HHS Public Access

Author manuscript

Phys Med Biol. Author manuscript; available in PMC 2018 January 21.

Published in final edited form as:

Phys Med Biol. 2017 January 21; 62(2): 539–559. doi:10.1088/1361-6560/aa52b8.

Multi-resolution statistical image reconstruction for mitigation of truncation effects: application to cone-beam CT of the head

Hao Dang[1], **J. Webster Stayman**[1], **Alejandro Sisniega**[1], **Wojciech Zbijewski**[1], **Jennifer Xu**[1], **Xiaohui Wang**[2], **David H. Foos**[2], **Nafi Aygun**[3], **Vassillis E. Koliatsos**[4], and **Jeffrey H. Siewerdsen**[1,3]

[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 MD 21205, USA

[4]Department of Neurology, Johns Hopkins University, Baltimore MD 21205, USA

Abstract

A prototype cone-beam CT (CBCT) head scanner featuring model-based iterative reconstruction (MBIR) has been recently developed and demonstrated the potential for reliable detection of acute intracranial hemorrhage (ICH), which is vital to diagnosis of traumatic brain injury and hemorrhagic stroke. However, data truncation (e.g., due to the head holder) can result in artifacts that reduce image uniformity and challenge ICH detection. We propose a multi-resolution MBIR method with an extended reconstruction field of view (RFOV) to mitigate truncation effects in CBCT of the head. The image volume includes a fine voxel size in the (inner) nontruncated region and a coarse voxel size in the (outer) truncated region. This multi-resolution scheme allows extension of the RFOV to mitigate truncation effects while introducing minimal increase in computational complexity. The multi-resolution method was incorporated in a penalized weighted least-squares (PWLS) reconstruction framework previously developed for CBCT of the head. Experiments involving an anthropomorphic head phantom with truncation due to a carbon-fiber holder were shown to result in severe artifacts in conventional single-resolution PWLS, whereas extending the RFOV within the multi-resolution framework strongly reduced truncation artifacts. For the same extended RFOV, the multi-resolution approach reduced computation time compared to the single-resolution approach (viz. time reduced by 40.7%, 83.0%, and over 95% for an image volume of 600^3, 800^3, 1000^3 voxels). Algorithm parameters (e.g., regularization strength, the ratio of the fine and coarse voxel size, and RFOV size) were investigated to guide reliable parameter selection. The findings provide a promising method for truncation artifact reduction in CBCT and may be useful for other MBIR methods and applications for which truncation is a challenge.

Keywords

intracranial hemorrhage; cone-beam computed tomography; iterative reconstruction; projection truncation; truncation correction; scatter correction

1. Introduction

Accurate and reliable detection of intracranial hemorrhage (ICH) is essential to the diagnosis of a number of neurological pathologies, including traumatic brain injury, hypertensive intracerebral hemorrhage, hemorrhagic stroke, ruptured aneurysm, and cerebral amyloid angiopathy (Kidwell and Wintermark 2008, Parizel *et al* 2001). Non-contrast multi-detector CT (MDCT) is the current frontline imaging modality for detecting acute ICH (40–80 HU contrast of fresh blood to brain) (Kidwell and Wintermark 2008, Dublin *et al* 1977). For patients in the intensive care unit (ICU) or neurological critical care unit (NCCU), transport to a MDCT suite requires time and dedicated personnel, often with risk to the patient. For example, a mean round-trip time of 50–80 min outside the critical care environment and a 71% incidence of adverse events have been reported for such transport (Masaryk *et al* 2008, Smith *et al* 1990, Indeck *et al* 1988).

Cone-beam CT (CBCT) has a number of characteristics that make it potentially well suited to imaging acute ICH at the point-of-care. These include small footprint, open geometry, portability, low cost, and volumetric acquisition from a single rotation. In the past decade, CBCT imaging technology has been developed for a variety of applications, including dedicated musculoskeletal imaging (Koskinen *et al* 2013, Carrino *et al* 2014), breast imaging (Boone *et al* 2001, Yang *et al* 2007), maxillofacial imaging (Penninger *et al* 2011, Xu *et al* 2012), and interventional C-arms (Siewerdsen *et al* 2009, Navab *et al* 2010, Schafer *et al* 2011, Dang *et al* 2012). However, the detection of acute ICH requires a high level of contrast resolution (40–80 HU), spatial resolution (0.5–10 mm), and image uniformity, which poses major challenges to the current generation of CBCT systems.

Recent work has led to the development of a prototype CBCT head scanner suitable for detection of acute ICH. The scanner prototype was designed and optimized according to a task-based image quality model (Xu *et al* 2016c, 2016b). A fairly comprehensive artifact correction framework was developed to mitigate effects of x-ray scatter, beam hardening, detector lag, and veiling glare (Sisniega et al 2015). Furthermore, a model-based iterative reconstruction (MBIR) method was developed with statistical weights modified to account for the change in variance following artifact corrections, providing improved noise-resolution tradeoffs compared to analytical and conventional iterative reconstruction methods (Dang *et al* 2015). The method was recently extended to include spatially varying penalty strength that maximizes task-based detectability throughout the brain (Dang *et al* 2016).

While the aforementioned studies have shown great promise for acute ICH detection in CBCT, a practical challenge arises in translating this technology to clinical use. Specifically, a head holder – typically a U-shaped carbon-fiber support as shown in Fig. 1(a) – is typically used to support the head and minimize motion during the scan. However, the head holder can be partially truncated in the projection data (even for the fairly large – 43 x 43 cm^2 – detector employed on the prototype), and the amount of truncation varies depending on the separation (e.g., a pillow) between the head and the holder. In this respect, the head holder is truncated in the axial plane. In the z direction, the head holder may or may not be longitudinally truncated at its superior extent (top of the head) and is certainly truncated at

the inferior extent (below the neck); however, such longitudinal truncation is not studied in the current work. Axial truncation introduces artifacts in the reconstructed image as shown in Fig. 1, giving rise to nonuniformity that could hinder ICH detection. MBIR methods tend to be sensitive to truncation, since they attempt to solve for an image estimate that best matches all of the measurements.

A variety of strategies to mitigate truncation effects have been investigated. For example, lateral extrapolation of the projection data prior to MBIR has been proposed, including symmetric mirroring (Ohnesorge *et al* 2000), approximation as a scalable water cylinder (Hsieh *et al* 2004), elliptical fitting (Kolditz *et al* 2010), and using scout images to constrain anatomical boundaries (Xia *et al* 2014). These methods have demonstrated reduction of truncation effects to various extents but usually assume the main source of truncation is the patient, and the missing projection data are treated as a continuous extension of the projection of the patient at the edge of the detector. These assumptions may not hold well when the truncation is primarily due to patient support – e.g., the head holder as in Fig. 1. Alternatively, truncation effects can be mitigated by increasing the reconstruction field-of-view (RFOV). A large RFOV provides space for "tomosynthesis-like" reconstruction of the truncated object outside the scan field-of-view (SFOV) – i.e., the central (untruncated) region – and reduces bias within the SFOV. One advantage of this method is that it does not require additional processing of the projection data (e.g., extrapolation). However, simply increasing the RFOV increases the computational cost of MBIR. In cases where the truncated object is relatively far from the patient (e.g., a thick pillow inserted between the patient and the head holder), a RFOV much larger than the SFOV may be needed to mitigate truncation effects, posing a significant burden to computation time and memory.

We propose a method to mitigate truncation effects by a multi-resolution reconstruction approach, thereby extending the RFOV without major increase in computational burden. Specifically, an image volume is defined to contain two regions: 1) a fine interior region containing the region of interest (i.e., the head) with voxel size appropriate to the diagnostic task; and 2) a coarse outer region that can be extended as much as needed to mitigate truncation and coarser voxel size to reduce computational load. Multi-resolution MBIR has been studied previously in a 2D digital phantom by Hamelin *et al*. (Hamelin *et al* 2007) and applied to region-of-interest (ROI) reconstruction of high-resolution bone morphology by Cao *et al*. (Cao *et al* 2015). In this work, the multi-resolution approach is incorporated into the penalized weighted least-squares (PWLS) framework previously developed for high-quality CBCT of the head (Dang *et al* 2015). Accordingly, the previously reported scatter correction method was also modified to account for the presence of the head holder. The method was evaluated in CBCT scans of an anthropomorphic head phantom with varying degrees of realistic data truncation by a carbon-fiber head support.

2. Methods

2.1 Multi-resolution PWLS for high-quality head imaging

Conventional PWLS methods (Sauer and Bouman 1993) usually model an image volume μ as a 3D region containing voxels with a fixed voxel size. In this work, we model μ as a combination of an inner 3D rectangular region with a fine voxel size (referred to as the "fine

region" or μ_F) and an outer 3D rectangular shell with a coarser voxel size (referred to as the "coarse region" or μ_C). Figure 2 illustrates the two regions in the multi-resolution method in imaging of the head. The fine region is defined to cover the SFOV of the CBCT system, while the coarse region is defined to cover objects that are outside the SFOV and subject to truncation. The combination of both regions defines the RFOV. For CBCT of the head, the anatomy is entirely within the SFOV (i.e., the fine region), while the head holder spans the fine and/or coarse regions of the RFOV. The resulting boundary between the fine and coarse regions is outside the cranium (in air, presumably not of diagnostic interest), so downsampling / upsampling voxels in the other region is not considered when calculating neighboring voxel differences in the subsequent image reconstruction. In the current work, we investigate an implementation specifically with two voxel sizes (coarse and fine), but the term "multi-resolution" (c.f., "dual resolution") is used for consistency with previous work (Delaney and Bresler 1995, Cao *et al* 2016) and for generality in anticipation of future work in which voxel size is more continuously varied from a fine value in the SFOV to progressively coarser values outside the SFOV.

Following Cao *et al.* (Cao *et al* 2015), one can write the forward model for multi-resolution PWLS reconstruction as follows, assuming independent measurements:

$$\bar{y}=\mathbf{D}\left(g\right)\exp\left(-\tilde{\mathbf{A}}\mu\right)=\mathbf{D}\left(g\right)\exp\left(-\begin{bmatrix}\mathbf{A}_F & \mathbf{A}_C\end{bmatrix}\begin{bmatrix}\mu_F\\ \mu_C\end{bmatrix}\right) \quad (1)$$

where the mean measurements are modeled by $\bar{y}$ (a $N_y \times 1$ vector), g is a $N_y \times 1$ vector of measurement-dependent gains, and $\mathbf{D}(\cdot)$ is an operator that places a vector on the main diagonal of a matrix. The notation $\tilde{\mathbf{A}}$ denotes a system matrix representing the linear projection operation (and $\tilde{\mathbf{A}}^{\mathrm{T}}$ denotes the matched backprojection operation), which consists a $N_y \times N_{\mu F}$ system matrix $\mathbf{A}_F$ for the fine region μ_F and a $N_y \times N_{\mu C}$ system matrix $\mathbf{A}_C$ for the coarse region μ_C. The resulting line integral estimate $\tilde{\mathbf{A}}\mu$ is thus a sum of the line integral estimate from the fine region (i.e., $\mathbf{A}_F\mu_F$) and that from the coarse region (i.e., $\mathbf{A}_C\mu_C$).

The objective function for multi-resolution PWLS reconstruction can be written:

$$\hat{\mu}=\arg\min_{\mu\geq 0}\frac{1}{2}\|\tilde{\mathbf{A}}\mu-l\|_{\mathbf{W}}^{2}+\beta_F R_F\left(\mu\right)+\beta_C R_C\left(\mu\right) \quad (2)$$

where l denotes a vector of line integrals, and $\mathbf{W}$ is a diagonal weighting matrix with the i^{th} diagonal element $\mathbf{W}_i$ representing the fidelity of the i^{th} measurement. The terms R_F (R_C) and β_F (β_C) are the regularization term and regularization parameter for the fine region (the coarse region).

The two regularization terms enforce image smoothness in the fine and coarse regions respectively, which can be defined as:

$$R_F(\mu) = \frac{1}{2} \sum_{k \in K_F} H\left((\Psi_F \mu_F)_k\right) \quad (3)$$

$$R_C(\mu) = \frac{1}{2} \sum_{k \in K_C} H\left((\Psi_C \mu_C)_k\right) \quad (4)$$

where K_F (K_C) denotes the indices of voxels in the fine region (coarse region), Ψ_F (Ψ_C) is an operator that computes first-order neighborhood differences in the fine region (coarse region), and $H(\cdot)$ is Huber penalty function (Huber 1981) which is quadratic within a neighborhood of $[-\delta, \delta]$ and linear for larger differences as in (Wang *et al* 2014). Separate regularization terms for the fine and coarse regions allow independent control of the regularization strength. Calculation of neighborhood differences for voxels near the boundary between fine and coarse regions downsamples (or upsamples) neighboring voxels in the other region. This downsampling / upsampling operation is especially important when the boundary contains anatomy of interest (e.g., bone morphology in (Cao *et al* 2015)); however, in the scenario considered here, the boundary is outside the cranium (in air), so downsampling / upsampling at the boundary was not considered in the current work.

In the data fidelity term of the multi-resolution PWLS objective, the line integrals in l are typically derived from raw measurements y through a number of steps. Such steps include a log transformation to convert from the measurement domain to the line integral domain, and in many situations, also include correction for artifacts and/or processing to reduce noise in the measurements (Li *et al* 2004). Such steps can potentially lead to changes in the noise characteristics of the measurements, which need to be accommodated into the PWLS weighting terms. Previous work (Dang *et al* 2015) modeled the processing of the measured data as a generic function f as:

$$l_i = f(y_i) \quad (5)$$

and derived the variance following data processing using first-order Taylor expansion of f:

$$\mathrm{var}(l_i) \approx \left[\dot{f}(\bar{y}_i)\right]^2 \mathrm{var}(y_i) \quad (6)$$

where $\dot{f}$ denotes the derivative. In CBCT of the head, scatter and beam hardening corrections represent the dominant corrections in the artifact correction framework corrections (Sisniega *et al* 2015). The function f in this case thereby corresponds to scatter correction in the measurement domain, followed by log transformation, and then beam hardening correction in the line integral domain. The variance following this particular

function f using Eq. (6) implies modification of the statistical weights as in (Dang *et al* 2015):

$$\mathbf{W}_i = \frac{1}{\mathrm{var}\,(l_i)} \approx y_i \cdot \left(\frac{y_i - \bar{S}_i}{y_i}\right)^2 \cdot \frac{1}{\eta_w\left(y_i - \bar{S}_i\right)} \cdot \frac{1}{\eta_b\left(y_i - \bar{S}_i\right)} \quad (7)$$

The term y_i corresponds to the weights used in conventional PWLS methods that model data processing simply as a log transformation (Sauer and Bouman 1993). The term $((y_i - \bar{S}_i)/y_i)^2$ corresponds to the variance changes following scatter correction, where $\bar{S}_i$ denotes the mean scatter for the i^{th} measurement. For the Joseph-Spital beam-hardening correction method (Joseph and Spital 1978), the terms $1/\eta_w(y_i - \bar{S}_i)$ and $1/\eta_b(y_i - \bar{S}_i)$ correspond to the variance changes following water correction and bone correction, respectively, defined in Eq. (16b) of (Dang *et al* 2015). The statistical weights in Eq. (7) were used in all PWLS reconstructions in this work.

2.2 Optimization approach for multi-resolution PWLS

The multi-resolution PWLS objective in Eq. (2) was solved using the separable quadratic surrogate (SQS) algorithm with ordered subsets (OS) (Erdoğan and Fessler 1999). The OS-SQS algorithm facilitates fast convergence not only via ordered subsets (nominally 10 subsets from 360 projection data; see below) but also via parallelizable image updates allowing parallel implementation on GPU. Previous work (Dang *et al* 2015) adapted OS-SQS to the single-resolution PWLS objective with modified statistical weights. For the multi-resolution case, for every subset of projections in every iteration, the image update was computed and applied to the fine and coarse regions separately. Moreover, since the optimal curvature c_j of the data fidelity term is constant in PWLS objective, the term d in the image update can be precomputed (d_{Fj} and d_{Cj} for the fine and coarse regions, respectively).

Table 1 shows pseudocode of the OS-SQS solution of the multi-resolution PWLS objective. The pseudocode is similar to that in (Cao *et al* 2016), updated with respect to notation and detector pixel model. The notation $[\cdot]_+$ denotes the nonnegativity constraint, γ_i is the i^{th} projection of an image of all ones, a_{Fij} and a_{Cij} are the $(i, j)^{\text{th}}$ element of the matrix $\mathbf{A}_F$ and $\mathbf{A}_C$ respectively, n_{iter} is the maximum number of iterations, M is the number of subsets, $\hat{l}_i$ is the sum of the projection of the current image estimate $\hat{\mu}_F$ and $\hat{\mu}_C$, and S_m denotes all the projections in the m^{th} subset. In the regularization part, taking the fine region as an example, K_F is the number of neighboring voxels in the fine region, and $\dot{H}$ and ω are the gradients and curvatures of the Huber penalty function H, respectively. While the pseudocode in (Cao *et al* 2016) used a small detector pixel size for the projection of a high-resolution region-of-interest and a large detector pixel size for the rest of the projection data, the pseudocode here used a single detector pixel size.

2.3 Experimental studies

The method was tested in phantom experiments performed on the CBCT test-bench shown in Fig. 3(a). The bench includes an x-ray source (RAD13, Dunlee, Aurora IL) and flat-panel detector (PaxScan 4343R, Varian, Palo Alto CA) in geometry equivalent to that of the

prototype head scanner (Xu *et al* 2016c): 550 mm source-to-axis distance (SAD) and 1000 mm source-to-detector distance (SDD). Scans were acquired at 90 kVp, 0.8 mAs per projection, with 360 projections (1° angular steps), and a 0.556 × 0.556 mm^2 pixel size (after 2 × 2 binning). The radiation dose was measured using a Farmer chamber in an extended length CTDI phantom of 16 mm diameter, weighting the central (D_o) and mean peripheral (D_p) dose according to $D_w = (1/3)D_o + (2/3)D_p$. The dose measured with no head holder in place was 26.8 mGy. Adding the carbon-fiber head holder reduced the dose slightly to 25.8 mGy, which was independent of Locations 1, 2, 3 (Fig. 3b) within 1%. This dose is comparable to that for scan protocols used in clinical studies using the prototype head scanner, 22.8 mGy (Xu *et al* 2016a). The anthropomorphic phantom and head holder emulated a typical clinical setup in which the head was fully covered by the (23.7 × 23.7 × 23.7 cm^3) SFOV, but the head holder was truncated to varying extent. The head phantom (The Phantom Laboratory, Greenwich NY) included a natural skull and tissue-equivalent plastic (Rando™). The head holder (Siemens AG, Forchheim, Germany) was a carbon fiber (~150 HU) unit identical to that used on routine head CT exams. The phantom was scanned with the head holder placed at three locations as illustrated in Fig. 3, increasing in anterior-posterior distance from the head in increments of 2.54 cm in a manner that emulated a broad range of clinically realistic setup (e.g., varying amount of padding beneath the head). A scan was also acquired without the head holder to provide a truncation-free dataset.

All projection data were first offset-corrected and gain-normalized by mean dark and flood field calibrations. Scatter correction involved a fast Monte Carlo method integrated with beam hardening correction using the Joseph and Spital approach. Previous work (Sisniega *et al* 2015) validated the scatter and beam hardening correction without a head holder. We added a head holder model to the Monte Carlo scatter simulation to estimate the scatter from the head holder in addition to the scatter from the head. The head holder model (3D map of attenuation coefficient) was obtained from a separate CT scan of the head holder using a diagnostic CT scanner (SOMATOM Definition, Siemens Healthineers, Erlangen, Germany) with a SFOV sufficient to cover the entire head holder (i.e., without truncation). This separate scan yielded an accurate attenuation map of the head holder. In the Monte Carlo scatter simulation, the head holder was added to the system geometry based on its position as evident in the projection data. In the current work, the position of the head holder model was manually adjusted for each scan, but was subsequently automated by detecting the long edges of the holder in the scan data and computing a rigid 3D-2D registration.

Projection data were reconstructed using both conventional single-resolution PWLS and the proposed multi-resolution PWLS method. Both methods used matched separable footprint projectors and backprojectors (Long *et al* 2010) and 10 ordered subsets. A total of 50 iterations was found sufficient for convergence for both PWLS methods. The voxel size for single-resolution PWLS was 0.5 × 0.5 × 0.5 mm^3. For multi-resolution PWLS, the voxel size for the fine region was also 0.5 × 0.5 × 0.5 mm^3, and the voxel size for the coarse region was varied as described in the next section. A RFOV of 400 × 480 × 480 voxels (at isotropic 0.5 mm voxel size) was sufficient to cover the SFOV and was defined as the *basic* RFOV. For multi-resolution PWLS, the fine region was set to the basic RFOV, and the coarse region was varied as described in the next section (equivalent to varying the relative ratio of areas between the fine and coarse regions). This study investigates how extension of the coarse

region of reconstruction outside the head (and the head holder) allows more accurate reconstruction of attenuation coefficient within the SFOV, essentially distributing bias from axial truncation outside regions of interest. In the current work, the boundary between the fine and coarse regions is a fixed value determined by the system geometry (simply equal to the SFOV) and is not a parameter that needs to be manually defined. The water attenuation coefficient was 0.0216 mm^{-1}, and the Huber parameter δ was set to 10^{-4} mm^{-1}, which enforced a degree of edge-preservation for features such as ICH and ventricles without causing an overly patchy appearance to the images (Dang *et al* 2015).

2.4 Multi-resolution PWLS: parameter selection

Key parameters affecting the performance of multi-resolution PWLS were investigated, including regularization strength in the fine and coarse regions (β_F and β_C), the voxel size in fine and coarse regions (related by the downsampling factor, DS), and the size of the RFOV. For example, previous work (Cao *et al* 2015) in extremity orthopaedic imaging showed that using a coarse region voxel size four times larger than the fine region voxel size yielded accurate ROI reconstruction.

1. Regularization parameter. The parameter β_F controls the noise-resolution tradeoff in the fine region in a similar manner to β in single-resolution PWLS (Wang *et al* 2014, Dang *et al* 2015). The parameter β_C, however, affects the fine region indirectly, and its effect on image quality was investigated as a function of downsampling factor, RFOV, and location of the head holder.
2. Downsampling factor. The ratio of the voxel size in the coarse region to that in the fine region defined the downsampling factor (DS), which is expected to control the amount of speedup in multi-resolution PWLS. In the studies presented below, multi-resolution PWLS reconstructions were performed with the fine voxel size fixed at 0.5 mm, and DS was varied from 1 to 40.
3. Reconstruction field-of-view (RFOV). Extending the RFOV is expected to reduce truncation effects but increase reconstruction time. In the work reported below, multi-resolution PWLS images were reconstructed for RFOV ranging from the basic SFOV to a much larger RFOV, and the impact on image quality and reconstruction time were evaluated.

2.5 Figures of merit and computational complexity

The accuracy of image reconstruction was defined as the root mean square difference (RMSD) from a “truth” image, restricted to a region of the image within the cranium (i.e., in the brain). The “truth” image was defined as a single-resolution PWLS image reconstructed from the “no-holder” dataset (i.e., free of truncation effects). Spatial resolution was also

assessed as in (Wang *et al* 2014) in terms of the width ε (mm) of the edge spread function (ESF) of a low-contrast sphere within the brain [see Fig. 3(f)]. Contrast and contrast-to-noise ratio (CNR) were evaluated with respect to a 50 HU sphere and nearby uniform ROI [see Fig. 3(f)].

The computational complexity of both single-resolution and multi-resolution PWLS methods are primarily determined by the total number of projection operations (including forward projection and backprojection). For single-resolution PWLS, one iteration of OS-SQS algorithm as shown in Table 1 in (Dang *et al* 2015) requires two projection operations (one forward projection and one backprojection) for the entire RFOV, which can be written as:

$$T_{single} = (2T_{RFOV}) \cdot M \quad (8)$$

where T_{RFOV} denotes the time for one projection operation (for one forward projection and one backprojection) for the entire RFOV, and M is the number of subsets. For multi-resolution PWLS, one iteration of the OS-SQS algorithm requires two projection operations for both the fine and coarse regions, giving:

$$T_{multi} = (2T_F + 2T_C) \cdot M \quad (9)$$

where T_F and T_C denote the time for one projection operation for the fine and coarse region, respectively. Assuming the same RFOV, multi-resolution PWLS is expected to require less computation time than single-resolution PWLS, since projection operations at the fine voxel size are performed only for the fine region ($2T_F$) for multi-resolution PWLS, but are performed for the entire RFOV ($2T_{RFOV}$) for single-resolution PWLS. Although multi-resolution PWLS requires two additional projection operations for the coarse region, the time associated with these two operations ($2T_C$) is expected to be small.

Both PWLS methods were implemented in Matlab (The Mathworks, Natick MA), with projection operations executed on GPU using CUDA-based libraries. All image reconstructions were performed on a workstation equipped with a GeForce GTX TITAN (Nvidia, Santa Clara CA) graphics card.

3. Results

3.1 Scatter correction with a head holder model

The previously developed Monte Carlo scatter correction method was modified to include a model of the head holder, with results summarized in Fig. 4. The head holder was truncated in the three scans at Location 1 to 3 in Fig. 4. Here, we focus on the evaluation of scatter artifacts (not truncation artifacts), so the results show FBP reconstructions for simplicity which appears to be somewhat less sensitive to truncation than PWLS reconstructions. As shown in Fig. 4(a–c), ignoring the head holder in the scatter correction model resulted in residual artifacts – evident primarily as shading, streaks, and overall underestimation in HU

as shown in Fig. 4(d–f). In comparison, the results in Fig. 4(g–i) show that including the head holder model in scatter correction yielded images with improved uniformity for all three locations of the head holder. The "truth" image (FBP image with the head holder removed during the scan) is shown in Fig. 4(m). Including the head holder in the scatter correction method reduced the RMSD from the "truth" image from 15.1×10^{-4} mm^{-1} to 10.2×10^{-4} mm^{-1} at location 1, 13.5×10^{-4} mm^{-1} to 9.4×10^{-4} mm^{-1} at location 2, and 11.7×10^{-4} mm^{-1} to 9.1×10^{-4} mm^{-1} at location 3. In subsequent results reported below, the head holder model was always included in the scatter correction.

3.2 Single-resolution PWLS

The influence of truncation on the image quality of single-resolution PWLS reconstructions was first investigated. First, the nominal β value suitable for CBCT of the head was selected using the "no-holder" dataset. Figure 5(a) plots the ESF width and CNR measured as a function of β, showing a steep increase in CNR for $\beta > 10^2$, owing to the Huber penalty as shown in previous work (Wang *et al* 2014). A nominal value of $\beta = 10^{2.4}$ was selected as balancing noise reduction and edge preservation without overly patchy image appearance, giving CNR = 25.4 and ESF width = 0.77 mm for the ROIs shown in Fig. 3. The resulting image [Fig. 5(b)] was taken as "reference" / "truth" in subsequent results.

Next, single-resolution PWLS reconstructions were computed with the head holder at three locations as shown in Fig. 6. Severe artifacts – including both positive and negative bias – are evident throughout the head, attributable to truncation by the head holder (not to x-ray scatter, which was corrected with the head holder model as summarized in the previous section and Fig. 4). The magnitude of truncation artifacts is seen to depend on the position of the head holder with respect to the head – i.e., somewhat stronger artifacts for location 1, and reduced for location 3. These artifacts appear to be associated with the truncated anterior edges of the head holder, giving rise to shading and streaks in the anterior part of the head as shown in the difference image of Fig. 6. The streaks are strongest for location 1 (where the edges are closest to the head) and reduced as the head holder was positioned toward the posterior of the head (location 3). Because the streaks appear to arise from the edges of the holder, they are most severe for location 1, even though location 3 involves a greater bulk of material attenuation farther from the SFOV. The RMSD from "truth" was 11.0×10^{-4} mm^{-1} at location 1, 10.2×10^{-4} mm^{-1} at location 2, and 8.2×10^{-4} mm^{-1} at location 3.

3.3 Multi-resolution PWLS

The sections below systematically evaluate the performance of multi-resolution PWLS in the presence of truncation. First, a very large RFOV was chosen (1000 × 1000 × 1000 voxels, with 0.5 mm isotropic voxel size), and the effects of regularization parameter β_C (Section 3.3.1) and downsampling factor DS (Section 3.3.2) were studied. Based on that analysis, nominal values of β_C and DS were selected, and the dependence of image quality on RFOV (Section 3.3.3) was investigated.

3.3.1 Regularization parameters—Figure 7 plots the RMSD of multi-resolution PWLS as a function of β_C and DS for the three locations of the head holder. The value of β_F was fixed at $10^{2.4}$ for each case. The reconstruction accuracy exhibited low dependence on β_C as

long as β_C was below an upper limit in regularization strength, but quickly reduced when β_C exceeded this limit. This was observed for all locations of the head holder and all DS levels. The rapid degradation in reconstruction accuracy beyond a regularization limit was also observed in previous work (Cao *et al* 2015). Note that for any level of DS, the regularization limit was the same for different locations of the head holder, suggesting that in practice the exact location of the head holder does not affect selection of β_C.

Figure 8 shows the fine (a–c) and coarse (d–f) regions of the multi-resolution PWLS images for three values of β_C. In each figure (a–c), the left half is the PWLS image, and the right half is the difference from "truth". For the cases in Fig. 8, the head holder was at location 2, and the DS was set to 4. As evident in Fig. 8(a) and 8(b), PWLS exhibited fairly accurate reconstruction for a broad range of β_C below or near the regularization limit (~$10^{6.4}$), but performance degraded markedly for β_C above the limit (Fig. 8(c)). Figures 8(d–f) show the amount of smoothing in the coarse region (outside cyan box) for the three β_C values. It can be seen that the use of β_C beyond the regularization limit resulted in over-smoothing and low intensity error throughout the air region, which was a possible cause of the reduced accuracy in the fine region.

3.3.2 Downsampling factor—Figure 9 shows that the reconstruction accuracy was robust as DS was increased from 4 to 20, beyond which we observe moderate degradation (e.g., at DS = 40). In the cases shown, the head holder was at location 3, and the β_C value was selected to achieve the lowest RMSD for each DS (fairly insensitive to selection as shown in Fig. 7). A small increase in streak artifacts is evident with increasing DS. Figures 9(e–h) show the coarse regions for various DS levels. Taking DS = 40 as an example, the coarse region involves a very large voxel size (20 × 20 × 20 mm^3), which led to coarse reconstruction of the head holder and likely led to the streaks observed in the fine region. Despite the small reduction in accuracy as DS increased, the reconstruction accuracy at all four DS levels was still much better than the single-resolution PWLS image (Fig. 6).

3.3.3 Reconstruction field-of-view—The RFOV of multi-resolution PWLS was varied as summarized in Fig. 10, which effectively varies the size of the coarse region while keeping the size of the fine region equal to the basic RFOV. Specifically, the RFOV was varied from (400 × 480 × 480 voxels) to ~10 times as large (1000^3 voxels). In the cases shown in Fig. 10, the head holder was at location 2, the DS was 4, and β_C was selected to minimize RMSD for each RFOV. The noise-resolution tradeoff in the fine region was found to exhibit small changes for varying the RFOV (which was not observed when varying β_C and DS), so the β_F value was selected to give the same noise-resolution performance for each RFOV. Images reconstructed using any of the three RFOV in Fig. 10 exhibited similarly high reconstruction accuracy compared to the basic RFOV (single-resolution PWLS) shown in Fig. 6. This again shows the benefit of increasing RFOV to mitigate truncation effects and suggests that one could freely choose a RFOV (> ~600 × 600 × 600 voxels) to mitigate truncation effects.

3.4 Computation time

The computation time between single-resolution and multi-resolution PWLS reconstruction is summarized in Fig. 11 in terms of the measured time per projection operation (averaged over one forward projection and one backprojection) and memory usage as a function of RFOV (for isotropic voxel size of 0.5 mm). As shown in Fig. 11(a), both time and memory usage increase dramatically if the RFOV increases from the basic RFOV (denoted by the dashed line), suggesting that simply increasing the RFOV in single-resolution PWLS is computationally expensive and likely impractical. Figure 11(b) plots the measured time per iteration as a function of RFOV for various PWLS reconstruction methods. Single-resolution PWLS exhibited a steep increase in time per iteration with larger RFOV, consistent with the steep increase in Fig. 11(a). By comparison, the time per iteration was much reduced for the multi-resolution approach at DS = 2 and was further reduced at DS = 4. The time per iteration was reduced from the single-resolution approach (DS = 4) by 40.7% for a RFOV of 600^3 voxels, 83.0% for a RFOV of 800^3 voxels, and over 95% for a RFOV of 1000^3 voxels. Moreover, the time per iteration became almost independent of RFOV when DS increased to 4 or larger in the multi-resolution approach. This suggests that one could increase the RFOV as much as needed to mitigate truncation effects in the multi-resolution approach without corresponding increase in computational complexity. The results clearly demonstrate the advantages of multi-resolution reconstruction, especially in situations where a large RFOV is needed. Because an increase in DS larger than 4 gradually reduced the reconstruction accuracy as shown in Section 3.3.2 (but does not correspondingly reduce the computation time), DS = 4 was selected as the nominal / optimal DS.

3.5 Comparison of reconstruction methods

Figure 12 shows a single-resolution PWLS reconstruction using (a–b) the basic RFOV and (c–d) an extended RFOV in comparison to (e–f) multi-resolution PWLS reconstruction (with the same extended RFOV). In this case, the head holder was at location 2. The multi-resolution PWLS reconstruction parameters were: (1) $\beta_C = 10^{6.4}$, which is near the upper range of stable regularization identified in Section 3.3.1; (2) DS = 4, as indicated in Sections 3.3.2 and 3.4; and (3) an extended RFOV of 600^3 voxels (assuming isotropic voxel size of 0.5 mm), which the results of Section 3.3.3 identify as the smallest RFOV providing good mitigation of truncation artifacts. Single-resolution PWLS using the basic RFOV exhibits severe artifacts due to truncation, whereas single-resolution PWLS with an extended RFOV substantially reduced such artifacts, but doubled computation time. Multi-resolution PWLS using the same extended RFOV exhibited visually and quantifiably similar reduction of truncation effects and only increased computation time by 12% (109 sec/iter vs. 97 sec/iter) compared to reconstruction with the basic RFOV.

4. Conclusions and Discussion

We have proposed a multi-resolution MBIR method to mitigate truncation effects and applied to CBCT of the head in which the main source of truncation is the patient support / head holder. While conventional reconstruction methods employ a fixed voxel size throughout the image, the multi-resolution method uses a fine voxel size within the untruncated region (i.e., inside the SFOV encompassing the area of interest – in this case, the

head) and a coarse voxel size in the truncated region outside the SFOV (i.e., outside the area of interest). The approach was implemented in a PWLS reconstruction framework and evaluated in experiments involving a head phantom imaged on a CBCT test-bench with varying levels of truncation using a commercially available carbon-fiber head holder. The multi-resolution method demonstrated substantial mitigation of truncation effects and major reduction in computational cost compared to single-resolution reconstruction with an extended RFOV.

Investigation of the main algorithm parameters suggest that: (1) reconstruction accuracy in the fine region (the head) does not depend strongly on the regularization parameter in the coarse region β_C as long as the parameter is below a regularization "limit," which in turn was found not to depend on the location of the head holder and can therefore be held fixed; (2) use of a larger voxel size in the coarse region (larger DS) reduces computational complexity but slightly reduces reconstruction accuracy, suggesting an optimal DS such that the voxel size in the coarse region was ~4 times that in the fine region; and (3) reconstruction accuracy improved with a larger RFOV up to a certain extent (600^3 voxels assuming isotropic voxel size in this work) beyond which accuracy was modestly improved. In the current work, truncation was due solely to the head holder (which varied in location but not in size or mass), and more severe truncation (i.e., greater mass of attenuation outside the SFOV) may require larger RFOV. In summary, the method presents a promising means to mitigate truncation effects in CBCT of the head and supports translation of a newly developed CBCT head scanner in point-of-care imaging applications.

A variety of alternative methods to managing truncation artifacts have been reported. For example, some methods treat missing projection data as a continuous extension of the projection at the edge of the detector and extrapolate the missing data before image reconstruction (Ohnesorge *et al* 2000, Hsieh *et al* 2004, Kolditz *et al* 2010, Xia *et al* 2014). These methods have demonstrated reduction of truncation effects to various extents but the assumption on the continuous extension of the projection may not hold well when the truncation is primarily due to patient support. Other methods attempt to directly reconstruct a ROI inside the patient anatomy that has not been truncated during the scan (Noo *et al* 2004, Defrise *et al* 2006, Kudo *et al* 2008, Zou and Pan 2004, Pan *et al* 2005, Yu *et al* 2006, Yu and Wang 2009). For example, a widely recognized approach in ROI reconstruction is to backproject the derivative of the projection data and apply Hilbert filtering along certain lines covering the ROI (Noo *et al* 2004, Zou and Pan 2004). These ROI reconstruction methods have demonstrated substantial reduction of truncation artifacts in the ROI, but as analytical methods, they usually do not enjoy the noise-resolution benefits exhibited by MBIR. The method proposed in this work allows more general treatment of the source of truncation than extrapolation-based methods and therefore can be used to manage truncation effects that do not arise from the patient (e.g., due instead to the patient support). Moreover, the proposed method is formulated within a MBIR framework, which allows the use of advanced system models and regularization techniques.

An alternative method to mitigate truncation effects in CBCT of the head is to include a model of the head holder within the image reconstruction process. For example, previous work (Stayman *et al* 2012) reported a known-component reconstruction (KCR) approach

that could be extended to include the known shape of the head holder. This could yield even better agreement with the measured projection data and mitigate truncation artifacts in the image. Moreover, the multi-resolution approach proposed above could be combined with the KCR approach to improve computational efficiency.

The multi-resolution method presents a more efficient means to recover attenuation information from truncated objects than simple extension of the RFOV. This is a particularly important consideration in MBIR, which can be sensitive to truncation effects not only in terms of artifacts and accuracy of reconstruction but also in the speed and stability of convergence. MBIR also carries a fairly high computational burden, and straightforward extension of the RFOV could lead to impractical reconstruction time. The current work focused on a particular form of data truncation encountered in CBCT of the head but offers a potential general solution for other scenarios in CT or CBCT. In C-arm CBCT for interventional imaging, for example, the patient periphery, interventional tools, and operating table are often truncated due to the limited SFOV. Moreover, in diagnostic imaging, truncation can occur for obese patients or (purposeful or inadvertent) setup of the patient off center. Such scenarios are the subject of future work, where the proposed method may offer a means to mitigate truncation effects without major increase in computational cost.

Acknowledgments

This work was supported by academic-industry partnership with Carestream Health (Rochester, NY) and by the National Institutes of Health Grant R01-EB-018896. The authors thank Mr. Qian Cao for useful discussion of multi-resolution image reconstruction.

References

Boone JM, Nelson TR, Lindfors KK, Seibert JA. Dedicated Breast CT: Radiation Dose and Image Quality Evaluation. Radiology. 2001; 221:657–67. [PubMed: 11719660]

Cao, Q.; Zbijewski, W.; Yorkston, J.; Siewerdsen, JH.; Stayman, JW. Multi-Resolution Penalized Weighted Least-Squares Reconstruction for Quantitative Cone-Beam CT Imaging of Bone Morphology. International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine; 2015. p. 452-5.

Cao Q, Zbijweski W, Sisniega A, Yorkston J, Siewerdsen JH, Stayman JW. Multiresolution iterative reconstruction in high-resolution extremity cone-beam CT. Phys Med Biol. 2016; 61:7263–81. [PubMed: 27694701]

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; 270:816–24. [PubMed: 24475803]

Dang H, Otake Y, Schafer S, Stayman JW, Kleinszig G, Siewerdsen JH. Robust methods for automatic image-to-world registration in cone-beam CT interventional guidance. Med Phys. 2012; 39:6484–98. [PubMed: 23039683]

Dang H, Stayman JW, Sisniega A, Xu J, Zbijewski W, Wang X, Foos DH, Aygun N, Koliatsos VE, Siewerdsen JH. Statistical reconstruction for cone-beam CT with a post-artifact-correction noise model: application to high-quality head imaging. Phys Med Biol. 2015; 60:6153–75. [PubMed: 26225912]

Dang, H.; Stayman, JW.; Xu, J.; Sisniega, A.; Zbijewski, W.; Wang, X.; Foos, DH.; Aygun, N.; Koliatsos, VE.; Siewerdsen, JH. Task-based regularization design for detection of intracranial hemorrhage in cone-beam CT. The Fourth International Conference on Image Formation in X-ray Computed Tomography; 2016. p. 557-60.

Defrise M, Noo F, Clackdoyle R, Kudo H. Truncated Hilbert transform and image reconstruction from limited tomographic data. Inverse Probl. 2006; 22:1037–53.

Delaney AH, Bresler Y. Multiresolution tomographic reconstruction using wavelets. IEEE Trans Image Process. 1995; 4:799–813. [PubMed: 18290029]

Dublin AB, French BN, Rennick JM. Computed tomography in head trauma. Radiology. 1977; 122:365–9. [PubMed: 299951]

Erdoğan H, Fessler JA. Ordered subsets algorithms for transmission tomography. Phys Med Biol. 1999; 44:2835–51. [PubMed: 10588288]

Hamelin, B.; Goussard, Y.; Dussault, JP. Penalized-likelihood region-of-interest CT reconstruction by local object supersampling. Proceedings of the 29th Annual International Conference of the IEEE EMBS; 2007. p. 739-42.

Hsieh J, Chao E, Thibault J, Grekowicz B, Horst A, McOlash S, Myers TJ. A novel reconstruction algorithm to extend the CT scan field-of-view. Med Phys. 2004; 31:2385–91. [PubMed: 15487717]

Huber, PJ. Robust Statistics. New York: Wiley; 1981.

Indeck M, Peterson S, Smith J, Brotman S. Risk, cost, and benefit of transporting ICU patients for special studies. J Trauma. 1988; 28:1020–5. [PubMed: 3135417]

Joseph PM, Spital RD. A method for correcting bone induced artifacts in computed tomography scanners. J Comput Assist Tomogr. 1978; 2:100–8. [PubMed: 670461]

Kidwell CS, Wintermark M. Imaging of intracranial haemorrhage. Lancet Neurol. 2008; 7:256–67. [PubMed: 18275927]

Kolditz D, Meyer M, Kyriakou Y, Kalender WA. Comparison of extended field-of-view reconstructions in C-arm flat-detector CT using patient size, shape or attenuation information. Phys Med Biol. 2010; 56:39–56. [PubMed: 21119229]

Koskinen SK, Haapamäki VV, Salo J, Lindfors NC, Kortesniemi M, Seppälä L, Mattila KT. CT arthrography of the wrist using a novel, mobile, dedicated extremity cone-beam CT (CBCT). Skeletal Radiol. 2013; 42:649–57. [PubMed: 22990597]

Kudo H, Courdurier M, Noo F, Defrise M. Tiny a priori knowledge solves the interior problem in computed tomography. Phys Med Biol. 2008; 53:2207–31. [PubMed: 18401067]

Li T, Li X, Jing W, Wen J, Lu H, Hsieh J, Liang Z. Nonlinear sinogram smoothing for low-dose X-ray CT. IEEE Trans Nucl Sci. 2004; 51:2505–13.

Long Y, Fessler JA, Balter JM. 3D forward and back-projection for X-ray CT using separable footprints. IEEE Trans Image Process. 2010; 29:1839–50.

Masaryk TJ, Kolonick R, Painter T, Weinreb DB. Benefits of Portable Head / Neck CT Imaging in the Intensive Care Unit. Radiol Manage. 2008; 30:50–4.

Navab N, Heining S-M, Traub J. Camera augmented mobile C-arm (CAMC): Calibration, accuracy study, and clinical applications. IEEE Trans Med Imaging. 2010; 29:1412–23. [PubMed: 20659830]

Noo F, Clackdoyle R, Pack JD. A two-step Hilbert transform method for 2D image reconstruction. Phys Med Biol. 2004; 49:3903–23. [PubMed: 15470913]

Ohnesorge B, Flohr T, Schwarz K, Heiken JP, Bae KT. Efficient correction for CT image artifacts caused by objects extending outside the scan field of view. Med Phys. 2000; 27:39–46. [PubMed: 10659736]

Pan X, Zou Y, Xia D. Image reconstruction in peripheral and central regions-of-interest and data redundancy. Med Phys. 2005; 32:673–84. [PubMed: 15839339]

Parizel PM, Makkat S, Van Miert E, Van Goethem JW, Van den Hauwe L, De Schepper AM. Intracranial hemorrhage: Principles of CT and MRI interpretation. Eur Radiol. 2001; 11:1770–83. [PubMed: 11511901]

Penninger RT, Tavassolie TS, Carey JP. Cone-beam volumetric tomography for applications in the temporal bone. Otol Neurotol. 2011; 32:453–60. [PubMed: 21307814]

Sauer K, Bouman C. A local update strategy for iterative reconstruction from projections. IEEE Trans Signal Process. 1993; 41:534–48.

Schafer S, Nithiananthan S, Mirota DJ, Uneri A, Stayman JW, Zbijewski W, Schmidgunst C, Kleinszig G, Khanna AJ, Siewerdsen JH. Mobile C-arm cone-beam CT for guidance of spine surgery: Image quality, radiation dose, and integration with interventional guidance. Med Phys. 2011; 38:4563–74. [PubMed: 21928628]

Siewerdsen JH, Daly MJ, Chan H, Nithiananthan S, Hamming N, Brock KK, Irish JC. High-performance intraoperative cone-beam CT on a mobile C-arm: an integrated system for guidance of head and neck surgery. SPIE Medical Imaging. 2009; 7261:72610J.

Sisniega A, Zbijewski W, Xu J, Dang H, Stayman JW, Yorkston J, Aygun N, Koliatsos V, Siewerdsen JH. High-fidelity artifact correction for cone-beam CT imaging of the brain. Phys Med Biol. 2015; 60:1415–39. [PubMed: 25611041]

Smith I, Fleming S, Cernaianu A. Mishaps during transport from the intensive care unit. Crit Care Med. 1990; 18:278–81. [PubMed: 2302952]

Stayman JW, Otake Y, Prince JL, Khanna AJ, Siewerdsen JH. Model-based tomographic reconstruction of objects containing known components. IEEE Trans Med Imaging. 2012; 31:1837–48. [PubMed: 22614574]

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; 59:1005–26. [PubMed: 24504126]

Xia Y, Hofmann H, Dennerlein F, Mueller K, Schwemmer C, Bauer S, Chintalapani G, Chinnadurai P, Hornegger J, Maier A. Towards clinical application of a laplace operator-based region of interest reconstruction algorithm in c-Arm Ct. IEEE Trans Med Imaging. 2014; 33:593–606. [PubMed: 24595336]

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; 39:4932–42. [PubMed: 22894419]

Xu J, Sisniega A, Zbijewski W, Dang H, Stayman JW, Mow M, Wang X, Foos DH, Koliatsos VE, Aygun N, Siewerdsen JH. Technical assessment of a prototype cone-beam CT system for imaging of acute intracranial hemorrhage. Med Phys. 2016a; 43:5745–57. [PubMed: 27782694]

Xu J, Sisniega A, Zbijewski W, Dang H, Stayman JW, Wang X, Foos DH, Aygun N, Koliatsos VE, Siewerdsen JH. Evaluation of detector readout gain mode and bowtie filters for cone-beam CT imaging of the head. Phys Med Biol. 2016b; 61:5973–92. [PubMed: 27435162]

Xu J, Sisniega A, Zbijewski W, Dang H, Stayman JW, Wang X, Foos DH, Aygun N, Koliatsos VE, Siewerdsen JH. Modeling and design of a cone-beam CT head scanner using task-based imaging performance optimization. Phys Med Biol. 2016c; 61:3180–207. [PubMed: 27025783]

Yang WT, Carkaci S, Chen L, Lai C, Sahin A, Gary J, Shaw CC. Dedicated cone-beam breast CT: feasibility study with surgical mastectomy specimens. Am J Roentgenol. 2007; 189:1312–5. [PubMed: 18029864]

Yu H, Wang G. Compressed sensing based interior tomography. Phys Med Biol. 2009; 54:2791–805. [PubMed: 19369711]

Yu H, Ye Y, Wang G. Practical cone-beam lambda tomography. Med Phys. 2006; 33:3640–6. [PubMed: 17089830]

Zou Y, Pan X. Exact image reconstruction on PI-lines from minimum data in helical cone-beam CT. Phys Med Biol. 2004; 49:941–59. [PubMed: 15104318]

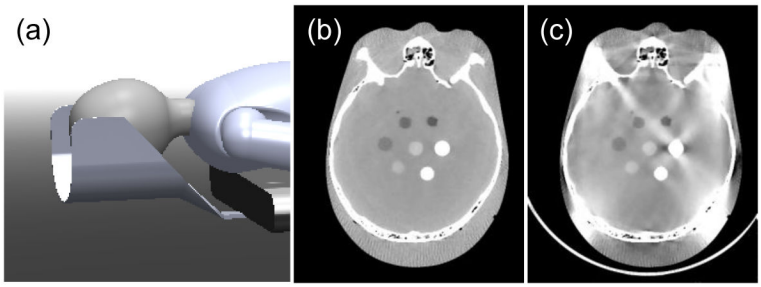

Figure 1.

Artifacts caused by truncation of the head support in CBCT of the head. (a) CAD drawing of a patient with head supported by a carbon-fiber head holder during a CBCT scan. (b) PWLS image of an anthropomorphic head phantom without a head holder. The circular inserts within the central region of the cranium span a range of contrast including that of ICH. (c) PWLS image of the same, with a U-shaped carbon-fiber head holder in place during the scan (evident beneath the posterior of the head).

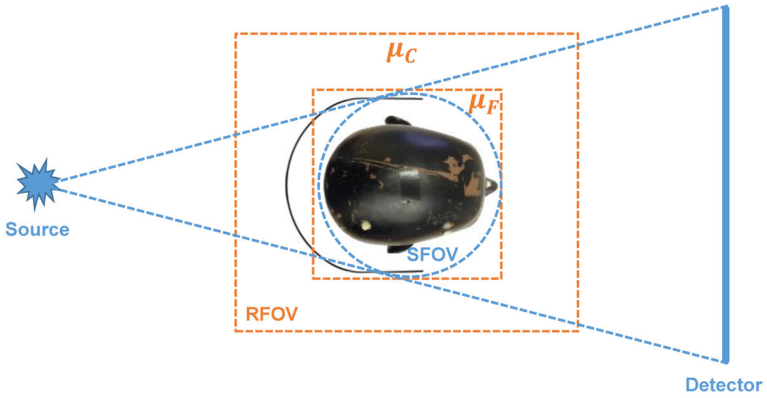

Figure 2.

Illustration of fine and coarse regions in multi-resolution reconstruction (only x-y plane shown here). The dashed circle denotes the SFOV. The fine region (μ_F) is a 3D rectangle that contains the head, and the coarse region (μ_C) is the space outside the fine region that contains the head holder. The volume encompassing both the fine and coarse regions is the RFOV.

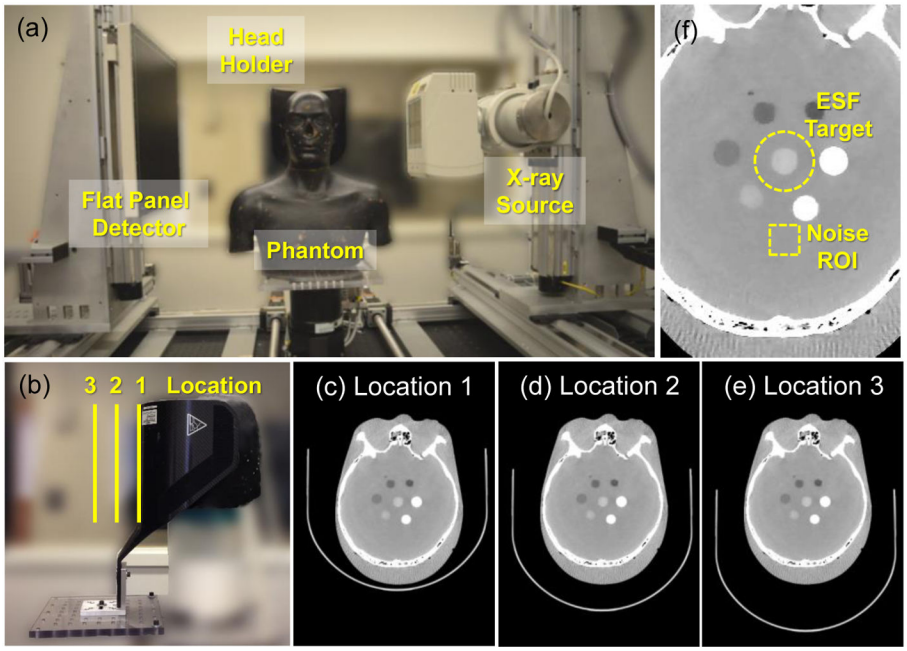

Figure 3.

Experimental setup. (a) Photograph of the CBCT test-bench, head phantom, and head holder. (b) Illustration of three locations at which the head holder was positioned during the experiments. (c–e) Axial images superimposed with a representation of the head holder at each location. (f) Axial image illustrating structures and ROI used for image quality assessment. The central circular insert was used to compute spatial resolution (edge spread function) and contrast, and the nearby rectangular ROI was used to compute noise.

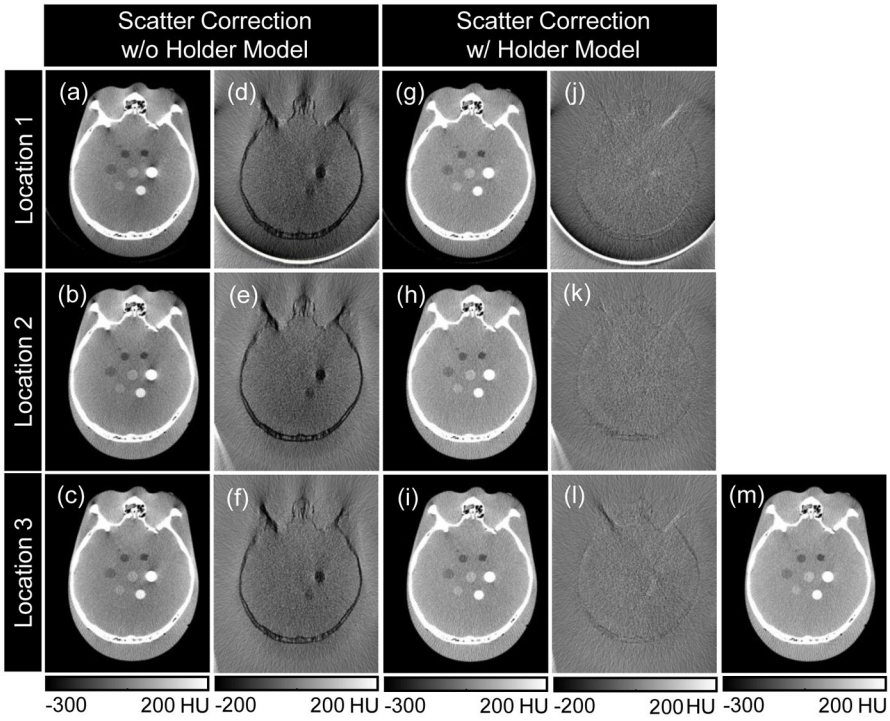

Figure 4.

Scatter correction (a–f) without and (g–l) with the head holder included in the Monte Carlo model. (a–c) FBP reconstructions without a head holder model exhibit shading and streaks in the (d–f) difference images from (m) “truth”. Including the head holder in scatter correction reduces such residual errors as shown in (g–i) and difference images (j–l).

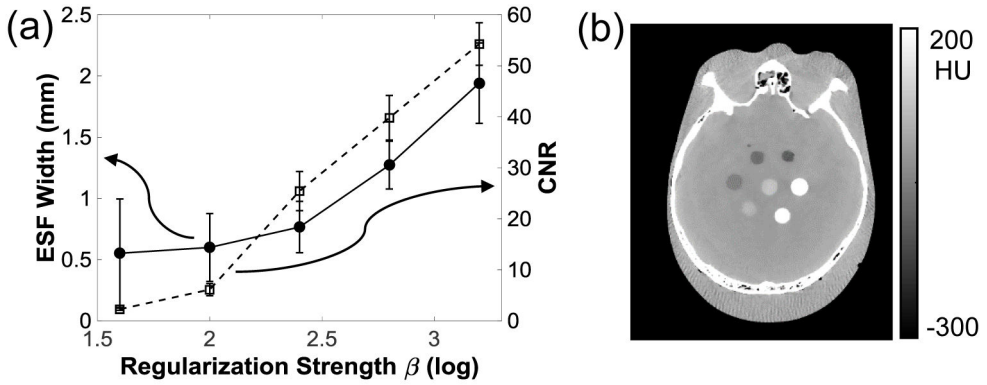

Figure 5.

Nominal parameter selection for single-resolution PWLS. (a) ESF and CNR as a function of the regularization strength β (in the absence of truncation). (b) Axial slice of a single-resolution PWLS image using $\beta = 10^{2.4}$, exhibiting a reasonable balance between ESF and CNR and taken as the "truth" image for subsequent PWLS reconstructions.

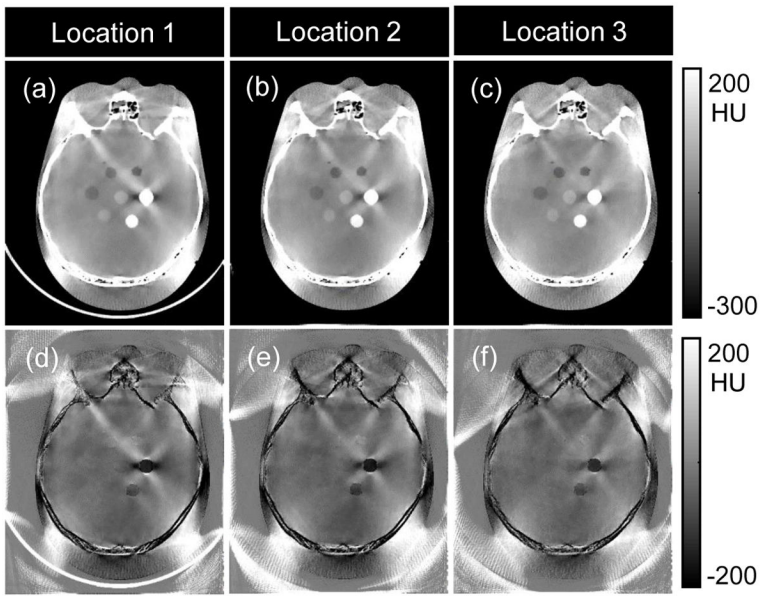

Figure 6.

(a–c) Single-resolution PWLS with a carbon-fiber head holder positioned at three locations posterior to the head. RFOV is (400 × 480 × 480 voxels). (d–f) Difference images between (a–c) and the “truth” image of Fig. 5(b).

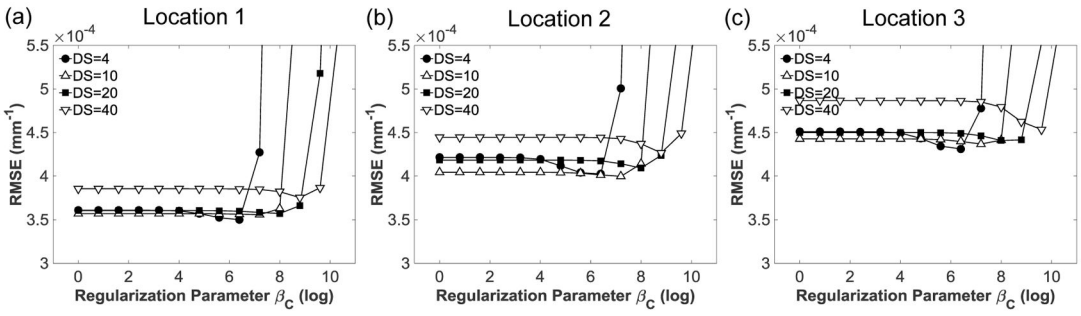

Figure 7.

Accuracy of multi-resolution PWLS reconstructions as a function of coarse region regularization strength β_C and downsampling factor (DS) at three locations.

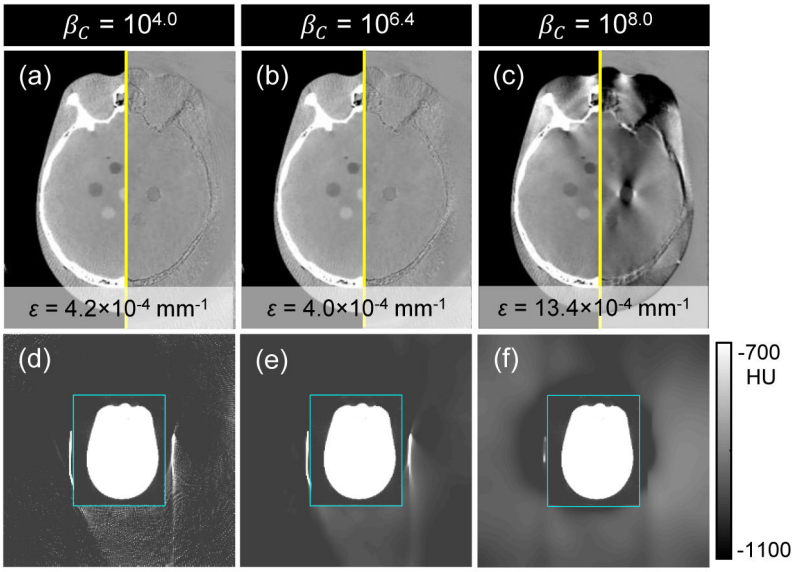

Figure 8.

Multi-resolution PWLS reconstruction for various choices of coarse region regularization strength, β_C. Images (a–c) show the fine region, with the left half showing the PWLS image (grayscale window: [−300, 200] HU) and the right half showing the difference from truth (grayscale window: [−1100 −700] HU). Images (d–f) show the coarse region (outside cyan box). The head holder was at location 2, the DS was 4, and the RFOV was 1000^3 voxels with 0.5 mm isotropic voxel size.

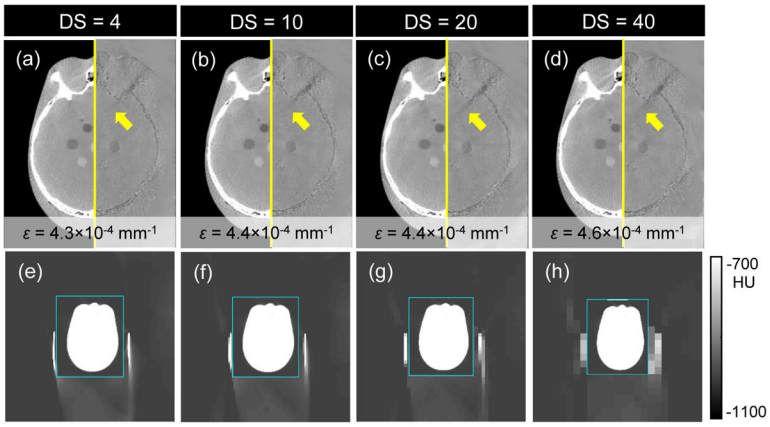

Figure 9.

Multi-resolution PWLS reconstruction for various choices of coarse region voxel size, characterized by DS. Images (a–c) show the fine region, with the left / right presentation of the PWLS (grayscale window: [−300, 200] HU) and difference image (grayscale window: [−1100 −700] HU) as in Figure 8. The head holder was at location 3, the RFOV was 1000^3 voxels (with fine region voxel size = 0.5 mm isotropic), and β_C was chosen to minimize RMSD for each DS.

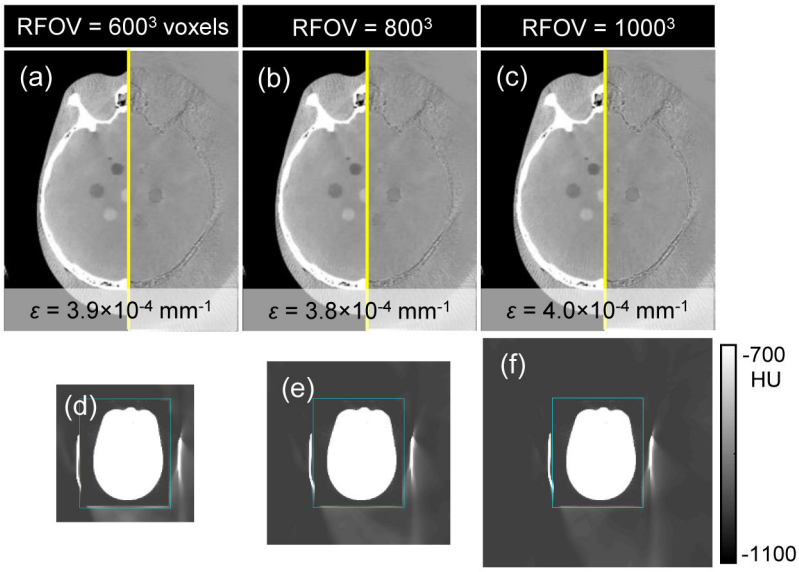

Figure 10.

Multi-resolution PWLS reconstruction for various choices of RFOV. Images (a–c) show the fine region, with the left / right presentation of the PWLS (grayscale window: [−300, 200] HU) and difference image (grayscale window: [−1100 −700] HU) as in Figs. 8–9. The head holder was at location 2, the DS was set to 4, and β_F was adjusted slightly to maintain constant noise-resolution performance: (a) $\beta_F = 10^{2.35}$, (b) $\beta_F = 10^{2.38}$, and (c) $\beta_F = 10^{2.40}$. The parameter β_C was chosen to minimize RMSD for each RFOV (after selecting β_F).

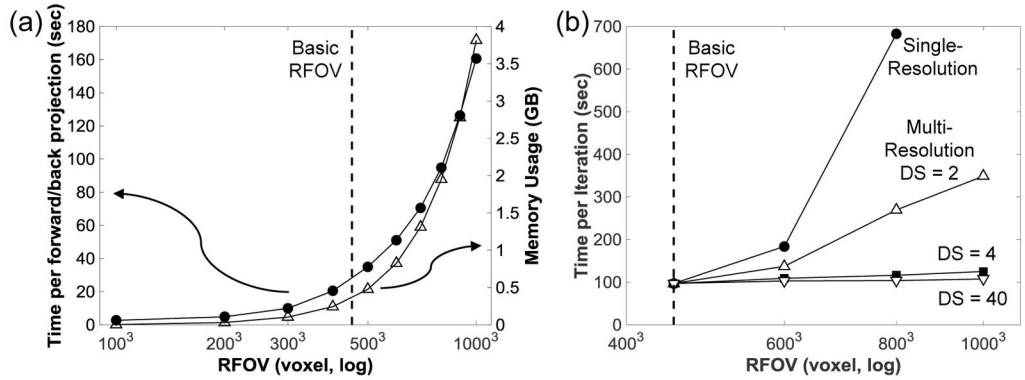

Figure 11.

(a) Measured computation time (averaged over one forward projection and one backprojection) and measured memory usage as a function of reconstruction field-of-view (RFOV). (b) Time per iteration (i.e., for all subsets) as a function of RFOV for single-resolution and multi-resolution PWLS reconstruction at different DS levels. Substantial speedup can be seen compared to the single-resolution approach at DS = 2 and to multi-resolution approach at DS = 4.

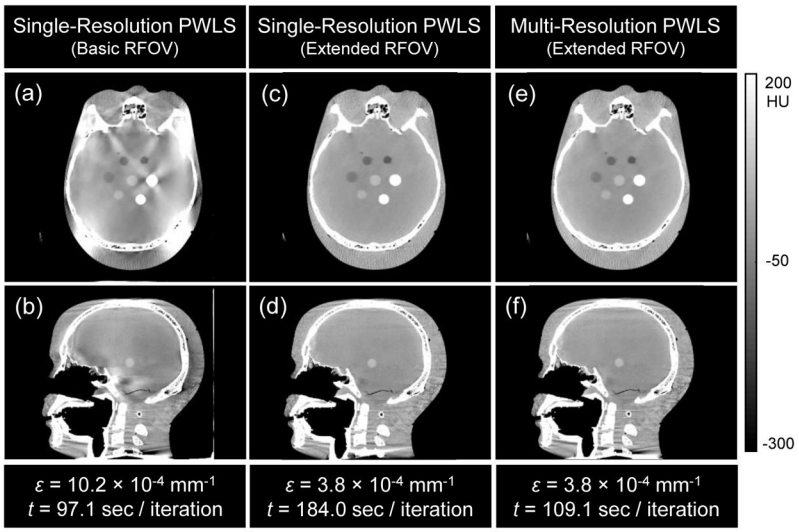

Figure 12.

Comparison of single-resolution and multi-resolution PWLS reconstruction. (a–b) Single-resolution PWLS reconstruction using the basic RFOV (400 × 480 × 480 voxels). (c–d) Single-resolution PWLS reconstruction using an extended RFOV (600^3 voxels assuming an isotropic voxel size of 0.5 mm). (e–f) Multi-resolution PWLS reconstruction using the same extended RFOV. The symbols ε and t denote RMSD and computation time, respectively, quantifying the reduction in artifact using an extended RFOV and the benefit to computation time using the multi-resolution method.

Table 1

Pseudocode for solving the multi-resolution PWLS reconstruction using OS-SQS.

Precompute optimal curvatures $c_i(\hat{l}_i) = [\mathbf{W}_i]_+ \; \hat{l}_i \geq 0$

Precompute approximate d term for the fine region $$d_{F_j} = \sum_{i=1}^{N_y} a_{F_{ij}} \gamma_i c_i \left(\hat{l}_i\right)$$

Precompute approximate d term for the coarse region $$d_{C_j} = \sum_{i=1}^{N_y} a_{C_{ij}} \gamma_i c_i \left(\hat{l}_i\right)$$

for each iteration $n = 1, \ldots, n_{iter}$

for each subset $m = 1, \ldots, M$

$$\hat{l}_i = \sum_{j=1}^{N_{\mu_F}} a_{F_{ij}} \hat{\mu}_{F_j} + \sum_{j=1}^{N_{\mu_C}} a_{C_{ij}} \hat{\mu}_{C_j} \quad \forall i \in S_m$$

$$\dot{h}_i = w_i \left(l_i - \hat{l}_i\right) \quad \forall i \in S_m$$

for voxel in the fine region $j = 1, \ldots, N_{\mu F}$

$$\dot{L}_{F_j} = M \sum_{i \in S_m} a_{F_{ij}} \dot{h}_i, \quad \hat{\mu}_{F_j} = \left[\hat{\mu}_{F_j} + \frac{\dot{L}_{F_j} - \beta_F \sum_{k=1}^{K_F} \left[\mathbf{\Psi}_F\right]_{kj} \dot{H}\left(\left[\mathbf{\Psi}_F \hat{\mu}_F\right]_k\right)}{d_{F_j} + 2\beta_F \sum_{k=1}^{K_F} \left[\mathbf{\Psi}_F\right]_{kj}^2 \omega\left(\left[\mathbf{\Psi}_F \hat{\mu}_F\right]_k\right)} \right]_+$$

End

for voxel in the coarse region $j = 1, \ldots, N_{\mu C}$

$$\dot{L}_{C_j} = M \sum_{i \in S_m} a_{C_{ij}} \dot{h}_i, \quad \hat{\mu}_{C_j} = \left[\hat{\mu}_{C_j} + \frac{\dot{L}_{C_j} - \beta_C \sum_{k=1}^{K_C} \left[\mathbf{\Psi}_C\right]_{kj} \dot{H}\left(\left[\mathbf{\Psi}_C \hat{\mu}_C\right]_k\right)}{d_{C_j} + 2\beta_C \sum_{k=1}^{K_C} \left[\mathbf{\Psi}_C\right]_{kj}^2 \omega\left(\left[\mathbf{\Psi}_C \hat{\mu}_C\right]_k\right)} \right]_+$$

End

End

End