A digitally reconstructed radiograph algorithm calculated from first principles

David Staub^{a)} and Martin J. Murphy

Department of Radiation Oncology, Virginia Commonwealth University, Richmond, Virginia 23298

(Received 16 July 2012; revised 8 November 2012; accepted for publication 9 November 2012; published 14 December 2012)

Purpose: To develop an algorithm for computing realistic digitally reconstructed radiographs (DRRs) that match real cone-beam CT (CBCT) projections with no artificial adjustments.

Methods: The authors used measured attenuation data from cone-beam CT projection radiographs of different materials to obtain a function to convert CT number to linear attenuation coefficient (LAC). The effects of scatter, beam hardening, and veiling glare were first removed from the attenuation data. Using this conversion function the authors calculated the line integral of LAC through a CT along rays connecting the radiation source and detector pixels with a ray-tracing algorithm, producing raw DRRs. The effects of scatter, beam hardening, and veiling glare were then included in the DRRs through postprocessing.

Results: The authors compared actual CBCT projections to DRRs produced with all corrections (scatter, beam hardening, and veiling glare) and to uncorrected DRRs. Algorithm accuracy was assessed through visual comparison of projections and DRRs, pixel intensity comparisons, intensity histogram comparisons, and correlation plots of DRR-to-projection pixel intensities. In general, the fully corrected algorithm provided a small but nontrivial improvement in accuracy over the uncorrected algorithm. The authors also investigated both measurement- and computation-based methods for determining the beam hardening correction, and found the computation-based method to be superior, as it accounted for nonuniform bowtie filter thickness. The authors benchmarked the algorithm for speed and found that it produced DRRs in about 0.35 s for full detector and CT resolution at a ray step-size of 0.5 mm.

Conclusions: The authors have demonstrated a DRR algorithm calculated from first principles that accounts for scatter, beam hardening, and veiling glare in order to produce accurate DRRs. The algorithm is computationally efficient, making it a good candidate for iterative CT reconstruction techniques that require a data fidelity term based on the matching of DRRs and projections.

© 2013 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4769413>]

Key words: digitally reconstructed radiograph, DRR, scatter, beam hardening, veiling glare

I. INTRODUCTION

We are developing a forward iterative projection matching (FIPM) algorithm for deformable 2D/3D registration of raw cone-beam computed tomography (CBCT) projections and fan-beam computed tomography (FBCT) images.¹⁻⁴ This process requires the calculation of digitally reconstructed radiographs (DRRs) from the FBCT for comparison to actual cone-beam radiographic projections. The anatomical features that must be registered, as well as the quantitative manner in which the DRRs and the projections are compared (i.e., the matching metric, or merit function), determine the accuracy with which the two images must match.

For example, rigid 2D/3D registration for radiotherapy applications typically focuses on high-contrast bony features, while a deformable registration algorithm must be sensitive to low contrast soft tissue features. Among other applications of DRR-to-projection matching we take notice of simultaneous algebraic reconstruction techniques and prior-image-constrained remote sensing.^{5,6} Each of these FIPM applications has its own requirements for the DRR quality.

The problem of rigid 2D/3D registration via projection matching to DRRs has a long history of study, during which many different image matching metrics have been investigated.⁷⁻⁹ Though our interest is in deformable 2D/3D registration, insights from rigid registration can provide us with useful guiding principles with regards to such optimization parameters. The lengthy experience documented in the literature demonstrates that each metric has merits and weaknesses. The sum-of-squared-intensity-differences (SSD) has the advantage of being robust with respect to noise and is easy to formulate for a gradient-driven iterative search process, but requires a close correspondence of the absolute image intensities in the DRRs and the projections.^{4,10,11} The many formulations of mutual information allow much greater variability in the comparative image intensities, but most studies of its effectiveness have involved 3D images.¹⁰ Mutual information has been shown to be less robust for 2D image matching, owing to the relatively small number of available data samples compared to 3D/3D registration.^{7,10} This weakness is compounded when there is noise. It also has a smaller and more sharply delimited capture range, which makes it less robust in a gradient-driven iterative search. Formulating gradients of

the iterative cost function for mutual information is also more difficult than for SSD.¹¹

In our development of deformable FIPM we have preferred a gradient search and therefore have opted to use the sum-of-squared differences as our image matching metric. We have already shown via simulations that this formulation is sensitive to contrast mismatch between the DRRs and the projection images, but is relatively insensitive to noise.⁴ This has motivated the present work. Our goal here is to calculate DRRs solely from fundamental physical principles underlying image formation to accurately replicate the absolute image intensities in a CBCT projection; i.e., to calculate DRRs from first principles. We avoid recourse to intensity normalization or any other artificial numerical intensity adjustment, which could bias the registration toward an inaccurate deformation. It is not our purpose to maximize the contrast in the DRRs, but to reproduce what is present in the projections. The effects of noise on FIPM are minor and relatively easy to incorporate, and therefore are not considered here.

Many investigators have described algorithms covering various aspects of the DRR generation process. In particular, much attention has been paid to the x-ray volume rendering component, with regards to geometric realism, spatial resolution, and computational efficiency. Siddon introduced a widely used algorithm for efficiently calculating the exact radiological path through a computed tomography (CT) volume.¹² Russakoff *et al.* developed the concept of “attenuation fields” for fast DRR rendering.¹³ Moore *et al.* used pencil beams rather than rays as projectors to increase DRR realism and avoid aliasing effects.¹⁴ Birkfellner *et al.* used a “wobbled splatting” technique for fast DRR generation.¹⁵ Other techniques including shear-warp volume rendering, frequency domain volume rendering, and Monte Carlo integration based methods have also been investigated.^{16–18} We note that much of this prior work is tailored to rigid 2D/3D registration for patient setup, which emphasizes precise geometrical reproduction of the anatomical pose and benefits from the high contrast of skeletal features. In addition, many authors have used a graphics processing unit (GPU) to speed up DRR computation.^{19–23} Our algorithm uses a standard ray-casting approach with GPU acceleration.²³

The basic ray-tracing component of a DRR algorithm models only the attenuation of primary photons, whereas projection image formation is influenced by additional physical effects. One such effect is scatter, which can have a large presence in CBCT projections.²⁴ Various methods of estimating CBCT scatter exist, including beam stop arrays, scatter kernel methods, and Monte Carlo simulation. Ruhrschoepf and Klingensbeck give a comprehensive summary of scatter estimation and correction techniques.²⁵ These solutions are generally framed in the context of scatter mitigation in CBCT projections, rather than scatter modeling for DRRs. Moore *et al.* applied a scatter correction specifically aimed at improving DRRs using the beam stop array method.¹⁴ Our approach is to use an in-house Monte Carlo algorithm simulating the CBCT system to obtain scatter estimates.²⁶ In contrast to the beam stop array method, this approach does not require any special equipment and more impor-

tantly avoids any extra scans, which deliver extra dose to a patient.

A second effect influencing projection image formation is beam hardening. Again, there are many papers describing beam hardening and beam hardening corrections, but nearly all are applied to improving CT reconstruction. Two popular methods assume that the subject is composed entirely of water or bone.^{27,28} Yan *et al.* describe a method for correcting single and dual energy CT using an iterative method by representing each voxel in the CT as a sum of basis materials.²⁹ In their DRR algorithm Moore *et al.* explicitly modeled a poly-energetic beam rather than applying a correction. For our DRR algorithm we evaluated both an analytical and measurement approach to a beam hardening correction implemented in postprocessing.

A final effect is veiling glare, or leakage of signal from a detector pixel to its neighbors.^{30–32} At our institution we have measured the effect of veiling glare on the flat panel detector used in this study.

We maximize the realism in our DRRs by using an efficient ray-tracing routine combined with scatter, beam hardening, and veiling glare corrections. In this paper we provide a step by step account of the measurements and calculations needed by our algorithm. Though many descriptions of these effects have been described separately, few works exist that consolidate them into one algorithm for the purpose of computing realistic DRRs.

II. METHOD AND MATERIALS

II.A. Projection definition

We begin by defining a CBCT projection through an object,

$$P_{\text{CBCT}}(\mathbf{x}_i) = A_1 \Phi_{\text{CBCT}}(\mathbf{x}_i) \alpha_{\text{CBCT}}(\mathbf{x}_i) R_{\text{CBCT}}(\mathbf{x}_i) + S_{\text{CBCT}}(\mathbf{x}_i) \quad (1)$$

where $\mathbf{x}_i$ is the location of detector pixel i (all coordinates are in the world coordinate frame), $P_{\text{CBCT}}(\mathbf{x}_i)$ is the signal recorded by detector pixel i , A_1 is a tube output factor corresponding to mAs, $\Phi_{\text{CBCT}}(\mathbf{x}_i)$ is the source fluence distribution aimed at the detector pixel, $\alpha_{\text{CBCT}}(\mathbf{x}_i)$ is the transmission, or the fraction of primary particles that remain when the beam reaches the detector pixel, $R_{\text{CBCT}}(\mathbf{x}_i)$ is a pixel dependent detector response function, and $S_{\text{CBCT}}(\mathbf{x}_i)$ is the contribution of scattered photons to the detector signal. We further define the transmission as

$$\alpha_{\text{CBCT}}(\mathbf{x}_i) = \exp[-M(\mathbf{x}_i)] = \exp\left(-\int_{\mathbf{x}_0}^{\mathbf{x}_i} \mu(\mathbf{x}) d\mathbf{x}\right), \quad (2)$$

where $\mu(\mathbf{x})$ is the linear attenuation coefficient (LAC) at location $\mathbf{x}$ within the object being imaged, $\mathbf{x}_0$ is the location of the source, and $M(\mathbf{x}_i)$ is a the line integral of linear attenuation coefficients (cumulative LAC) along the ray connecting $\mathbf{x}_0$ to $\mathbf{x}_i$. We can expand the scatter term as

$$S_{\text{CBCT}}(\mathbf{x}_i) = A_1 \sigma_{\text{CBCT}}(\mathbf{x}_i) R_{\text{CBCT}}(\mathbf{x}_i), \quad (3)$$

where $\sigma_{\text{CBCT}}(\mathbf{x}_i)$ is the object scatter distribution. In general, the scatter distribution is a complex function of the object being imaged, the source fluence distribution, and the detector pixel location.

Finally, it is useful to define a special case of Eq. (1) termed a flood projection. A flood projection is taken with no object in the beam, thus transmission is 1 and scatter is 0 across the detector surface. This leads to the following simplified mathematical description:

$$F_{\text{CBCT}}(\mathbf{x}_i) = A_0 \Phi_{\text{CBCT}}(\mathbf{x}_i) R_{\text{CBCT}}(\mathbf{x}_i), \quad (4)$$

where $F_{\text{CBCT}}(\mathbf{x}_i)$ is the flood projection signal recorded by detector pixel i . The flood projection empirically captures the source fluence distribution and detector response function.

II.B. DRR challenges

A CT image represents the x-ray attenuation coefficients of materials in the imaged object, relative to the attenuation coefficient of water, for the particular energy spectrum of the incident x-rays. The goal of our algorithm is to use a CT image of an object to simulate an x-ray radiograph (projection) of the object as defined in Eq. (1). We can divide the physical processes that produce a projection into primary and secondary effects.

The primary process responsible for image formation is the attenuation of primary photons by the object. This process is governed by the object's LAC distribution, $\mu(\mathbf{x})$, in the following manner:

$$I_i = I_0 \exp\left(-\int_{x_0}^{x_i} c \mu(\mathbf{x}) dx\right), \quad (5)$$

where I_0 is the initial intensity of the photon beam, I_i is the intensity of primary photons remaining after the beam passes through the object, and the integral runs along the line connecting $\mathbf{x}_0$ to $\mathbf{x}_i$. We can use Eq. (5) to calculate the transmission factor $\alpha(\mathbf{x}_i)$ in Eq. (1). The line integral in Eq. (5) can be estimated using a ray-tracing routine; however, the CT through which we make projections has dimensionless units of CT number (attenuation relative to water), not absolute LAC. Therefore, our first requirement is to convert CT number to absolute LAC in order to render the CT in physically meaningful units. This is complicated by the fact that the absolute LAC depends on the incident x-ray energy spectrum. Therefore, one must measure the x-ray attenuation through objects of known composition using the precise x-ray spectrum of the CBCT system whose projections one wishes to emulate.

The secondary processes that we need to model are scatter, beam hardening, and veiling glare. Scatter refers to particles that have been deflected by the object but are still absorbed by the detector. Beam hardening refers to the effect whereby the energy distribution of a poly-energetic beam is shifted toward higher energies as it passes through an attenuating medium. Veiling glare denotes the leakage of signal from the detector pixel in which it is generated to its neighbors. These effects must be accounted for in two ways: they must be removed from the measurements used to determine the CT number to LAC conversion function, and they must be

folded into the raw DRR produced by the ray-tracing routine. This separation ensures that scatter, beam hardening, and veiling glare are explicitly and precisely accounted for in each individual DRR rather than entangled within the CT number conversion function.

II.C. Secondary physics effects

II.C.1. Beam hardening

As a poly-energetic beam passes through a material the lower energy photons are preferentially attenuated, causing the average beam energy to increase. This effect is known as beam hardening. The LAC is an energy dependent quantity, so for a given material with a uniform CT number, beam hardening causes the beam to “see” a progressively lower effective LAC as the beam traverses the material.

In order to account for this effect we require a function, $h(M_{bh1})$, that will convert hardened cumulative LAC from measured attenuation data to unhardened cumulative LAC in order to eliminate beam hardening from the data used to calculate the CT number to LAC conversion function. In addition, we require a function $g(M_{bh0})$ that does the reverse in order to add the beam hardening effect to the unhardened cumulative LAC integrated by our ray tracing routine. The subscripts “bh1” and “bh0” refer to the presence and absence of beam hardening, respectively. We can accomplish both goals using the same data—a set of hardened cumulative LAC values, M_{bh1} , and a corresponding set of unhardened cumulative LAC values, M_{bh0} , calculated through different thicknesses of a test material. Fitting this data-set with M_{bh0} as the independent variable produces the beam hardening function, and fitting the data with M_{bh1} as the independent variable generates the beam unhardening function. In this paper we present both measurement- and computation-based approaches for obtaining this data-set.

The raw data for the measurement approach is a set of 32 CBCT projections taken through solid water blocks ranging in thickness from 2 mm to 180 mm. The solid water slabs are positioned so that the couch does not interfere with the measurement, and a minimal volume of the slabs is exposed to the beam to minimize scatter. In addition, we acquire a flood projection at the same CT technique. We begin by choosing a position on the detector that lies beneath the shadow of the solid water slabs as our measurement point and calculate the transmission for each slab thickness. Assuming scatter to be negligible, Eqs. (1) and (4) can be used to calculate transmission as follows:

$$\alpha_{\text{CBCT}}(\mathbf{x}_i) = \frac{A_0 P_{\text{CBCT}}(\mathbf{x}_i)}{A_1 F_{\text{CBCT}}(\mathbf{x}_i)} \quad (6)$$

In general A_0/A_1 can be estimated as the ratio of the mAs of the flood projection to that of the object projection, or calculated as the ratio of the flood projection to the object projection in an area where there is no attenuation in the object projection, i.e., $\alpha_{\text{CBCT}}(\mathbf{x}_j) = 1$. The transmission for each thickness l of solid water can be converted to the hardened cumulative LAC using Eq. (2),

$$M_{bh1}(l) = -\ln[\alpha_{\text{CBCT}}(l)]. \quad (7)$$

After making these measurements for each thickness we estimate the corresponding set of unhardened cumulative LAC values as follows:

$$M_{bh0}(l) = \mu_0 l \quad (8)$$

where μ_0 is the average LAC of the material in the unhardened beam spectrum and l is the thickness of the solid water slab. While we know l for each slab we must calculate μ_0 . To do this we start by fitting an 8th order polynomial function to our previously determined hardened $M_{bh1}(l)$ data plotted as a function of l , generating a function $y(l)$ such that

$$M_{bh1}(l) \approx y(l) = al + bl^2 + \dots + hl^8. \quad (9)$$

As thickness l approaches zero we know that $M_{bh1}(l)$ and $M_{bh0}(l)$ will converge

$$\lim_{l \rightarrow 0} M_{bh1}(l) = M_{bh0}(l) = \mu_0 l. \quad (10)$$

Thus we can estimate μ_0 as the linear coefficient of our polynomial function from Eq. (9),

$$\mu_0 \approx a. \quad (11)$$

From this we can calculate $M_{bh0}(\mathbf{x}_i)$ for each slab thickness using Eq. (8).

Our second approach is to analytically calculate the attenuation process through slabs of a test material to obtain the same data-set as above. This requires three items of input: the beam spectrum of our CBCT source, $W_0(E)$, the LAC distribution of the bowtie filter material (aluminum), $\mu_{Al}(E)$, and the LAC distribution of a test material, $\mu_{mat}(E)$, through which we calculate the passage of the beam. We use liquid water as our test material, though any material whose attenuation coefficient is known is a possible candidate. We begin by propagating the beam through the middle of the bowtie filter, where it is at its thinnest,

$$W_{BT,min}(E) = W_0(E) \exp[-\mu_{Al}(E)l_{BT,min}]. \quad (12)$$

For all practical purposes this is our true unhardened spectrum, as images are never acquired without the bowtie filter and this is the area of the beam used to calibrate our CT number to LAC conversion function, described in Sec. II.D. We can now calculate the μ_0 value for our test material using an average of $\mu_{mat}(E)$ weighted by $W_{BT,min}(E)$:

$$\mu_0 = \frac{\int W_{BT,min}(E) \mu_{mat}(E) E dE}{\int W_{BT,min}(E) E dE}. \quad (13)$$

We then proceed on a pixel by pixel basis, propagating the spectrum through the thickness of bowtie filter that each pixel “sees.” For detector pixel i this new spectrum is calculated as

$$W_{BT,i}(E) = W_0(E) \exp[-\mu_{Al}(E)l_{BT,i}]. \quad (14)$$

Next, the spectrum for each pixel is filtered through thicknesses of the test material ranging from 0 mm to 150 mm in 5 mm increments. For a given thickness of test material $l_{mat,j}$ this newly filtered spectrum is

$$W_{mat,i,j}(E) = W_{BT,i}(E) \exp[-\mu_{mat}(E)l_{mat,j}]. \quad (15)$$

From here the cumulative hardened LAC corresponding to thickness $l_{mat,j}$ of material for detector pixel i is calculated as

$$M_{bh1}(\mathbf{x}_i) = -\ln \left(\frac{\int W_{mat,i,j}(E) E dE}{\int W_{BT,min}(E) E dE} \right), \quad (16)$$

and the corresponding unhardened cumulative LAC is calculated as

$$M_{bh0}(\mathbf{x}_i) = \mu_0 l_{mat,j}. \quad (17)$$

In practice the unhardened cumulative LAC values need only be calculated once since they have no detector pixel dependence.

Having obtained our data-set of M_{bh0} and corresponding M_{bh1} values, either by measurement or calculation, we are prepared to calculate our beam hardening and unhardening functions. The beam hardening function $g(M_{bh0})$ is obtained by fitting an 8th order polynomial to the M_{bh1} values plotted as a function of M_{bh0} . The input to the function is an unhardened cumulative LAC value and the output is a hardening cumulative LAC value, as follows:

$$M_{bh1} = g(M_{bh0}) = aM_{bh0} + bM_{bh0}^2 + \dots + hM_{bh0}^8. \quad (18)$$

Likewise, the beam unhardening function $h(M_{bh1})$ is determined by fitting an 8th order polynomial to the M_{bh0} values plotted as a function of M_{bh1} . The input to this function is a hardened cumulative LAC value and the output is an unhardened cumulative LAC value,

$$M_{bh0} = h(M_{bh1}) = iM_{bh1} + jM_{bh1}^2 + \dots + pM_{bh1}^8. \quad (19)$$

If the raw beam hardening data are obtained by the measurement method only one of each type of function will be produced. If the data are obtained via calculation unique versions of both functions will be produced for each pixel.

II.C.2. Scatter

Scatter refers to those photons reaching the detector that have been previously deflected by the object. For our purposes we must be able to remove and add the effects of scatter. To do this we used an in-house Monte Carlo (MC) algorithm capable of simulating scatter-only projections of an object.^{26,33} A Monte Carlo scatter projection can be represented as

$$S_{MC}(\mathbf{x}_i) = A_2 \sigma_{MC}(\mathbf{x}_i) R_{MC}(\mathbf{x}_i). \quad (20)$$

The Monte Carlo algorithm can also simulate a flood projection,

$$F_{MC}(\mathbf{x}_i) = A_2 \Phi_{MC}(\mathbf{x}_i) R_{MC}(\mathbf{x}_i). \quad (21)$$

Dividing Eq. (20) by Eq. (21) produces

$$\frac{S_{MC}(\mathbf{x}_i)}{F_{MC}(\mathbf{x}_i)} = \frac{\sigma_{MC}(\mathbf{x}_i)}{\Phi_{MC}(\mathbf{x}_i)}. \quad (22)$$

Making use of our CBCT flood projection we multiply Eq. (22) by Eq. (4) to yield

$$\begin{aligned} F_{CBCT}(\mathbf{x}_i) \frac{S_{MC}(\mathbf{x}_i)}{F_{MC}(\mathbf{x}_i)} \\ = A_0 \Phi_{CBCT}(\mathbf{x}_i) R_{CBCT}(\mathbf{x}_i) \frac{\sigma_{MC}(\mathbf{x}_i)}{\Phi_{MC}(\mathbf{x}_i)}. \end{aligned} \quad (23)$$

Finally, multiplying Eq. (23) by the inverse of the output factor from Eq. (6), $\frac{A_1}{A_0}$, and assuming that $\Phi_{MC}(\mathbf{x}_i)$ is approximately equal to $\Phi_{CBCT}(\mathbf{x}_i)$, we obtain an estimate of the scatter described in Eq. (3),

$$\tilde{S}_{CBCT}(\mathbf{x}_i) = A_1 \sigma_{MC}(\mathbf{x}_i) R_{CBCT}(\mathbf{x}_i) \quad (24)$$

This scatter projection can be added or subtracted from another projection to add or remove scatter effects.

II.C.3. Veiling glare

Veiling glare refers to the leakage of signal from its detector pixel of origin to surrounding detector pixels. We model this effect using a point spread function (PSF) that characterizes this signal spread over an area up to 10 cm distant from the signal origin. The veiling glare PSF is modeled by fitting a parametric function to the edge response function obtained from half-blocked projection images, which were obtained by acquiring projections with a lead plate positioned over half of the detector surface. The parametric function is a weighted sum of Gaussian and exponential terms used to describe the short and long range effects of veiling glare, respectively. The veiling glare effect can be introduced to a projection by convolving it with the PSF,

$$P_{vg1} = P_{vg0} \otimes \text{PSF}, \quad (25)$$

where the subscript “vg1” indicates that veiling glare is present, “vg0” indicates that veiling glare is absent, and $\otimes$ is the convolution operator. Likewise veiling glare can be removed by deconvolving the PSF,

$$P_{vg0} = P_{vg1} \otimes^{-1} \text{PSF}, \quad (26)$$

where $\otimes^{-1}$ denotes the deconvolution operator.

II.D. CT number to LAC conversion function

In order to evaluate the integral in Eq. (5) we must convert the CT numbers to units of linear attenuation coefficient. This conversion depends on the energy spectrum (and thus on the kVp) at which one wishes to emulate a CBCT projection. Our method is the same general procedure as that used by Moore *et al.*¹⁴ We begin by collecting cylindrical plugs composed of different materials. The plugs are inserts of known dimensions taken from a Gammex electron density calibration phantom. We first acquire a FBCT of the phantom with plugs in place to determine the CT number of each plug. Next we separately capture CBCT projections at a given technique for each plug. After removing beam hardening, scatter, and veiling glare from these projections, LAC values for each plug material can be calculated using Eqs. (2) and (8). Once this is accomplished, a 2nd order polynomial fit to the LACs plotted as a function of CT number yields our desired CT number to LAC conversion function, $f_{LAC}[CT(\mathbf{x})]$, where $CT(\mathbf{x})$ is the CT number at location $\mathbf{x}$. First, veiling glare is removed from the plug projections,

$$P_{CBCT,plug,vg0} = P_{CBCT,plug} \otimes^{-1} \text{PSF}. \quad (27)$$

FIG. 1. Plot of LAC vs CT number data points and 2nd order polynomial fits generated from attenuation measurements of different plug materials. Attenuation measurements are either uncorrected or fully corrected for scatter, beam hardening, and veiling glare effects.

Next scatter is subtracted

$$P_{CBCT,plug,vg0sc0}(\mathbf{x}_i) = P_{CBCT,plug,vg0}(\mathbf{x}_i) - \tilde{S}_{CBCT,plug}(\mathbf{x}_i). \quad (28)$$

Then the projection is reduced to cumulative LAC

$$M_{CBCT,plug,vg0sc0}(\mathbf{x}_i) = -\ln \left(\frac{A_0 P_{CBCT}(\mathbf{x}_i)}{A_1 F_{CBCT}(\mathbf{x}_i)} \right), \quad (29)$$

and a correction is applied to remove beam hardening

$$M_{CBCT,plug,vg0sc0bh0}(\mathbf{x}_i) = h(M_{CBCT,plug,vg0sc0}(\mathbf{x}_i)). \quad (30)$$

Finally, LAC is calculated from the cumulative unhardened LAC

$$\mu_{CBCT,plug,vg0sc0bh0} = \frac{M_{CBCT,plug,vg0sc0bh0}(\mathbf{x}_c)}{t}, \quad (31)$$

where t is the thickness of the plugs, 7 cm, and $\mathbf{x}_c$ is the coordinate of the detector center over which the plugs are positioned during imaging. Plotting $\mu_{CBCT,plug,vg0sc0bh0}$ as a function of CT number and fitting with a polynomial yields our conversion function, $f_{LAC}[CT(\mathbf{x})]$. (As previously mentioned, this conversion function is specific to the x-ray tube technique at which the CBCT projections were made.) To illustrate the effect of the scatter, beam hardening, and veiling glare corrections on the final conversion function, the data points used to build the function and their polynomial fits are plotted in Fig. 1 for both uncorrected and fully corrected data.

II.E. DRR workflow

The workflow for producing a fully corrected DRR is similar to the reverse of that for processing the plug projections. First, a raw DRR is produced using an x-ray volume rendering routine. We use a ray-tracing based approach that

approximates the line integral in Eq. (5) as the following:

$$\int_{x_0}^{x_i} \mu_{bh0}(\mathbf{x}) dx \approx M_{\text{DRR},bh0}(\mathbf{x}_i) = \sum_{k=0}^K f_{\text{LAC}}[\text{CT}(\mathbf{x}_0 + k\Delta\mathbf{x})] \Delta x, \quad (32)$$

where K is the number of steps that fit between $\mathbf{x}_0$ and $\mathbf{x}_i$, $\Delta\mathbf{x}$ is the ray-step vector, and Δx is the ray-step length. The CT is sampled at location $\mathbf{x}_0 + k\Delta\mathbf{x}$ using trilinear interpolation. This algorithm is fully supported by the hardware implementation of ray tracing and interpolation on a GPU platform.²³ After the raw DRR is generated, beam hardening effects are introduced,

$$M_{\text{DRR},bh1}(\mathbf{x}_i) = g(M_{\text{DRR},bh0}(\mathbf{x}_i)). \quad (33)$$

Using this result and Eqs. (1) and (2) a beam hardening corrected projection image is produced,

$$P_{\text{DRR},bh1}(\mathbf{x}_i) = \frac{A_1}{A_0} F_{\text{CBCT}}(\mathbf{x}_i) \exp[-M_{\text{DRR},bh1}(\mathbf{x}_i)]. \quad (34)$$

The flood projection allows us to reproduce the CBCT source fluence distribution in the DRR, and the output factor allows us to recover the beam output of the target projection. This image is then corrected for scatter,

$$P_{\text{DRR},bh1sc1}(\mathbf{x}_i) = P_{\text{DRR},bh1}(\mathbf{x}_i) + \tilde{S}_{\text{CBCT},obj}(\mathbf{x}_i), \quad (35)$$

where the scatter projection has been calculated using the method of Sec. II.C.2. Lastly, the veiling glare effect is introduced and the DRR is fully corrected,

$$P_{\text{DRR},bh1sc1vg1} = P_{\text{DRR},bh1sc1} \otimes \text{PSF}. \quad (36)$$

II.F. Projection acquisition

For verification, DRRs were computed through object FBCTs and compared to actual CBCT projections of the same object. Two different objects were tested in this study, a CatphanTM CT phantom (The Phantom Laboratory, Salem, NY), and an actual patient. In each case projections were acquired at 20 ms, 20 mA, and 125 kVp. CatphanTM projections were acquired in full-fan mode while patient projections were captured in half-fan mode, each using their respective aluminum bow-tie filters. All images were acquired on the OBITM CBCT system (Varian Medical Systems, Palo Alto, CA) which utilizes a G242 kilovoltage x-ray source and a PAXSCAN 4030CB flat panel detector mounted to the CLINAC. In addition a 10:1 antiscatter grid is positioned on top of the detector system, which is protected by a 0.5 mm aluminum cover.

III. RESULTS

We tested the accuracy of the DRR algorithm using comparisons between our DRRs and the actual CBCT projections they were intended to match. We calculated two different sets of DRRs—one set using just the measured linear attenuation through the medium (uncorrected) and a second set using the full algorithm with corrections for beam hardening,

scatter, and veiling glare in both the LAC measurements and the DRR projections. This displays the relative importance of the corrections. We used two different objects as test cases—a CatphanTM CT phantom, and an actual patient. For both sets of DRRs, a CBCT flood projection was used to measure the output and source fluence distribution of the CBCT system when the raw DRRs were postprocessed. We emphasize that these DRRs are absolute calculations of detected signal intensity, without any normalization factor or other numerical adjustment to match to the CBCT projections. For qualitative validation we focus on visual comparison, comparison of pixel intensity profiles, and comparison of image intensity histograms. For quantitative validation we correlate the DRR pixel intensities with the actual projection pixel intensities.

III.A. Visual comparison

Figure 2 displays projections cast through the CT phantom, and Figs. 3 and 4 show projections cast through an actual patient at different angles. The images in each figure are set to the same window and level for comparison. In Fig. 2 the bottom of each DRR was contaminated by metal artifacts in the CatphanTM FBCT. This has no bearing on the efficacy of our method and is subsequently ignored.

The goal of this evaluation is not to provide an expert analysis of DRR image quality, as might be provided by a radiologist, but to make some basic visual observations on the output of our algorithm. We begin by noting that both uncorrected and corrected DRRs bear a close resemblance to the CBCT projections, with most obvious features clearly reproduced. In general, the uncorrected and corrected DRRs look similar; both look less crisp than the CBCT projections due to the coarse resolution of the FBCT relative to the size of the detector pixels. This effect causes the loss of some detail in both corrected and uncorrected DRRs, such as the spine and vertebral disks in Fig. 3.

However, in the patient projections in particular the corrected DRRs show more realistic contrast than the uncorrected DRRs with respect to the actual CBCT projections. In Fig. 3(a), the uncorrected DRR, the upper right side of the patient is essentially saturated in black while the actual projection, Fig. 3(c) shows considerable contrast. In the corrected DRR, Fig. 3(b), some of this lost contrast has been visibly recovered. Similarly, Fig. 4 shows areas, particularly at the top of the image above the crest of the lung, where the uncorrected DRR is saturated in black, while both the corrected DRR and actual CBCT projection maintain visible contrast.

III.B. Pixel intensity profiles

Figures 5–7 show pixel intensity profiles sampled along the horizontal direction from each set of uncorrected, fully corrected, and actual projections. The locations at which these profiles were sampled are marked on the actual CBCT projections in Figs. 2–4. For the most part, both the curves for the fully corrected and uncorrected DRRs track with the actual CBCT projection, although the uncorrected curves tend to consistently underestimate the pixel intensity. This is particularly visible in Fig. 7. Figures 6 and 7 compare intensity

FIG. 2. Projections through a CatphanTM CT phantom at a projection angle of 0° . (a) uncorrected DRR, (b) fully corrected DRR, (c) actual CBCT projection. Red lines in (c) indicate locations where line profiles are sampled.

profiles for real patient data. The FBCT through which the DRRs were cast was taken on a different day than the CBCT projections to which the DRRs are compared. For this reason it is impossible to perfectly match patient CBCT projections and DRRs, as the patient setup will invariably differ and

FIG. 3. Projections through a patient at a projection angle of 92.67° . (a) Uncorrected DRR, (b) fully corrected DRR, and (c) actual CBCT projection. Red lines in (c) indicate locations where line profiles are sampled.

some internal anatomic deformation will be unavoidable. In addition, the CBCT and FBCT systems have different couches, and in Figs. 3 and 4 couch structures that are plainly visible in the actual CBCT projections are absent from the DRRs. Areas in the line profiles of Figs. 6 and 7 where the DRRs and CBCT projections radically disagree reflect these differences.

FIG. 4. Projections through a patient at a projection angle of 178.93° . (a) Uncorrected DRR, (b) fully corrected DRR, and (c) actual CBCT projection. Red lines in (c) indicate locations where line profiles are sampled.

III.C. Intensity histograms

Figure 8 plots intensity histograms of the actual CBCT projections alongside histograms of uncorrected and fully corrected DRRs. Areas contaminated by the treatment couch, in the case of the patient projections, or metal artifact, in the case of the CatphanTM projections, were excluded from

FIG. 5. Comparison of line profiles sampled from an actual CBCT projection, a DRR uncorrected for beam hardening, scatter, or veiling glare, and a fully corrected DRR. Projections are taken through a CatphanTM CT phantom at a projection angle of 0° . Line profiles are sampled at (a) row 48, (b) row 284, and (c) row 582.

the histograms. For all cases the fully corrected DRR histograms exhibit a closer match to the CBCT projections than the uncorrected DRRs. Computing the Dice coefficient between CBCT and DRR histograms yields a measure of the

FIG. 6. Comparison of line profiles sampled from an actual CBCT projection, a DRR uncorrected for beam hardening, scatter, or veiling glare, and a fully corrected DRR. Projections are taken through an actual patient at a projection angle of 92.67°. Line profiles are sampled at (a) row 145, (b) 360, and (c) row 630.

overlap between the two distributions. The Dice coefficient ranges from 0 (no overlap) to 1 (perfect overlap) and is defined as

$$D = \frac{2(H_{\text{CBCT}} \cap H_{\text{DRR}})}{H_{\text{CBCT}} + H_{\text{DRR}}}. \quad (37)$$

FIG. 7. Comparison of line profiles sampled from an actual CBCT projection, a DRR uncorrected for beam hardening, scatter, or veiling glare, and a fully corrected DRR. Projections are taken through an actual patient at a projection angle of 178.93°. Line profiles are sampled at (a) row 145, (b) 360, and (c) row 630.

For the CatphanTM histograms in Fig. 8(c) the Dice coefficients for the fully corrected and uncorrected DRR histograms when measured against the CBCT histogram were 0.87 and 0.76, respectively. For the patient histograms in Fig. 8(b) corrected and uncorrected values were 0.84 and 0.80, and for the

FIG. 8. Histograms comparing the intensity distributions of actual CBCT projections, DRRs uncorrected for beam hardening, scatter, or veiling glare, and fully corrected DRRs. Comparisons are for projections cast through (a) a patient at a projection angle of 92.67° , (b) a patient at a projection angle of 178.93° , and (c) a CatphanTM CT phantom at a projection angle of 0° . In (a) and (b) areas of the projections contaminated by the treatment couch are omitted from the distributions. In (c) areas of the projections contaminated by metal artifacts in the FBCT image are omitted from the distributions.

patient histograms in Fig. 8(a) corrected and uncorrected values were 0.73 and 0.59.

III.D. DRR-to-projection intensity correlation

The scatter plots in Fig. 10 display the correspondence between intensity values recorded by actual detector pixels and

their counterparts in both corrected and uncorrected DRRs. A perfect match of DRR to projection is represented by the (dashed) line of slope one. Intensity mismatch is represented by the scatter around this line as well as any systematic offset in slope or intercept.

We obtain a quantitative measure of these correlations by computing the mean percent disagreement between CBCT projection and DRR pixel intensities averaged over all pixels, along with the standard deviation about this mean. The mean value describes the extent to which a scatter distribution is systematically offset from the line of perfect agreement, while the standard deviation describes how tightly grouped the distribution is. A positive mean difference indicates that on average the DRR overestimated the true pixel intensity, while a negative difference signifies that on average the DRR underestimated the true pixel intensity. For the patient data plotted in Fig. 10(a) the percent disagreements for the corrected and uncorrected DRRs were $34.22\% \pm 33.88\%$ and $-46.61\% \pm 41.19\%$, respectively. For the patient data plotted in Fig. 10(b) percent disagreements for corrected and uncorrected DRRs were $6.52\% \pm 6.46\%$ and $-22.09\% \pm 21.87\%$, respectively. For the CatphanTM data plotted in Fig. 10(c) percent disagreements for corrected and uncorrected DRRs were $4.01\% \pm 3.97\%$ and $-6.58\% \pm 6.52\%$, respectively.

III.E. Beam hardening comparison

Figure 9 displays intensity profiles comparing the measured and calculated beam hardening corrections. For reference, the actual CBCT projection intensity profiles are also included. In the center of the profiles the two beam hardening methods exhibit the best agreement, while the two curves begin to diverge closer to the edges. This is expected, as the data for the measured correction is collected at the center of the detector. At the sides of the detector the calculated method, which accounts for the heterogeneous bowtie filter thickness, increasingly deviates from the measurement-based method, which does not.

IV. DISCUSSION

Before discussing our results it is important to reiterate our motivation for developing a DRR algorithm from first principles for the purpose of FIPM based deformable 2D/3D registration. As our goal is to minimize the discrepancy between DRRs and CBCT projections in absolute units we must have a DRR algorithm that produces realistic image intensities. However, our DRR algorithm must also be unbiased, i.e., it must not conceal or misrepresent differences between DRRs and CBCT projections that correspond to actual deformation features. Any scheme that artificially scales or adjusts the DRRs to the actual CBCT projections runs the risk of compromising those features. By grounding our algorithm in the fundamental physics of image formation our method meets the dual goals of closely reproducing CBCT projection intensities while avoiding bias that could affect the integrity of our image registration process.

FIG. 9. Line profiles sampled from projections cast through Catphan CT phantom for comparison of two different beam hardening correction methods. Line profiles are sampled at (a) row 48, (b) row 284, and (c) row 582.

IV.A. Qualitative and quantitative algorithm assessments

In testing our algorithm we calculated DRRs both with and without physics corrections and compared the results to actual CBCT projections. This allowed us to compare our DRRs to a “ground truth” in order to determine the absolute accuracy

FIG. 10. Scatter plots showing the correspondence between pixel intensities in actual CBCT projections and those of both corrected (black) and uncorrected (green) DRRs. The pink line indicates the track along which perfect correspondence is achieved. Scatter plots are plotted for projections cast through (a) a patient at a projection angle of 92.67° , (b) a patient at a projection angle of 178.93° , and (c) a CatphanTM CT phantom at a projection angle of 0° .

achieved by our algorithm, as well as to examine the relative effect of our physics based corrections. In our assessment of absolute DRR and CBCT projection agreement there were two potentially confounding factors. The DRRs were

manually aligned to the CBCT system geometry, with small misalignments possibly resulting in large discrepancies between pixel intensities. Additionally, for patient data some deformation inevitably occurred in the time between FBCT and CBCT projection acquisition. This likely caused differences between DRRs and CBCT projections unrelated to the DRR algorithm. Due to these factors our quantitative assessments likely indicate an upper bound on the accuracy we can expect from our algorithm.

With respect to our visual comparison of DRRs to the actual CBCT projections, both corrected and uncorrected methods appear to produce images that look similar to each other and to the actual projection. Differences mainly appear in low intensity areas where the uncorrected DRR became saturated in black while the corrected DRR and actual projection maintained some visible contrast. The line plots in Figs. 5–7 show that in general both corrected and uncorrected DRRs follow the general shape of the actual CBCT projection. It is observed that the uncorrected DRRs tend to underestimate the true CBCT projection intensity, while the corrected DRRs slightly overestimate the CBCT projection intensity. A similar effect is observed in the intensity histogram plots of Fig. 8, with histograms of corrected DRRs shifted to the right and those of uncorrected DRRs shifted to the left. In each of the three cases shown the Dice coefficient for the corrected DRRs was larger than that of the uncorrected DRRs, indicating that the physics corrections consistently provide a quantifiable improvement in DRR quality. Assessment of the scatter plots in Fig. 10 confirms this trend. Scatter distributions from corrected DRRs all fall somewhat above the line of perfect agreement indicating overestimation of pixel intensity, while distributions for uncorrected DRRs consistently fall below the line, indicating underestimation. In all cases the mean error for the corrected DRRs is less than the uncorrected DRRs. The same goes for the standard deviations of the errors, indicating a tighter distribution for the corrected DRRs. The smallest error and error standard deviation of $4.01\% \pm 3.97\%$ occurs for the corrected DRR cast through the CatphanTM. We expect the CatphanTM to give us the smallest errors since it is a rigid object and will thus avoid the deformations that can occur in patient images. For this reason it is our best estimate of the true error of our algorithm. We choose to quantitatively evaluate our algorithm using percent mean error of the absolute DRR intensities since this is the most relevant quantity with regards to the SSD merit function. This measure is a broad indicator, in the sense that a small percent error in absolute intensity will usually be indicative of a small error in relative metrics such as contrast, while the reverse may not be true.

In a previous study we quantitatively examined the effect of intensity mismatch on registration quality by introducing an artificial, nonlinear variation between projection and DRR intensities.⁴ The magnitude of disagreement required to induce appreciable degradation of the registration results was in the neighborhood of 1%, which is below our present best value of 4.01%. However, looking at Fig. 10(c), the variations in our CatphanTM DRR appear to follow a straight line, and therefore may have less of an effect on the FIPM algorithm

than those previously considered. We are now in a position to repeat our test of the sensitivity of 2D/3D FIPM to DRR accuracy using actual rather than contrived contrast mismatch.

The corrected DRRs consistently overestimate true projection intensities, while our uncorrected DRRs consistently underestimate them. In the uncorrected DRRs the effects of scatter, beam hardening, and veiling glare are included in the CT number to LAC conversion function to the degree that they are present in the plug projections. If the sum of those effects is greater in the plug projections than in the actual object projections the DRR will overestimate the true CBCT projection intensity. If the reverse is true the DRR will underestimate the true CBCT projection intensity. Averaged over all 12 plug materials, the absence of the scatter, beam hardening, and veiling glare effects reduces the image intensity within the plug shadows by 9%, 5%, and 19% respectively when each is individually removed from the plug projections. Likewise, when scatter, beam hardening, and veiling glare effects are not accounted for in DRR postprocessing the average DRR pixel intensities are decreased by 32%, 19%, and 1%, respectively, using the CatphanTM DRRs as an example. From this we see that the veiling glare effect is overrepresented in the plug projections as compared to the object projections, while the beam hardening and scatter effects are under-represented. This suggests that the two large underestimations with regards to scatter and beam hardening outweigh the overestimation of veiling glare, leading to a net underestimation of the total effect, which in turn produces the observed underestimation of CBCT projection intensity in the uncorrected DRRs. Analyzing the source of the consistent overestimation of CBCT pixel intensity in the corrected DRRs is somewhat trickier, since our whole effort has been to accurately mimic the physics behind image formation. The most likely source is an overly aggressive beam hardening correction, which assumes that the object is composed of a single material (solid water in the case of the measured correction, liquid water in the case of the calculated correction). Using our calculation based framework we can test many different materials and examine the beam hardening functions they produce. For example, using polystyrene or adipose as the test material produces a less aggressive beam hardening function than liquid water, using muscle or lung produces a nearly identical correction, and using bone produces a more aggressive correction.

In addition to the systematic intensity offset present in the DRRs, the scatter distributions in Fig. 10 are also spread about the mean. This is likely a result of geometric misalignment of the FBCT scan, deformation of the object between FBCT and CBCT projection acquisition (if applicable), and the volume averaging effect caused by the relatively coarse resolution of the FBCT with respect to the size of the detector pixels. Indeed, the CatphanTM scatter plot, where the second of these factors is absent, shows the tightest distribution.

Of additional interest is our comparison of the calculated and measured beam hardening corrections. While the measurement approach has the benefit of using the actual x-ray technique and detector used to form the CBCT projections, the calculation method intrinsically avoids scatter and veiling glare, is uncontaminated by noise, and can be calculated

separately for each detector pixel to account for the nonuniform thickness of the bowtie filter. In addition, any desired test material can be used to simulate the beam hardening effect, opening up the possibility of a material-sensitive correction. The two corrections tend to agree well in the center of projection images, where we expect the measured approach to perform well since the beam hardening data are recorded there. They diverge toward the edges of the projections, with the calculation method providing a better match to the actual CBCT projections as increased bowtie filter thickness becomes a factor. For this reason we chose the calculation based correction as our method of choice, and all corrected DRRs other than Fig. 9 were postprocessed using calculated beam hardening.

IV.B. Comparison to methods in the literature

The x-ray volume rendering approach that we have employed has various benefits over other methods for the purpose of FIPM. Splatting algorithms are fast but produce inferior quality DRRs compared to ray-tracing methods, especially with respect to image sharpness.^{34,35} Shear-warp methods produce similarly blurry images.³⁵ Fourier volume rendering does not support perspective projection and thus is of limited use.¹⁸ Exact path length ray-tracing methods are known to suffer from artifacts due to aliasing.^{14,35} Our method avoids such artifacts because the trilinear interpolation used to sample the CT serves as an antialiasing (smoothing) filter. Moore *et al.* modified an exact path length algorithm to use a pencil beam rather than a ray as the projector to avoid artifacts.¹⁴ This approach provides a similar smoothing effect but is not as easily ported to run on the GPU as our method, which takes advantage of hardware-implemented trilinear interpolation.²³ Light field methods provide a means of interpolating between precomputed DRRs to generate DRRs from arbitrary angles.¹³ However, FIPM algorithms for deformable registration must calculate projections through continuously changing CT volumes, thus any precomputed projection data become useless as the algorithm iterates. Likewise, Monte Carlo methods rely on a precomputed sample of points within the CT space, limiting their utility for our purposes.¹⁸

While many different approaches to the x-ray volume rendering problem exist in the literature, references that construct realistic DRRs based on complete image formation physics are difficult to find. Moore *et al.* provide one such work, with the goal of creating a tool to aid in the optimization of computed radiography (CR) imaging parameters for chest radiographs that avoids additional patient exposures. Their algorithm contains many of the same elements as ours, though implemented in different ways. Their ray-tracing routine uses a volume based projector to sample the CT, whereas we use trilinear interpolation. Both algorithms are likely to produce similar results, but our approach is attractive since it is easily ported to run on the GPU. Beam hardening is another area of differentiation. Moore *et al.* explicitly account for beam hardening by casting many projections through the CT at different simulated energies, and then computing an average projection weighted by the measured x-ray spectrum of their machine. In

contrast, our approach computes a single DRR at the average beam energy and then applies a beam hardening correction as a postprocessing step. While their approach may provide a realistic simulation they pay a massive price in speed, with reported times of 45–90 min per DRR. Our correction incurs minimal additional computational overhead, a necessity for an application such as FIPM. In addition, Moore *et al.* use a beam stop array method to measure scatter, as opposed to our Monte Carlo approach. Measurement-based approaches have the benefit of intrinsically including the characteristics of the x-ray source/detector system. However, in order to avoid giving patients extra dose Moore *et al.* were forced to adapt scatter distributions measured from a RANDO phantom to their patient geometries, limiting them to patients with approximately the same proportions as the phantom. Using Monte Carlo, independent scatter distributions can be calculated for any patient, which is preferable for a registration algorithm meant to be generally applicable. Finally, Moore *et al.* do not account for the veiling glare effect. As discussed in Sec. IV.A, this effect has the greatest impact in the plug projections used to make our CT number to LAC conversion function. Moore *et al.* create a similar conversion function; however instead of directly measuring the LAC of each plug insert they calculate it from reported density and elemental makeup values. This approach has the drawback of depending on manufacturer-reported values for a critical aspect of the overall DRR calibration process, but renders the absence of a veiling glare correction more or less harmless. Overall, Moore *et al.* report CBCT and DRR intensity histograms that look similar, as well as comparable signal to noise ratios (SNRs) and tissue to rib ratios (TRRs, their measure of contrast) for different tissues within their real and calculated projections. The differences between their algorithm and our own do not necessarily reflect better or worse ways of producing DRRs, but instead relate to the different applications for which the DRRs are meant to be applied, particularly with regards to speed.

IV.C. Computational efficiency

Computational efficiency is an area of concern for the application of FIPM. Our test scenario used a 1024×768 detector matrix, a $512 \times 512 \times 281$ CT with voxels of approximately $1 \times 1 \times 1$ mm, and a ray step-size of 0.5 mm. Our workstation ran Debian Linux on an Intel Core i7 CPU clocked at 2.8 GHz with 12 GB of memory. Execution time scales linearly with number of detector pixels and ray step-size, and approximately linearly with the length of the projected rays. Under these conditions, uncorrected DRRs were calculated in about 0.35 s on average using a GeForce GTX 480 GPU. Over 95% of the computation time was spent performing trilinear interpolation on the CT image, which was carried out by the GPU hardware using special texture memory. Once the GPU finished ray-tracing, it handed the raw DRR back to the CPU where the scatter, beam hardening, and veiling glare postprocessing corrections were applied. The scatter and beam hardening corrections were performed in a matter of milliseconds, however the veiling glare correction, which requires a convolution between a kernel sampled

on a 600×600 grid and the 1024×768 pixel projection, generally took around 15 s. Of the three postprocessing corrections the veiling glare correction is by far the least important. As noted in Sec. IV.A, the veiling glare effect produces only a small change in image intensities of large object projections, such as CatphanTM images. Therefore in practice it will likely be acceptable to omit this step for the sake of efficiency. The veiling glare effect was mainly apparent when a small highly attenuating object was imaged in an otherwise unobstructed x-ray fluence. In this scenario a large proportion of the signal that appears in the object's shadow could be attributed to veiling glare. As demonstrated in Sec. IV.A the plug projections used for our CT number calibration curve fall into this category, thus it remains important to correct for veiling glare in the attenuation measurement data.

Finally, we note that the postprocessing corrections applied to our DRRs can just as easily be applied in reverse to "clean up" actual CBCT projections. This is likely the preferable route in the context of a FIPM algorithm, as it can be done once, offline, before the iterative reconstruction process is initiated. Even though this type of backwards correction is done offline, the veiling glare effect may still be too costly to include, as in reverse it requires a deconvolution taking 15–20 min using an implementation of the Richardson-Lucy algorithm.³⁶ However, it is likely that a large speedup can be obtained using the GPU.

V. CONCLUSIONS

We have demonstrated a step by step recipe for producing a physically realistic DRR that matches an actual CBCT projection with no artificial scaling, renormalization, or other adjustment. We have analyzed the effects of beam hardening, scatter, and veiling glare and have found that they play a small but non-negligible role in the agreement between DRR and CBCT projection. In addition, we have compared beam hardening corrections generated from measured and simulated data and determined simulation to be the superior method. Finally, we have benchmarked our algorithm for speed and found it to be sufficiently fast for application to FIPM type CT reconstruction methods.

ACKNOWLEDGMENTS

The authors would like to recognize prior work by our colleagues Dimitrios Lazos, Dandan Zheng, and Jeffrey Williamson that supported some aspects of this study. In addition, we would like to acknowledge Andrew Sampson for generously donating his time to help facilitate Monte Carlo simulations. This research was supported in part by National Institutes of Health (NIH) (Grant No. R01CA123299). The authors report no conflicts of interest in conducting the research.

^{a)}Author to whom correspondence should be addressed. Electronic mail: staubda@mymail.vcu.edu

¹A. Docef, M. J. Murphy, P. Keall, J. Siebers, and J. Williamson, "Forward CT reconstruction from limited projection data," in *Proceedings of the 19th*

Conference on Computer-Assisted Radiology and Surgery (Elsevier, Berlin, Germany, 2005), pp. 104–108.

²A. Docef and M. J. Murphy, "Reconstruction of a deformed 4D CT for moving anatomy," *Int. J. CARS* **3**, 591–598 (2008).

³D. Staub, A. Docef, R. S. Brock, C. Vaman, and M. J. Murphy, "4D cone-beam CT reconstruction using a motion model based on principal component analysis," *Med. Phys.* **38**, 6697–6709 (2011).

⁴R. S. Brock, A. Docef, and M. J. Murphy, "Reconstruction of a cone-beam CT image via forward iterative projection matching," *Med. Phys.* **37**, 6212–6220 (2010).

⁵A. H. Andersen and A. C. Kak, "Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm," *Ultrason. Imaging* **6**, 81–94 (1984).

⁶S. Leng, J. Tang, J. Zambelli, B. Nett, R. Tolakanahalli, and G. H. Chen, "High temporal resolution and streak-free four-dimensional cone-beam computed tomography," *Phys. Med. Biol.* **53**, 5653–5673 (2008).

⁷G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2D-3D medical image registration," *IEEE Trans. Med. Imaging* **17**, 586–595 (1998).

⁸D. Skerl, D. Tomazevic, B. Likar, and F. Pernus, "Evaluation of similarity measures for reconstruction-based registration in image-guided radiotherapy and surgery," *Int. J. Radiat. Oncol., Biol., Phys.* **65**, 943–953 (2006).

⁹D. Skerl, B. Likar, J. M. Fitzpatrick, and F. Pernus, "Comparative evaluation of similarity measures for the rigid registration of multi-modal head images," *Phys. Med. Biol.* **52**, 5587–5601 (2007).

¹⁰J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Mutual-information-based registration of medical images: A survey," *IEEE Trans. Med. Imaging* **22**, 986–1004 (2003).

¹¹J. Wu and M. J. Murphy, "Assessing the intrinsic accuracy of 3D/3D rigid image registration results for patient setup in the absence of a ground truth," *Med. Phys.* **37**, 2501–2508 (2010).

¹²R. L. Siddon, "Fast calculation of the exact radiological path for a 3-dimensional CT array," *Med. Phys.* **12**, 252–255 (1985).

¹³D. B. Russakoff, T. Rohlfing, K. Mori, D. Rueckert, A. Ho, J. R. Adler, and C. R. Maurer, "Fast generation of digitally reconstructed radiographs using attenuation fields with application to 2D-3D image registration," *IEEE Trans. Med. Imaging* **24**, 1441–1454 (2005).

¹⁴C. S. Moore, G. P. Liney, A. W. Beavis, and J. R. Saunderson, "A method to produce and validate a digitally reconstructed radiograph-based computer simulation for optimisation of chest radiographs acquired with a computed radiography imaging system," *Br. J. Radiol.* **84**, 890–902 (2011).

¹⁵W. Birkfellner, R. Seemann, M. Figl, J. Hummel, C. Ede, P. Homolka, X. H. Yang, P. Niederer, and H. Bergmann, "Wobbled splatting—A fast perspective volume rendering method for simulation of x-ray images from CT," *Phys. Med. Biol.* **50**, N73–N84 (2005).

¹⁶P. Lacroute and M. Levoy, "Fast volume rendering using a shear-warp factorization of the viewing transformation," in *Proceedings of the 21st Annual Conference on Computer Graphics And Interactive Techniques* (ACM, 1994), pp. 451–458.

¹⁷T. Totsuka and M. Levoy, "Frequency domain volume rendering," in *Proceedings of the 20th Annual Conference on Computer Graphics And Interactive Techniques* (ACM, Anaheim, CA, 1993), pp. 271–278.

¹⁸X. L. Li, J. Yang, and Y. M. Zhu, "Digitally reconstructed radiograph generation by an adaptive Monte Carlo method," *Phys. Med. Biol.* **51**, 2745–2752 (2006).

¹⁹J. Spoerk, H. Bergmann, F. Wanschitz, S. Dong, and W. Birkfellner, "Fast DRR splat rendering using common consumer graphics hardware," *Med. Phys.* **34**, 4302–4308 (2007).

²⁰S. Mori, M. Kobayashi, M. Kumagai, and S. Minohara, "Development of a GPU-based multithreaded software application to calculate digitally reconstructed radiographs for radiotherapy," *Radiol. Phys. Technol.* **2**, 40–45 (2009).

²¹C. L. Yang, B. D. Yang, and J. L. Wang, "Fast and high precision digitally reconstructed radiograph generation on graphic processing unit," *Biomed. Eng. Appl. Basis Commun.* **24**, 9–15 (2012).

²²H. Hong, K. Kim, and S. Park, "Fast 2D-3D point-based registration using GPU-based preprocessing for image-guided surgery," *Lect. Notes Comput. Sc.* **4225**, 218–226 (2006).

²³G. J. Tornai, G. Cserey, and I. Pappas, "Fast DRR generation for 2D to 3D registration on GPUs," *Med. Phys.* **39**, 4795–4799 (2012).

²⁴J. T. Bushberg, J. A. Seibert, E. M. Leidholdt, Jr., and J. M. Boone, *The Essential Physics of Medical Imaging*, 2nd ed. (Lippincott/Williams and Wilkins, Philadelphia, 2002).

- ²⁵E. P. Ruhrnschopf and K. Klingensbeck, "A general framework and review of scatter correction methods in x-ray cone-beam computerized tomography. Part 1: Scatter compensation approaches," *Med. Phys.* **38**, 4296–4311 (2011).
- ²⁶D. Lazos, D. Pokhrel, Z. Su, J. Lu, and J. F. Williamson, "Experimental validation of a Monte Carlo-based kV x-ray projection model for the Varian linac-mounted cone-beam CT imaging system," *Proc. SPIE* **6913**, 69133Q (2008).
- ²⁷O. Nalcioglu and R. Y. Lou, "Post-reconstruction method for beam hardening in computerized tomography," *Phys. Med. Biol.* **24**, 330–340 (1979).
- ²⁸P. M. Joseph and R. D. Spital, "A method for correcting bone induced artifacts in computed tomography scanners," *J. Comput. Assist. Tomogr.* **2**, 100–108 (1978).
- ²⁹C. H. Yan, R. T. Whalen, G. S. Beaupre, S. Y. Yen, and S. Napel, "Reconstruction algorithm for polychromatic CT imaging: Application to beam hardening correction," *IEEE Trans. Med. Imaging* **19**, 1–11 (2000).
- ³⁰A.-K. Carton, R. Acciavatti, J. Kuo, and A. D. A. Maidment, "The effect of scatter and glare on image quality in contrast-enhanced breast imaging using an *a*-Si/CsI(Tl) full-field flat panel detector," *Med. Phys.* **36**, 920–928 (2009).
- ³¹R. B. Benitez, R. Ning, and D. Conover, "Cone beam CT image quality measurements: PSF de-convolution," in *Medical Imaging 2007: Physics of Medical Imaging* (San Diego, CA, 2007), Vol. 6510, pp. U1925–U1936.
- ³²G. Poludniowski, P. M. Evans, A. Kavanagh, and S. Webb, "Removal and effects of scatter-glare in cone-beam CT with an amorphous-silicon flat-panel detector," *Phys. Med. Biol.* **56**, 1837–1851 (2011).
- ³³D. Lazos and J. F. Williamson, "Monte Carlo evaluation of scatter mitigation strategies in cone-beam CT," *Med. Phys.* **37**, 5456–5470 (2010).
- ³⁴N. Milickovic, D. Baltast, S. Giannouli, M. Lahanas, and N. Zamboglou, "CT imaging based digitally reconstructed radiographs and their application in brachytherapy," *Phys. Med. Biol.* **45**, 2787–2800 (2000).
- ³⁵W. Cai and G. Sakas, "DRR volume rendering using splatting in shear-warp context," in *Nuclear Science Symposium Conference Record, 2000 IEEE, Vol. 3* (IEEE Conference Publications, Lyon, France, 2000), pp. 19/12–19/17.
- ³⁶L. B. Lucy, "An iterative technique for the rectification of observed distributions," *Astron. J.* **79**, 745–754 (1974).