Dose reduction using a dynamic, piecewise-linear attenuator

Scott S. Hsieh^{a)}

Department of Radiology, Stanford University, Stanford, California 94305 and Department of Electrical Engineering, Stanford University, Stanford, California 94305

Dominik Fleischmann

Department of Radiology, Stanford University, Stanford, California 94305

Norbert J. Pelc

Department of Radiology, Stanford University, Stanford, California 94305 and Department of Bioengineering, Stanford University, Stanford, California 94305

(Received 21 June 2013; revised 11 November 2013; accepted for publication 1 January 2014; published 30 January 2014)

Purpose: The authors recently proposed a dynamic, prepatient x-ray attenuator capable of producing a piecewise-linear attenuation profile customized to each patient and viewing angle. This attenuator was intended to reduce scatter-to-primary ratio (SPR), dynamic range, and dose by redistributing flux. In this work the authors tested the ability of the attenuator to reduce dose and SPR in simulations.

Methods: The authors selected four clinical applications, including routine full field-of-view scans of the thorax and abdomen, and targeted reconstruction tasks for an abdominal aortic aneurysm and the pancreas. Raw data were estimated by forward projection of the image volume datasets. The dynamic attenuator was controlled to reduce dose while maintaining peak variance by solving a convex optimization problem, assuming *a priori* knowledge of the patient anatomy. In targeted reconstruction tasks, the noise in specific regions was given increased weighting. A system with a standard attenuator (or “bowtie filter”) was used as a reference, and used either convex optimized tube current modulation (TCM) or a standard TCM heuristic. The noise of the scan was determined analytically while the dose was estimated using Monte Carlo simulations. Scatter was also estimated using Monte Carlo simulations. The sensitivity of the dynamic attenuator to patient centering was also examined by shifting the abdomen in 2 cm intervals.

Results: Compared to a reference system with optimized TCM, use of the dynamic attenuator reduced dose by about 30% in routine scans and 50% in targeted scans. Compared to the TCM heuristics which are typically used without *a priori* knowledge, the dose reduction is about 50% for routine scans. The dynamic attenuator gives the ability to redistribute noise and variance and produces more uniform noise profiles than systems with a conventional bowtie filter. The SPR was also modestly reduced by 10% in the thorax and 24% in the abdomen. Imaging with the dynamic attenuator was relatively insensitive to patient centering, showing a 17% increase in peak variance for a 6 cm shift of the abdomen, instead of an 82% increase in peak variance for a fixed bowtie filter.

Conclusions: A dynamic prepatient x-ray attenuator consisting of multiple wedges is capable of achieving substantial dose reductions and modest SPR reductions. © 2014 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4862079>]

Key words: dynamic bowtie, dose reduction, fluence field modulation

1. INTRODUCTION

The ability to control the x-ray illumination field incident on the patient in CT improves image quality. The line integrals measured by CT vary drastically in their attenuation, and if no equalization is used, the variance of different measurements can vary by orders of magnitude because of differences in quantum statistics. The redistribution of flux can greatly reduce the uncertainty in the noisiest measurements without increasing total patient dose. Control over this flux distribution can additionally be used to avoid radiosensitive organs or to perform targeted scanning, with improved noise characteristics in a limited volume of interest for the same or lower patient dose.

Several techniques exist that provide limited control over the flux distribution. Tube current modulation (TCM), also

called mA modulation, controls the intensity of the x-ray source as it rotates about the patient, allowing flux to increase when photons must travel through the highly attenuating long axis of the patient.^{1,2} Beam shaping filters,³ such as the “bowtie filter” (named for its bowtie-like shape), are prepatient attenuators which control the flux distribution as a function of the fan angle. Typically, bowtie filters are minimally attenuating at the center and increase in attenuation as the fan angle increases, in order to partially compensate for the expected attenuation differences in a generic, centered, circularly symmetric object. However, it is not practical for CT scanners to be equipped with more than a small number of bowtie filters, so the filter cannot be personalized for a specific patient or optimized for a specific task. Beam shaping filters additionally reduce the dynamic range of the detector and reduce the scatter-to-primary ratio (SPR).

More sophisticated methods exist which provide greater control over the flux distribution, although none of these methods is currently used in clinical practice. Inverse geometry CT uses a small detector which is illuminated by a distributed source array.⁴⁻⁶ By modulating the mA used in each individual source, the flux distribution can be highly customized in a patient-specific and view-specific manner. Inverse geometry CT can, therefore, implement a *virtual bowtie filter* which is more flexible than a traditional, fixed bowtie filter. Investigations of this type of detailed flux control have previously shown that a dose reduction of 39%–52% is possible without increasing peak variance in one study,⁷ and by up to 83% in another, more favorable case in comparison to a system with a fixed bowtie and sinusoidal tube current modulation.⁸

Recently, we proposed a dynamic beam shaping filter (or dynamic attenuator) which is able to introduce a time-dependent piecewise-linear attenuation function into the beam and, therefore, achieves a level of flux control comparable to inverse geometry CT.⁹ (In this work, “attenuation” will refer to the integral of the linear attenuation coefficient of the ray measured, i.e., $\int \mu(x)dx$ in the Beer-Lambert law $I = I_0 e^{-\int \mu(x)dx}$.)

A related, piecewise-constant dynamic attenuator has also been independently proposed and demonstrated.¹⁰⁻¹² This attenuator is quite similar conceptually but produces a piecewise-constant attenuation function instead of a piecewise-linear function. This piecewise-constant function is discontinuous at certain angles, which may present challenges during reconstruction but should otherwise provide a similar level of fluence control.

The goal of our study was to evaluate the ability of the piecewise-linear dynamic attenuator to reduce radiation dose. Preliminary studies have shown the benefits of dynamic flux control are substantial,^{7,8} but generally simplified the problem, ignoring polychromatic and volumetric effects, and did not study the specific piecewise-linear attenuator that we have proposed.

The paper is organized as follows. In Sec. 2, we begin by summarizing the dynamic attenuator concept. We then describe the other tools used in our simulation framework. In Sec. 3, we describe the studies we conducted. While the primary focus of this work is studying dose reduction from the dynamic attenuator, we also carried out secondary experiments on patient centering effects, scatter-to-primary ratio, and kVp choice. We show results in Sec. 4, with discussion and concluding remarks presented in Sec. 5.

2. SIMULATION FRAMEWORK

2.A. Piecewise-linear dynamic attenuator

As context for the reader, we first briefly review the piecewise-linear dynamic attenuator. For details, the reader is directed towards Ref. 9.

To build up the piecewise-linear attenuation, the dynamic attenuator comprises N wedges. For ease of notation let N be odd with $N = 2M + 1$ for a positive integer M ; a general-

ization to N even would be straightforward. Each wedge of the dynamic attenuator casts a shadow on the detector, and let the total angular width of the shadow of one wedge be γ_b . In simulations, we assumed the wedges to be composed of iron.

We design each wedge such that the cross section of each attenuating wedge is a triangle, and we place each wedge so that the shadows they cast on the detector are $\frac{\gamma_b}{2}$ apart. If the k th triangular wedge has maximum thickness c_k , then the cumulative attenuation provided by our dynamic attenuator, as a function of fan angle γ , takes the form

$$\mu_{\text{DA}}(\gamma) = \sum_{k=-M}^M c_k \Lambda\left(\frac{2\gamma}{\gamma_b} - k\right), \quad \Lambda(x) = \max(0, 1 - |x|).$$

Here, $\Lambda(x)$ is simply the triangle function. The sum of all the triangular wedges is, therefore, a piecewise-linear function. The c_k are dynamically adjusted during the scan to provide the flexibility that we desire. The adjustment is performed by translation of each wedge in the axial direction (also called the transverse, slice, or z-direction). The wedges are shaped such that each cross section is triangular with varying thicknesses, and that by moving a wedge in the axial direction, the thickness c_k of the wedge is changed. Two layers of wedges are used, with the layers offset by half of the triangle base. Figure 1(a) shows the wedges in relation to the fan beam, and Fig. 1(b) shows a cross section with the individual triangle pieces visible. Figure 2 shows example piecewise-linear attenuation profiles for a stylized thorax in both the lateral and A/P directions. In the thorax, the dynamic attenuator has sufficient flexibility to compensate for the lung fields.

The piecewise-constant dynamic attenuator¹⁰⁻¹² is quite similar to the piecewise-linear attenuator, but uses differently shaped wedges. Specifically, the wedges produce rectangle function cross sections instead of triangle function cross sections.

The dynamic attenuator is designed for single or multi-slice systems of reasonable thickness (e.g., ≤ 4 cm), not wide cone beam systems. The attenuation profile can be flat in the cone (or axial) direction with the addition of a compensating wedge.

FIG. 1. The piecewise-linear dynamic attenuator, consisting of a series of wedges attached to actuators (not shown). The fan beam is also shown, and the x-ray focal spot sits at the apex of the fan. (a) The set of N wedges. Two of the wedges have been shaded and translated. (b) A cross section of the wedges, showing the triangular cross sections in the plane of the fan beam. The wedge which was translated downward introduces only a small amount of attenuation, while the wedge that was translated upward introduces more attenuation.

FIG. 2. Example attenuation functions for the thorax. In the lateral direction, the dynamic attenuator is able to modulate its shape to become more narrow. In the anterior/posterior or A/P direction, the dynamic attenuator can expand its shape and additionally compensate for the effect of the lung field.

2.B. Raw data

Raw data from CT scans were not available to us. We estimated raw data by forward projection of deidentified DICOM files. This retrospective use of human data was approved by our Institutional Review Board.

To estimate the raw data, the DICOM files were segmented into mixtures of bone, water, and air on a voxel-by-voxel basis and forward projected. Detector and source blurring were not explicitly modeled. All voxels with value less than -700 HU were rounded to -1000 HU and assumed to be in air. If the voxel value was between -700 and 200 HU, it was assumed to be composed of water with density $\frac{x}{1000} + 1$ g/cc, where x is the observed HU value. If the voxel value was greater than 200 HU, the voxel was assumed to be composed of the mixture of water and cortical bone which would cause the voxel to

be reconstructed at the observed HU value. We assumed that the CT number of pure cortical bone was 2000 HU and that its density was 1.92 g/cc. We note that it may be more appropriate to model certain voxels as mixtures of water and iodine instead of water and bone, but for the purpose of estimating dose and noise and especially since the same numerical object was used in both the standard bowtie and the dynamic attenuator, we assumed that these errors were negligible. Additionally, in our scans a small section of the table is segmented into water and included in the dose calculations, and we assumed that this error was also negligible.

To simulate polychromatic effects, a 120 kVp spectrum was discretized into 10 keV intervals and the energy weighted signal at each 10 keV interval was estimated, scaled by the number of photons in the interval, and summed.

Table I summarizes the system parameters used in this work. A helical trajectory with a pitch of 1 was assumed. In order to decrease the computation necessary for the simulations, the resolution of the system used and the raw data matrix were both significantly downsampled. We believed that this downsampling is acceptable because the variance of a scan has little content at high spatial frequencies, and the dose is an aggregate quantity which should not depend greatly on resolution. A sensitivity analysis was later conducted to confirm our assumption, using a different resolution, to show that the results are not significantly impacted by the choice of system parameters.

2.C. Variance calculations

Many algorithms exist for reconstructing from a helical trajectory. We used an algorithm that suppressed helical artifacts by three-dimensional ray weighting.¹³ The reconstruction algorithm contains a tunable value kh which may be adjusted to reduced helical artifacts, and in this work we chose $kh = 2$, as suggested in Ref. 13 for scans with a pitch of approximately 1 . Note that, because of the decrease in system resolution (Table I), reconstructing the images from the estimated raw data is not useful. The limited number of views and detector channels would result in limited resolution and streaking artifacts. We limit ourselves to reconstructing estimates of the variance only.

TABLE I. System parameters.

Source-isocenter distance	54.1 cm
Detector-isocenter distance	40.8 cm
Reconstructed field of view	40 cm
Length of scan	2160° for AAA 1800° for pancreas 1080° for other datasets
Helical pitch	1
Detector size	1040 × 70 mm
Detector pixel size	5 × 5 mm
Views per revolution	64
Detector type	Equiangular
X-ray spectrum	120 kVp
Number of wedges	15
Wedge material	Iron
Gantry rotation time	300 ms
Attenuator modulation rate at wedge center	1.94% per millisecond
Reconstruction voxel side length	1.56 mm
Reconstruction voxel side length in optimization step	6.25 mm
Views per revolution in optimization step	16
Voxel side length in Monte Carlo dose simulations	8 mm

The variance of each voxel in the image was estimated analytically by a weighted, unfiltered backprojection of the variance of each ray. Unfiltered backprojection is derived in Ref. 14 for variance estimation in parallel-beam scans but applies to any FBP-type reconstruction. Reference 14 shows that unfiltered backprojection is equivalent to a full propagation of uncertainty from the sinogram into the reconstruction, assuming that the ramp filter is approximated as a local operation. The analytic calculation of the noise agreed with Monte Carlo simulations using multiple noise realizations. Briefly, in reconstructing a function $f(r)$ from sinogram data $s(t, \theta)$, with t being the signed distance from the origin, FBP-type algorithms can be written in the form

$$f(r) = \sum_j \sum_i w(r, i, j) s_{\text{ramp}}(i, j),$$

$$s_{\text{ramp}}(t, \theta) = s(t, \theta) * h(t),$$

where $h(t)$ is the ramp filter and $w(r, i, j)$ includes weighting which depends on the specific reconstruction algorithm. The variance of the reconstruction is calculated using propagation of uncertainty as

$$\sigma_f^2(r) = \sum_j \sum_i w^2(r, i, j) \sigma_{s_{\text{ramp}}}^2(i, j).$$

We make the approximation that $\sigma_{s_{\text{ramp}}}^2(i, j) \propto \sigma_s^2(i, j)$. For most datasets, this is justified because the variance in the sinogram is smoothly varying and because $h(t)$ is very local, and decreases as $\frac{1}{t^2}$.¹⁵ In calculating the variance, the weights are again squared so the contribution to variance goes as $\frac{1}{t^4}$.

To adapt unfiltered backprojection to a particular analytic reconstruction algorithm, the weighting term $w^2(r, i, j)$ must be accurately modeled. The three-dimensional ray weighting algorithm we used recalculates the weights of the rays for each voxel.¹³ The variance of each ray must further be adjusted to include a water beam hardening correction. The water beam hardening was implemented by interpolation of a lookup table which estimates the length of water penetrated from the signal received at the detector, after adjusting for the known thickness of the bowtie (iron for the dynamic attenuator, water for the traditional bowtie). The square of the local slope of the water beam hardening correction was used to adjust the variance of the ray.

We assumed Poisson statistics only in estimating the variance of each ray, and neglected potential electronic noise.

2.D. Dose calculations

Monte Carlo software (GEANT4) was used to estimate the dose of the scan.¹⁶ To estimate dose, the input dataset was downsampled into cubic voxels consisting of mixtures of bone, water, and air. The dataset was extruded by 10 cm in each direction so that dose resulting from scatter would be included. A polychromatic spectrum was used in the Monte Carlo simulations.

Because a piecewise-linear attenuation function for a diverging beam requires the modeling of nonpolygonal shapes, the bowtie was approximated in GEANT4 using a series of

wedges of constant thickness. The bowtie thickness was sampled at 161 equally spaced locations. At each location, a piecewise-constant wedge was placed in the GEANT4 simulation with the sampled thickness. The GEANT4 simulation includes contributions for photons which scatter from the dynamic attenuator or bowtie. About 5000 photons were used in each view to calculate dose. The calculated dose in each view was then scaled by the actual number of photons in that view in the simulation. For each photon, a random location on the detector was chosen and the photon was sent toward that detector location.

2.E. Systems compared and control methods

We compared five system configurations: (1) the reference bowtie with heuristic TCM, (2) without any TCM, (3) with optimized TCM, (4) the dynamic attenuator with optimized control, and (5) the dynamic attenuator with heuristic control. We include three different control methods for TCM because there is no standard mechanism for determining the TCM profile. We do not include the combination of dynamic attenuator and TCM because the effect of TCM can be approximated by a uniform shift of all the dynamic attenuator wedges. In implementations, however, the use of TCM will be important to reduce the thermal load on the anode.

The optimized control methods provide best-case estimates of the performance of both the dynamic attenuator and the reference system. We assumed knowledge of the patient anatomy, so that the control of the dynamic attenuator or tube current modulation can be tailored specifically to the patient. Possible sources of knowledge include a low-dose prescan prior to the diagnostic scan, a patient atlas, scout scans, or a prior CT scan. The dose of the prescan could be many times less than that of the diagnostic scan because its purpose is only to estimate the major features of the anatomy. In this work, we assume a noiseless prescan and neglect the dose associated with the prescan, but we do not expect that noise in the prescan will materially impact our results. The clinician or an algorithm could segment regions of interest for targeted scans using the prescan, or segment radiosensitive organs which would be protected in the diagnostic scan.

In many applications a prescan of the patient is not practical, and the optimized performance we report here would not be realized. In general, the dose reduction that can be achieved depends on the control method used. Control methods which do not require knowledge of patient anatomy are an important subject of active research by our group. As described below, we used a well-known heuristic control algorithm for TCM,¹ which has theoretical justification, but no such control algorithm exists for the dynamic attenuator. We have adapted the TCM heuristic to the dynamic attenuator as one possible control algorithm, but we emphasize that better control algorithms which do not require *a priori* knowledge may be possible.

2.E.1. The reference bowtie

The reference bowtie was assumed to be composed of a material equivalent to water. Although real bowtie filters

FIG. 3. Bowtie filter used in this study. An example of a possible attenuation function delivered by the piecewise-linear dynamic bowtie is also provided.

are not composed of water, the materials used are similar to water because this choice reduces beam hardening artifacts. In contrast, the material used for the dynamic attenuator is iron because it is more highly attenuating. This could exacerbate beam hardening artifacts, but previous work has showed that corrections can be used to compensate.^{9,17}

The shape of a reference bowtie was chosen to be similar to a bowtie in clinical systems. Figure 3 shows the reference bowtie used, in addition to an example attenuation function from the dynamic attenuator.

2.E.2. Heuristic TCM

The combination of reference bowtie and square-root heuristic TCM approximates clinical systems which are used today. The tube current is modulated in each view proportional to $\frac{I_0}{\sqrt{I}}$, where I_0 is the intensity of the photons incident on the patient and I is the intensity of the photons detected for the central channel of the detector. The rationale for this method is described in Ref. 1, but it can be shown under certain assumptions to minimize the variance for the center pixel of a scan for a fixed dose. We assumed for ease of modeling that the source current could change instantaneously, although in actual implementations the response speed of the x-ray source may be more limited. We emphasize that, according to the assumptions in Ref. 1, the square-root control method minimizes the variance of the center pixel for a given dose constraint, but in some cases, the relevant variance to minimize is far from the center, and in these cases disabling the square-root heuristic TCM may be beneficial.

2.E.3. No TCM

As a second point of comparison, we include a system without TCM.

2.E.4. Optimized control of TCM

If the anatomy is known, the TCM can be determined using optimization. Because the reconstruction is a linear process, the variance of the reconstruction can be decomposed into the variance of individual views. Let v_i be a column vector corresponding to the i th view. The j th entry of v_i is the variance of the j th voxel in the image if the tube current for

this view is set to unity and no other views contribute to the noise. If the system has N views, then let $V = [v_1 v_2 v_3 \dots v_N]$, and let x be a column vector with x_k being the tube current for the k th view. With these assumptions, $V^T x^{-1}$ is a vector describing the variance of the reconstructed image. x^{-1} is the elementwise inverse of the tube current vector x . The variance of a line integral measurement scales inversely with the number of photons detected, and so the variance scales linearly with x^{-1} if the only source of noise is photon statistics, as we have assumed. V itself may be constructed using simulation by considering the variance contribution of each view separately, assuming unit source current.

Finally, let d be a column vector with d_i being the average dose in the imaged volume contributed by the i th view if the tube current is set to unit intensity, so that $d^T x$ is the total dose delivered by the scan. The dose vector d is determined using Monte Carlo simulations.

The optimization problem can now be written as

$$\text{minimize } \|V^T x^{-1}\|_{\infty}$$

$$\text{subject to } x > 0$$

$$d^T x \leq d_{total}$$

$$|x_k - x_{k-1}| < c_1 \text{mean}(x) \forall k > 2$$

$$|x_k + x_{k-2} - 2x_{k-1}| < c_2 \text{mean}(x) \forall k > 3$$

This problem is convex and optimization can proceed with guarantees on convergence and correctness. A dedicated convex optimization package (CVX) was used to find the solution.^{18,19}

$\|V^T x^{-1}\|_{\infty}$ is the infinity-norm or the *peak variance* of the reconstruction. We note that we could instead optimize $\|V^T x^{-1}\|_{\alpha}$ with $\alpha = 1$ to optimize the mean variance, or we could choose $1 < \alpha < \infty$ to find a compromise between the mean and peak variance, while still preserving the convexity of the problem. d_{total} is a constant which determines the maximum permissible dose of the scan.

For targeted reconstructions, the rows of V corresponding to voxels in the region of interest are multiplied by a weighting factor which depends on the specific task. We used a relative weighting factor of 10. As previously stated, CT systems cannot instantaneously change the source current. The regularization parameters c_1 and c_2 reflect this fact by constraining both the first and second derivatives of the source current over time. For our study, the values of c_1 and c_2 were chosen by comparison to a unit amplitude sinusoid which undergoes two periods per revolution. For R views per revolution, such a sinusoid takes the form

$$x(t) = 1 + \sin\left(\frac{4\pi t}{R}\right), \quad |x'(t)| \leq \frac{4\pi}{R}, \quad |x''(t)| \leq \left(\frac{4\pi}{R}\right)^2.$$

We, therefore, choose $c_1 = \frac{4\pi}{NR}$ and $c_2 = \left(\frac{4\pi}{NR}\right)^2$. To make this choice invariant to a scaling of $x(it)$, a $\text{mean}(x)$ is included in the optimization constraints. If the actual generator limitations are known, they can readily be incorporated into our formalism.

In order to make the computation tractable, the system was downsampled (Table I). We also assumed a monochromatic

spectrum, which is not necessary for the optimization of TCM but which simplifies the dynamic attenuator control optimization in Sec. 2.E.5. The optimized tube current vector x is, therefore, an approximate optimum. Note that while monochromatic x rays were assumed for the calculation of the TCM or dynamic attenuator control, the computation of image variance and noise used the polychromatic spectrum.

2.E.5. Optimized control of dynamic attenuator

Control of the dynamic attenuator was also posed as a convex optimization problem. For a dynamic attenuator with N wedges, let x_{MNi+j} correspond to the number of photons transmitted through the center of the j th attenuating element for the i th view. We note that x corresponds to photons transmitted, not the positions of the attenuating elements themselves, which are proportional to attenuation or thickness rather than flux transmitted. We formulate the optimization problem as

$$\begin{aligned} & \text{minimize } \|V^T x^{-1}\|_{\infty} \\ & \text{subject to } x > 0 \\ & \quad d^T x \leq d_{total} \\ & \quad x_k > c x_{k-N} \forall k > N \\ & \quad x_k > \frac{1}{c} x_{k-N} \forall k > N. \end{aligned}$$

The position of the attenuating elements is constrained by the finite speed of the actuators used. Our specification is that the system should be able to modulate the photon transmission at a rate of 2.5 orders of magnitude per 300 ms gantry rotation. This is roughly equivalent to a change of 1.94%/ms, or a doubling of intensity every 28 ms (or 34° of gantry rotation). For a spacing between adjacent views of d degrees, we choose $c = (1 + 1.94\%)^d$. Physically, the maximum c possible depends on the speed of the actuator and on geometry of the wedges. For an iron wedge with a maximum thickness of 8 mm and an axial length of 80 mm, an actuator which translates at 25 cm/scan meet our specification as it introduces 0.025 mm of additional iron into the beam every millisecond, or approximately 2.1% additional attenuation at 60 keV. With a polychromatic beam the rate of change of the intensity possible will vary as the beam hardens, but based on our experience with small actuators we believe that our current choice of c is practical. If desired, second derivatives to prescribe force limits could also be included in the convex formulation, and a minimum thickness of attenuator could also be included.

The optimization problem was formulated assuming that for any given view, the dynamic attenuator transmits a piecewise-linear *flux* profile instead of a piecewise-linear *attenuation* profile (with attenuation proportional to the thickness of the attenuator). As an example of the difference, consider two adjacent attenuating elements, with the first attenuating element transmitting NP photons at the center and the second attenuating element transmitting $2NP$ photons at the center. Consider the ray midway between the two attenuating elements. Under the piecewise-linear flux assumption, $\frac{3NP}{2}$ photons are transmitted, but with the physically realizable piecewise-linear attenuation assumption, $\sqrt{2}NP$ photons are transmitted in this ray. While convex optimization

with piecewise-linear attenuation is possible, the formulation is more cumbersome and it is computationally more intensive. Again, even though the optimization used this simplifying assumption, the dose and noise testing used the correct physical model.

Assuming piecewise-linear flux, the elements of the dose vector can be calculated by simulating a triangle function of flux for each attenuator. In calculating v_{MNi+j} , the variance contributed by the j th attenuating element in the i th view, we used a triangle function with unit amplitude (corresponding to unit flux) in reciprocal flux space, while all other attenuating elements transmitted infinite intensity flux (and, therefore, zero variance).

The x-ray source is assumed to be monochromatic in order to eliminate the nonlinear effect of beam hardening on the variance of the ray. This would make the problem much more complicated and possibly nonconvex. The dose constraints also do not include dose resulting from scatter from the attenuator.

The optimization problem, therefore, includes a number of simplifications, and we only claim to find an approximate optimum. After the control of the wedges was determined, the dose was re-estimated using a polychromatic spectrum, using wedges obeying piecewise-linear attenuation instead of piecewise-linear flux, and with the attenuator directly modeled in GEANT4. In converting the control scheme from monochromatic to polychromatic, the positions of the attenuating wedges at each view were modified so that the fraction of the energy transmitted was unchanged. We emphasize that these simplifications are only used to determine an approximate optimum for attenuator control, and the final dose reduction values we report do not make these simplifications.

The total running time of the algorithm, including constructing the V matrix and the d vector, simulating all possible control methods and conducting Monte Carlo simulations for each of them, was about 3 h for a 20-cm thick dataset. The convex optimization itself was quite fast, typically requiring less than a few minutes per control method.

2.E.6. Heuristic control of dynamic attenuator

While optimal control of the dynamic attenuator establishes an upper bound on performance, *a priori* information of the patient anatomy is not always available. We, therefore, included an alternative control mechanism that positions the wedges using real-time feedback from the previous view.

In analogy with the heuristic control for TCM, we modulate the position of each wedge such that the intensity of the radiation incident on the patient is scaled as the $\sqrt{\frac{I_0(\gamma_c)}{I(\gamma_c)}}$, where γ_c is the fan angle of a representative ray for each wedge, or the ray which passes through the center of the wedge, and $I_0(\gamma_c)$ is the intensity of the radiation after transmission through the dynamic attenuator but prior to the patient. Limitations on the wedge motion are still respected, and the wedge is moved in a greedy fashion. If the wedge cannot be moved to the target location in a given view, it is moved as far as the velocity constraint permits. We emphasize that this

heuristic control method was chosen for its simplicity and does not have the theoretically desirable properties of the square-root law for TCM. Better heuristic control methods are a subject for future work.

3. STUDIES

3.A. Dose reduction study

The primary goal of our work was to estimate the dose reduction possible with the dynamic attenuator in clinical settings. To capture the possible clinical use cases of the dynamic attenuator, we identified four separate clinical applications, and obtained deidentified DICOM image datasets from either our clinic or from an online source.²⁰

To assess the value of the dynamic attenuator in routine imaging, an abdominal/pelvic exam and a chest (or thorax) study were used. Abdominal and pelvic CT scans are among the most common types of CT exams and one study of potential cancer risk identified abdominal/pelvic CT as being the source of 48% of all future potential radiation risk from CT.²¹ Typical body bowtie filters perform well for the abdomen or pelvis. Chest CT is also common but the thorax has highly nonuniform attenuation and the bowtie filter may not perform as well here. In particular, the use of a standard bowtie will lead to increased dose and reduced noise in the lung field. This is inefficient from a clinical perspective because structures in the lung field are intrinsically high in contrast, the lung is fairly sensitive to radiation, and for many studies it would be preferential to increase SNR in the mediastinum, not in the lung field.

The dynamic attenuator is able to perform targeted studies. Therefore, datasets from a pancreatic study and a scan of an abdominal aortic aneurysm (AAA) were used as our third and fourth datasets. In these datasets, a small volume of interest was defined and the noise characteristics in this volume of interest were preserved. The noise elsewhere was allowed to increase to reduce dose.

Figure 4 shows a representative cross-sectional image from each of these 3D datasets.

The control of the wedge motion or tube current was determined as described in Sec. 2.E. Once an approximately optimal control trajectory was determined, the dose reduction was calculated without the simplifying assumptions used in the optimization except reduced resolution, which was used throughout.

To aid in visualization, noisy image realizations were generated in individual axial slices. To generate these noisy image realizations, we created noise-only images that depended on the simulated flux and added them to the original DICOM images. We emphasize that this was not a reconstruction using forward projected data with noise, which would have created images with highly visible artifacts. To make the noise more realistic, the number of views in the noise sinogram was upsampled by a factor of 4, and the detector spacing was upsampled by a factor of 5 (1 mm detector elements). The entries of the noise sinogram were realizations of Gaussian random variables, assuming that Poisson statistics could be approximated by normal variables. The variance of each Gaussian random variable was the variance of the corresponding ray in the raw data. This noise realization was then filtered and backprojected and added to the original, full-resolution image to produce a noisy reconstruction. Note that by adding a noise image to the original DICOM image, we avoid introducing artifacts that would result from forward projection and then backprojection with a small number of views and detector elements.

Although similar to existing noise insertion methods, our noise estimates contain a number of simplifications because the raw data of the scan was not available to us.²² We do not have estimates of the noise of the original raw data, so we scaled the number of photons in the simulations so that the inserted noise dominates the noise already present in the original image. Also, the helical geometry of our forward projected raw data and the apodization filter used in the reconstruction do not match the original DICOM image. The noise realizations we show are estimates only, with the goal of providing examples of the noise texture distribution throughout the image and a comparison between the dynamic attenuator and a standard bowtie.

3.B. Patient centering study

Because the bowtie filter is designed for a centered patient, the dose efficiency and image quality of a conventional system suffer for off-centered patients.^{23,24} Although careful attention or automated algorithms can often improve patient centering, this introduces an additional step to the workflow, and for some patients (for example, those with scoliosis) proper positioning is very difficult. We hypothesized that the dose inefficiency resulting from off-centered patients should be much less pronounced on systems with a dynamic bowtie. A

FIG. 4. Datasets used in this study, consisting of (from left to right) an abdomen, thorax, pancreas, and abdominal aortic aneurysm (AAA). The pancreas and AAA are targeted scans, and the targeted region is displayed at regular intensity while the nontargeted region is darkened.

FIG. 5. (Left column) Sinograms, (middle column) TCM profiles, and (right column) example piecewise-linear attenuation profiles for the (top row) abdomen, (top middle row) thorax, (middle row) AAA, (bottom middle row) pancreas, and (bottom row) abdomen shifted by 4 cm. The sinograms have three horizontal lines drawn in them, corresponding to different views for which the example piecewise-linear attenuation profiles are plotted. The bowtie filter is plotted as a reference to the example piecewise-linear attenuation profiles, with its attenuation converted to an equivalent thickness of iron at 60 keV. The conformity of the dynamic attenuator to the patient attenuation can be seen in several locations, including the middle example for the shifted abdomen.

secondary study was, therefore, conducted to compare the performance of the reference system to the dynamic attenuator system by shifting the abdomen dataset in image space prior to processing.

We studied shifts of 0 (i.e., centered patient), 2, 4, and 6 cm. For each, the systems were controlled as described in the dose reduction study. We report the relative dose neces-

sary to achieve the same peak variance as in the centered patient.

3.C. Scatter-to-primary ratio study

In addition to calculating dose, the Monte Carlo software can also calculate scatter. The SPR varies across the field of

view. In this study, we estimated it in a region 10×7 cm at the center of the detector (corresponding to 5.7×4 cm at isocenter). We did not model an antiscatter grid in our simulations. The SPR calculation requires better statistics than the dose calculation, because many photons are attenuated in the patient and do not reach the detector. For the SPR calculation, we simulated 400 000 photons per view, and only sampled 64 views distributed uniformly throughout the scan.

The SPR is reported for both the standard bowtie and the optimized dynamic attenuator. Note that the SPR is independent of tube current modulation, because a rescaling of flux in any view scales both primary and scatter together. We examined the SPR for three different datasets: the abdomen, thorax, and the targeted AAA scan.

3.D. Resolution and kVp sensitivity analyses

To ensure that our results were not an artifact of the limited resolution of the forward and backprojection, we doubled the resolution for the variance calculations. In Table I, we used 128 views per resolution instead of 64 and doubled the detector sampling density so that pixels of 2.5×2.5 mm were used.

It is possible that the optimization strategy or dose reduction depends on the kVp of the scan. Therefore, we conducted the dose reduction study both at 80 and 120 kVp.

For both sensitivity analyses, only the abdomen dataset was studied.

4. RESULTS

4.A. Systems compared and control methods

Sample tube current modulation profiles, sinograms, and piecewise-linear attenuation profiles are shown in Fig. 5. While the square root heuristic for TCM produces a complex, rapidly changing mA profile that may not be possible with real x-ray sources, the convex optimized profile has been regularized to be reasonably smooth. The piecewise-linear profiles can be seen to compensate for the attenuation of the patient.

4.B. Dose reduction study

Table II summarizes the relative dose for each of the datasets while maintaining the same peak variance. In the abdomen, the dose reduction for the dynamic attenuator with optimized control was 33% relative to a standard system with optimized TCM ($\frac{42}{63} = 0.67$) and 48% relative to the square root heuristic TCM. For the thorax, the dose reduction from the dynamic attenuator was 25% relative to optimized TCM and 54% relative to fixed mA (which outperformed the square root TCM on this particular dataset). In targeted scans, the potential gains are 56% and 48% for the AAA and the pancreas, respectively, compared to optimized TCM. A comparison to the square root heuristic TCM is not as sensible for targeted scans because the organ of interest does not span the entire volume. In all cases, the process of optimizing the TCM

TABLE II. Dose reduction found on different datasets and with different system options.

Dataset	System configuration	Dose for equal peak variance (%)
Abdomen	Standard bowtie, no TCM	100
	Standard bowtie, heuristic TCM	80
	Standard bowtie, optimized TCM	63
	Dynamic attenuator, heuristic	55
	Dynamic attenuator, optimized	42
Thorax	Standard bowtie, no TCM	100
	Standard bowtie, heuristic TCM	159
	Standard bowtie, optimized TCM	61
	Dynamic attenuator, heuristic	59
	Dynamic attenuator, optimized	46
Abdominal aortic aneurysm (targeted)	Standard bowtie, no TCM	100
	Standard bowtie, heuristic TCM	85
	Standard bowtie, optimized TCM	72
	Dynamic attenuator, optimized	32
Pancreas (targeted)	Standard bowtie, no TCM	100
	Standard bowtie, heuristic TCM	84
	Standard bowtie, optimized TCM	42
	Dynamic attenuator, optimized	22

produces substantial dose savings over the use of the square-root TCM heuristic, consistent with previously reported results.⁸ When *a priori* knowledge is not available, heuristic control of the dynamic attenuator still performs well.

For the targeted AAA study, strong discontinuities occur in the weighting function used to calculate peak variance and which controls the targeted nature of the scan. As indicated above, the optimization of the TCM is performed at a reduced resolution, which leads to potential errors at the boundaries of the targeted volume of interest. We also found that errors occurred at the very end of the datasets and we discarded two slices at one of the ends (2 slices out of a total of 128) in the calculation of the peak variance. Otherwise, the optimized TCM produces worse results (87% instead of 72%) and the optimized dynamic attenuator also suffers somewhat (35% instead of 31%). The performance of TCM with the square-root heuristic is unchanged.

Noisy realizations with the dynamic attenuator and the standard bowtie filter are shown in Figs. 6 and 7. Figure 6 shows low-dose noise realizations at constant peak variance, and Fig. 7 shows them at constant dose. The noise realizations are generated by adding noise images to the original DICOM datasets. In selected rectangular ROIs, we indicate the standard deviation of the noise images, not the sum of the noise images with the DICOM datasets. This gives us the flexibility of selecting ROIs with nonuniform image signal.

Figure 8 shows variance maps through the central axial slice of the reconstruction. This provides a complementary way of visualizing the noise distribution in the reconstructed images. In order to illustrate the benefits provided by the dynamic attenuator, these variance maps show the variance at constant dose in each dataset (not constant peak variance). In most cases, the dynamic attenuator redistributes photons to

FIG. 6. Noise realizations using (top row) the standard bowtie filter with optimized TCM and (bottom row) the dynamic bowtie with optimized control. The datasets are (far left) centered abdomen, (center left) thorax, (center right) targeted AAA, and (far right) targeted pancreas. Within each dataset, the peak variance is held constant. This corresponds to dose reductions of 33%, 25%, 56%, 47% for the abdomen, thorax, AAA, and pancreas, respectively. The standard deviation of the noise is calculated in certain ROIs and is indicated on the corner of each image. Note that the ROI standard deviation is calculated from the noise image only, and does not include the variation from the anatomy. See text for details. In the ROI names, “left” and “right” refer to the location on the page, and not in the patient.

make the noise distribution more homogeneous. In targeted imaging studies, the region of interest has superior noise characteristics.

Figure 9 shows coronal views of these variance maps. We believe the banded nature of noise arose from the particular

helical reconstruction algorithm used, which causes rays with a very shallow cone angle to have a weight approaching 1, while rays with intermediate cone angles often have weights of close to 0.5 because a given ray and its conjugate ray contribute equally to the reconstruction.¹³ Reducing

FIG. 7. Same noise realizations as Fig. 6 but at constant dose. In this application, the dynamic attenuator is used for reduced noise instead of reduced dose.

FIG. 8. Axial cross sections of variance maps using (a) the reference bowtie with square-root heuristic TCM, (b) the reference bowtie with optimized TCM, and (c) the dynamic attenuator with optimized control. Within each dataset, the variance maps are individually windowed. Variance maps represent analytically predicted variance with the mA tuned so that the average dose delivered is the same for all methods in each dataset.

the kh parameter in the reconstruction may reduce these bands.

In the thorax, the square root heuristic TCM does a poor job at minimizing the peak variance. The heuristic TCM is designed to minimize the variance at the center pixel for a given total mAs, but in the thorax the peak variance is not at the center pixel. The heuristic TCM decreases the variance in the mediastinum at the expense of increasing the variance in the back and in the spine, including the region containing the peak variance. As shown also in Fig. 6, the dynamic attenuator decreases the variance in the spine while increasing the variance to the lungs and to the heart. In many applications, peak variance is not an appropriate measure of image quality, and in these applications, optimizing

variance in selected organs of interest will provide superior images.

Figure 10 shows dose maps, integrated across the entire volume in the z direction.

4.C. Patient centering study

Table III shows the performance of the dynamic attenuator on the shifted abdomen. As expected, image quality with the bowtie filter is very dependent on patient centering. By contrast, the dynamic attenuator is much less sensitive. One reason for the remaining sensitivity to patient centering is the speed constraints on the actuators. The noise visualizations

FIG. 9. Coronal reformats of Fig. 8.

FIG. 10. Dose for each dataset, with (left column) square-root TCM, (middle column) optimized TCM, and (right column) dynamic attenuator. The bottom row (40 mm shift) is the shifted abdomen. The dose has been integrated in the z direction, and is shown normalized to constant peak variance within each dataset.

for the 4 cm offset case are shown in Figs. 8 and 9, and the dose is shown in Fig. 10.

4.D. SPR study

Table IV shows the average SPR (before any antiscatter grid) of a system with the dynamic attenuator and a system using a standard bowtie filter on three of our datasets. The average SPR is the average of the SPRs calculated in each of 64 views distributed uniformly throughout the scan. For the AAA dataset, we discarded the final 19% of the views where the region of interest was no longer being scanned. A modest

TABLE III. Relative dose required to achieve the same peak variance as the abdomen dataset is shifted off-center in 2 cm intervals. The optimized TCM and dynamic attenuator systems achieve a large dose benefit even with no shift, but this gap increases as the abdomen is shifted further off center.

Abdomen shift amount (cm)	Square root TCM (%)	Optimized TCM (%)	Optimized dynamic attenuator (%)
0	100	79	53
2	107	78	54
4	167	108	58
6	280	144	62

TABLE IV. SPR comparison between a system using a standard bowtie and a system with a dynamic attenuator.

Dataset	Average SPR for standard bowtie (%)	Average SPR for dynamic attenuator	SPR reduction
Abdomen	52	41	21
Thorax	35	31	12
AAA (Targeted region only)	125	43	66

SPR reduction is seen in the abdomen and the thorax, and a more significant reduction in SPR is seen in the AAA. The more dramatic improvement in the AAA can be traced to the targeted nature of the scan, in which the scatter was reduced by decreasing the flux to the outer parts of the patient.

The AAA study has a much higher average SPR with the standard bowtie because the patient is larger and, therefore, more attenuating (about 38 cm wide instead of 29 cm wide).

4.E. Resolution and kVp sensitivity analyses

We found little difference between the original resolution and improved resolution, with the values changing by only 1% or 2%. For example, the dose reduction with optimized dynamic attenuator calculated with improved resolution was 69% compared to the standard bowtie without TCM, instead of 68% with the original resolution.

On the abdomen dataset, we found that there was minimal difference in dose reduction between using a 120 kVp spectrum and an 80 kVp spectrum. At 120 kVp, the dynamic attenuator showed a 33% dose reduction compared to optimized TCM, while at 80 kVp, the dynamic attenuator showed a 34% dose reduction.

5. DISCUSSION

In agreement with previous, preliminary studies,^{7,8} the proposed, piecewise-linear dynamic attenuator is able to achieve significant dose reductions without increasing peak variance. Dose reductions of about 30% are observed relative to optimized TCM for the routine scans, and dose reductions of 50% are observed for targeted scans. For the routine scans with a standard bowtie, optimizing the TCM delivers an additional 30% dose reduction over the square root heuristic, which itself may be superior to implemented tube current modulation schemes which are constrained by generator limitations. When the dynamic attenuator with optimized control is compared against the standard bowtie filter with heuristics for controlling TCM, the total dose reduction is about 50% even in full-field-of-view scanning.

Peak variance may not always be the best image quality metric. In the thorax, the dynamic attenuator decreased variance near the spine, which is important for some applications but may be unimportant if the features of interest are localized near the mediastinum or the lungs. However, absent more information, we believe that peak variance is a reasonable quantity to minimize. Optimizing peak variance enforces a

minimum image quality everywhere. Using the mean variance metric, in contrast, could compromise a clinically important region of the image in order to produce image quality that is higher than necessary elsewhere. When the clinically important region is known before the scan, the dynamic attenuator could be used to perform semitargeted scans, with additional image quality and reduced noise in organs of interest. In general, the image quality metric may be task dependent.

The process of optimizing the control for the dynamic attenuator through the use of a prescan is limiting. A simple heuristic control algorithm has been suggested for the dynamic attenuator which eliminates the prescan, but we emphasize that more sophisticated heuristic control methods are possible and will be the subject of future work. The performance of more sophisticated heuristic control methods can be bounded above by the optimized control of the attenuator and bounded below by our simple square-root heuristic. These heuristics would benefit from input from the operator to specify regions of interest or radiosensitive organs. This would add a step to the workflow, but on the other hand, the relative insensitivity of the dynamic attenuator to exact patient centering could eliminate a step in the clinical workflow by increasing the acceptable margin of error in patient positioning.

The SPR at the center of the detector is modestly reduced by the dynamic attenuator. By reducing the flux incident on the outer portions of the patient, which are less attenuating, the overall level of scatter is reduced. This decreases the SPR in the center but may increase the SPR somewhat for those rays far from the center of the patient, although the primary signal in those rays is larger. The dynamic attenuator itself is a source of scatter, but the net impact of the dynamic attenuator for SPR still appears to be positive. A larger SPR reduction should be possible if the dynamic attenuator is controlled to minimize SPR and not dose. The piecewise-constant dynamic attenuator has been demonstrated experimentally to show a fourfold reduction in SPR compared to flat field imaging (without a bowtie filter).¹¹

We assumed that a linear reconstruction algorithm was used. Recently, iterative reconstruction algorithms have become popular.²⁵ We expect that the dynamic attenuator would still provide dose reduction when used with iterative reconstruction, for the same reason that TCM is still beneficial with iterative reconstruction. However, with a dynamic attenuator, the benefit of the iterative reconstruction may be diminished because iterative reconstruction improves image quality by effectively decreasing the weight of noisy rays. If the noise distribution of the rays becomes more homogeneous, as is the case with the dynamic attenuator, iterative reconstruction may provide less added value. However, estimates of dose reduction possible with iterative reconstruction remain difficult because traditional metrics such as standard deviation are not meaningful for these nonlinear reconstruction algorithms, and quantitative judgments on dose reduction ultimately depend on clinical observer studies.

The dose reductions observed here with a piecewise-linear dynamic attenuator are highly encouraging. Compared to other dose reduction strategies being pursued, such as iterative reconstruction or photon-counting detectors,^{25,26} the

dynamic attenuator is relatively simple in concept and inexpensive while still offering compelling dose reduction. The dynamic attenuator could be used synergistically with the photon-counting detector by reducing the maximum transmitted flux and thereby mitigating pulse pile-up effects. The primary challenges in the adoption of the dynamic attenuator may be mechanical in nature, because the clinical scanner will impose demanding requirements on the speed and precision of the device. However, the piecewise-linear nature of the attenuator may be smooth enough so that artifacts resulting from imperfect precision remain acceptable. A prototype system is currently under construction to test this hypothesis.

ACKNOWLEDGMENT

The authors acknowledge the Coulter Foundation and the Lucas Foundation for funding, and Lior Molvin and Mark Peng for their assistance.

^{a)} Author to whom correspondence should be addressed. Electronic mail: sshsieh@stanford.edu

¹ M. Gies, W. A. Kalender, H. Wolf, C. Suess, and M. T. Madsen, "Dose reduction in CT by anatomically adapted tube current modulation. I. Simulation studies," *Med. Phys.* **26**, 2235–2247 (1999).

² W. A. Kalender, H. Wolf, and C. Suess, "Dose reduction in CT by anatomically adapted tube current modulation. II. Phantom measurements," *Med. Phys.* **26**, 2248–2253 (1999).

³ J. Hsieh, *Computed Tomography: Principles, Design, Artifacts, and Recent Advances* (Society of Photo Optical, 2003).

⁴ T. G. Schmidt, R. Fahrig, N. J. Pelc, and E. G. Solomon, "An inverse-geometry volumetric CT system with a large-area scanned source: A feasibility study," *Med. Phys.* **31**, 2623–2627 (2004).

⁵ T. G. Schmidt, J. Star-Lack, N. R. Bennett, S. R. Mazin, E. G. Solomon, R. Fahrig, and N. J. Pelc, "A prototype table-top inverse-geometry volumetric CT system," *Med. Phys.* **33**, 1867–1878 (2006).

⁶ S. R. Mazin, J. Star-Lack, N. R. Bennett, and N. J. Pelc, "Inverse-geometry volumetric CT system with multiple detector arrays for wide field-of-view imaging," *Med. Phys.* **34**, 2133–2142 (2007).

⁷ S. Bartolac, S. Graham, J. Siewerdsen, and D. Jaffray, "Fluence field optimization for noise and dose objectives in CT," *Med. Phys.* **38**, S2–S17 (2011).

⁸ J. Sperl, D. Beque, B. Claus, B. De Man, B. Senzig, and M. Brokate, "Computer-assisted scan protocol and reconstruction (CASPAR)—Reduction of image noise and patient dose," *IEEE Trans. Med. Imaging* **29**(3), 724–732 (2010).

⁹ S. S. Hsieh and N. J. Pelc, "The feasibility of a piecewise-linear dynamic bowtie filter," *Med. Phys.* **40**(3), 031910 (12pp.) (2013).

¹⁰ T. P. Szczykutowicz and C. A. Mistretta, "Design of a digital beam attenuation system for computed tomography: Part I. System design and simulation framework," *Med. Phys.* **40**, 021905 (12pp.) (2013).

¹¹ T. P. Szczykutowicz and C. A. Mistretta, "Design of a digital beam attenuation system for computed tomography: Part II. Performance study and initial results," *Med. Phys.* **40**, 021906 (9pp.) (2013).

¹² T. Szczykutowicz and C. Mistretta, "Practical considerations for intensity modulated CT," *Proc. SPIE* **8313**, 83134E-1–83134E-11 (2012).

¹³ X. Tang, J. Hsieh, R. A. Nilsen, S. Dutta, D. Samsonov, and A. Hagiwara, "A three-dimensional-weighted cone beam filtered backprojection (CB-FBP) algorithm for image reconstruction in volumetric CT—Helical scanning," *Phys. Med. Biol.* **51**, 855–874 (2006).

¹⁴ D. A. Chesler, S. J. Riederer, and N. J. Pelc, "Noise due to photon counting statistics in computed x-ray tomography," *J. Comput. Assist. Tomogr.* **1**(1), 64–74 (1977).

¹⁵ A. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging* (SIAM, Philadelphia, 1988).

- ¹⁶S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, and G. Barrand, "Geant4-a simulation toolkit," *Nucl. Instrum. Methods Phys. Res.: Sect. A* **506**(3), 250–303 (2003).
- ¹⁷P. M. Joseph and R. D. Spital, "A method for correcting bone induced artifacts in computed tomography scanners," *J. Comput. Assist. Tomogr.* **2**(1), 100–108 (1978).
- ¹⁸M. Grant and S. Boyd, CVX: Matlab software for disciplined convex programming, version 1.21 (2011).
- ¹⁹M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, edited by V. Blondel, S. Boyd, and H. Kimura (Springer-Verlag Limited, 2008), pp. 95–110.
- ²⁰CASIMAGE radiology teaching files database available from <http://pubimage.hcuge.ch>.
- ²¹A. Berrington de Gonzalez, M. Mahesh, K. P. Kim, M. Bhargavan, R. Lewis, F. Mettler, and C. Land, "Projected cancer risks from computed tomographic scans performed in the United States in 2007," *Arch. Intern. Med.* **169**(22), 2071–2077 (2009).
- ²²J. R. Mayo, K. P. Whittall, A. N. Leung, T. E. Hartman, C. S. Park, S. L. Primack, G. K. Chambers, M. K. Limkeman, T. L. Toth, and S. H. Fox, "Simulated dose reduction in conventional chest CT: Validation study," *Radiology* **202**(2), 453–457 (1997).
- ²³J. Li, U. K. Udayasankar, T. L. Toth, J. Seamans, W. C. Small, and M. K. Kalra, "Automatic patient centering for MDCT: Effect on radiation dose," *Am. J. Roentgenol.* **188**(2), 547–552 (2007).
- ²⁴T. Toth, Z. Ge, and M. P. Daly, "The influence of patient centering on CT dose and image noise," *Med. Phys.* **34**, 3093–3101 (2007).
- ²⁵J. B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multislice helical CT," *Med. Phys.* **34**, 4526–4544 (2007).
- ²⁶K. Taguchi, S. Srivastava, H. Kudo, and W. C. Barber, in *Nuclear Science Symposium Conference Record on Enabling Photon Counting Clinical X-ray CT* (IEEE, 2009), pp. 3581–3585 (2009).