of Leiden, Leiden, The Netherlands
{Tristan.van.Leeuwen,kbatenbu}@cwi.nl

Abstract. Discrete tomography has proven itself as a powerful approach to image reconstruction from limited data. In recent years, algebraic reconstruction methods have been applied successfully to a range of experimental data sets. However, the computational cost of such reconstruction techniques currently prevents routine application to large data-sets. In this paper we investigate the use of adaptive refinement on QuadTree grids to reduce the number of pixels (or voxels) needed to represent an image. Such locally refined grids match well with the domain of discrete tomography as they are optimally suited for representing images containing large homogeneous regions. Reducing the number of pixels ultimately promises a reduction in both the computation time of discrete algebraic reconstruction techniques as well as reduced memory requirements. At the same time, a reduction of the number of unknowns can reduce the influence of noise on the reconstruction. The resulting refined grid can be used directly for further post-processing (such as segmentation, feature extraction or metrology). The proposed approach can also be used in a non-adaptive manner for region-of-interest tomography. We present a computational approach for automatic determination of the locations where the grid must be defined. We demonstrate how algebraic discrete tomography algorithms can be constructed based on the QuadTree data structure, resulting in reconstruction methods that are fast, accurate and memory efficient.

Keywords: Tomography, adaptive refinement, QuadTree grids, algebraic reconstruction techniques.

1 Introduction

We consider a linear tomography problem

$$\mathbf{p} = W\mathbf{x} + \mathbf{n}, \tag{1}$$

where $W \in \mathbb{R}^{M \times N}$ is the projection matrix, $\mathbf{x} \in \mathbb{R}^N$ is the image, $\mathbf{p} \in \mathbb{R}^M$ are the projections and $\mathbf{n}$ is additive noise. The goal is to retrieve $\mathbf{x}$ from the noisy projections $\mathbf{p}$, which is typically done by solving a least-squares problem:

$$\min_{\mathbf{x}} ||W\mathbf{x} - \mathbf{p}||_2^2. \tag{2}$$

The system of equations is often underdetermined (i.e., $M < N$) due to the limited number of measurements (small M) and the demand for high-resolution images (large N). The resulting non-uniqueness can be partially mitigated by adding prior knowledge, either in the form of a regularization penalty or by employing a tailored reconstruction algorithm that enforces the prior.

In discrete tomography, the prior is particularly strong – the object consists of only a few different materials – and the reconstruction problem can be formulated as a discrete optimization problem [1,2]. Solving such problems exactly is not feasible for large-scale problems due to their combinatorial nature and often not desirable due to noise. Many heuristic reconstruction have been developed over the years, which fall into two basic classes: methods that aim directly at solving the discrete optimization problem [3,4] and methods that solve (a series) of continuous optimization problems [5,6]. Even state-of-the-art iterative algorithms such as DART [6,7] are computationally very expensive as they are based on iterative reconstruction algorithms. Not only the costs of forward and backward projection and the memory usage scale linearly with N, the number of iterations required is also expected to scale linearly with N. To reduce the computational costs and the required number of iterations of iterative reconstruction methods, many authors have considered multi-scale or multi-grid methods for general, continuous tomography problems. [8] proposes a multi-level strategy that coarsens the projection images by averaging or subsampling the detector pixels. The use of classical multi-grid algorithms is discussed by [9,10,11]. Algorithms of a more heuristic nature are discussed by [12], who propose coarsening in both the image and data space, and [13], who develop a two-level approach. The closest in spirit to the current work is [14] who present an adaptive refinement strategy using QuadTree grids but use a much simpler refinement criterion.

Multi-scale reconstruction approaches aimed specifically at binary tomography have also been proposed. In [3] the authors use a simulated annealing approach in conjunction with uniform refinement. The use of QT grids in a similar context is explored by [4], who proposes refinement of the edges of the object.

In this paper, we investigate the use of QuadTree (QT) grids for iterative image reconstruction in discrete tomography. If the original object consists of large homogeneous regions, each consisting of a single material, QuadTree grids can strongly reduce the number of pixels needed to represent the image. The use of QT grids serves a double purpose in the case of discrete tomography; it can help to regularize the problem and to reduce the computational cost. For the interior of the object, coarse grid pixels can be used, thereby implicitly enforcing the discrete tomography constraint of constant gray levels for these interior regions. As a consequence, even when algorithms for continuous tomography (i.e. allowing all Gray values) are applied to the QT representation, the resulting reconstructions will contain large homogeneous regions and therefore this choice of image representation allows standard iterative methods from continuous tomography to be applied to DT problems successfully.

To illustrate the main ideas we consider the following toy example. A binary phantom and its corresponding (optimal) QT grid are shown in figure 1 (a-b). In this case the phantom consists of $128^2 = 16384$ pixels, while the QT grid allows us to represent the same image with only 25 pixels. To illustrate the potential benefits, we assume for the moment that we know the optimal QT grid and use it for reconstruction. We consider three scenarios: *i)* A benchmark reconstruction with 32 angles between 0 and 180 degrees and no noise; *ii)* a reconstruction with 32 angles and 20% Gaussian noise and finally, *iii)* a reconstruction with only 5 angles and no noise. The results are shown in figure 2. These examples clearly illustrate the potential benefits of reducing the number of unknowns; the results are more stable with respect to noise and it allows us to recover from severely limited data *using a conventional algebraic reconstruction algorithm.* Moreover, since the computational cost of the forward projection and the required memory is proportional to the number of pixels, using QT grids may also lead to significant computational savings.

Of course, we do not know the optimal QT grid a-priori in practice. Therefore, we propose adaptive refinement strategy that allows us to construct a QT grid as part of the iterative reconstruction. By starting from a coarse grid and refining only in areas of high variability, we never introduce more unknowns than needed and are able to construct an efficient representation of the reconstruction directly from the projection data. We apply the proposed method on 3 (binary) phantoms.

The outline of the paper is as follows. First, we discuss multi-level reconstruction and adaptive refinement in section 3. Numerical experiments where we apply the standard SIRT algorithm for continuous tomography to discrete image reconstruction on a QT grid are presented in section 4. Finally, we present conclusions and discuss possible future extensions as well as open questions in section 6.
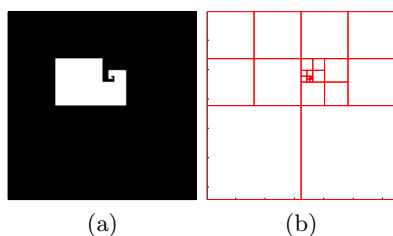

(a)   (b)

Fig. 1. (a) Spiral phantom and (b) corresponding QuadTree grid

2   Algorithm

We represent the image as a piece-wise constant function on a QuadTree grid. An example of a QT grid is shown in figure 3. A QT grid is represented by a collection of triples (i, j, s) which store the location of the upper-left corner of each cell as well as its size, both w.r.t. an underlying fine grid. An image on this

32 angles, no noise – 32 angles, 20% Gaussian noise – 5 angles, no noise
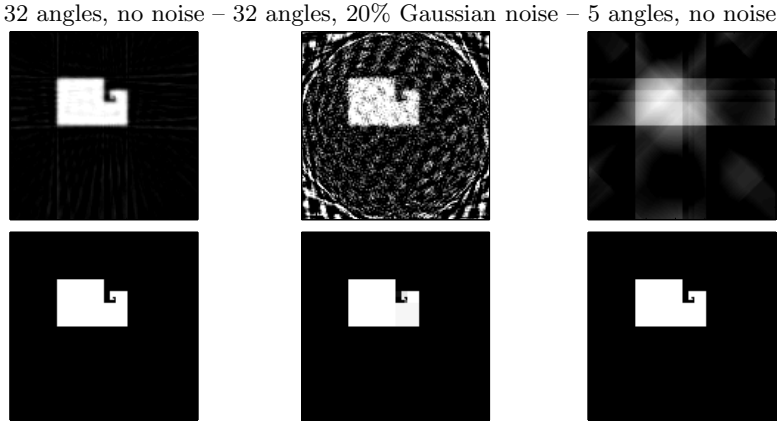
Fig. 2. Reconstructions of the phantom depicted in figure 1 (a) for different scenarios. The top row shows reconstructions on a fine grid with 16384 pixels while the bottom row the results for the reconstructions on the QuadTree grid with 25 pixels (cf. figure 1 (b)).

grid is represented with a single number for each cell. Although it is in principle possible to work with the image directly in this representation, is often more convenient to work with images that are represented on a uniform fine grid. For this purpose, we introduce the *mapping* matrix V that maps from a given QT grid to the underlying fine grid. For the QT grid depicted in figure 3 this matrix is given by

$$V^T = \frac{1}{2} \begin{pmatrix} 1\,1\, \cdot\, \cdot\, 1\,1\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot \\ \cdot\, \cdot\, 1\,1\, \cdot\, \cdot\, 1\,1\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot \\ \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, 2\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot \\ \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, 2\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot \\ \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, 1\,1\, \cdot\, \cdot\, 1\,1 \\ \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, 2\, \cdot\, \cdot\, \cdot \\ \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, \cdot\, 2\, \cdot\, \cdot \end{pmatrix}. \tag{3}$$

Here, the underlying fine grid has 4×4 pixels and dictates the finest level of the QT grid. The columns of V represent the cells of the QT grid and couple the corresponding cells in the fine grid. Note that these matrices are normalized such that for any given image $\mathbf{x}$ on the fine grid $V^T\mathbf{x}$ is the best approximation in the euclidean norm of this image on the QT grid.

For a given QT grid, we can pose the reconstruction problem as

$$\min_{\mathbf{x}} ||WV\mathbf{x} - \mathbf{p}||_2^2,$$

where $\mathbf{x}$ represents the image on the QT grid defined by V. The resulting reconstruction problem can be solved with any conventional reconstruction algorithm,

such as SIRT or ART. If the QT grid has fewer cells than measurements, this reconstruction problem is overdetermined and much better posed than the original reconstruction problem.
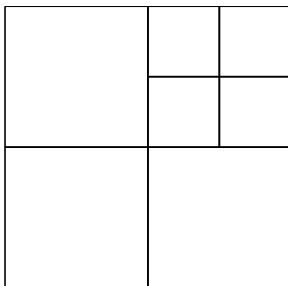
Fig. 3. Example of a QT grid

2.1 Adaptive Refinement

To construct a QT representation of a given image with as few cells as possible we propose the following refinement procedure, starting from an initial coarse grid:

1. refine all cells on the finest level;
2. compute the error between the current reconstruction and the reconstruction on the refined grid;
3. for each refined cell, keep the refinement if the local error is bigger than some threshold, coarsen otherwise;

This approach is different from traditional adaptive refinement strategies which typically refine after the fact based on local image gradients [14]. A problem with such approaches is that the image gradient needs to be estimated on the current (coarse) grid, making it difficult to detect features that are not properly resolved on this grid. Our procedure circumvents this problem by refining before measuring the error and reverting back to the coarse grid if the difference is small.

A more detailed description this procedure is shown in Algorithm 1. Here, $\mathsf{refine}(V_k)$ refines all cells on the finest level by splitting them into 4. $\mathsf{refine}(V, \mathcal{I})$ refines only the cells in the index set $\mathcal{I}$. By refining only the cells on the finest level, we avoid having to refine the same cells over and over again. The algorithm automatically terminates when we have reached the finest level corresponding to the underlying fine grid. So, for an underlying fine grid with $N = n^2$ pixels, we have a total of $K = \log_2(n)$ levels.

We compute the difference between the reconstructions on the coarse and refined grids V_c and V_f as

$$\Delta \mathbf{x}_f = \mathbf{x}_f - V_f^T V_c \mathbf{x}_c.$$

Algorithm 1. Adaptive refinement algorithm

Require:
   x - input image
   V_0 - basis for initial subspace
   δ - tolerance for refinement
Ensure:
   $\mathbf{x}_K$ - final representation on QT grid
   V_K - corresponding matrix

   $\mathbf{x}_0 = V_0^T \mathbf{x}$ {initial reconstruction}
   for $k = 0$ **to** $K - 1$ **do**
      $\widetilde{V} = \mathsf{refine}(V_k, \mathcal{I}_k)$ {refinement proposal}
      $\widetilde{\mathbf{x}} = \widetilde{V}^T \mathbf{x}$ {map onto refined grid}
      $\widetilde{\mathbf{e}} = \mathsf{error}(\widetilde{\mathbf{x}} - \widetilde{V}^T V_k \mathbf{x}_k)$ {compute error}
      $\widetilde{\mathcal{I}} = \{i \,|\, \widetilde{e}_i > \delta\}$
      $V_{k+1} = \mathsf{refine}(V_k, \widetilde{\mathcal{I}})$ {update}
      $\mathbf{x}_{k+1} = V_{k+1}^T \widetilde{V} \widetilde{\mathbf{x}}$
   end for

We then define a quantity $\mathbf{e}_c = \mathsf{error}(\Delta \mathbf{x}_f)$ on the coarse grid that contains the accumulated contributions for each grid cell such that

$$\mathbf{e}_c^T \mathbf{1} = \|\Delta \mathbf{x}_f\|_2^2. \tag{4}$$

All cells i on the coarse grid for which $e_{c,i} > \delta$ are subsequently refined. Thus, when the algorithm terminates (i.e., when $e_{c,i} \leq \delta \,\forall i$) we have $\|\Delta \mathbf{x}_f\|_2^2 \leq \delta N_c$ where N_c is the number of cells in the coarse grid.

An example of a series of adaptively refined grids for the Shepp-Logan phantom is shown in figure 4. On the finest level, we perfectly reconstruct the original image with only 1948 cells (compared to 16384 for the original image).

2.2   Reconstruction

We can adapt the above described refinement algorithm for reconstruction by replacing the mapping of the true image onto the refined QT grids in Algorithm 1 by a mapping of the projection data onto the refined QT grid. This can be achieved by

$$\widetilde{\mathbf{x}} = \left(W\widetilde{V}\right)^\dagger \mathbf{p},$$

where † denotes the pseudo-inverse. In practice, we never compute the pseudo inverse explicitly, but instead perform a tomographic reconstruction on the refined grid using the previous iterate $\mathbf{x}_k$ as initial guess. The resulting algorithm is stated in Algorithm 2. Here, $\mathbf{x}_1 = \mathsf{reconstruction}(W, \mathbf{p}, \mathbf{x}_0, L, \epsilon)$ performs up to L iterations of an iterative reconstruction technique starting from initial guess $\mathbf{x}_0$ with stopping criterion $\|W\mathbf{x}_1 - \mathbf{p}\|_2 \leq \epsilon \|\mathbf{p}\|_2$.
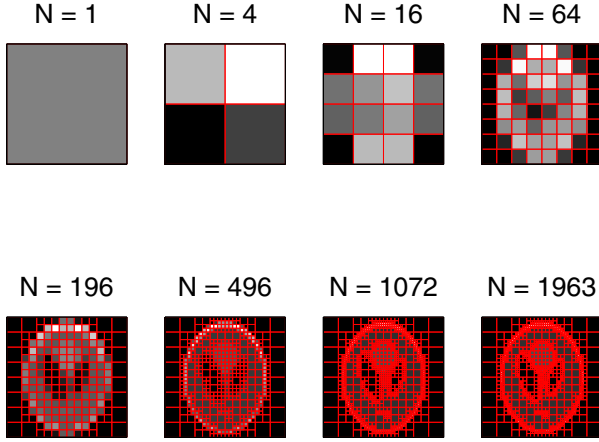
Fig. 4. Example of adaptive refinement

Algorithm 2. Adaptive multi-scale reconstruction algorithm

Require:
W - projection operator
p - projection data
V_0 - basis for initial subspace
L - iteration count for iterative reconstruction
ϵ - tolerance for iterative reconstruction
δ - tolerance for refinement

Ensure:
$\mathbf{x}_K$ - final reconstruction

$\mathbf{x}_0 = \mathsf{reconstruction}(WV_0, \mathbf{p}, 0, L, \epsilon)\{\text{initial reconstruction}\}$
for $k = 0$ **to** $K - 1$ **do**
$\quad \widetilde{V} = \mathsf{refine}(V_k, \mathcal{I}_k)$ {refinement proposal}
$\quad \widetilde{\mathbf{x}} = \mathsf{reconstruction}(W\widetilde{V}, \mathbf{p}, \widetilde{V}^T V_k \mathbf{x}_k, L, \epsilon)$ {new reconstruction}
$\quad \widetilde{\mathbf{e}} = \mathsf{error}(\widetilde{\mathbf{x}} - \widetilde{V}^T V_k \mathbf{x}_k)$ {compute error}
$\quad \widetilde{\mathcal{I}} = \{i \,|\, \widetilde{e}_i > \delta\}$
$\quad V_{k+1} = \mathsf{refine}(V_k, \widetilde{\mathcal{I}})$ {update}
$\quad \mathbf{x}_{k+1} = V_{k+1}^T \widetilde{V} \widetilde{\mathbf{x}}$
end for

---

3   Numerical Results

We conduct numerical experiments on three phantoms. For the phantoms, the projection data is generated on a 256×256 grid with 128 detectors and 64 projections. The reconstruction is done on an underlying fine grid of 128×128 in order to avoid the *inverse crime*. We use the ASTRA toolbox to compute the forward and backward projections [15]. For the adaptive method the mapping matrices as discussed above are used to map to and from the QT grids to the

(a)

n = 16384

n = 2203
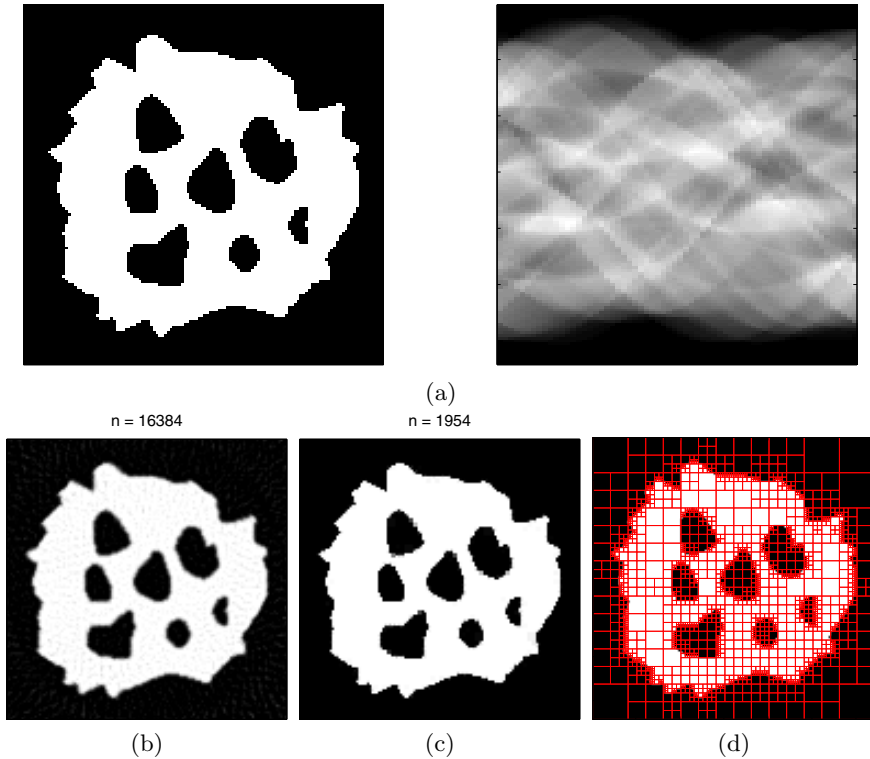

(b)   (c)   (d)
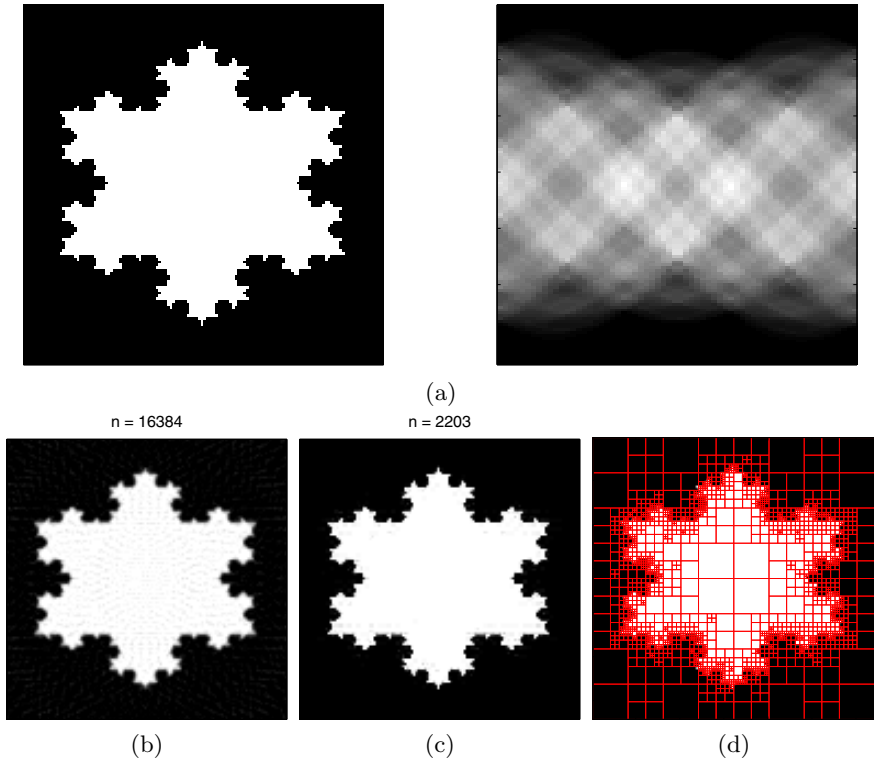
Fig. 5. (a) Ground truth and corresponding projection data, (b) SIRT reconstruction, (c) multi-scale reconstruction and (d) corresponding QT grid overlaying the ground truth

n = 16384   n = 2605

(a)

(b)   (c)   (d)
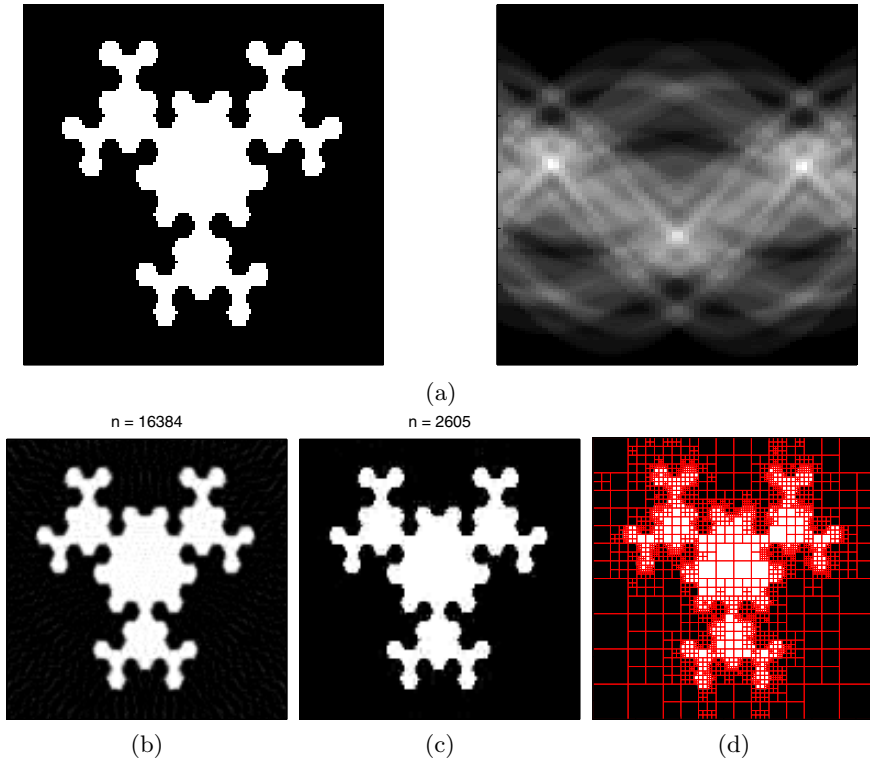
Fig. 6. (a) Ground truth and corresponding projection data, (b) SIRT reconstruction, (c) multi-scale reconstruction and (d) corresponding QT grid overlaying the ground truth

(a)

n = 16384   n = 1954

(b)   (c)   (d)

Fig. 7. (a) Ground truth and corresponding projection data, (b) SIRT reconstruction, (c) multi-scale reconstruction and (d) corresponding QT grid overlaying the ground truth

underlying fine grid. As reconstruction algorithm, we use SIRT with $L = 200$ and $\epsilon = 10^{-3}$. For the refinement we use a tolerance of $\delta = 0.2$. For comparison we also show the result obtained when applying SIRT directly on the finest grid.

The results on the phantoms are shown in figures 5, 6 and 7. These results show that the proposed refinement successfully detects areas of high variability and is able to capture most of the fine detail in the phantoms while using large cells in homogeneous areas. The resulting reconstructions are (almost) binary, showing the regularizing properties of the QT grid.

4   Conclusions and Discussion

We have presented an adaptive refinement strategy for tomographic reconstruction on QuadTree grids. The algorithm starts from a coarse grid and adaptively refines those cells where the reconstruction error is above some threshold. If we combine the QT grid approach with the standard SIRT algorithm for continuous tomography, and apply it to discrete images, the resulting grid represents the reconstructed image with only a fraction of the number of pixels otherwise

required. We expect this approach to be useful in a wide range of applications where high resolution is required and where the images are characterized by large homogeneous regions, which is typically the case in discrete tomography. We also envision that QuadTree grids will be useful for region-of-interest tomography. In this case, the QT grids can be reconstructed a-priori based on a simple FBP reconstruction to identify regions of interest.

To optimally benefit from the reduction of the number of pixels when using QuadTree grids, the projection operator will have to compute the projections directly based on the QuadTree representation of the image, without mapping to an underlying fine grid first (as we did in this paper). Since the cost of forward projection is proportional to the number of pixels, this would directly reduce the computations by an order of magnitude. For reconstruction, we expect that having less unknowns will lead to less iterations, promising another reduction of the computational cost. An extension of the proposed algorithm to 3D reconstruction using OcTree grids is straightforward. Future research is aimed at including regularization to explicitly enforce desirable properties (such as discreteness) on the reconstruction and application to continuous tomography. For the latter, we expect that moving away from the piece-wise constant representation (e.g., by using linear basis functions on the QT grid) will be beneficial.

References

1. Herman, G.T., Kuba, A.: Discrete Tomography: Foundations, Algorithms, and Applications. Birkhäuser (1999)
2. Herman, G.T., Kuba, A.: Advances in discrete tomography and its applications. Birkhäuser (2007)
3. Ruskó, L., Kuba, A.: Multi-resolution method for binary tomography. Electronic Notes in Discrete Mathematics 20, 299–311 (2005)
4. Gerard, Y.: Elementary algorithms for multiresolution geometric tomography with strip model of projections. In: 2013 8th International Symposium on Image and Signal Processing and Analysis (ISPA), pp. 600–605. University of Trieste and University of Zagreb (2013)
5. Schule, T., Schnorr, C., Weber, S., Hornegger, J.: Discrete tomography by convex–concave regularization and D.C. programming. Discrete Applied Mathematics 151(1-3), 229–243 (2005)
6. Batenburg, K.J., Sijbers, J.: DART: a practical reconstruction algorithm for discrete tomography. IEEE Transactions on Image Processing 20(9), 2542–2553 (2011)
7. Batenburg, K., Sijbers, J.: Dart: A Fast Heuristic Algebraic Reconstruction Algorithm for Discrete Tomography. In: IEEE International Conference on Image Processing, pp. IV-133–IV-136. IEEE (2007)
8. Herman, G.T., Levkowitz, H., Tuy, H.: Multilevel Image Reconstruction. In: Rosenfeld, A. (ed.) Multiresolution Image Processing and Analysis. Springer Series in Information Sciences, vol. 12, pp. 121–135. Springer, Heidelberg (1984)

9. Henson, V.E., Limber, M.A., McCormick, S.F., Robinson, B.T.: Multilevel Image Reconstruction with Natural Pixels. SIAM Journal on Scientific Computing 17(1), 193–216 (1996)
10. Kostler, H., Popa, C., Ummer, M., Rude, U.: Towards an algebraic multigrid method for tomographic image reconstruction-improving convergence of ART. In: European Conference on Computational Fluid Dynamics, pp. 1–12 (2006)
11. Cools, S., Ghysels, P., van Aarle, W., Vanroose, W.: A multilevel preconditioned Krylov method for algebraic tomographic reconstruction. arXiv, 26 (2013)
12. Bouman, C., Webb, K.: Multigrid tomographic inversion with variable resolution data and image spaces. IEEE Transactions on Image Processing 15(9), 2805–2819 (2006)
13. De Witte, Y., Vlassenbroeck, J., Van Hoorebeke, L.: A multiresolution approach to iterative reconstruction algorithms in X-ray computed tomography. IEEE Transactions on Image Processing: A Publication of the IEEE Signal Processing Society 19(9), 2419–2427 (2010)
14. Schumacher, H., Heldmann, S., Haber, E., Fischer, B.: Iterative Reconstruction of SPECT Images Using Adaptive Multi-level Refinement. In: Tolxdorff, T., Braun, J., Deserno, T.M., Horsch, A., Handels, H., Meinzer, H.P. (eds.) Bildverarbeitung für die Medizin 2008, pp. 318–322. Springer, Heidelberg (2008)
15. Palenstijn, W.J., Batenburg, K.J., Sijbers, J.: Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs). Journal of Structural Biology 176(2), 250–253 (2011)