Published in final edited form as:
Phys Med Biol. 2016 October 21; 61(20): 7263–7281. doi:10.1088/0031-9155/61/20/7263.

Multiresolution Iterative Reconstruction in High-Resolution Extremity Cone-Beam CT

Qian Cao[1], **Wojciech Zbijewski**[1], **Alejandro Sisniega**[1], **John Yorkston**[2], **Jeffrey H Siewerdsen**[1,3], and **J Webster Stayman**[1]

J Webster Stayman: web.stayman@jhu.edu

[1]Department of Biomedical Engineering, Johns Hopkins University, Baltimore MD, 21205, USA

[2]Carestream Health, Rochester, NY, 14608, USA

[3]Russell H Morgan Department of Radiology, Johns Hopkins University, Baltimore, 21205, USA

Abstract

Application of model-based iterative reconstruction (MBIR) to high resolution cone-beam CT (CBCT) is computationally challenging because of the very fine discretization (voxel size <100 μm) of the reconstructed volume. Moreover, standard MBIR techniques require that the complete transaxial support for the acquired projections is reconstructed, thus precluding acceleration by restricting the reconstruction to a region-of-interest. To reduce the computational burden of high resolution MBIR, we propose a multiresolution Penalized-Weighted Least Squares (PWLS) algorithm, where the volume is parameterized as a union of fine and coarse voxel grids as well as selective binning of detector pixels. We introduce a penalty function designed to regularize across the boundaries between the two grids. The algorithm was evaluated in simulation studies emulating an extremity CBCT system and in a physical study on a test-bench. Artifacts arising from the mismatched discretization of the fine and coarse sub-volumes were investigated. The fine grid region was parameterized using 0.15 mm voxels and the voxel size in the coarse grid region was varied by changing a downsampling factor. No significant artifacts were found in either of the regions for downsampling factors of up to 4×. For a typical extremities CBCT volume size, this downsampling corresponds to an acceleration of the reconstruction that is more than five times faster than a brute force solution that applies fine voxel parameterization to the entire volume. For certain configurations of the coarse and fine grid regions, in particular when the boundary between the regions does not cross high attenuation gradients, downsampling factors as high as 10× can be used without introducing artifacts, yielding a ~50× speedup in PWLS. The proposed multiresolution algorithm significantly reduces the computational burden of high resolution iterative CBCT reconstruction and can be extended to other applications of MBIR where computationally expensive, high-fidelity forward models are applied only to a sub-region of the field-of-view.

1. Introduction

Modern cone-beam CT (CBCT) systems utilize a variety of digital x-ray detectors, such as amorphous silicon flat panel detectors (a-Si FPDs), charge-coupled devices (CCDs) and complementary metal-oxide-semiconductors (CMOS). These sensors offer very fine pixel sizes (150 μm and less), enabling high resolution clinical and preclinical imaging at a level of detail not accessible to conventional CT. Examples of emerging high resolution applications include: (i) small animal imaging (Clark and Badea 2014, Boerckel *et al* 2014), (ii) breast CBCT (Boone *et al* 2001, Yang *et al* 2007, O'Connell *et al* 2010, Shen *et al* 2013), in particular for detection of micro-calcifications, (iii) otolaryngology, dental, and head and neck imaging (De Vos *et al* 2009, Hodez *et al* 2011, Xu *et al* 2012), and (iv) extremities imaging (Tuominen *et al* 2013, Nishiyama and Shane 2013, Carrino *et al* 2014, Marinetto *et al* 2016), with potential use in quantitative assessment of bone microarchitecture.

The small size of the sensing elements in high resolution detectors conventionally necessitates an increase in imaging dose to maintain the contrast-to-noise ratio sufficient for clinical tasks. This increase in patient exposure may be mitigated by model-based iterative reconstruction (MBIR), which was previously shown to yield improved resolution-noise tradeoffs as compared to conventional FBP at similar or even lower dose levels (Thibault *et al* 2007, Wang *et al* 2014). Application of MBIR to high resolution CBCT imaging is however challenged by the computational burden of iterative algorithms. Whereas analytical methods typically require only one backprojection operation, each iteration of MBIR generally requires at least one forward and one backprojection. Furthermore, since MBIR relies on a data fidelity term that matches image estimates to the measured data, a complete transaxial region of support for the image volume must be reconstructed to capture all contributions from the object for each ray path. Thus, region-of-interest (ROI) reconstructions that are often straightforward to implement using analytical methods cannot be realized in standard iterative algorithms. The requirement to reconstruct the entire field of view (FOV), typically as large as $20\times20\times20\ \mathrm{cm}^3$, using very fine voxels ($\leq 75\ \mu m$), can make the application of MBIR to high resolution imaging prohibitively slow because the speed of the projection and backprojection operators is generally proportional to the number of voxels.

Projection operators can be accelerated by optimized implementation on parallel hardware using graphical processing units (GPUs) (Keck and Hofmann 2009, Brokish *et al* 2010, Papenhausen *et al* 2011, Pratx and Xing 2011) or distributed computing (Liu *et al* 2013). However, for large-FOV high resolution CBCT data, both approaches suffer from latency of data transfer – either between host computer and GPU or between computational nodes in a distributed network. Furthermore, while such methods improve the baseline speed of the projection operator, the performance gains remain dependent on the size of the FOV and resolution. The benefits of parallelization diminish when the number of required processes exceeds the number of available computational nodes due to e.g. very fine discretization of the FOV.

One approach to overcome the decreased performance of MBIR in high resolution applications relies on the observation that while iterative reconstruction of an isolated high resolution ROI is not possible, the model of the volume can be altered so that a fine parameterization is only used in the ROI. Some versions of such multiresolution schemes forego the traditional square voxel basis functions and employ a sparse representation of the volume using heterogeneous voxels, wavelets (Frese *et al* 2002, Degirmenci *et al* 2015) or meshes (Brankov *et al* 2004, Sitek *et al* 2006). Here we will focus on traditional voxel bases, where multiresolution reconstruction can be realized by using voxels of different size to parameterize different subvolumes (i.e. coarse grid regions and fine grid regions) in the image. These bases are straightforward and map well to physical memory in computing hardware.

Early examples of multiresolution methods include those that enable application of the iterative algorithm to a limited internal high resolution ROI and use an analytical method for the surrounding volume. The full FOV is first reconstructed with an analytical method. The ROI subvolume is set to zero and the surrounding volume is reprojected. The resulting sinogram is subsequently subtracted from the measured data to isolate the line integrals through the ROI. The line integrals of the ROI are then iteratively reconstructed (Ziegler *et al* 2008, Stearns *et al* 2007). An additional step consisting of smoothing of the reprojected coarse grid line integrals was found to yield reduced noise in the ROI reconstruction (Rashed *et al* 2007). However, artifacts attributable to edge degradation and beam-hardening artifacts in the initial FBP were found in the iteratively reconstructed ROI obtained using this general approach (Hamelin *et al* 2010).

A multiresolution approach where subvolumes with different levels of discretization are reconstructed using an iterative algorithm was proposed in (La Rivière 2007). The fine and coarse grids are jointly reconstructed using a maximum-likelihood (ML) algorithm applied in an alternating scheme, updating one grid while holding the other fixed. In contrast to this alternating optimization, simultaneous optimization of both coarse and fine grids using general purpose optimization algorithms was developed in (Hamelin *et al* 2007, Yu *et al* 2009). Similar to our work, this approach relies on a natural factorization of the line integrals into the projections of the fine and coarse grids, yielding a joint objective function for the entire volume. This formulation was also applied to reduce metal artifacts by applying high fidelity forward models (finer parameterization of metallic edges and polychromatic beam modeling) only in ROIs containing metallic objects (Van Slambrouck and Nuyts 2010, 2012).

This work investigates the application of multiresolution reconstruction in the context of extremity CBCT, in particular for visualization and quantitative assessment of bone microarchitecture. In this application, the fine trabecular detail (≤ 75 μ*m*) needs to be reconstructed at very high resolution only over a relatively small ROI, covering e.g. the subchondral bone in the knee joint. The surrounding anatomy (cartilage, ligaments, muscle, fat, etc.) can be reconstructed at lower resolution over the larger FOV. We introduce a new multiresolution MBIR algorithm based on the penalized-weighted least squares (PWLS) framework (Sauer and Bouman 1993) with a separable paraboloidal surrogates (SPS) optimizer (Erdoğan and Fessler 1999b, Elbakri and Fessler 2002). We propose a penalty function specifically designed to provide regularization across the boundaries between fine and coarse voxel grids. We extend the multiresolution framework to include the possibility of regional binning of the projection images, resulting in variable pixel size in addition to the variable voxel size. We focus our evaluation on the tradeoffs between artifacts in the fine grid ROI and voxel size and regularization applied in the surrounding coarse grid region.

2. Methods

2.1. Multiresolution PWLS

The conventional forward model for a transmission tomography system, assuming a voxelized object representation, is given by:

$$\bar{y}_l = g_i \exp(-[\boldsymbol{A}\mu]_{\boldsymbol{i}}) \quad (1)$$

where $\bar{y}_i$ denotes the mean of the i^{th} measurement, g_i represents the system gain (including bare beam fluence and detector sensitivity for measurement i), μ is a vector of voxel attenuation values, and $\boldsymbol{A}$ is the forward projection operator. To formulate an objective function for model-based reconstruction, a noise model needs to be chosen. For quantum-limited projection data with negligible contribution of electronic noise, Poisson noise is often assumed. For systems with non-negligible contributions of electronic noise (such as the FPDs used in extremity CBCT) and when the projection data is post-processed with artifact correction algorithms prior to reconstruction (e.g. beam hardening and scatter correction), a more general Gaussian noise model is often more appropriate, leading to a weighted-least-squares objective. We adopt the Gaussian model and PWLS in this work as better suited for the statistics of FPD CBCT projection data. However, the multiresolution methodology presented here could be incorporated into any forward model and easily generalized to MBIR algorithms based on the Poisson noise model. PWLS reconstruction (Sauer and Bouman 1993) employs the following objective function:

$$\hat{\mu}_{PWLS} = \underset{\mu>0}{\operatorname{argmin}} \frac{1}{2} \|\boldsymbol{A}\mu - l\|_{\boldsymbol{W}}^2 + \beta R(\mu) = \underset{\mu>0}{\operatorname{argmin}} \frac{1}{2} \|\boldsymbol{A}\mu - l\|_{\boldsymbol{W}}^2 + \beta \sum_{k \in K} \psi([\boldsymbol{C}\mu]_k) \quad (2)$$

where l denotes the vector of line integrals $l_i = -log\, y_i / g_i$. The traditional choice for $\boldsymbol{W}$ is a diagonal weighting by the inverse variance, approximated using the measurements: $w_i = y_i$. The second term (R) is a roughness penalty specified by the finite differencing operator $\boldsymbol{C}$ (Fessler 1997) and the potential function ψ. Here, $\boldsymbol{C}$ is a matrix and applies the 6 nearest-neighbor pairwise voxel difference (excluding diagonal neighbors). K is the total number of pairwise differences in the volume (K ~6 × number of voxels) and ψ is the quadratic penalty.

The above system model and reconstruction objective is general and permits many options for the parameterization of the image volume. We will use this framework and define forward projection and differencing operators for a multiresolution representation of the object. Specifically, we choose to represent the total volume μ as a union of a set of fine grid voxels μ_F and a set of coarse grid voxels μ_C:

$$\mu = \begin{bmatrix} \mu_C \\ \mu_F \end{bmatrix} \quad (3)$$

The ratio of the coarse grid voxel size to the fine grid voxel size defines a downsampling factor η_μ. The linearity of the forward projection operator permits definition of the following multiresolution projector $\boldsymbol{A}$, consisting of coarse and fine forward projectors $\boldsymbol{A}_C$ and $\boldsymbol{A}_F$:

$$\boldsymbol{A}\mu=[\boldsymbol{A}_C \quad \boldsymbol{A}_F]\begin{bmatrix}\mu_C\\ \mu_F\end{bmatrix} \quad (4)$$

Integrating (3) and (4) into (1), we may rewrite the forward model as:

$$\bar{y}=g\circ\exp(-\boldsymbol{A\mu})=g\circ\exp(-\boldsymbol{A}_C\mu_C)\circ\exp(-\boldsymbol{A}_F\mu_F) \quad (5)$$

where operator ◦ denotes the Hadamard (element-wise) matrix product. This forward model provides a convenient mathematical form for the development of iterative reconstruction algorithms that treat the reconstructed volume as a sum of non-overlapping regions, each with its own distinct forward model. In this work, we focus on forward models with different discretization of the image volume; however, this kind of decomposition has also been used in reconstruction methods that parameterize the object into known foreground components and an unknown background (Stayman *et al* 2012) for reconstruction in the presence of objects known to be in the FOV (eg, high-density surgical implants).

In addition to the large memory footprint of the reconstructed volume, projection datasets acquired with high-resolution flat-panel detectors can also be prohibitively large (up to ~3000×3000 pixels and ~1000 projections). Thus, we extend the multiresolution approach to include multiresolution binning of projection data. The binning ratio η_y is defined as the ratio of the coarse pixel size to fine pixel size. As illustrated in Figure 1, we assume that the binning is applied (via appropriate pre-processing) in such a manner that the fine grid regions of the volume are projected entirely onto the fine, natively sampled (unbinned) regions of the projection image. The forward projector now comprises $\boldsymbol{A}_{FN}$ (fine grid voxels to native pixels), $\boldsymbol{A}_{CN}$ (coarse gird voxels to native pixels) and $\boldsymbol{A}_{CB}$ (coarse grid voxels to binned pixels):

$$\hat{l}=\begin{bmatrix}\hat{l}_B\\ \hat{l}_N\end{bmatrix}=\boldsymbol{A}\hat{\mu}=\begin{bmatrix}\boldsymbol{A}_{CB} & 0\\ \boldsymbol{A}_{CN} & \boldsymbol{A}_{FN}\end{bmatrix}\begin{bmatrix}\hat{\mu}_C\\ \hat{\mu}_F\end{bmatrix} \quad (6)$$

where l_N and l_B denote the line integrals associated with native and binned pixels, respectively.

Using the above multiresolution definitions, the data fidelity term in equation (2) is nearly complete. It remains to define the weighting matrix:

$$\boldsymbol{W}=\begin{bmatrix}\boldsymbol{W}_B & 0\\ 0 & \boldsymbol{W}_N\end{bmatrix} \quad (7)$$

with the corresponding diagonal inverse variance weights, $\boldsymbol{W}_B$ and $\boldsymbol{W}_N$, which represent the measurement weights of the binned and native pixels, respectively. Such weighting can accommodate, for example, reduced noise in the binned measurements.

Implementation of the regularization term in the multiresolution PWLS of equation (2) requires specification of how the differencing operator $\boldsymbol{C}$ acts across the boundary between fine and coarse grid voxels. We propose the following general multiresolution penalty:

$$\tilde{\mu}_c = \begin{bmatrix} \mu_C \\ \boldsymbol{D}\mu_F \end{bmatrix} \quad \tilde{\mu}_F \begin{bmatrix} \boldsymbol{U}\mu_C \\ \mu_F \end{bmatrix}$$

$$R(\mu)=\beta_C\sum_{k\in K_C}\psi([\boldsymbol{C_C}\tilde{\mu}_c]_k)+\beta_F\sum_{k\in K_F}\psi([\boldsymbol{C_F}\tilde{\mu}_F]_k)=\beta_C R_C+\beta_F R_F \qquad (8)$$

where we have defined augmented coarse and fine image vectors, $\tilde{\mu}_c$ and $\tilde{\mu}_F$, respectively. These vectors include boundary voxel values from the adjacent ROI that have been resampled to the appropriate voxel grid. That is, we use $\boldsymbol{D}$, an interpolating operator that downsamples fine grid voxels neighboring the μ_F/μ_C boundary to coarse grid voxels, to allow application of a coarse pairwise voxel difference operator (Fessler 1997), $\boldsymbol{C_c}$, across the boundary of the coarse and fine ROIs (Figure 2). Alternately, we use $\boldsymbol{U}$, an operator that upsamples the coarse grid voxels on the μ_F/μ_C interface, to apply a fine pairwise voxel difference operator, $\boldsymbol{C_F}$, across the ROI boundary in the opposite direction. In this work, we choose $\boldsymbol{D}$ to be a binning operation that computes the mean over fine voxels, and $\boldsymbol{U}$ to upsample from coarse to fine voxels using linear interpolation. We have two regularization parameters, β_F and β_C, that control the penalty strengths for the fine and coarse grid ROIs, respectively. The sets, K_C and K_F, are rows of $\widehat{\widetilde{C}\mu}$ which enumerate all pairwise voxel differences contributing to the penalty terms for the coarse and fine ROIs, respectively.

Eq. (8) amounts to applying the conventional penalty to a region (fine or coarse grid) that was virtually expanded to include boundary voxels from the other region, which were appropriately resampled using interpolation. This procedure is general and can be applied to arbitrary boundaries, except perhaps for highly degenerate cases were interpolation is not applicable (e.g. ROIs consisting of a single voxel).

The objective function in (2) can be optimized using the SPS approach (Erdoğan and Fessler 1999b, Elbakri and Fessler 2002). We note that the same data-fit surrogate as in Eq. 17 of Elbakri and Fessler (2002) can be used here, with the standard projector operator replaced by the multiresolution projector. Crucially, the surrogate in the SPS algorithm is separable across voxels, i.e. the update for each voxel depends on the surrounding volume only through the line integrals of the previous iterate of the reconstruction. This yields a multiresolution algorithm that allows independent updates of the fine- and coarse-grid ROIs, as detailed in Table 1. Similarly, the separable surrogate for the penalty of (Erdoğan and Fessler 1999a) can be applied directly to the multiresolution objective, as the construction of the surrogate is independent of discretization. The algorithm presented in Table 1 can be easily generalized to an ordered subset form in a manner analogous to (Erdoğan and Fessler 1999b).

2.2. Experimental Setup for Simulation Studies

The proposed multiresolution approach was evaluated in a simulation study. The primary goal of the study was to evaluate the impact of different coarse voxel downsampling factors and the effect of regularization strength in the coarse and fine grid volumes on image artifacts in the fine grid region. Fig. 3 shows the digital phantom used in simulations. The phantom included soft tissues (emulating fat and muscle) and two bone regions. The bones consisted of a cortical shell and fine trabecular features based on a segmented micro-CT image. The attenuation values of the simulated tissues corresponded to effective beam energy of 60 keV, which approximates that of the 90 kVp (4 mm Al, 0.3mm Cu added filtration) spectrum of a dedicated extremity CBCT at our institution (Siewerdsen *et al* 2004, Zbijewski *et al* 2011). The digital phantom used 0.075 mm voxels and simulated projections were generated on a 0.194 mm pixel grid. A bare beam fluence of 10^5 photons per detector pixel was assumed and Poisson noise was added to the projections. Simulated detector pixels had 100% efficiency. System geometry emulated the extremities CBCT scanner: source-axis distance (SAD) of 436 mm and source-detector distance (SDD) of 560 mm.

In the multiresolution reconstruction, the regularization strength of the interior fine grid ROI β_F was fixed at 10^3 based on visual assessment of spatial resolution. 60 subsets were used. The reconstructed voxel size in the fine grid ROI was 0.15 mm, whereas the voxel size in the coarse grid region was varied by changing the downsampling factor η_μ from 1 (i.e., the same voxel size for μ_F and μ_C) to 10 (i.e., 10× larger voxels for μ_C). In cases where detector binning was considered in addition to a multiresolution volume representation, the binning factor (η_y) applied to the simulated projections was kept equal to the volume downsampling factor. The forward and back-projection operators in the simulation and reconstruction used an in-house GPU implementation of separable footprints algorithm using trapezoidal functions (Long *et al* 2010). The separable footprint algorithm is voxel-driven and thus can easily accommodate arbitrary ROI shapes. However, our current implementation of the projector relies on GPU texture memory and thus only handles rectangular volumes. Non-rectangular fine grid ROIs were simulated using a minimum bounding box and masking operations.

In addition to varying the coarse voxel size, the regularization strength β_C in the coarse region was also varied. Since image downsampling imparts a degree of smoothing, the effect of a specific β_C value will differ between reconstructions with different coarse voxel sizes. To facilitate comparison of the effects of regularization across a range of voxel sizes (downsampling factors), a normalization that accounts for the effects of sampling was applied to the quadratic penalty (Yu *et al* 2013). Following the derivation by Yu, for quadratic penalties:

$$\psi=\frac{t^2}{2}=\frac{(\mu_i-\mu_k)^2}{2}\approx\frac{[\nabla\mu\cdot\Delta_{ki}]^2}{2}=\Delta_{ki}^2\frac{\nabla\mu^2}{2} \quad (9)$$

where Δ_{ki} is the distance between neighboring voxels *k* and *i*. Eq. (9) implies that within a neighborhood where the change in underlying attenuation volume is smooth, the action of the penalty (for a fixed penalty strength) scales approximately as square of the voxel size or, equivalently, as square of η_μ. Henceforth, reconstructions are compared across a range of η_μ in a manner that accounts for additional blur due to downsampling by using the following normalized penalty strength, denoted as β:

$$\beta=(1/\eta_\mu^2)\beta_C \quad (10)$$

In our studies we varied β from 0 (no regularization) to 10^8 for each value of η_μ.

To assess the artifacts in the fine grid region, a small (3 mm × 3.8 mm) ROI, μ_{art}, was selected in a uniform area of soft tissue adjacent to two bone edges (Fig. 3), where initial evaluation indicated pronounced streaking in reconstructions with high downsampling and regularization of the coarse grid region. To quantify artifacts, reconstruction of noiseless projection data were generated for each set of parameters of the multiresolution reconstruction. Root mean squared error (RMSE) was calculated in the ROI between the noiseless multiresolution reconstruction and a truth image that was obtained by downsampling the digital phantom (0.075 mm voxels) to the voxel size of the fine grid region (0.15 mm):

$$RMSE=\sqrt{\frac{\sum_j\left(\hat{\mu}_{\mathrm{art}_j}-\mu_{\mathrm{art}_j}\right)^2}{N}} \quad (11)$$

where $\hat{\mu}_{art_j}$ are voxels in the artifact ROI of the noiseless multiresolution reconstruction, μ_{art_j} are the corresponding voxels in the downsampled digital phantom and *N* is the total number of voxels in μ_{art}.

Increase in regularization strength and increase in voxel downsampling both impart resolution loss and noise reduction in the coarse grid region. The noise in μ_C (denoted as SD_c) was used as a metric quantifying the cumulative effect of coarse discretization and regularization. The noise was measured as RMSE between a flat region in the coarse grid subvolume of a noisy reconstruction, denoted as μ_{noise} (Fig. 3), and the same region in the truth image.

2.3. Physical Study using an Anthropomorphic Knee Phantom on a Benchtop CBCT

An anthropomorphic lower extremity phantom was scanned on a CBCT test bench (Zhao *et al* 2014) simulating a dedicated extremity CBCT (SAD=435.7 mm, SDD=559.2 mm). The detector was a PaxScan4030CB (Varian, Palo Alto, CA) operated at 1.5 fps and 0.194 mm pixel size. The x-ray source was a rotating anode DU694 x-ray tube with 14° anode angle, enclosed in EA10 housing (Dunlee, Aurora, IL). The scan consisted of 360 projections acquired over 360° at 90 kVp (+0.2 mm Cu, 2 mm Al), with 0.4 mAs per projection.

The total reconstructed volume was 18×18×18 cm^3. The irregularly-shaped fine grid ROI contained a region of subchondral bone identified using an initial Feldkamp-Davis-Kress (FDK) analytical reconstruction and consisted of 31,623,168 voxels of 0.15 mm size. There were 26,505,888 voxels of 0.6 mm size in the coarse grid region ($\eta_\mu = 4$). Multiresolution detector binning with $\eta_y = 4$ was applied, with the region of fine pixels in each view matching a forward projection of the subchondral bone mask. Regularization strengths were set to $\beta_F = 10^{-0.5}$ and $\beta_C = 10^4$.

2.4. Analysis of Runtime and Memory Footprint

To estimate the savings in computation time and memory consumption provided by the multiresolution method, benchmarking was performed for a 12×12×12 cm^3 volume. This volume is smaller than the one used in the anthropomorphic phantom experiments of Sec. 2.3. This is because a fine grid reconstruction of the entire grid was needed for benchmarking. The memory available on the GPU used in the experiments was not sufficient to fit a fine voxel grid larger than 12×12×12 cm^3. Note however that this size of the volume is sufficient to cover a typical knee joint on the extremity CBCT system.

Fine grid ROIs with 0.15 mm voxels and varying volume (5×5×5 cm^3 to 12×12×12 cm^3) were considered. For each fine grid ROI, the remainder of the FOV was parameterized using a coarse grid over a range of upsampling factors ($\eta_\mu = 1,2,4$). Since computation times for projection and backprojection vary and can be dependent on implementation, we considered the computational cost associated with projection-backprojection pairs. Average total projection-backprojection times were obtained separately for each fine grid volume (denoted as t_f) and for coarse grid volumes covering the entire FOV (t_c) at the various coarse grid voxel sizes. Furthermore, for each fine grid ROI, the projection-backprojection time for a volume equal in size to the fine grid ROI, but parameterized using coarse voxels was measured (t_Δ). The coarse grid projection-backprojection times were computed with and without projection binning. The estimate of projection-backprojection time for a given multiresolution parameterization (defined by a combination of fine grid volume size and upsampling factor) is $t_f + t_c - t_\Delta$, assuming that projection and backprojection times are linear with the number of voxels. The validity of this assumption was confirmed for the range of volume discretizations considered here for both projection and backprojection operations. The reconstruction speedup was computed as the ratio of the average measured time for a projection-backprojection pair of our multiresolution implementation to the same average measured time using only the fine grid ($\eta_\mu = 1$).

Each PWLS iteration requires storing the following variables: the current image μ, derivatives of data fit and regularization surrogates, $\dot{L}$ and $\dot{R}$, the curvatures of data fit and regularization surrogates, d_j and r, and the measured projections. (Here we ignored potential memory savings due to ordered subsets that may be offset by increased memory transfer times.) Memory footprint was thus calculated as the size of the projection dataset plus five times the size of the image volume, assuming all volumes are kept in memory. The memory footprint could be reduced by reusing the arrays associated with some of the variables, however at the expense of increased data transfer latency.

3. Results

3.1. Digital Phantom Study

Figure 4 illustrates the convergence properties of multiresolution PWLS reconstruction as a function of the downsampling factor. For each value of η_μ, we computed root-mean-square difference (RMSD) between the volume estimate at the n^{th} iteration and a nearly converged estimate (approximated by the solution after 200 iterations). RMSD computed within the fine grid ROI (dashed line) and within the coarse grid ROI (solid lines) are shown separately. 50 iterations of multiresolution PWLS were sufficient to achieve an RMSD of less than 10^{-4} mm^{-1} with this stable solution in both ROIs and across all values of η_μ. The RMSD behavior in μ_F was the same across all downsampling factors. Convergence was generally faster in μ_C than in μ_F, even at $\eta_\mu = 1$ (equal voxel size in both grids), likely because the phantom contains fewer high-frequency structures in the coarse grid ROI. The convergence rate in μ_C increases with higher downsampling. This behavior is intuitive since coarser grids (with fewer parameters to estimate) generally have better condition numbers, leading to faster convergence. Overall, the number of iterations needed to arrive at a stable solution in multiresolution PWLS is determined by the convergence rate in fine grid ROI and is thus independent of η_μ. The crucial advantage of using higher downsampling is in improved time per iteration. Based on this analysis, 50 iterations of multiresolution PWLS were used throughout this study.

Figure 5 shows a comparison of multiresolution PWLS reconstructions [without projection binning in (a) and with projection binning in (b)] across a range of downsampling factors and normalized regularization strengths. Note that reconstructions with matched β but unmatched η_μ use different β_C, as given by Eq. 12. The case of $\eta_\mu = 1$ corresponds to the same voxel size in μ_F and μ_C, but different values of β_F and β_C. A zoomed ROI is shown for each reconstruction that covers the fine grid ROI. Note that an increase in coarse grid regularization increases streak artifacts in the fine grid region. The artifacts become pronounced at large values of the downsampling factor. Additionally, when detector pixels are binned [Fig. 5 (b)], noticeable streak artifacts are apparent in the coarse grid near the ROI boundary in the case where $\eta_\mu = 10$. The right-most column in Fig. 5 (a) and (b) shows a difference image computed between the case of $\eta_\mu = 10$ and a reference image at $\eta_\mu = 1$ (for this computation, the coarse grid voxels for $\eta_\mu = 10$ were upsampled using nearest neighbor interpolation). Even for this high downsampling factor, there is no distortion in the trabecular features contained in the fine grid ROI μ_F (note that the same fine grid regularization strength β_F is applied in the downsampled and reference reconstructions).

Figure 6a investigates the trade-off between the strength of artifacts in μ_F, measured using equation (11), and the cumulative effect of regularization and downsampling in μ_C, represented by the level of noise in the coarse grid region. The results are very similar for cases with and without detector binning; for clarity, results are shown only for reconstructions with detector binning.

For all downsampling levels, the artifact metric exhibits a plateau accross a range of low-to-moderate noise levels (standard deviations of 10^{-4} to 10^{-3} mm^{-1} compared to muscle attenuation value of ~2×10^{-2} mm^{-1}). In this regime, regularization strength in μ_C can be adjusted without introducing significant artifacts in the fine grid ROI. The RMSE rapidly increases at lower noise levels, corresponding to increased smoothing in μ_C. The graphs in Figure 6 (a) are largely overlapping, indicating that the magnitude of the artifacts is a function of the overall blur in the coarse grid ROI, regardless of whether the blur results from downasampling or regularization. Consequently, when a lower value of the downsampling factor is used, there is more room to adjust the regularization in the corase grid ROI, i.e. β_C can be varied over a broader range of values without introducing artifacts in the fine grid region. If the corase grid region is of clinical interest, this ablity to tune the resolution-noise tradeoff in μ_C when the downsapling factor is relatively low may be relevant for optimizing the performance of the reconstruction.

Figure 6 (b) illustrates the artifacts in the fine grid region using a zoom on the artifact ROI (μ_{art}). As the normalized penalty strength (β) and η_μ increase, streaking becomes apparent. The more downsampling, the lower the value of normalized β where the streaking becomes pronounced.

Further examination of Figure 6 (a) reveals a slight increase in RMSE at the lowest regularizations (highest noise levels) in the presence of downsampling ($\eta_\mu \geq 4$). This is attributed to bias in the forward projection of μ_C introduced by sharp edges in the coarse discretization of fat-muscle interfaces of the phantom. This effect is diminished when sufficiently high levels of regularization are applied in the coarse grid region, blurring the discretization-induced sharp intensity transitions. The artifact due to sharp edges in weakly regularized μ_C is visible as a pronounced cross-hatch pattern at $\eta_\mu \geq 4$ and low β values in Figure 6 (b). Note that a similar, but much less conspicuous pattern is visible at $\eta_\mu = 1$, where it is likely to represent a combination of various discretization artifacts that are often present in noiseless reconstruction of digital phantoms with sharp edges.

3.2. Effect of the Location of the Fine Grid ROI

Examination of Fig. 5 suggests that the most pronounced artifacts in the fine grid ROI emerge from areas where the interface between μ_F and μ_C crosses regions of high contrast and sharp intensity transitions. The appearance and magnitude of the artifacts will depend on the location of the fine and coarse grid ROIs, as illustrated in Figure 7 for a case with both voxel downsampling and projection binning. Here, the boundaries of μ_F (marked with an orange dashed line) were designed to avoid cortical bone boundaries and other high contrast gradients. The downsampling factor was set to 10, which leads to severe artifacts in the reconstructions of Fig. 5 (5th column). With the fine grid ROI conforming to the bone boundaries, however, no streaking artifacts are visible in the space between the two bones. A strong mismatch in the magnitude of noise in μ_C and μ_F is apparent at higher values of normalized regularization strength. Similarly to Fig. 5, the right-most column of Fig. 7 investigates a difference image between the downsampled case and the reference image at $\eta_\mu = 1$ (equal voxel size in the coarse and fine grids). No visible distortions in the trabecular structure were introduced using multiresolution reconstruction.

3.3. Multiresolution CBCT Reconstruction of an Anthropomorphic Knee Phantom

Figure 8 shows a reconstruction of an anthropomorphic knee phantom acquired on the CBCT test bench. Fig. 8 (a) illustrates the boundaries of the fine grid ROI. Four-fold downsampling of the volume and projections was employed. The multiresolution PWLS used 3.36% of the number of voxels and 17.7% of the number of detector pixels compared to a full FOV fine grid reconstruction. No artifacts arise from downsampling, and high-resolution trabecular features are visible in the fine grid ROI [outlined with orange dashed line in Fig. 8 (b)–(d)]. The resulting multiresolution image can be stored in less than 0.5 GB, whereas for the full-field fine grid image, approximately 7 GB would be required.

3.4. Computational Cost of Multiresolution PWLS

Figure 9 summarizes the investigation of computational savings provided by multiresolution PWLS. As expected, the speedup factor (compared to reconstruction using fine grid voxels throughout the entire FOV) can be as high as a 5× – 10× when the fine grid ROI is a relatively small fraction of the total size of the FOV. Most of the speedup is from image downsampling, rather than detector binning, which is to be expected with the voxel-driven forward projection method used here.

Fig. 9 (b) examines the reduction in memory footprint provided by multiresolution PWLS. As anticipated, the memory savings are enhanced when projection binning is used in conjunction with volume downsampling. For a fine grid ROI of a similar size as the one used in the anthropomorphic knee phantom of Sec. 3.3, the memory footprint reduction is approximately 20% compared to reconstruction using only the fine grid voxels.

4. Discussion and Conclusions

We developed a PWLS reconstruction algorithm implementing a multiresolution voxelized parameterization of the object. The algorithm was evaluated for application in accelerated iterative reconstruction of finely sampled ROIs (voxel sizes <100 μm) from high resolution extremity CBCT projection data. It was assumed that the fine grid region is of the primary clinical interest and thus the investigation was focused on mitigation of artifacts in the fine grid region and quantification of the reconstruction speedup.

Streaking artifacts found in the fine grid ROI are likely caused by inconsistencies between the simulated forward projections of the coarsely and finely sampled ROIs. Artifact magnitude is thus a function of the cumulative blur due to downsampling and regularization of the coarse grid region. For example, downsampling factors as high as 10× were used without introducing visible streaking in the fine grid region by applying a relatively weak regularization in the coarsely sampled sub-volume. However, the tradeoff was that the μ_C exhibited a “blocky” appearance, which could limit the diagnostic utility of the coarsely sampled region. This may however be acceptable if the clinical interest is primarily in the fine grid ROI and achieving maximal speedup is a priority (and the coarsely sampled region may be reconstructed separately by other means using parameters suitable to soft-tissue visualization).

Additional memory savings can be attained when the multiresolution representation is augmented with binning of projection data. This work evaluated a scenario where projection binning was applied to pixels receiving the line integrals that traversed only through μ_C. This approach results in similar performance with respect to artifacts in the fine grid region as the method using only volume downsampling. In the coarse grid region, however, streaks tangential to the ROI boundary were found when using high levels of projection binning (η_Y = 10) in conjunction with volume downsampling. The artifacts are hypothesized to emerge from mismatches in simulated projections of some of the coarse grid voxels located close to the ROI boundary. Such voxels are projected both on the binned and native pixels, depending on whether a given line integral traverses through μ_F for each projection angle. Since this inconsistency is not present for μ_F voxels, the fine grid ROI was unaffected.

The simulation study of Sec. 3.1 involves an ROI boundary that crosses through high contrast bone regions, resulting in a challenging scenario for multiresolution PWLS. Sec. 3.2 illustrates that both types of streaking artifacts discussed above can be effectively mitigated using a fine grid ROI that conforms to high contrast edges.

Overall, the optimal value of the downsampling factor will depend on the location of the fine and coarse grid regions and on the clinical application (i.e. whether both ROIs are of clinical interest, or only the finely sampled sub-volume). The results indicate that downsampling factors of ~4× are possible without perceptible artifacts in either of the ROIs. At this level of η_μ, the regularization strength in the coarsely sampled sub-volume can be adjusted over a relatively broad range without adversely affecting the fine grid region. For a typical knee volume, this downsampling corresponds to more than 5-fold acceleration of the iterative reconstruction. The current execution time for the multiresolution reconstruction of the knee phantom of Figure 8, obtained with $\eta_\mu = 4$, is ~2 min per iteration (or ~100 min for the complete reconstruction of 50 iterations) using an un-optimized MATLAB-CUDA implementation. A naïve approach utilizing fine voxels throughout the whole volume would take ~10 hours for the reconstruction to complete and require ~15× more memory, and thus the multiresolution algorithm provided a significant step towards achieving clinically practical runtimes for iterative reconstruction of the trabecular detail in human studies. We anticipate that an additional 5× acceleration over the execution time stated here will be possible with an optimized implementation using a compiled executable with improved memory management on a multi-GPU workstation.

For certain configurations of the coarse and fine grid regions, downsampling factors as high as 10× can be used. For the typical knee CBCT (e.g. the sample reconstruction in figure 8), the 10× downsampling should yield ~52× speedup in PWLS, assuming a projection time that scales linearly with number of voxels. The acceleration factor will increase if the high resolution ROI is a smaller fraction of the total volume.

Similarly to the selection of the downsampling factor, the location, shape and the procedure for delineating the fine grid ROI will depend on the clinical application. The knee phantom reconstruction in Figure 8 represents a scenario where the scan parameters were selected to yield projection data that supports high-resolution reconstruction of the entire knee joint. In this case, a high fidelity FDK reconstruction could be obtained using the same projection data (thus at no additional dose), enabling precise definition of the fine grid ROI. However, such close adherence to bony boundaries is not always necessary. As shown in Figure 7, image artifacts due to downsampling will be minimized as long as the boundaries of the fine grid ROI do not cross high contrast edges. Detection of such high-contrast edges will typically be possible in an initial FDK reconstruction even if the projection data is noisy or sparsely sampled, as in low-dose imaging protocols intended for use with iterative reconstruction algorithms. In such cases, the initial analytical reconstruction may be non-diagnostic, but will likely be sufficient to delineate the fine grid ROI for the iterative algorithm. Furthermore, while the use of precisely defined fine grid ROIs (conforming with the object of interest) optimizes reconstruction time by minimizing the volume parameterized with fine voxels, the low dose applications are unlikely to yield data supporting reconstructions on very fine voxels. The use of sub-optimal fine grid ROIs (e.g. slightly larger than the object of interest) for low dose data will thus not be as detrimental for reconstruction performance as in the case of high resolution imaging.

Other factors to be considered in choosing the shape of the fine grid ROI involve computational burden and ease of implementation. In principle, the shape of the ROIs is only limited by the requirement that the fine grid voxel size must be an integer fraction of the coarse voxel size to ensure complete coverage of the volume. A fine grid ROI that is tightly matched to the shape of area of high resolution reconstruction has the advantage of minimizing the memory usage. At the same time, the execution speed may be hampered by non-contiguous memory access patterns that are likely to emerge when e.g. the fine grid ROI consists of multiple disjoint patches. Rectangular ROIs are easier to map to contiguous memory blocks and may thus yield improved computation speed compared to more complex ROI shapes when executed using general purpose software. In practice, the use of regular, rectangular or cylindrical fine grid ROIs loosely following the boundaries of the objects of interest (e.g. bones) and avoiding sharp tissue boundaries, combined with downsampling factors of 4–5, is likely to yield robust performance in typical clinical applications of multiresolution PWLS.

In conclusion, the proposed multiresolution algorithm for PWLS was tested in application to extremity CBCT. Artifact-free reconstructions of finely sampled ROIs were achieved at computation times that are 5× – 10× shorter compared to a brute force solution that applies fine voxel parameterization to the entire volume. The multiresolution framework can be further expanded to support other applications where computationally expensive forward models (e.g. modeling of detector blur (Tilley II *et al* 2015)) are applied only to a sub-region of the field-of-view. The algorithm permits application of different regularization strengths in the regions of coarse and fine grid, enabling reconstructions in which the sampling and regularization are individually optimized depending on the clinical task in each image region.

Acknowledgments

Research supported by National Institutes of Health Grant No. R01EB018896. Qian Cao is supported by the Howard Hughes Medical Institute (HHMI) International Student Research Fellowship. The authors thank Ali Uneri (Department of Computer Science, Johns Hopkins University) for valuable discussion regarding computer implementation of multiresolution projectors.

References

Boerckel JD, Mason DE, McDermott AM, Alsberg E. Microcomputed tomography: approaches and applications in bioengineering. Stem Cell Res. Ther. 2014; 5:144. [PubMed: 25689288]

Boone JM, Nelson TR, Lindfors KK, Seibert JA. Dedicated Breast CT: Radiation Dose and Image Quality Evaluation. Radiology. 2001; 221:657–667. [PubMed: 11719660]

Brankov JG, Yang Y, Wernick MN. Tomographic Image Reconstruction Based on a Content-Adaptive Mesh Model. IEEE Trans. Med. Imaging. 2004; 23:202–212. [PubMed: 14964565]

Brokish J, Sack P, Bresler Y. Combined algorithmic and GPU acceleration for ultra-fast circular conebeam backprojection. Proc. SPIE Medical Imaging 2010. 2010:762256.

Carrino JA, Al Muhit A, Zbijewski W, Thawait GK, Stayman JW, Packard N, Senn R, Yang D, Foos DH, Yorkston J, Siewerdsen JH. Dedicated cone-beam CT system for extremity imaging. Radiology. 2014; 270:816–824. [PubMed: 24475803]

Clark DP, Badea CT. Micro-CT of rodents: State-of-the-art and future perspectives. Phys. Medica. 2014; 30:619–634.

Degirmenci S, Politte DG, Bosch C, Tricha N, Sullivan JAO, Drive B, Louis S. Acceleration of iterative image reconstruction for X-ray imaging for security applications. Proc. SPIE Computational Imaging XIII. 2015:94010C.

Elbakri IA, Fessler JA. Statistical image reconstruction for polyenergetic X-ray computed tomography. IEEE Trans. Med. Imaging. 2002; 21:828–831.

Erdoğan H, Fessler Ja. Monotonic algorithms for transmission tomography. IEEE Trans. Med. Imaging. 1999a; 18:801–814. [PubMed: 10571385]

Erdoğan H, Fessler JA. Ordered subsets algorithms for transmission tomography. Phys. Med. Biol. 1999b; 44:2835–2851. [PubMed: 10588288]

Fessler JA. Grouped Coordinate Descent Algorithms for Robust Edge-Preserving Image Restoration. Proc. SPIE Image Reconstruction and Restoration. 1997; 3170:184–194.

Frese T, Bouman CA, Sauer K. Adaptive wavelet graph model for Bayesian tomographic reconstruction. IEEE Trans. Image Process. 2002; 11:756–770. [PubMed: 18244672]

Hamelin, B.; Goussard, Y.; Dussault, JP. Penalized-likelihood region-of-interest CT reconstruction by local object supersampling; Annual International Conference of the IEEE Engineering in Medicine and Biology; 2007. p. 739-742.

Hamelin B, Goussard Y, Dussault J-P, Cloutier G, Beaudoin G, Soulez G. Design of iterative ROI transmission tomography reconstruction procedures and image quality analysis. Med. Phys. 2010; 37:4577–4589. [PubMed: 20964175]

Hodez C, Griffaton-Taillandier C, Bensimon I. Cone-beam imaging: Applications in ENT. Eur. Ann. Otorhinolaryngol. Head Neck Dis. 2011; 128:65–78. [PubMed: 21459066]

Keck B, Hofmann H. GPU-accelerated SART reconstruction using the CUDA programming environment. Proc. SPIE Medical Imaging. 2009:72582B.

Liu X, Boons S, Sasov A. Strategies for GPU-based cone-beam CT reconstruction for very large data volumes. Proc. Int. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med. 2013:529–532.

Long Y, Fessler Ja, Balter JM. 3D forward and back-projection for X-ray CT using separable footprints. IEEE Trans. Med. Imaging. 2010; 29:1839–1850. [PubMed: 20529732]

Marinetto E, Brehler M, Sisniega A, Cao Q, Stayman W, Yorkston J, Siewerdsen J, Zbijewski W. Quantification of bone microarchitecture in ultra-high resolution extremities cone-beam CT with a CMOS detector and compensation of patient motion. Computer Assisted Radiology and Surgery: 30th International Congress and Exhibition. 2016

Nishiyama KK, Shane E. Clinical imaging of bone microarchitecture with HR-pQCT. Curr. Osteoporos. Rep. 2013; 11:147–155. [PubMed: 23504496]

O'Connell A, Conover DL, Zhang Y, Seifert P, Logan-Young W, Lin CFL, Sahler L, Ning R. Cone-beam CT for breast imaging: Radiation dose, breast coverage, and image quality. Am. J. Roentgenol. 2010; 195:496–509. [PubMed: 20651210]

Papenhausen E, Zheng Z, Mueller K. GPU-accelerated back-projection revisited : squeezing performance by careful tuning. Proc. Int. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med. 2011:19–22.

Pratx G, Xing L. GPU computing in medical physics: a review. Med. Phys. 2011; 38:2685–2697. [PubMed: 21776805]

Rashed EA, Kudo H, Zeniya T, Iida H. Practical statistical models for region-of-interest tomographic reconstruction and long object problem. IEEE Nucl. Sci. Symp. Conf. Rec. 2007; 5:3505–3511.

La Rivière PJ. Monotonic iterative reconstruction algorithms for targeted reconstruction in emission and transmission computed tomography. IEEE Nucl. Sci. Symp. Conf. Rec. 2007; 5:2924–2928.

Sauer K, Bouman C. A local update strategy for iterative reconstruction from projections. IEEE Trans. Signal Process. 1993; 41:534–548.

Shen Y, Zhong Y, Lai C-J, Wang T, Shaw CC. Cone beam breast CT with a high pitch (75 μm), thick (500 μm) scintillator CMOS flat panel detector: Visibility of simulated microcalcifications. Med. Phys. 2013; 40:101915. [PubMed: 24089917]

Siewerdsen JH, Waese AM, Moseley DJ, Richard S, Jaffray Da. Spektr: A computational tool for x-ray spectral analysis and imaging system optimization. Med. Phys. 2004; 31:3057–3067. [PubMed: 15587659]

Sitek A, Huesman RH, Gullberg GT. Tomographic reconstruction using an adaptive tetrahedral mesh defined by a point cloud. IEEE Trans. Med. Imaging. 2006; 25:1172–1179. [PubMed: 16967802]

Van Slambrouck, K.; Nuyts, J. A patchwork (back)projector to accelerate artifact reduction in CT reconstruction; IEEE Nucl. Sci. Symp. Med. Imaging Conf; 2010. p. 2625-2629.

Van Slambrouck K, Nuyts J. Metal artifact reduction in computed tomography using local models in an image block-iterative scheme. Med. Phys. 2012; 39:7080–7093. [PubMed: 23127099]

Stayman JW, Otake Y, Prince JL, Khanna aJ, Siewerdsen JH. Model-based tomographic reconstruction of objects containing known components. IEEE Trans. Med. Imaging. 2012; 31:1837–1848. [PubMed: 22614574]

Stearns CW, Manjeshwar RM, Wollenweber SD. An efficient algorithm for targeted reconstruction of tomographic data. IEEE Nucl. Sci. Symp. Conf. Rec. 2007; 5:2808–2811.

Thibault J-B, Sauer KD, Bouman Ca, Hsieh J. A three-dimensional statistical approach to improved image quality for multislice helical CT. Med. Phys. 2007; 34:4526. [PubMed: 18072519]

Tilley S II, Siewerdsen JH, Stayman JW. Model-based iterative reconstruction for flat-panel cone-beam CT with focal spot blur, detector blur, and correlated noise. Phys. Med. Biol. 2015; 296:296.

Tuominen EKJ, Kankare J, Koskinen SK, Mattila KT. Weight-bearing CT imaging of the lower extremity. AJR. Am. J. Roentgenol. 2013; 200:146–148. [PubMed: 23255755]

De Vos W, Casselman J, Swennen GRJ. Cone-beam computerized tomography (CBCT) imaging of the oral and maxillofacial region: A systematic review of the literature. Int. J. Oral Maxillofac. Surg. 2009; 38:609–625. [PubMed: 19464146]

Wang AS, Stayman JW, Otake Y, Kleinszig G, Vogt S, Gallia GL, Khanna AJ, Siewerdsen JH. Soft-tissue imaging with C-arm cone-beam CT using statistical reconstruction. Phys. Med. Biol. 2014; 59:1005–1026. [PubMed: 24504126]

Xu J, Reh DD, Carey JP, Mahesh M, Siewerdsen JH. Technical assessment of a cone-beam CT scanner for otolaryngology imaging: Image quality, dose, and technique protocols. Med. Phys. 2012; 39:4932. [PubMed: 22894419]

Yang WT, Carkaci S, Chen L, Lai CJ, Sahin A, Whitman GJ, Shaw CC. Dedicated cone-beam breast CT: Feasibility study with surgical mastectomy specimens. Am. J. Roentgenol. 2007; 189:1312–1315. [PubMed: 18029864]

Yu Z, Bouman CA, Thibault J, Sauer KD. Image Grid Invariant Regularization for Iterative Reconstruction. Proc. Int. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med. 2013:517–520.

Yu Z, Thibault J-B, Bouman C, Sauer K, Hsieh J. Edge-localized iterative reconstruction for computed tomography. Proc. Int. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med. 2009:255–258.

Zbijewski W, De Jean P, Prakash P, Ding Y, Stayman JW, Packard N, Senn R, Yang D, Yorkston J, Machado A, Carrino JA, Siewerdsen JH. Design and Optimization of a Dedicated Cone-Beam CT System for Musculoskeletal Extremities Imaging. Med. Phys. 2011; 38:4700. [PubMed: 21928644]

Zhao Z, Gang GJ, Siewerdsen JH. Noise, sampling, and the number of projections in cone-beam CT with a flat-panel detector. Med. Phys. 2014; 41:061909. [PubMed: 24877820]

Ziegler A, Nielsen T, Grass M. Iterative reconstruction of a region of interest for transmission tomography. Med. Phys. 2008; 35:1317–1327. [PubMed: 18491526]

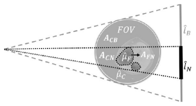

Figure 1.

Schematic of multiresolution forward projection with detector binning. The estimated bone region parameterized using fine grid voxels ($\hat{\mu}_F$) is marked with a black dashed line. This region is projected onto native detector pixels (marked with $\hat{l}_N$) using projection operator $\boldsymbol{A_{FN}}$. The line integrals captured by the native projection pixels also include contributions from a subset of coarse voxels ($\hat{\mu}_C$), denoted by projection operator $\boldsymbol{A_{CN}}$. Line integrals for binned projection pixels (marked as $\hat{l}_B$) traverse only through coarse grid voxels, computed using projector operator $\boldsymbol{A_{CB}}$.

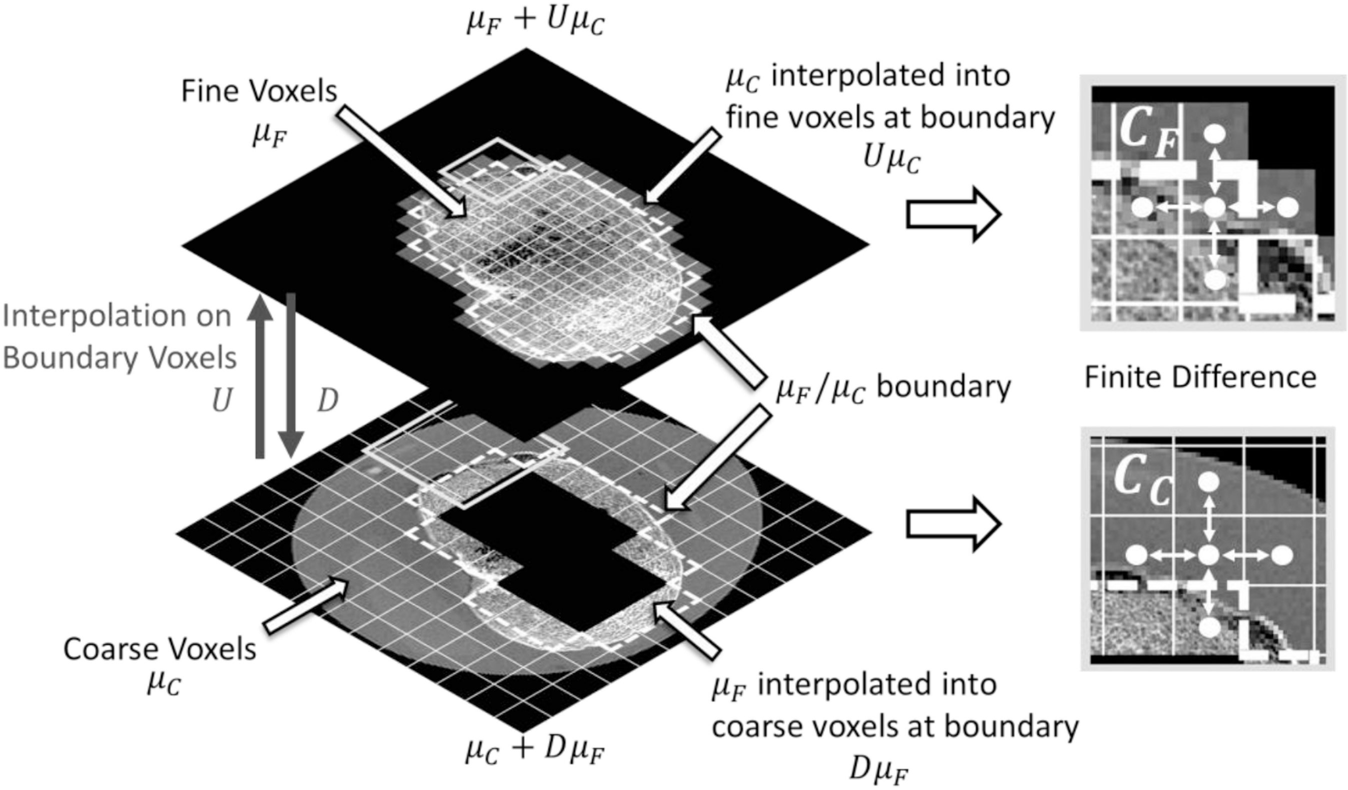

Figure 2.

Schematic of multiresolution regularization scheme illustrating how the regularization over boundaries between the fine and coarse voxel grids is performed using interpolation operators.

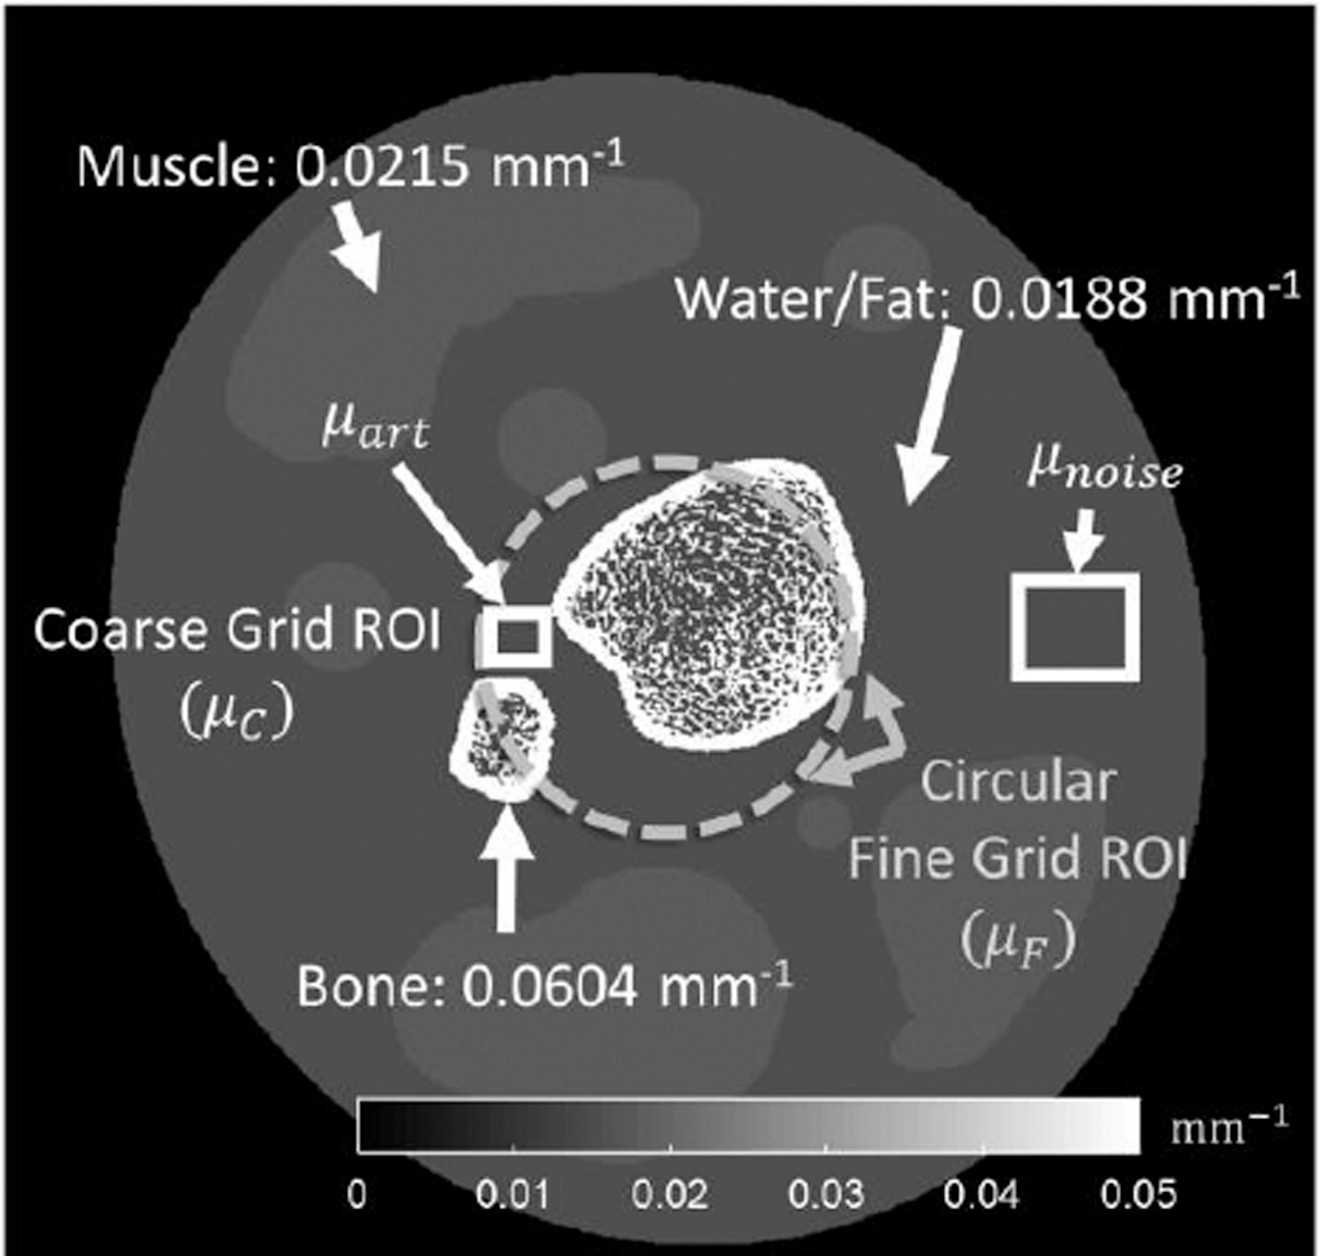

Figure 3.

Schematic of digital bone phantom with the boundary of a central fine voxel grid region marked with a thick dashed line. The ROIs used for measurement of artifacts in the fine grid region (μ_{art}) and noise (μ_{noise}) in the coarse grid region are indicated with boxes.

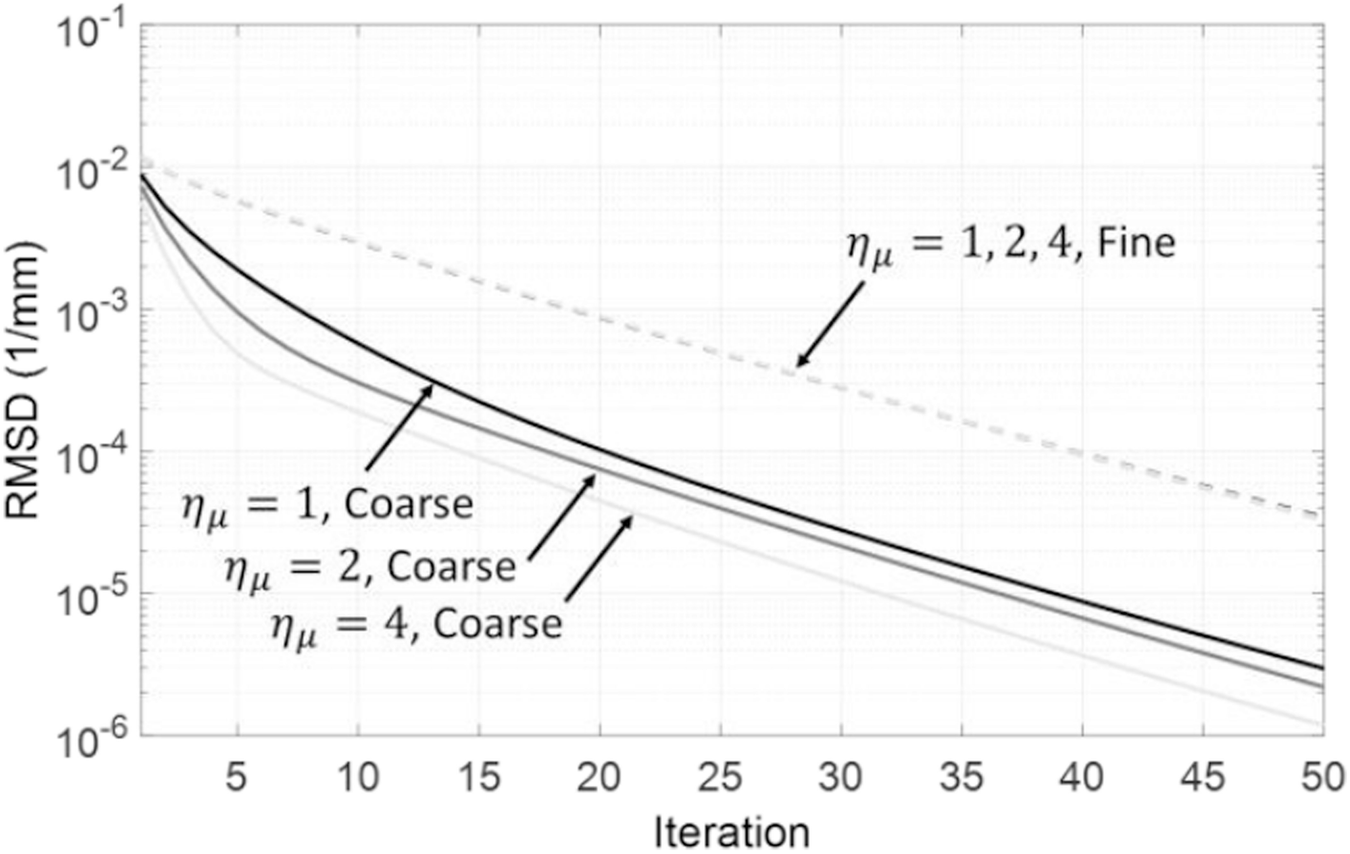

Figure 4.

Convergence profile measured in reconstructions of the digital phantom. The estimate at the current iteration is compared with a converged image at 200 iterations. $\beta_F = 10^{3.5}$; $\beta_C = 10^{4.5}$. The RMSD curves for the fie grid region (dashed line) overlap for all values of η_μ.

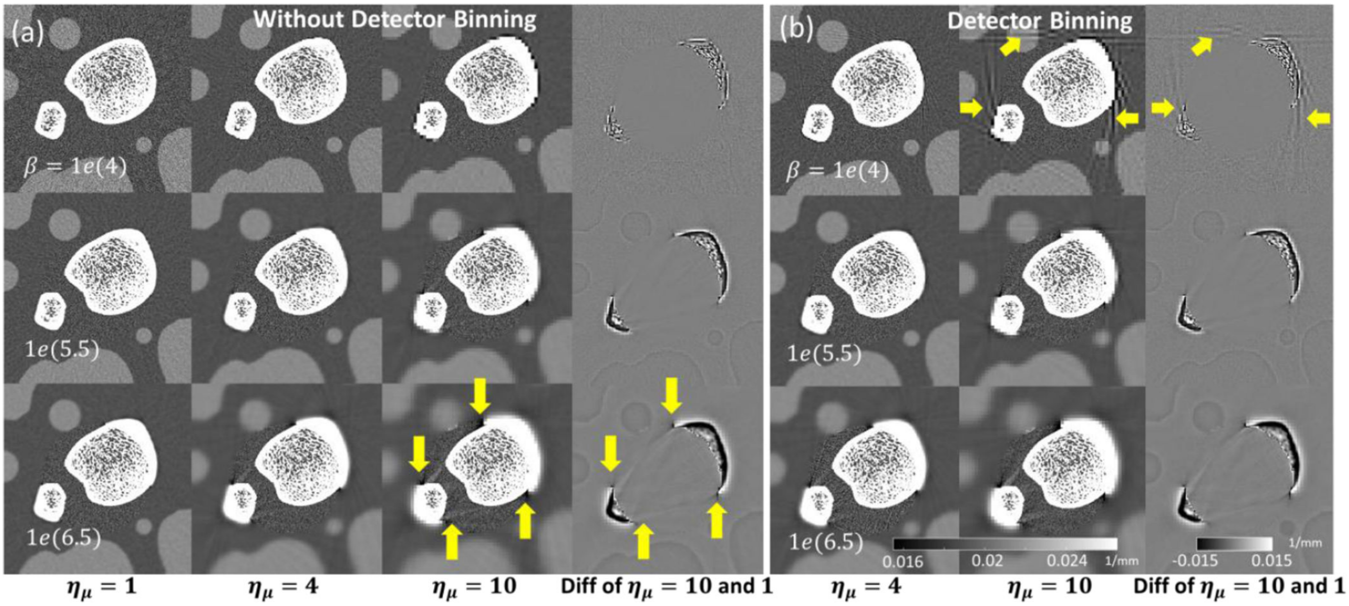

Figure 5.

Multiresolution PWLS reconstructions for different downsampling factors and coarse grid regularization values. A central region of the digital phantom is shown for PWLS without detector binning (a) and with detector binning (b). The last column in each subfigure is a difference image between the case with maximal downsampling ($\eta_\mu = 10$) and the case of no downsampling ($\eta_\mu = 1$). Arrows indicate artifacts due to strong coarse grid downsampling and regularization.

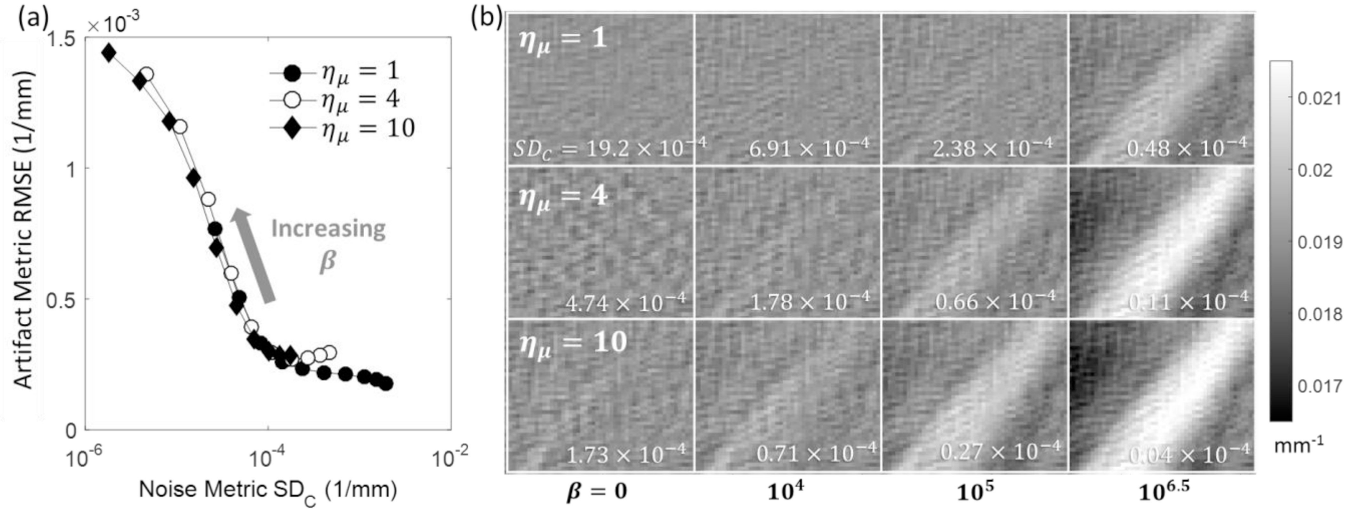

Figure 6.

(a) Magnitude of the artifact in fine grid ROI (RMSE of μ_{art} in Fig. 3, given by Eq. 11) as a function of noise in the coarse grid region. The noise metric quantifies the cumulative effect of regularization and voxel downsampling in μ_C. (b) The artifact ROI (μ_{art}) in the fine grid region for a range of downsampling factors (rows) and normalized coarse grid regularization strengths β (columns).

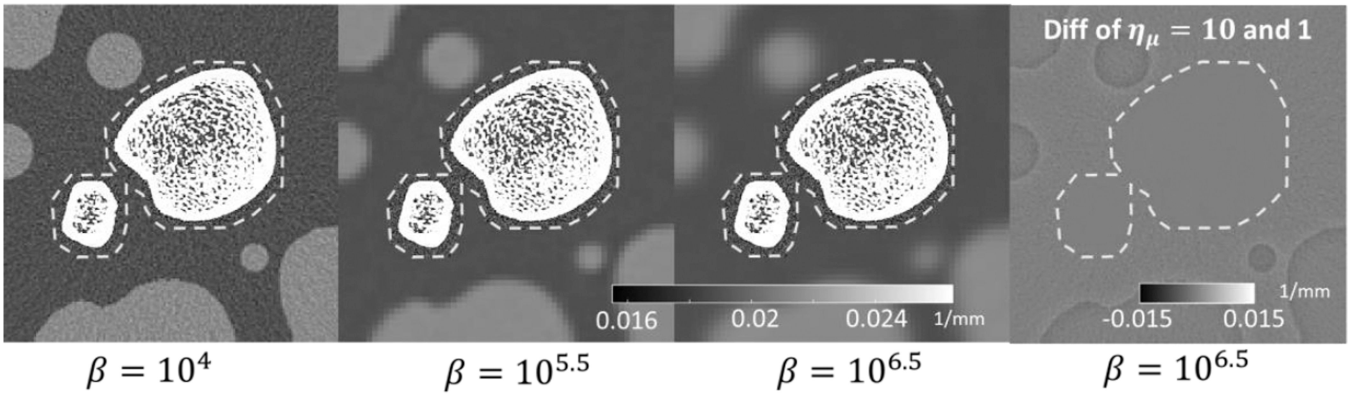

Figure 7.

Multiresolution PWLS with fine grid ROI conforming to bone boundaries (marked with dashed line in the leftmost mage) for three values of the normalized penalty strength β. The downsampling factor η_μ is set to 10. The right-most subfigure shows a difference image between the downsampled case ($\beta = 10^{6.5}$) and the reference reconstruction with no downsampling. Projection binning is included.

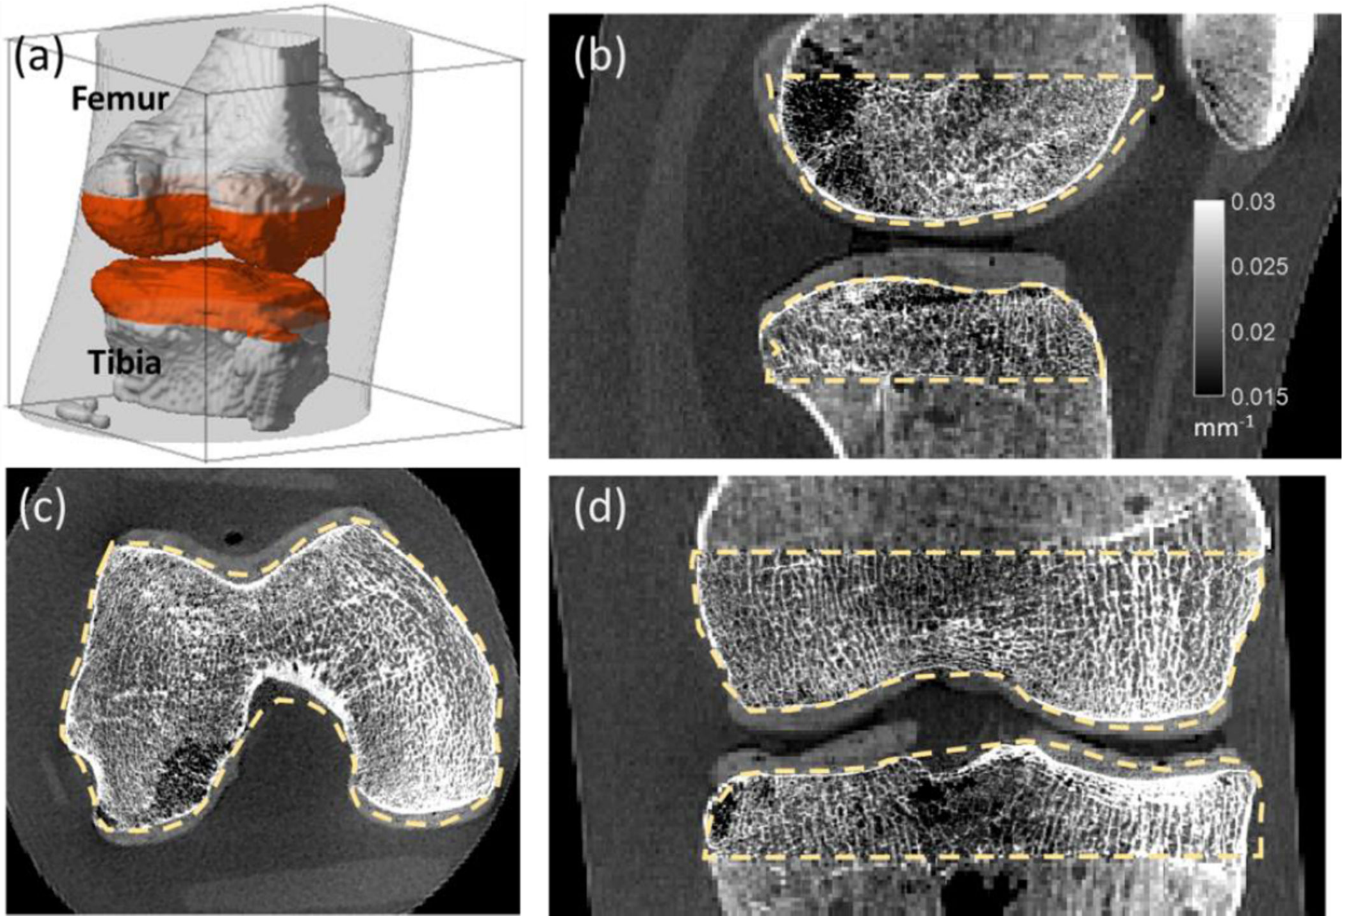

Figure 8.

Multiresolution PWLS reconstruction of an anthropomorphic knee phantom acquired on a CBCT benchtop. Fine and coarse grid regions are delineated in (a). (b)–(d) shows details of the high resolution trabecular ROI in the sagittal, axial, and coronal planes. Boundaries of the ROI are marked with a dashed line.

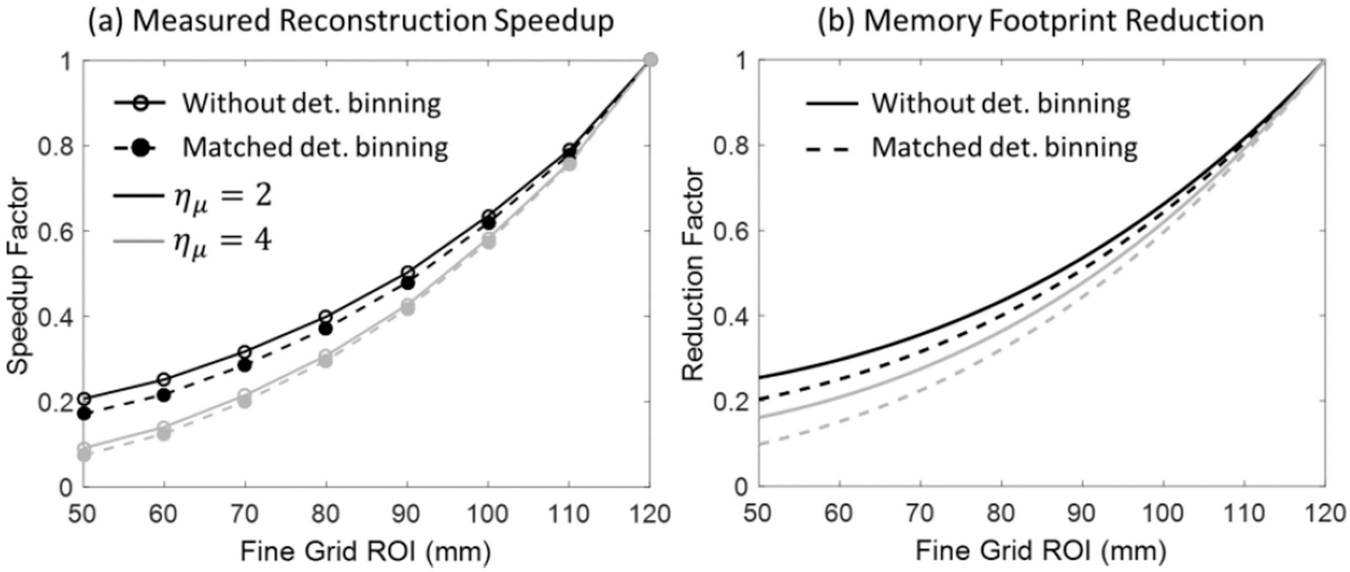

Figure 9.

(A) Measured reconstruction speedup of multiresolution PWLS as a function of the size of fine grid FOV (for total volume size of 120×120×120 mm^3). (B) Total memory consumption of downsampled cases compared to fine grid cases.

Table 1

Glossary of terms and pseudocode of SPS optimization method for multiresolution PWLS. Ordered-subsets are omitted for simplification.

Definitions	
$\hat{l}_B, \hat{l}_N$	forward-projected line integrals onto binned and native detector pixels
l_B, l_N	binned and native preprocessed line integral data
w_B, w_N	diagonal entries in weighting matrices $\boldsymbol{W}_B$ and $\boldsymbol{W}_N$.
c	precomputed PWLS curvature, $c = [\mathbf{W}]_+$
$\mathbf{1}$	volume covering the entire reconstruction FOV of all 1's.
$\dot{L}_C, \dot{L}_F$	derivatives of data-fit surrogates for voxels of $\hat{\mu}_C$ and $\hat{\mu}_F$.
d_C, d_F	curvatures of data-fit surrogates for voxels of $\hat{\mu}_C$ and $\hat{\mu}_F$.
$\dot{\psi}, \omega_{\psi k}$	gradient and curvature of the penalty function ψ.
$\dot{R}_C, \dot{R}_F$	derivatives of penalty surrogates for voxels of $\hat{\mu}_C$ and $\hat{\mu}_F$.
r_C, r_F	curvatures of penalty surrogates for voxels of $\hat{\mu}_C$ and $\hat{\mu}_F$.

Algorithm

```
Precompute $d_C = A_{CB}^T [A\mathbf{1} \circ c]_B + A_{CN}^T [A\mathbf{1} \circ c]_N$    $d_F = A_{FN}^T [A\mathbf{1} \circ c]_N$
for iteration n=1, …, N
    Coarse Grid Update
        $\hat{l}_B = A_{CB}\hat{\mu}_C$
        $\dot{h}_B = w_B \circ (\hat{l}_B - l_B)$
        $\dot{L}_C = A_{CB}^T \dot{h}_B + A_{CN}^T \dot{h}_N$
        $\dot{R}_{C_j} = \Sigma_{k \in K_C} c_{kj} \dot{\psi}([\tilde{C}\hat{\mu}]_k)$, $j \in$ coarse voxels
        $r_{C_j} = \Sigma_{k \in K_C} c_{kj}^2 \omega_{\psi_k}(\left[\tilde{C}\hat{\mu}\right]_k)$, $j \in$ coarse voxels
        $\hat{u}_C = \max(\hat{u}_C - \frac{\dot{L}_C + \beta_C \dot{R}_C}{d_C + \beta_C r_C}, 0)$
    Fine Grid Update
        $\hat{l}_N = A_{CN}\hat{\mu}_C + A_{FN}\hat{\mu}_F$
        $\dot{h}_N = w_N \circ (\hat{l}_N - l_N)$
        $\dot{L}_F = A_{FN}^T \dot{h}_N$
        $\dot{R}_{F_i} = \Sigma_{k \in K_F} c_{ki} \dot{\psi}([\tilde{C}\hat{\mu}]_k)$, $i \in$ fine voxels
        $r_{F_i} = \Sigma_{k \in K_F} c_{ki}^2 \omega_{\psi_k}(\left[\tilde{C}\hat{\mu}\right]_k)$, $i \in$ fine voxels
        $\hat{u}_F = \max(\hat{u}_F - \frac{\dot{L}_F + \beta_F \dot{R}_F}{d_F + \beta_F r_F}, 0)$
end
```