

Accelerating Monte Carlo simulations of radiation therapy dose distributions using wavelet threshold de-noising

Joseph O. Deasy^{a)}

Department of Radiation Oncology, Mallinckrodt Institute of Radiology, Washington University School of Medicine, St. Louis, Missouri 63110

M. Victor Wickerhauser

Department of Mathematics, Washington University, St. Louis, Missouri 63110

Mathieu Picard

Ecole Polytechnique, Palaiseau, France

(Received 30 August 2001; accepted for publication 15 July 2002; published 30 September 2002)

The Monte Carlo dose calculation method works by simulating individual energetic photons or electrons as they traverse a digital representation of the patient anatomy. However, Monte Carlo results fluctuate until a large number of particles are simulated. We propose wavelet threshold de-noising as a postprocessing step to accelerate convergence of Monte Carlo dose calculations. A sampled rough function (such as Monte Carlo noise) gives wavelet transform coefficients which are more nearly equal in amplitude than those of a sampled smooth function. Wavelet hard-threshold de-noising sets to zero those wavelet coefficients which fall below a threshold; the image is then reconstructed. We implemented the computationally efficient 9,7-biorthogonal filters in the C language. Transform results were averaged over transform origin selections to reduce artifacts. A method for selecting best threshold values is described. The algorithm requires about 336 floating point arithmetic operations per dose grid point. We applied wavelet threshold de-noising to two two-dimensional dose distributions: a dose distribution generated by 10 MeV electrons incident on a water phantom with a step-heterogeneity, and a slice from a lung heterogeneity phantom. Dose distributions were simulated using the Integrated Tiger Series Monte Carlo code. We studied threshold selection, resulting dose image smoothness, and resulting dose image accuracy as a function of the number of source particles. For both phantoms, with a suitable value of the threshold parameter, voxel-to-voxel noise was suppressed with little introduction of bias. The roughness of wavelet de-noised dose distributions (according to a Laplacian metric) was nearly independent of the number of source electrons, though the accuracy of the de-noised dose image improved with increasing numbers of source electrons. We conclude that wavelet shrinkage de-noising is a promising method for effectively accelerating Monte Carlo dose calculations by factors of 2 or more. © 2002 American Association of Physicists in Medicine. [DOI: 10.1118/1.1508112]

Key words: Monte Carlo, wavelet threshold de-noising

I. INTRODUCTION

Radiation therapy utilizes directed beams of radiation to sterilize tumor cells via DNA damage. Accurate pretreatment patient-specific dose calculations are essential tools for optimizing radiotherapy field shapes and intensities to maximize the probability of sterilizing the tumor while minimizing normal tissue damage. Monte Carlo (MC) dose calculation methods, which can accurately model radiation fields in the presence of complex anatomy, treatment aids, and radiation source geometries, are clearly the most general and accurate class of dose calculation engines developed to date.¹⁻⁵ MC methods for photon beams work by tracing the paths of source photons, scattered photons, and secondary electrons through computed-tomography-derived patient anatomy representations according to known scattering and energy deposition probabilities. The energy deposition events are binned into cuboid voxels, the energy deposited per unit mass being taken as the estimate of the dose which would be deposited under actual irradiation conditions.

Several manufacturers of radiation therapy treatment planning systems have recently announced plans to develop MC-based dose calculation algorithms. However, MC algorithms are typically much slower than the currently available but less-accurate algorithms. MC run-times are determined by the need to produce smooth averages of the energy deposition events for treatment plan evaluation. In this article, we show how MC calculations can be effectively accelerated using a wavelet threshold de-noising-based algorithm. This is a general technique which can potentially be applied to any MC dose calculation algorithm.

De-noising of MC electron beam dose distributions has previously been discussed using digital filtering techniques.⁶ The basic rationale for de-noising is that radiation transport is essentially a diffusive process which can be said to operate on sharp images (field shapes), thereby producing dose images [two-dimensional (2D) dose cross sections from full three-dimensional (3D) distributions] which are smoother. Therefore a postcalculation method of removing statistical “noise” from dose images, but without introducing unac-




FIG. 1. The digital electron beam step-heterogeneity phantom used as input to the ITS MC calculations. Source electrons of energy 10 MeV are transported through air to the water surface and form a 10-cm-diam radiation field. Dose is scored in a plane transverse to the beam axis, in voxels of cross section $0.2\text{ cm} \times 0.2\text{ cm}$ and thickness 0.1 mm , at a depth of 2.55 cm .

ceptably large bias (defined as deviation from the true dose), is feasible and desirable. The potential of producing clinically usable dose distributions, i.e., with low enough noise, with much shorter calculation times, makes MC de-noising attractive for clinical implementation.

We apply wavelet threshold de-noising techniques to the MC de-noising problem. Unlike traditional digital filtering techniques, wavelets adapt to the sharpness of local features, and are potentially less likely to introduce unacceptable distortions of the true underlying dose distribution.

II. METHODS

A. Monte Carlo simulations

We used Monte Carlo simulations produced with the ITS (Integrated Tiger System) package.⁷ Two-dimensional slices of dose are extracted from the resulting 3D dose distributions for de-noising.

Two test geometries are used here, one for electron beam tests and one for photon beam tests. A “step-heterogeneity” phantom, with a 1-cm-thick slab of water on the surface of a water medium (shown in Fig. 1) was simulated with 10 MeV electrons started from a source 100 cm in air from the surface of the phantom. The source particles are emitted uniformly over a truncated cone resulting in a circular field of diameter 10 cm at the phantom surface. Dose was scored in voxels which are 0.2 cm wide in the scoring plan and 0.1 cm thick in the depth direction, at a depth of 2.55 cm .

To test the use of de-noising with photon beams, a challenging lung-like test phantom⁸ is used which consists of: slabs of water (depth $0\text{--}3\text{ cm}$), aluminum ($3\text{--}5\text{ cm}$), homogeneous lung-like media ($5\text{--}12\text{ cm}$), and water ($12\text{--}32\text{ cm}$). The 6 MV spectrum specified for this test⁸ was used by the ITS code. The beam was collimated to $5\text{ cm} \times 5\text{ cm}$ by an upstream lead collimator. Photons, scattered photons, and secondary electrons were transported through vacuum to the phantom surface. Secondary electrons were transported until kinetic energy dropped below 100 keV , whereupon energy was deposited on the spot. Dose was scored in cubic voxels

0.2 cm on each side. The dose distribution was padded with zeros before wavelet de-noising to avoid edge distortions.

B. Wavelet de-noising

To reduce Monte Carlo noise fluctuations, we use a modified *wavelet threshold* de-noising algorithm⁹ with the so-called 9,7-biorthogonal filters. We suppose that the simulated dose distribution d consists of two parts, $d = s + n$, where s is the smooth function we would obtain in the limit by running the simulation forever, and n is the rough function giving the error in our short simulation. The data d represent the dose sums, giving an array of values that is then linearly transformed into its discrete wavelet coefficients, which we call $W(d) = W(s + n) = W(s) + W(n)$.

A key property of the wavelet transform is that a sampled rough function n gives values $W(n)$ which are more nearly equal in amplitude than those of a sampled smooth function s . If n is relatively small compared to s , then any sufficiently small wavelet coefficient is more likely to be part of n than s . By picking a positive threshold ϵ and setting $W(d)_i = 0$ if $|W(d)_i| < \epsilon$, we preferentially attenuate the n , or noise component of d . The array reconstructed from the surviving coefficients is therefore a closer approximation to s than was $d = s + n$. For the key property to hold, it is necessary that the wavelet transform use a wavelet that is about as smooth as s , and thus smoother than n .

Various wavelet bases can be used. We present results here using the so-called 9,7-biorthogonal filters. The biorthogonal family has attractive computational properties (discussed in the following) which makes it very efficient to implement. The 9,7-biorthogonal filter was chosen as the basis of the JPEG2000 image compression standard for continuous value still images.¹⁰ As one might expect, compression and de-noising performance are intimately related, as both depend on separating essential image features from small rough details. “Spin-cycling,” which is not in the JPEG2000 standard, refers to choosing nine different nearest-neighbor points in the image as computational centers for the wavelet transform, and averaging the results (this is not the same as averaging dose distributions). Preliminary tests (not shown) indicated that the use of spin-cycling (used for all results presented in this paper) reduces small artifacts introduced by the wavelet transform.

C. Approximations of dose distributions

Monte Carlo radiotherapy simulations produce sampled dose distributions on a regular grid. A simulated two-dimensional slice of a three-dimensional dose distribution on an $M \times N$ grid is given by a non-negative function $f = f(m, n)$, where $0 \leq m < M$ and $0 \leq n < N$ are integer indices. Such functions may be approximated by distributions $\phi(x, y)$ defined on a continuum $0 \leq x, y < 1$ as a superposition of basic distributions $\phi_{mn} = \phi_{mn}(x, y)$ concentrated near $(m/M, n/N)$, given by

$$\phi(x, y) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) \phi_{mn}(x, y).$$




FIG. 2. 9,7-biorthogonal wavelets. Examples of 9,7-biorthogonal symmetric wavelets on 120 gridpoints. The wavelets differ in their central location and width.

The represented ϕ will be as smooth as its basic densities, regardless of f , so for this work we chose 9,7-biorthogonal basic densities, with two continuous derivatives, that are “almost” confined to the unit square.

D. Efficient coding of fluctuations as wavelets

Regions of nearly constant dosage might span many grid points, with relatively small fluctuations on various scales. The amplitude and scale of each fluctuation can be used to decide what is noise and what is signal.

We modeled fluctuations by almost independent wave functions of different sizes and positions. By approximating the original dose distribution as a superposition of biorthogonal filters, we can use the fast *discrete wavelet transform algorithm*,¹¹ with 9,7-biorthogonal symmetric wavelets, for computing the scale and amplitude of fluctuations. A representative few of these are graphed on the same axis in Fig. 2.

For computational efficiency, we used the *lifting scheme*.¹² To reduce the bias arising from the arbitrary choice of grid origin, we averaged the results of de-noising the original image and the eight images shifted by one grid point: left, right, up, down, and four diagonal shifts. This is referred to here as “spin-cycling.”

E. Discrete wavelet transforms

For two-dimensional and three-dimensional arrays, the transform is applied separately along each dimension, so only the one-dimensional transform will be described.¹¹ Let $x = x(k)$, $0 \leq k < K$ be the array of K values to be transformed. We generate a new array $F = F(k)$, $0 \leq k < K$, by the following rule:

$$F(k) = \begin{cases} \sum_{n=-4}^4 h(n)X(k-n), & \text{if } k \text{ is even} \\ \sum_{n=-4}^2 g(n)X(k-1-n), & \text{if } k \text{ is odd} \end{cases} \quad (1)$$

TABLE I. 9,7-biorthogonal analysis filters h, g and their inverse, or synthesis filters h', g' .

k	$h(k) = (-1)^k g'(k)$	k	$g(k) = -(-1)^k h'(k)$
-4, 4	0.037 828 455 506 99	-4, 2	0.064 538 882 628 938
-3, 3	-0.023 849 465 019 38	-3, 1	-0.040 689 417 609 558
-2, 2	-0.110 624 404 418 42	-2, 0	-0.418 092 273 222 212
-1, 1	0.377 402 855 612 65	-1	0.788 485 616 405 664
0	0.852 698 679 009 40		

Here $h = h(n)$ and $g = g(n)$ are the filter sequences defining the 9,7-biorthogonal wavelets. Table I gives their values. The index ranges $[-4, 4]$ and $[-4, 2]$ have 9 and 7 indices, respectively. Before summing, we extend the original signal x to a longer signal X by *whole-sample symmetric reflection*, as defined by

$$X(k) = \begin{cases} x(k) & \text{if } 0 \leq k < K \\ x(2K-2-k) & \text{if } K \leq k < 2K-1. \end{cases}$$

X is then *extended periodically* by defining $X(k \pm [2K-2]) = X(k)$ for all integer indices k . The result is $(2K-2)$ -periodic and symmetric with respect to reflection about indices 0 and $K-1$, with $X(k) = x(k)$ for the original indices $k = 0, 1, \dots, K-1$. The output sequence $F(k)$ is also $(2K-2)$ -periodic and defined at all integers k , but because the filters are symmetric, a complete set of output values may be found at just the indices $0, 1, \dots, K-1$. This allows computation of F in about $8K$ operations.

As described in Ref. 12, the lifting scheme reorganizes the computation of F to save operations. With x extended to X as before, F is computed sequentially in place as follows (the symbol “ \leftarrow ” means “takes the value of”):

1. $X(k) \leftarrow X(k) + \alpha[X(k-1) + X(k+1)]$, for all odd k in the range $0 < k < K$,
2. $X(k) \leftarrow X(k) + \beta[X(k-1) + X(k+1)]$, for all even k in the range $0 \leq k < K$,
3. $X(k) \leftarrow X(k) + \gamma[X(k-1) + X(k+1)]$, for all odd k in the range $0 < k < K$,
4. $X(k) \leftarrow X(k) + \delta[X(k-1) + X(k+1)]$, for all even k in the range $0 \leq k < K$,
5. $F(k) \leftarrow X(k)/\zeta$, for all odd k in the range $0 < k < K$,
6. $F(k) \leftarrow \zeta X(k)$, for all even k in the range $0 \leq k < K$.

The coefficients $\{\alpha, \beta, \gamma, \delta, \zeta\}$ are given in Table II. The resulting sequence F is the same as the one defined by Eq. (1), but it takes only $5K$ operations to compute, less intermediate storage, and fewer data exchanges.

TABLE II. Lifting coefficients for 9,7-biorthogonal analysis filters.

Coefficient	Value
α	-1.586 134 342 059 92
β	-0.052 980 118 572 96
γ	0.882 911 075 530 93
δ	0.443 506 852 043 97
ζ	1.149 604 398 860 24

To invert F , we use $\{-\alpha, -\beta, -\gamma, -\delta, 1/\zeta\}$ and apply the operations in reverse order:

1. $Y(k) \leftarrow Y(k)/\zeta$, for all even k in the range $0 \leq k < K$,
2. $Y(k) \leftarrow \zeta Y(k)$, for all odd k in the range $0 < k < K$,
3. $Y(k) \leftarrow Y(k) - \delta[Y(k-1) + Y(k+1)]$, for all even k in the range $0 \leq k < K$,
4. $Y(k) \leftarrow Y(k) - \gamma[Y(k-1) + Y(k+1)]$, for all odd k in the range $0 < k < K$,
5. $F^{-1}(k) \leftarrow Y(k) - \beta[Y(k-1) + Y(k+1)]$, for all even k in the range $0 \leq k < K$,
6. $F^{-1}(k) \leftarrow Y(k) - \alpha[Y(k-1) + Y(k+1)]$, for all odd k in the range $0 < k < K$.

The analog of Eq. (1) for F^{-1} is

$$F^{-1}(n) = \sum_{k \in K_0(n)} h'(2k-n)Y(2k) + \sum_{k \in K_1(n)} g'(2k-n)Y(2k+1), \quad (2)$$

where Y is the same extension of y as X is of x . The inverse filters h' , g' are given in Table I. The index ranges are $K_0(n) = \{k: -4 \leq 2k-n \leq 2\}$ and $K_1(n) = \{k: -4 \leq 2k-n \leq 4\}$, for $n=0,1,\dots,K-1$.

Each even-indexed element $F(2k)$ is a weighted average of x over nine grid points near $2k$, with the weights being the filter coefficients.

To compute a one-dimensional wavelet transform on a signal x , we apply F , retain the odd-indexed values, extract the even-indexed values to another array of about half the length, call that new array x , and repeat until x contains just a single number:

1. Let $j=1$.
2. Compute $y(k) = F(k)$, $k=0,1,\dots,K-1$.
3. Extract $D_j(k) = y(2k+1)$ for all indices k with $0 \leq 2k+1 < K$.
4. Replace $x(k) = y(2k)$ for all indices k with $0 \leq 2k < K$.
5. Replace $K \leftarrow K/2$ if K is even, or $K \leftarrow (K+1)/2$ if K is odd. This the number of elements written to x at step 4.
6. If $K > 1$, then increment $j \leftarrow j+1$ and go to step 2. Otherwise, stop.

At the termination of the above-given algorithm, the discrete wavelet transform of the original signal is stored in the arrays D_1, D_2, \dots, D_J , plus $y(0)$. The values in D_j are the amplitudes of fluctuations of the signal at scales of 2^j grid points; $y(0)$ contains the average value of the signal. J is the *depth* of the wavelet transform, the maximum value of j in the above-given algorithm, and is the least integer such that $2^J \geq K$.

As mentioned, 9,7-biorthogonal wavelets can be applied in a fully 3D manner. We are only reporting here, however, on the application to 3D dose distributions of a slice-by-slice 2D approach.

F. Approximation and de-noising

We define the *thresholding function* $T = T_\epsilon[x]$ for a function x on M by letting $a = \max\{|x(m)|:m \in M\}$ and then setting

$$T_\epsilon[x](m) = \begin{cases} 0 & \text{if } |x(m)| < \epsilon a \\ x(m) & \text{otherwise.} \end{cases}$$

Let ϵ be a fixed positive parameter, let $0 \in M$ be the origin in the grid M , and let $r \in M$ be some fixed grid point. Designate the wavelet transform of f by $W(f) = \{W(f)_m:m \in M\}$, with $W(f)_0$ being the single average coefficient, previously called $y(0)$. We perform the following algorithm:

1. Translate the function f to f_r , so that $f_r(0) = f(r)$.
2. Compute the wavelet transform $y_r \leftarrow W(f_r)$.
3. Apply ϵ -thresholding $z_r(m) \leftarrow T(y_r(m))$ for all $m \neq 0$.
4. Compute the inverse wavelet transform $x_r \leftarrow W^{-1}(z_r)$.
5. Translate the function x_r to x , so that $x(r) = x_r(0)$.

We now let O be the origin of M and its nearest neighbors, and then average together all the x 's obtained for different choices of $r \in O$ ("spin-cycling"). In the one-dimensional case, $O = \{-1, 0, +1\}$, while in the two-dimensional case O is the nine-point set produced by permuting such offsets in both directions, resulting in eight nearest neighbors and the center point. Denote by $E(f, \epsilon)$ the average of the x 's produced by all the shifts of the original f .

Suppose that $d = s + n$ is an approximation to the true dose distribution s . We may assume that s is known because we have run a Monte Carlo simulation long enough to approximate it adequately, referred to as the "target." We measure rms error as the square root of the mean difference between the dose values and the target (low noise) dose values, including only those voxels where the target dose value is greater than 50% of the maximum target dose. Variance error is defined as the square of the rms error. Dose values are variance stabilized (made closer to normally distributed) by taking the square root before wavelet de-noising and squaring after de-noising, as explained in Ref. 6. This gives an estimator for s from $E(\sqrt{d}, \epsilon)^2$. The threshold ϵ that minimizes the rms error $\|E(\sqrt{d}, \epsilon)^2 - s\|$ for each of the short segments d that comprise the long run is found by exhaustive search, as further discussed in the following. Exhaustive search is necessary due to multiple local minima. The average of these thresholds is an estimator for the best threshold to use for all MC simulations with similar noise and image characteristics.

Estimating noise reduction is not straightforward, because de-noising introduces nonrandom errors (bias). The standard deviation of voxels between batches is a misleading measurement of noise due to variations in bias. To quantify noise reduction, we use the local second derivative (Laplacian) of an image g , computed using the five-point formula:

$$\Delta(x, y) = g(x-1, y) + g(x+1, y) + g(x, y-1) + g(x, y+1) - 4g(x, y).$$

We define "roughness" as the square root of the median of the squared Laplacian values for all dose voxels above the

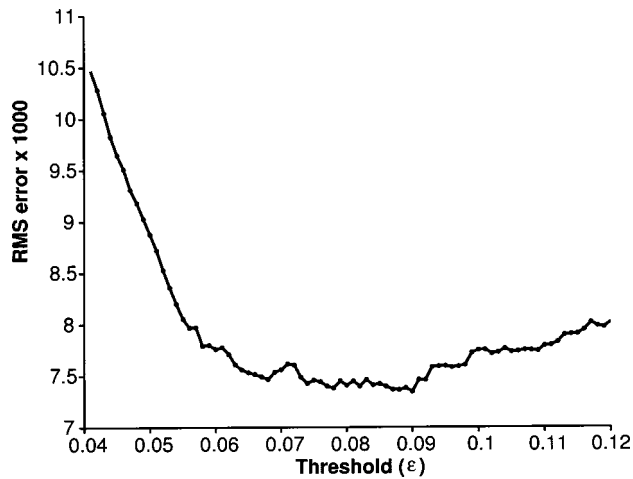


Fig. 3. Error vs threshold. The change in the rms error between a de-noised image of 4 mil source electrons and the target image as the threshold (ϵ) is changed.

50% maximum dose cutoff. The median operator makes our roughness metric nearly independent of the large values of the Laplacian which occur near edges or other regions of high second-derivative. Roughness is useful because it measures *local* roughness of the image, which is important for dose visualization. This differs from the rms error metric, which includes the effect of voxel-to-voxel fluctuations and bias introduced by the de-noising procedure. Roughness is not a dose accuracy measure, but is introduced primarily to demonstrate the effect of wavelet de-noising, as further discussed in the following.

G. Computational efficiency

The total number of arithmetic operations required to apply this de-noising algorithm in 2-D, including spin-cycling, is about 336 operations (multiplications or additions) per dose grid point. For a typical three-dimensional dose distribution of size less than 500 000 voxels, the total number of operations required for de-noising (without counting input/output operations) will be less than 200 mil arithmetic operations. Given that current desktop computers are capable of several hundred million arithmetic operations per second, de-noising computation time for this algorithm is expected to be negligible in a clinical setting.

H. Results

Figure 3 shows the improvement in the rms difference between a de-noised image of 4 mil source electrons and the target image as the threshold (ϵ) is changed, showing multiple local minima. In Fig. 4, we plot best thresholds ϵ and the rms error compared to a target with 192 mil source electrons. With more source electrons, the MC simulation becomes less noisy, weaker features in the signal become statistically significant, and the best threshold decreases.

The result of de-noising of the electron beam dose distribution with 4 mil source electrons is shown in Fig. 5. A color version is available via the EPAPS link to this document.¹³




Fig. 4. Threshold determination. Best thresholds (ϵ values) according to the minimum rms error criteria, determined by exhaustive search. The 192 mil source electron image is used as the reference. Eight simulations were run for increasing numbers of source electrons. The solid line connects the mean values of the best thresholds.

The standard deviation of un-de-noised dose voxels greater than 50% of the maximum dose voxel was 2.7%. The threshold was 0.0839. The low noise target result used 192 mil source electrons. Figure 6 shows the mean interbatch voxel standard deviation of dose for voxels above 50% maximum dose, determined from eight de-noised runs, and the ratio of mean interbatch dose variance for the raw and de-noised MC results. When the standard error of the de-noised MC is 1%, de-noising reduces variance by a factor of 8.5–9.5. Because variance is reduced in a non-de-noised MC run inversely proportional to the run time, it would require 8.5–9.5 as many source particles without de-noising to achieve the same interbatch consistency. However, this is not the same as




Fig. 5. Electron beam de-noising. De-noising of the electron beam dose distribution with 4 mil source electrons: (a) the low-noise 192 mil source electron result, (b) the high-noise 4 mil source electron result, (c) the de-noised high-noise result with a best threshold of 0.0839, and (d) the difference of the de-noised and low-noise dose distributions. The dose distributions were scaled such that the maximum low noise dose value is one.




FIG. 6. Dose calculation noise reduction using de-noising. The solid curve (left axis) gives the mean standard deviation of dose for voxels above 50% maximum dose, determined from eight runs. The dashed curve (right axis) is the ratio of dose variance for the raw and de-noised MC results. Again, the mean variance is for the same voxels and the same eight runs. When the standard error is 1% (perhaps the magnitude needed for clinical review) then the ratio of the variances is 8.5–9.5. This represents a computational acceleration of the same magnitude to reduce noise to a clinically acceptable level.

an improvement in dose accuracy. The ratio of variance error for the de-noised and un-de-noised slice dose distributions, for 4 mil source electrons, is 4.96.

Figure 7 plots three different performance metrics for MC de-noising applied to the electron beam phantom: (a) mean interbatch variance between batches of dose values (open circles, plotted as the ratio of values for the raw and de-noised results), (b) the mean squared error with respect to the “true” (low-noise) image (closed circles, also plotted as a ratio), and (c) the maximum dose deviation from the true

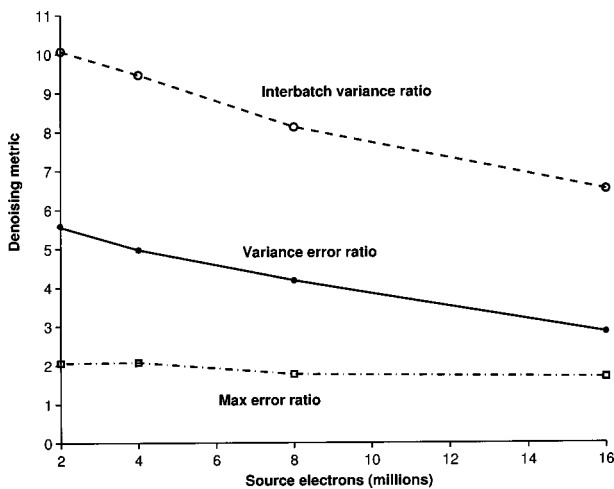


FIG. 7. De-noising performance, as measured by three different metrics: mean variance of dose values between batches (open circles); mean squared difference with the “true” (low-noise) image (closed circles); maximum deviation from the true image (open squares). All points are ratios: (raw image)/(de-noised image). All results are averages over all voxels greater than 50% maximum, for eight MC runs, with 2, 4, 8, or 16 mil source electrons.

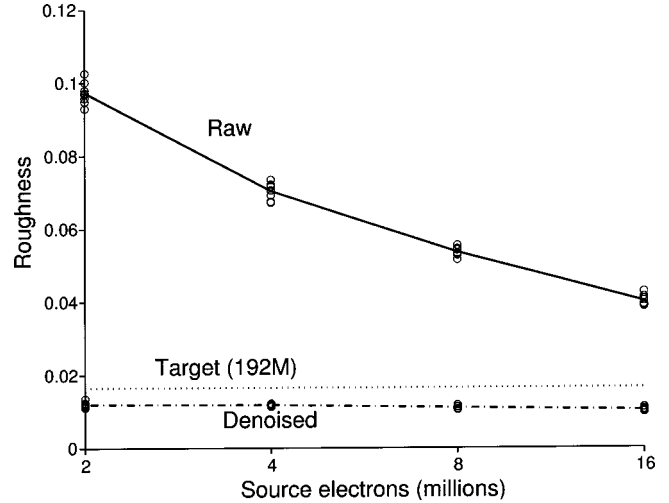


FIG. 8. Roughness of de-noised electron beam images. Roughness is plotted for raw images (solid line) and de-noised images (dashed line). The value for the 192 million source electron target image is shown as a dotted line. According to this metric, even the de-noised 2 mil source electron image is less rough than the target image. However, de-noised dosimetric accuracy decreases with decreasing source electrons.

image (open squares, as a ratio). All results are averages over all voxels greater than 50% maximum dose, for eight MC runs, with 2, 4, 8, or 16 mil source electrons.

The roughness metric for the electron beam phantom is plotted in Fig. 8. Roughness is defined primarily to understand the effect of wavelet de-noising. We note that the roughness of the de-noised images (using decreasing thresholds) is nearly constant, independent of the number of source electrons, and that it is similar to the target image roughness. Low-pass filtering techniques would typically result in roughness values which decrease with the number of source electrons. However, as the number of source electrons increases the de-noised image more closely approaches the true image.

De-noising performance for the lung-like photon beam phantom is also encouraging. We applied 2D de-noising in planes which contain the beam axis. The midplane dose distribution is shown in Fig. 9: (a) low noise (3.25 bil source photons), (b) high noise (250 mil source photons), (c) de-noised high noise, and (d) low noise minus de-noised. An average best threshold of 0.03275 was used. The mean standard deviation of voxels from one batch to the next was 2.4% of the dose maximum for the high noise runs, 0.66% for the low noise target, and 1.19% for the de-noised high noise run. The effective acceleration measured by reductions in rms fluctuations is therefore 4.0 for this example. The roughness metric is more representative of the differences in image smoothness: roughness was 0.048 for the high noise image, 0.018 for the low noise image, and 0.0096 for the de-noised image. The effective acceleration in terms of the variance error ratio is 1.97.

III. DISCUSSION

The Monte Carlo method for treatment planning is itself a statistical estimation technique, with random and systematic




FIG. 9. De-noising of the photon beam phantom dose for a beam-aligned midplane: (a) low-noise (3.25 bil source photons), (b) high-noise (250 mil source photons), (c) de-noised high-noise, and (d) low-noise minus de-noised. The dose distributions were scaled such that the maximum low noise dose value is one. A color version is available via the EPAPS link to this document.¹³

error components due to unknowns or errors in incident phase space specifications, computed tomography voxel values, material properties, etc. Because transport diffusion makes true dose distributions more smooth than MC voxel-to-voxel statistical fluctuations, de-noising of voxel sums of energy deposition events is superior to using the raw sums. De-noising of MC generated data, which can potentially be accomplished in a variety of ways, is a natural last step for the MC dose calculation technique, and represents a new type of variance reduction method.

As shown in Fig. 7, wavelet de-noising improves both the rms dose deviation and the maximum dose deviation from the “true” dose distribution. However, extensive dosimetric analysis of clinical examples will be required to determine the MC noise level which must be attained prior to de-noising to yield acceptable clinical accuracy with the least run time. Relatively smooth de-noised dose distributions can be obtained even with rather large initial dose standard deviations. This is reflected in the small and nearly constant roughness values plotted in Fig. 8. Wavelet de-noising of high-noise images is of course not perfectly accurate, and is most likely to disagree with target doses in regions of rapid dose variation, as shown in Figs. 5 and 9. However, compared to conventional digital filtering techniques,⁶ the fast wavelet de-noising algorithm is particularly good at preserving high-frequency features, such as the interface dose spikes in the photon beam example shown in Fig. 9.

Thresholds have been determined by direct comparison with “true” images in this report. We expect that reasonable thresholds could be set by using typical clinical dose distributions, computed to very high accuracy. Threshold values near the “best” value also perform well (see Fig. 3)

Monte Carlo de-noising is based on the relative statistical independence of voxels. However, this clearly is not strictly true, as neighboring voxels share many of the same particle trajectories. This may sometimes result in “clumping” of

statistical fluctuations of the same sign over a shared neighborhood of voxels. Clearly, the choice of de-noising planes affects noise removal in such “clumps.” More importantly, such clumps, when they do occur, are unlikely to be completely smoothed out by, say, doubling the number of Monte Carlo histories. New histories are unlikely to introduce statistical fluctuations which would have a structure which just so happens to be the same in shape but the opposite in magnitude as the original “clump.” Hence, “clumping” is an occasional hazard of MC-based treatment planning, and is not particularly a de-noising problem per se.

Wavelet de-noising, or other de-noising techniques, could potentially substantially reduce the cost and increase the efficiency of MC-based radiation therapy treatment planning. This efficiency gain could produce both lower operating costs and, perhaps more importantly, more time for potential improvements of the treatment plan or other important clinical activities.

Different metrics emphasize different aspects of de-noising. Accuracy is easy to measure in a mean-square sense (variance error), but could be significantly affected by high-frequency rolloff of voxels near steep gradients. De-noising or smoothing methods typically perform worst near edges or discontinuities, effectively introducing small geometric distortions. Wavelet de-noising appears superior to frequency-based de-noising in this respect,⁶ however, bias is still introduced preferentially near edges (see Fig. 9). Such slight geometric distortions may not be significant in treatment planning, as other geometric uncertainties may dominate.

The results show that the proposed wavelet threshold de-noising algorithm with spin-cycling decreases the computational time needed to compute smooth, low-noise MC dose distributions. We conclude that de-noising of MC-generated data based on wavelet threshold de-noising could potentially provide significant efficiency gains for clinical radiotherapy treatment planning. However, clinical validation will require much more extensive testing with a range of dose distributions based on individual patient anatomy as represented by computed tomography scans. The computationally efficient wavelet de-noising algorithm investigated here, based on the 9,7-biorthogonal filters, including a computationally efficient lifting scheme and spin-cycling, is attractive for further clinical testing, or as a basis for further algorithm improvements.

ACKNOWLEDGMENT

This research was supported by Grant No. R29 CA85181 from the National Cancer Institute. Washington University has applied for a patent on wavelet threshold de-noising applied to Monte Carlo dose distributions.

^aElectronic mail: deasy@radonc.wustl.edu

¹I. Kawrakow, M. Fippel, and K. Friedrich, “3D electron dose calculation using a Voxel based Monte Carlo algorithm (VMC),” *Med. Phys.* **23**, 445–457 (1996).

²J. S. Li, T. Pawlicki, J. Deng, S. B. Jiang, E. Mok, and C. M. Ma, “Validation of a Monte Carlo dose calculation tool for radiotherapy treatment planning,” *Phys. Med. Biol.* **45**, 2969–2985 (2000).

³H. Neuenschwander, W. Voken, D. Frei, C. Cris, E. Born, and R. Mini,

- “Macro Monte Carlo: Clinical implementation in a distributed computing environment,” paper presented at MC 2000: Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications, Lisbon, Portugal, 23–36 October 2000.
- ⁴C. M. Ma, E. Mok, A. Kapur, T. Pawlicki, D. Findley, S. Brain, K. Forster, and A. L. Boyer, “Clinical implementation of a Monte Carlo treatment planning system,” *Med. Phys.* **26**, 2133–2143 (1999).
- ⁵J. Sempau, S. J. Wilderman, and A. F. Bielajew, “DPM, a fast, accurate Monte Carlo code optimized for photon and electron radiotherapy treatment planning dose calculations,” *Phys. Med. Biol.* **45**, 2263–2291 (2000).
- ⁶J. O. Deasy, “De-noising of electron beam Monte Carlo dose distributions using digital filtering techniques,” *Phys. Med. Biol.* **45**, 1765–1779 (2000).
- ⁷J. A. Halbleib, R. P. Kensek, T. A. Mehlhorn, G. D. Valdez, S. M. Seltzer, and M. J. Berger, “Integrated TIGER series of coupled Electron/Photon Monte Carlo transport code system, version 3.0, RSICC code package CCC-467” (Distributed by the Radiation Shielding Information Center, <http://www-rsicc.ornl.gov/rsic.html>, 1994).
- ⁸D. W. O. Rogers and R. Mohan, “Questions for comparison of clinical Monte Carlo codes,” in *The Use of Computers in Radiation Therapy*, Proceedings of the 23rd International Conference on the Use of Computers in Radiation Therapy, Heidelberg, 22–25 May 2000, edited by W. Schlegel and T. Bortfeld (Springer, Berlin, 2000), pp. 120–122.
- ⁹D. Donoho, “Unconditional bases are optimal bases for data compression and for statistical estimation,” *Appl. Comp. Harmon. Anal.* **1**, 100–115 (1993).
- ¹⁰“Information technology-Digital compression and coding of continuous-tone still images: Requirements and guidelines” International Organization for Standardization (ISO) Technical Report No. ISO/IEC 10918-1, 2000.
- ¹¹M. V. Wickerhauser, *Adapted Wavelet Analysis from Theory to Software* (Peters, Natick, MA, 1994).
- ¹²I. Daubechies and W. Sweldens, “Factoring wavelet transforms into lifting steps,” *J. Fourier Anal. Appl.* **4**, 245–267 (1998).
- ¹³See EPAPS Document No. E-MPHYA6-29-010210 for color figures. A direct link to this document may be found in the online article’s HTML reference section. The document may also be reached via the EPAPS homepage (<http://www.aip.org/pubserve/epaps.html>) or from <ftp.aip.org> in the directory /epaps/. See the EPAPS homepage for more information.