A limit on dose reduction possible with CT reconstruction algorithms without prior knowledge of the scan subject

Scott S. Hsieh^{a)}

Department of Radiology, Stanford University, Stanford, California 94305

David A. Chesler

Department of Radiology, Massachusetts General Hospital, Charleston, Massachusetts 02114

Dominik Fleischmann

Department of Radiology, Stanford University, Stanford, California 94305

Norbert J. Pelc

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

(Received 12 October 2015; revised 5 January 2016; accepted for publication 2 February 2016; published 23 February 2016)

Purpose: To find an upper bound on the maximum dose reduction possible for any reconstruction algorithm, analytic or iterative, that result from the inclusion of the data statistics. The authors do not analyze noise reduction possible from prior knowledge or assumptions about the object.

Methods: The authors examined the task of estimating the density of a circular lesion in a cross section. Raw data were simulated by forward projection of existing images and numerical phantoms. To assess an upper bound on the achievable dose reduction by any algorithm, the authors assume that both the background and the shape of the lesion are completely known. Under these conditions, the best possible estimate of the density can be determined by solving a weighted least squares problem directly in the raw data domain. Any possible reconstruction algorithm that does not use prior knowledge or make assumptions about the object, including filtered backprojection (FBP) or iterative reconstruction methods with this constraint, must be no better than this least squares solution. The authors simulated 10 000 sets of noisy data and compared the variance in density from the least squares solution with those from FBP. Density was estimated from FBP images using either averaging within a ROI, or streak-adaptive averaging with better noise performance.

Results: The bound on the possible dose reduction depends on the degree to which the observer can read through the possibly streaky noise. For the described low contrast detection task with the signal shape and background known exactly, the average dose reduction possible compared to FBP with streak-adaptive averaging was 42% and it was 64% if only the ROI average is used with FBP. The exact amount of dose reduction also depends on the background anatomy, with statistically inhomogeneous backgrounds showing greater benefits.

Conclusions: The dose reductions from new, statistical reconstruction methods can be bounded. Larger dose reductions in the density estimation task studied here are only possible with the introduction of prior knowledge, which can introduce bias. © 2016 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4941954>]

Key words: dose reduction, iterative reconstruction, noise power spectrum, statistical reconstruction, low contrast detectability

1. INTRODUCTION

CT iterative reconstruction (IR) methods have regained widespread interest in recent years upon their incorporation into commercial CT platforms.^{1,2} These methods range from model-based IR algorithms that are sophisticated but slow, to fast hybrid algorithms that attain a compromise between traditional filtered backprojection (FBP) and model-based IR.³ Like FBP, model-based IR algorithms produce reconstructed images consistent with the measured data, but unlike FBP, IR algorithms are informed by the data statistics. A popular class of IR algorithms in addition uses data regularization to suppress rapid variations which appear to be noise, although alternative approaches have been explored.⁴ Very promis-

ing results have been reported in the literature. Investigators have claimed dose reductions between 50% and 80% using model-based IR algorithms,^{5–7} and even the simpler, hybrid techniques have claimed dose reductions of 40%.^{8,9} These dose reductions have been measured using an assortment of metrics, including linear systems metrics such as standard deviation or noise power spectrum, radiologist preference, and task-based metrics. Linear systems metrics are not strictly valid for nonlinear IR and may be misleading. Radiologist preference is an imperfect surrogate for diagnostic accuracy. For example, a new algorithm could be penalized simply because its noise texture is unusual when compared to existing reconstruction methods. Task-based metrics may be preferred for the evaluation of new reconstruction algorithms.¹⁰

It is often stated that by exploiting the statistical properties of the data, IR algorithms achieve far better noise performance than their analytic counterparts (e.g., FBP). It would seem reasonable to expect further dose reductions as this technology matures. For this reason, we analyzed the maximum noise reduction possible from optimal use of the measurement statistics, in contrast to noise reduction that arises from inclusion of *a priori* knowledge about the object. If *a priori* information is introduced, the amount of dose reduction depends on the amount and accuracy of the knowledge available. As an extreme, if we assumed that the object consisted of a single material with uniform density, the dose reduction could be extremely large, but the associated error would also be large if the single-material hypothesis was false. Iterative reconstruction frameworks in use today often assume smoothness in the ground truth, and this prior knowledge may lead to bias from excessive blurring in the reconstructed images if the reconstruction parameters are not correctly tuned.¹¹ By contrast, noise reduction from more efficient use of data statistics is free from these potential drawbacks. The purpose of our work was to find an upper bound on the maximum dose reduction possible for any reconstruction algorithm, analytic or iterative, which does not make *a priori* assumptions about the object, i.e., the dose reduction that comes from knowledge of the measurement statistics rather than that comes from knowledge or assumptions about the object.

2. MATERIALS AND METHODS

The dose reduction will vary depending on the clinical task. We restrict our analysis to one specific task, estimation of the average attenuation or CT number of a lesion, which we will refer to as simply being the “lesion density” for convenience. (Assuming a fixed chemical composition, the CT number of the lesion will be directly related to its density.) For low contrast lesions, the precision of lesion density estimation is closely related to lesion detection. All the methods considered here are linear, so the noise in the lesion density estimate is independent of the density value and will, together with the lesion contrast, largely determine the detectability.

The benefits of IR depend on the object being imaged. To examine this aspect, we superimpose the lesion on a range of objects, including two simulated phantoms and four anthropomorphic cross sections. Our approach is a generalization of the result reported in Ref. 13 that compared the noise variance in CT using FBP to that obtained if a pixel is measured in isolation. Precision of lesion density estimation can be viewed as a surrogate to low contrast detectability.

We compare three methods described below that estimate the density of the lesion. In all of them we assume that the location and shape of the lesion and the background are known exactly. Two of the methods compute the estimate from the FBP reconstructed image. While the estimators have the above knowledge, the reconstruction method, FBP, does not. The precision of these two estimates represent the performance of conventional reconstruction algorithms. In this special case of known background, while the density of the lesion can be estimated from reconstructed images, it can be better estimated

from the raw data directly. This raw data estimate will meet or exceed the precision of estimates derived from any reconstructed image made with no assumptions about the object, including those reconstructed from IR algorithms.

An important characteristic of all these estimators is that none of them assumes any knowledge of the density of the region or correlation between this value and the rest of the image. Therefore, the comparison between the direct raw data estimate and the estimates based on the FBP image represents an upper limit on the improvement possible over FBP for any reconstruction method that similarly does not assume or introduce correlations between the density of the lesion and the rest of the image.

2.A. Estimators

The estimate of the lesion density was obtained using three different methods: direct raw data estimation, ROI averaging, and streak-adaptive averaging. Figure 1 shows an overview of the three methods. ROI averaging and streak-adaptive averaging operate on images reconstructed using FBP, which serves as our baseline reconstruction method. Direct raw data estimation operates directly and optimally on the raw data.

In Fig. 1, all three estimation methods are described as weighted sums. A weighted sum of an image is calculated by first multiplying each pixel of the image with a weight map, and then adding the values together to produce a number which is an estimate of the lesion density. The weighted sum in the ROI estimate is simply an average within the ROI.

We will briefly describe the estimators in the remainder of this section; they are more formally derived in the Appendix.

2.A.1. Direct, raw data estimator

Because the background is known exactly, the raw data of the background can be subtracted from the measured projections, leaving the raw data corresponding to the lesion only (plus noise), shown in Fig. 1(b). Estimating the lesion density is now a classic weighted least squares problem. This least squares solution avoids image reconstruction altogether. It uses the statistical noise in the measurements and produces the unbiased estimate of the lesion density with the lowest possible variance—no other method can outperform it. Figure 1(c) is the weight map for the direct raw data estimator. The weights are high where the measurement noise is low, and vice versa. In Fig. 1(c), the weights appear to oscillate from high (bright white) to zero (black), but in fact the weights in the high-noise region are small but positive. Outside the projection of the lesion, the weights are zero, indicating that these regions are irrelevant to calculating the lesion density.

2.A.2. ROI averaging

Conceptually, the simplest method for estimating the density of the lesion is to reconstruct the data using FBP, draw a circular ROI around the lesion, and take the average value within the lesion. Our ROI averaging estimator is an improvement on this concept. First, the known background is

FIG. 1. A comparison of raw data estimation, ROI estimation, and streak-adaptive estimation of the lesion density. Starting with (a) the raw data, here of an elliptical 20 by 40 cm phantom, we extract (b) the difference raw data of the lesion alone by subtracting the background. The density of the lesion is exaggerated in (a) to aid in visibility. The (b) difference raw data are weighted by (c) the least squares solution and then summed to generate (d) the raw data estimate. The weights are negatively correlated with noise in the raw data, indicating that the raw data estimate is using the statistical properties of the measurements. To generate (g) the ROI estimate, the (b) difference raw data are reconstructed by FBP. This (e) FBP reconstruction is then weighted by (f) the ROI weights, which essentially averages the reconstruction within the boundary of the lesion. To generate (h) the streak-adaptive estimate, the (e) FBP reconstruction of the lesion is averaged with (h) the streak-adaptive (or matched filter) weights. In this example with asymmetric noise statistics, the streak-adaptive estimate significantly outperforms the ROI estimate.

subtracted, yielding an image of only the lesion, as shown in Fig. 1(e). Next, a ROI is selected. If the reconstructed lesion were a perfect circle, a perfectly circular ROI should be used. However, because of the finite spatial resolution of the system, the edge of the lesion is blurred. The best performance is obtained if a weighted average is used, with the pixels on the borderline given a reduced weight. The particular weighted average we use is in Fig. 1(f).

ROI averaging is a very intuitive method for calculating the CT number of the lesion, but it ignores noise texture (e.g., noise streaks). A human observer might account for these noise streaks, but the ROI averaging method does not. To improve the performance while still relying on the FBP reconstruction, we use a method which can adapt to the noise patterns.

2.A.3. Streak-adaptive averaging

Streak-adaptive averaging can account for noise texture. As with ROI averaging, the reconstructed image of the lesion alone in Fig. 1(e) is multiplied by a weighting function. Unlike ROI averaging, the shape of the weighting function Fig. 1(h) is not a simple circle but accommodates noise streaks and view-dependent statistics. This weighting function is closely related to the matched filter used in image processing. Under some conditions, streak-adaptive averaging is the best estimate possible from reconstructed images and represents the upper limit of what a human observer could extract from the FBP images. As described in the Appendix, it is equivalent to solving the weighted least squares problem but in frequency space and correctly handles the noise correlations among pixels.

2.B. Simulated objects

Raw data for the background were simulated by forward projection of either existing DICOM files or analytically defined shapes. The choice of background is important because it affects the statistics of the measurements. FBP performs best when all measurements have equivalent statistics (i.e., equal detected x-ray intensity in all rays), whereas statistical IR can adapt to inhomogeneous statistics. Seven backgrounds were used in our experiments.

2.B.1. Uniform background

The circular lesion was placed in the center of a circular water cylinder. This is the simplest scenario, although not very representative of clinical use. All the rays passing through the center of a circular water cylinder have uniform statistics. This is an advantageous situation for FBP.

2.B.2. Elliptical background

A circular lesion was placed in the center of an elliptical water cylinder. The dimensions of the water ellipse were 40 by 20 cm. In many clinical systems, tube current modulation is used to adjust the intensity of the x-ray beam to suppress variation in the measurement statistics and improve image quality.¹⁴ We disabled tube current modulation for this case in order to put FBP at the largest possible disadvantage. The noise variance of the rays along the short axis of the ellipse was 60 times better than the statistics along the long axis of the ellipse.

FIG. 2. Phantoms used in this study. The (top left) uniform circle is the most ideal case with very uniform statistics, with the lesion placed in the center of the field of view. The (top middle) ellipse, measuring 20 by 40 cm, leads to very nonuniform statistics and streaks in the noise pattern. The (top right) thorax, taken from a clinical dataset, has lesions inserted either in the mediastinum (marked “1”) or the lung (marked “2”). The (bottom left) shoulder, (bottom middle) pelvis, and (bottom right) abdomen are taken from another clinical dataset and have lesion locations marked by yellow arrows.

2.B.3. Thorax background

A deidentified scan of a thoracic section¹² was used for the third and fourth cases. Although only one thoracic section was used, we tested two different locations for the lesion: one in the mediastinum and one in the lung field.

2.B.4. Shoulder, pelvis, and liver backgrounds

Three slices were extracted from a deidentified volume and lesions were inserted between the shoulder blades, between the pelvic bones, and in the liver. Lesions in the liver represent a typical, low contrast detection task. Data for lesions between the shoulder blades or pelvic bones feature very high degrees of statistical inhomogeneity, with much better statistics along the short axis of the body than the long axis, and may represent the highest degree of improvement from statistical reconstruction.

2.C. CT simulations

Figure 2 shows all seven cases. Attenuation-dependent noise was added assuming enough photons were present that the noise statistics could be modeled as Gaussian instead of Poisson. For the human cross sections, we used the tube current

modulation scheme proposed in Ref. 14 and a bowtie filter to increase dose efficiency. The shape of the bowtie filter was similar to commercial scanners as previously reported in Ref. 15.

To decrease the computational burden, local reconstructions, 5 by 5 cm with 0.5 mm pixel size were used, with 256 views. Two lesion diameters were used, 1 and 10 mm. 10 000 random noise realizations were simulated for each object and each lesion size. For simplicity, we assumed a single-slice, parallel-beam system imaging with monochromatic x-rays, independent measurement noise in each ray, and we also assumed that the statistics were slowly varying in the detector direction so that the statistics through the lesion could be approximated as constant in each view. These assumptions could be relaxed, but we believe that the noise comparisons found in this simplified case are representative of the polychromatic, volumetric case. For each lesion and object, the mean and standard deviation across the 10 000 realizations were computed and compared.

3. RESULTS

Figure 3 shows examples of the lesion images in each of the seven datasets. Table I shows the normalized

FIG. 3. Example difference reconstructions (images of the lesion alone, after the background is subtracted) for 10 mm lesions. From left to right: mediastinum, lung field, circle phantom, ellipse phantom, shoulder, pelvis, and liver. The various reconstructions have differences in noise texture.

TABLE I. Relative standard deviations of the three estimators for different datasets and lesion sizes and normalized to that of the raw data estimate. Error bars are estimated from measured variances assuming Gaussian distributions. Numbers after the $\pm$ symbol represent standard deviation.

Background	Lesion (mm)	Raw data (%)	ROI averaging (%)	Streak-adaptive averaging (%)
Circle phantom	1	100 $\pm$ 1	133 $\pm$ 2	129 $\pm$ 2
Ellipse phantom	10	100 $\pm$ 1	128 $\pm$ 2	108 $\pm$ 1
Mediastinum	1	100 $\pm$ 1	291 $\pm$ 4	143 $\pm$ 2
	10	100 $\pm$ 1	286 $\pm$ 4	142 $\pm$ 2
Lung field	1	100 $\pm$ 1	143 $\pm$ 2	135 $\pm$ 2
	10	100 $\pm$ 1	139 $\pm$ 2	120 $\pm$ 2
Shoulder	1	100 $\pm$ 1	196 $\pm$ 3	142 $\pm$ 2
	10	100 $\pm$ 1	191 $\pm$ 3	144 $\pm$ 2
Pelvis	1	100 $\pm$ 1	212 $\pm$ 3	138 $\pm$ 2
	10	100 $\pm$ 1	206 $\pm$ 3	128 $\pm$ 2
Liver	1	100 $\pm$ 1	171 $\pm$ 2	134 $\pm$ 2
	10	100 $\pm$ 1	165 $\pm$ 2	118 $\pm$ 2
	1	100 $\pm$ 1	155 $\pm$ 2	146 $\pm$ 2
	10	100 $\pm$ 1	148 $\pm$ 2	133 $\pm$ 2

standard deviation of the lesion density estimate for each of the three methods in all seven backgrounds. The mean values are omitted from the table because all of them were accurate. The direct, raw data estimate provides the best performance possible and therefore provides an upper bound on the performance of any estimator based on reconstructed images that make no assumptions about the object. ROI averaging and streak-adaptive averaging are both calculated from FBP images. ROI averaging is a very simplistic method, whereas streak-adaptive averaging is able to extract more information from the FBP images, especially in the presence of asymmetric data statistics.

In the absence of electronic noise, dose and variance are inversely related. The maximum dose reduction possible depends on whether the observer is better modeled by ROI averaging or streak-adaptive averaging. Figure 4 is directly derived

from Table I and shows the maximum dose reduction possible from any reconstruction algorithm, including new iterative methods, for these two assumptions.

4. DISCUSSION

FBP is an imperfect reconstruction method in noisy situations because it ignores statistical information about the raw data and thereby produces reconstructions with reduced information content. New reconstruction algorithms could salvage some of this lost information, but are still limited by the information present in the original raw data. However, reports on the efficacy and dose saving potential of new IR algorithms have arguably been overoptimistic, and at times misleading. In this work, we analyzed the maximum dose reduction possible for any reconstruction algorithm over standard FBP methods for the evaluation of small circular lesions. We went directly to the raw data to find the best performance possible with any reconstruction algorithm. For our task, we found that estimates directly from the raw data had an average of 42% less variance than those derived from streak-adaptive averaging of the FBP images. Streak-adaptive averaging makes the best use of the information present in the FBP images. If we assume that the human reader is less sophisticated and uses simpler ROI averaging, then the raw data estimate performed 64% better. We believe these are limits that apply to any reconstruction method that does not make assumptions about the object.

The task studied here was that of measuring the average density of an inserted lesion. This is closely related to a low contrast detectability test. Our findings are consistent with some existing estimates in the literature on the detection of lesions.¹⁶ They are also reasonably consistent with the finding in Ref. 13 that under a different set of assumptions, the dose penalty of FBP over direct measurements of attenuation is 25% when the statistics are uniform.

Another use for iterative reconstruction could be to increase CNR and thereby the conspicuity of features which were already detectable. This could increase physician comfort or

FIG. 4. Maximum dose reduction possible with new reconstruction algorithms relative to filtered backprojection. Error bars represent 95% confidence intervals and were derived from measured variances assuming the lesion estimates are Gaussian distributed. The maximum dose reduction varies with the object, the size of the inserted lesion, and the model of the observer (simple ROI averaging or more sophisticated streak-adaptive averaging). The new reconstruction algorithm is not permitted to introduce *a priori* information or through-slice blurring.

the speed with which images are read but may not improve outcomes. For example, it is worth noting that in some instances, human observer performance diminishes even as objective metrics (such as standard deviation) and subjective metrics (image quality scores) improve.¹⁷

One major assumption in this work is that the envisioned reconstruction is not allowed to use *a priori* knowledge. Many IR algorithms do include implicit *a priori* knowledge (or assumptions). In fact, the regularizer in penalized-likelihood IR inherently assumes that small differences between neighboring pixel values are noise to be smoothed out whereas large differences are meaningful and should be retained. A very simplified model of the regularizer is that of edge-preserving smoothing. This disproportionately increases the CNR of already detectable lesions while potentially blurring out lesions below the detectability threshold, so that very low contrast lesions “vanish.”¹⁸ It has also been noted that if this *a priori* knowledge is eliminated and the regularizer behaves like conventional rather than edge-preserving smoothing, many of the image quality improvements in iterative reconstruction disappear.¹⁹ *A priori* knowledge could increase the noise reduction seen here but may also lead to biased images, such as the vanishing effect for very low contrast lesions. Stronger forms of *a priori* knowledge could be used but may be dangerous because when the *a priori* knowledge is incorrect, the final images will be further biased. Therefore, if noise reductions larger than those predicted here are observed, one should be cautious about possible introductions of bias and vanishing of small, low contrast features.

A separate, important assumption in our work is the study of a single-slice setting. This was done both to simplify the analysis and to avoid through-slice effects. Further apparent dose reductions are possible in multislice acquisitions if the iterative reconstruction algorithm blurs the images in the through-slice direction. This blurring ability is, in fact, built into modern regularizers. Larger dose reductions from IR have been reported²⁰ in multiple alternative, forced choice comparisons in which only a single slice from a multislice scanner is displayed to the observer. This is a context which could be most affected by through-slice blurring. It is perhaps telling that at the time of this writing, most dose reduction claims on IR algorithms cleared by the US FDA are explicitly restricted to thin-slice imaging, which is the setting most affected by hidden through-slice averaging. Thus, relatively large noise reductions may be seen when thin slices in a thick stack are viewed, but the observer should be aware that adjacent slices will be correlated, especially in regions of low contrast, essentially making the slices thicker than nominal in those regions.

We assumed that electronic noise was negligible in this study. In low dose scans, measurements of highly attenuated rays may be near the electronic noise floor. In this situation, unbiased estimation of the ray measurements becomes difficult, and the FBP reconstruction may also become biased. In addition, iterative reconstruction could leverage the statistics of the electronic noise and gain an additional performance advantage relative to FBP.

Some advantages of iterative reconstruction have been overlooked in this analysis. Improved modeling of the x-ray

source and detector aperture, present in iterative reconstruction using more sophisticated forward and backprojectors,^{21,22} can improve the performance of the system in high-resolution tasks. However, the impact of those improved models on detection of lesions several mm wide should be small. Some minor nonidealities of the system, such as decreased statistical efficiency arising from inefficient weighting functions used in volumetric reconstruction,²³ can be compensated for using statistical reconstruction but are not included here.

We have constrained our analysis to a task similar to the signal-known-exactly, background-known-exactly task. Real detection problems involve search for lesions of ambiguous shape and do not map directly to the task described here. It is possible that the bias that may be introduced by the regularizer, if limited, may not harm a detection task if the noise reduction from the regularizer increases conspicuity and human observer performance in a search task. On the other hand, bias does mean that noise alone, or CNR for large or higher contrast targets, cannot be used to predict the performance of the system in detection of small low contrast features. Similarly, through-plane filtering from 3D regularizers may improve performance; however, such an effect could also be achieved with conventional FBP methods followed by volumetric image processing. We do not rule out the possibility that dose reductions larger than the limits provided here are possible, but caution that, should these dose reductions be observed, they may be due to prior knowledge instead of more efficient use of the data statistics.

Historically, it has been observed that IR algorithms without regularization do not significantly improve on the FBP result. Maximum likelihood methods have existed for decades,²⁴ but these algorithms do not appear effective at reducing noise unless a regularizer is included. The regularizer stabilizes the objective function,¹ and a combination of statistical data weighting, in-plane regularization, and through-plane regularization is able to achieve noise and dose reductions. In this work, we have bounded the maximum benefit possible from statistical data weighting. Further dose reduction may be possible through in-plane regularization, which represents a weak form of prior knowledge. Still further dose reduction is possible in a slice-versus-slice comparison using through-plane regularization, at the cost of degrading the effective slice thickness.

The results of our study indicate that for low contrast detectability tasks, we should not expect substantial further reductions in dose from new reconstruction algorithms that do not use *a priori* knowledge. The amount of dose reduction varies with the anatomic setting, but assuming a 10 mm circular lesion and the streak-adaptive observer, the observed dose reduction varies between 15% and 52% in the cases studied here. Higher dose reductions are possible if the human observer does not account for noise texture, or if smaller lesions are assumed. If higher dose reductions are reported, then it may be due to the introduction of *a priori* knowledge or hidden through-slice blurring. In-plane blurring that is contrast and dose dependent has already been observed and is apparent in many images. These should not be considered a “free lunch”: while they can introduce further noise

reductions, they also lead to bias. The net effect may be beneficial clinically—i.e., the bias and through-slice blurring may be acceptable—but we suggest that when using reconstruction methods with claims of very large noise or dose reductions, the observer should be aware of these possible side effects.²⁵

ACKNOWLEDGMENTS

This work was funded in part by GE Healthcare and in part by the National Institutes of Health (No. U01 EB017140).

APPENDIX: MATHEMATICAL DESCRIPTION OF THE ESTIMATORS

1. Direct raw data estimator

Let $b_{\text{raw}}(t, \theta)$ be the noiseless raw data for the background, $\ell_{\text{raw}}(t, \theta)$ be raw data for the lesion with unit density above the background, α be the lesion density which we desire to estimate, $d_{\text{raw}}(t, \theta)$ be the measured raw data, and $\varepsilon(t, \theta)$ be a random variable representing the noise from the measurements with variance $\sigma^2(t, \theta)$. The t and θ variables correspond to detector index and view angle index, respectively. We have

$$d_{\text{raw}}(t, \theta) = b_{\text{raw}}(t, \theta) + \alpha \ell_{\text{raw}}(t, \theta) + \varepsilon(t, \theta).$$

Estimating α is a classic weighted least squares problem. The optimal estimator $\hat{\alpha}_{\text{bound}}$ for the lesion density is given by

$$\hat{\alpha}_{\text{bound}} = \sum_{t, \theta} c_{\text{bound}}(t, \theta) (d_{\text{raw}}(t, \theta) - b_{\text{raw}}(t, \theta)),$$

where we have

$$c_{\text{bound}}(t, \theta) = w_{\text{bound}}(t, \theta) \ell_{\text{raw}}(t, \theta)^{-1} \left(\sum_{t, \theta} w_{\text{bound}}(t, \theta) \right)^{-1},$$

$$w_{\text{bound}}(t, \theta) = \frac{\ell_{\text{raw}}(t, \theta)^2}{\sigma^2(t, \theta)}.$$

Note that this is in weighted sum format. The weight map $c_{\text{bound}}(t, \theta)$ is depicted in Fig. 1(c). The weight is larger when the projection of the lesion is larger or the measurement noise is low.

2. ROI averaging

Let $d_{\text{FBP}}(x, y)$ be the reconstructed image, and $b_{\text{FBP}}(x, y)$ be the reconstruction of the background alone. We seek an estimator of the form

$$\hat{\alpha}_{\text{ROI}} = \sum_{x, y} c_{\text{ROI}}(x, y) (d_{\text{FBP}}(x, y) - b_{\text{FBP}}(x, y)).$$

We choose $c_{\text{ROI}}(x, y) = \ell_{\text{FBP}}(x, y)$, the reconstructed image lesion including any blurring from the algorithm at unit density. $c_{\text{ROI}}(x, y)$ is shown in Fig. 1(f). This is equivalent to solving the least squares problem in image space if noise correlations between pixels are ignored.

3. Streak-adaptive averaging

To address the problems of noise correlation in FBP images, we move into the frequency space. Let $\mathcal{F}$ be the Fourier transform operator. The Fourier transform of the difference image, $\mathcal{F}\{d_{\text{FBP}}(x, y) - b_{\text{FBP}}(x, y)\}$, is equal to the Fourier transform of the unit density lesion shape, $\mathcal{F}\{\ell_{\text{FBP}}(x, y)\}$, multiplied by the lesion density (plus noise). Therefore, each point in frequency space can provide an independent estimate of the lesion density

$$\alpha(k_x, k_y) = \frac{\mathcal{F}\{d_{\text{FBP}}(x, y) - b_{\text{FBP}}(x, y)\}}{\mathcal{F}\{\ell_{\text{FBP}}(x, y)\}}.$$

Each of the estimates has an uncertainty that depends on the signal power and the noise power at that point in frequency space. A weighted average of $\alpha(k_x, k_y)$ is now calculated, with weights being proportional to the square of $\mathcal{F}\{\ell_{\text{FBP}}(x, y)\}$ divided by the noise power spectrum.

While this derivation was in frequency space, the linearity of the Fourier transform means that it can be implemented in the image domain. The matched filter, used commonly in image processing, has a similar derivation and is likewise applied multiplicatively in image space to detect (or measure the amplitude of) objects with a known shape in the presence of noise with known spectrum. Figure 1(h) shows the streak-adaptive weights as represented in image space.

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

¹J. 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).

²D. Fleischmann and F. E. Boas, "Computed tomography—Old ideas and new technology," *Eur. Radiol.* **21**(3), 510–517 (2011).

³A. K. Hara, R. G. Paden, A. C. Silva, J. L. Kujak, H. J. Lawder, and W. Pavlicek, "Iterative reconstruction technique for reducing body radiation dose at CT: Feasibility study," *Am. J. Roentgenol.* **193**(3), 764–771 (2009).

⁴Q. Xu, H. Yu, X. Mou, L. Zhang, J. Hsieh, and G. Wang, "Low-dose x-ray CT reconstruction via dictionary learning," *IEEE Trans. Med. Imaging* **31**(9), 1682–1697 (2012).

⁵M. Katsura, I. Matsuda, M. Akahane, J. Sato, H. Akai, K. Yasaka, A. Kunimatsu, and K. Ohtomo, "Model-based iterative reconstruction technique for radiation dose reduction in chest CT: Comparison with the adaptive statistical iterative reconstruction technique," *Eur. Radiol.* **22**(8), 1613–1623 (2012).

⁶S. Singh, M. K. Kalra, S. Do, J. B. Thibault, H. Pien, O. J. O'Connor, and M. A. Blake, "Comparison of hybrid and pure iterative reconstruction techniques with conventional filtered back projection: Dose reduction potential in the abdomen," *J. Comput. Assisted Tomogr.* **36**(3), 347–353 (2012).

⁷M. K. Kalra, M. Woisetschlager, N. Dahlstrom, S. Singh, M. Lindblom, G. Choy, P. Quick, B. Schmidt, M. Sedlmair, M. A. Blake, and A. Persson, "Radiation dose reduction with sinogram affirmed iterative reconstruction technique for abdominal computed tomography," *J. Comput. Assisted Tomogr.* **36**(3), 339–346 (2012).

⁸J. Leipsic, T. M. LaBounty, B. Heilbron, J. K. Min, G. J. Mancini, F. Y. Lin, C. Taylor, A. Dunning, and J. P. Earls, "Estimated radiation dose reduction using adaptive statistical iterative reconstruction in coronary CT angiography: The ERASIR study," *Am. J. Roentgenol.* **195**(3), 655–660 (2010).

⁹Y. Sagara, A. K. Hara, W. Pavlicek, A. C. Silva, R. G. Paden, and Q. Wu, "Abdominal CT: Comparison of low-dose CT with adaptive statistical iterative reconstruction and routine-dose CT with filtered back projection in 53 patients," *Am. J. Roentgenol.* **195**(3), 713–719 (2010).

¹⁰J. Vaishnav, W. Jung, L. Popescu, R. Zeng, and K. Myers, "Objective assessment of image quality and dose reduction in CT iterative reconstruction," *Med. Phys.* **41**(7), 071904 (12pp.) (2014).

- ¹¹J. Tang, B. E. Nett, and G. Chen, "Performance comparison between total variation (TV)-based compressed sensing and statistical iterative reconstruction algorithms," *Phys. Med. Biol.* **54**(19), 5781–5804 (2009).
- ¹²A. Rosset, H. Muller, M. Martins, N. Dfouni, J. Vallée, and O. Ratib, "Casimage project: A digital teaching files authoring environment," *J. Thorac. Imaging* **19**(2), 103–108 (2004).
- ¹³D. A. Chesler, S. J. Riederer, and N. J. Pelc, "Noise due to photon counting statistics in computed x-ray tomography," *J. Comput. Assisted Tomogr.* **1**(1), 64–74 (1977).
- ¹⁴M. 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).
- ¹⁵S. S. Hsieh and N. J. Pelc, "The feasibility of a piecewise-linear dynamic bowtie filter," *Med. Phys.* **40**(3), 031910 (12pp.) (2013).
- ¹⁶S. T. Schindera, D. Odedra, S. A. Raza, T. K. Kim, H. Jang, Z. Szucs-Farkas, and P. Rogalla, "Iterative reconstruction algorithm for CT: Can radiation dose be decreased while low-contrast detectability is preserved?," *Radiology* **269**(2), 511–518 (2013).
- ¹⁷P. J. Pickhardt, M. G. Lubner, D. H. Kim, J. Tang, J. A. Ruma, A. M. del Rio, and G. H. Chen, "Abdominal CT with model-based iterative reconstruction (MBIR): Initial results of a prospective trial comparing ultralow-dose with standard-dose imaging," *AJR, Am. J. Roentgenol.* **199**(6), 1266–1274 (2012).
- ¹⁸M. E. Baker, F. Dong, A. Primak, N. A. Obuchowski, D. Einstein, N. Gandhi, B. R. Herts, A. Purysko, E. Remer, and N. Vachani, "Contrast-to-noise ratio and low-contrast object resolution on full-and low-dose MDCT: SAFIRE versus filtered back projection in a low-contrast object phantom and in the liver," *Am. J. Roentgenol.* **199**(1), 8–18 (2012).
- ¹⁹A. S. Wang, J. W. Stayman, Y. Otake, G. Kleinszig, S. Vogt, G. L. Gallia, A. J. Khanna, and J. H. Siewerdsen, "Soft-tissue imaging with C-arm cone-beam CT using statistical reconstruction," *Phys. Med. Biol.* **59**(4), 1005–1026 (2014).
- ²⁰F. A. Miéville, F. Gudinchet, F. Brunelle, F. O. Bochud, and F. R. Verdun, "Iterative reconstruction methods in two different MDCT scanners: Physical metrics and 4-alternative forced-choice detectability experiments—A phantom approach," *Phys. Med.* **29**(1), 99–110 (2013).
- ²¹B. De Man and S. Basu, "Distance-driven projection and backprojection in three dimensions," *Phys. Med. Biol.* **49**(11), 2463–2475 (2004).
- ²²Y. Long, J. A. Fessler, and J. M. Balter, "3D forward and back-projection for x-ray CT using separable footprints," *IEEE Trans. Med. Imaging* **29**(11), 1839–1850 (2010).
- ²³T. Köhler, B. Brendel, and R. Proksa, "Beam shaper with optimized dose utility for helical cone-beam CT," *Med. Phys.* **38**, S76–S84 (2011).
- ²⁴L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. Med. Imaging* **1**(2), 113–122 (1982).
- ²⁵E. W. Olcott, L. K. Shin, G. Sommer, I. Chan, J. Rosenberg, F. L. Molvin, F. E. Boas, and D. Fleischmann, "Model-based iterative reconstruction compared to adaptive statistical iterative reconstruction and filtered back-projection in CT of the kidneys and the adjacent retroperitoneum," *Acad. Radiol.* **21**(6), 774–784 (2014).