Low radiation tomographic reconstruction with and without template information

Preeti Gopal^{a,b,c}, Sharat Chandran^a, Imants Svalbe^b, Ajit Rajwade^a

 $\{preetig, sharat, a jitvr\}$ @cse.iitb.ac.in, imants.svalbe@monash.edu

^aDepartment of Computer Science and Engineering, IIT Bombay ^bSchool of Physics and Astronomy, Monash University ^cIITB-Monash Research Academy

Abstract

Low-dose tomography is highly preferred in medical procedures for its reduced radiation risk when compared to standard-dose Computed Tomography (CT). However, the lower the intensity of X-rays, the higher the acquisition noise and hence the reconstructions suffer from artefacts. A large body of work has focussed on improving the algorithms to minimize these artefacts. In this work, we propose two new techniques, rescaled non-linear least squares and Poisson-Gaussian convolution, that reconstruct the underlying image making use of an accurate or near-accurate statistical model of the noise in the projections. We also propose a reconstruction method when prior knowledge of the underlying object is available in the form of templates. This is applicable to longitudinal studies wherein the same object is scanned multiple times to observe the changes that evolve in it over time. Our results on 3D data show that prior information can be used to compensate for the low-dose artefacts, and we demonstrate that it is possible to simultaneously prevent the prior from adversely biasing the reconstructions of new changes in the test object, by means of careful selection of a weights map, subsequently followed by a method called "re-irradiation". Additionally, we also present a technique for automated tuning of the regularization parameters for tomographic inversion.

Keywords: low-dose tomographic reconstruction, compressed sensing, priors, longitudinal studies.

1. Introduction

Reduction in radiation exposure is a critical goal, especially in CT of medical subjects [1] and biological specimens [2]. One of the ways to reduce this radiation is to acquire projections from fewer views. An alternate way, which is the focus of this work, is to lower the strength ('dose') of X-ray beam. The CT imaging model that incorporates the strength of X-rays, I_0 , is non-linear and non-deterministic and is given by:

$$\boldsymbol{y} \sim \text{Poisson}(I_0 \exp\{-\boldsymbol{\Phi} \boldsymbol{x}\}) + \boldsymbol{\eta}$$
 (1)

where η represents the zero mean additive Gaussian noise vector with a fixed signal-independent standard deviation σ , where Φ is the sensing matrix which represents the forward model for the tomographic projections, and x is the

- ⁵ underlying image representing the density values. The noise model for y is primarily Poisson in nature as this is a photon counting process [3], and the added Gaussian noise is due to the thermal effects [4]. This Poisson-Gaussian noise model is quite common in optical or X-ray based imaging systems, but we consider it here explicitly for tomography, where it induces a non-linear inversion
- ¹⁰ problem. Specifically, the i^{th} index (for bin number and projection angle) in the measurement vector \boldsymbol{y} is given as: $y_i \sim \text{Poisson}(I_0 \exp\{-\boldsymbol{\Phi}^i \boldsymbol{x}\}) + \eta_i$, where $\boldsymbol{\Phi}^i$ is the i^{th} row of the sensing matrix $\boldsymbol{\Phi}$. The major effect of low-dose acquisition is the large magnitude (relative to the signal) of Poisson noise due to the low strength of X-ray beam. This is because the Signal-to-Noise-Ratio (SNR) of Poisson noise with mean κ and variance κ is given by $\frac{\kappa}{\sqrt{\kappa}} = \sqrt{\kappa}$. Due to the

inherently low SNR, traditional low dose reconstructions are noisy.

2. Previous Work

Modelling of Poisson noise and recovery of images also finds applications in areas outside of CT. [5] recovered images from Poisson-noise corrupted and blurred images using alternating direction method of multipliers(ADMM). Lowdose imaging and reconstruction (with dense projection view sampling) has been more widely studied than the few-views imaging. This is probably because the former does not involve a strategy for selection of the set of view angles, which in itself is an active field of research [6, 7, 8]. For long, almost all of the commercial

²⁵ CT machines used FBP ¹ as the standard reconstruction technique [9]. Only recently are the iterative techniques being deployed for commercial use [10]. The power of iterative routines was reinforced by [11], where it was proved that iterative reconstructions from *ultra-low* dose ² CT are of similar quality to those of FBP reconstructions from *low*-dose CT. Here, a commercial forward ³⁰ projected model-based algorithm was deployed and compared with FBP.

Among the other iterative methods, [12] presented a technique that minimizes log-likelihood of the Poisson distribution and a patch-based spatially encoded non-local penalty. [13] used a smoothness prior along with data-fidelity constraint and solved using ADMM. In order to improve the reconstruction fur-

ther, various prior-based and learning-based methods have also been explored in literature. In these techniques, properties of available standard-dose CT images influence low-dose reconstruction of the test (i.e., the object which needs to be reconstructed from the current set of new tomographic projections). One such technique was described by [14], wherein the iterative reconstruction was

- ⁴⁰ formulated as a penalized weighted least squares problem with a pre-learned sparsifying transform. While the weights were set manually, the sparsifying transform was learned from a database of regular-dose CT images. Another technique presented by [15] clustered overlapping patches of previously scanned standard-dose CT images using Gaussian Mixture Model (GMM). The texture
- of the prior was learned for each cluster. Following this, patches from a pilot reconstruction of the test were classified using the learned GMM and depending on the class, the corresponding texture priors were imposed on patches of the reconstructed test image. The limitation here is- patches that correspond

¹Filtered Backprojection

 $^{^{2}}$ Typically, low-dose imaging is performed at 120 kVp and 30 mAs beam current, and ultra-low dose imaging is performed at 80-100 kVp and 20-30 mAs beam current settings.

to new changes between the test and the templates will also be influenced by

- some inappropriate texture of patches from prior. [16] solved a cost function with L1 norm for imposing similarity to a learned dictionary. They concluded that the number of measurements needed is progressively less for each of the four methods: Simultaneous Algebraic Reconstruction Technique (SART) [17], Adaptive Dictionary based Statistical Iterative Reconstruction (ADSIR) [18],
- Gradient Projection Barzilai Borwein (GBPP) [19] and their method (L1-DL), in the same order. [20] used edge-based priors to reconstruct normal-dose CT along with Compressed Sensing (CS) sparsity prior. An iterative method [21] in a related area (electrical impedance tomography) reconstructs using Split Bregman algorithm for L1 minimization. None of these methods explore optimizing
- a log-likelihood based cost-function that accurately reflects the Poisson-Gaussian noise statistics. In addition, they do not address the issue of the prior playing a role in the reconstruction of parts of the test that are dissimilar to the parts of the prior, which is undesirable. In contrast, this work focuses on applying a computationally fast global prior on only those regions of the test that are similar to the prior.

Given the noisy nature of tomographic projections under low radiation dosage, there are some techniques such as [22] that first seek to denoise these projections, possibly making use of the Poisson-Gaussian noise model, and subsequently reconstruct the final image from these cleaner projections. However, this can alter

- ⁷⁰ the noise statistics as all denoising techniques introduce their own 'method noise' which may introduce inconsistencies during reconstruction. Hence in this paper, we directly reconstruct the image from the noisy projections making use of the noise-model, and also use important prior information about the underlying image. Lately, artificial neural networks have also been designed for
- ⁷⁵ low-dose reconstruction. [23] proposed one such neural network to learn features of the image that is later imposed along with data-fidelity during iterative reconstruction. [24] showed that deep neural network based reconstructions are faster than iterative reconstructions for comparable reconstruction quality. All of these neural-network based techniques need large amount of data. This can

⁸⁰ be challenging in longitudinal studies where usually only a few of the previous scans of the same object are available. Hence, this paper focuses on analytical iterative techniques.

We also present a technique for parameter selection. Most techniques in literature tune the parameters omnisciently, i.e. by running the reconstruction

- ⁸⁵ algorithm for a wide range of parameters and choosing the result which is closest to the ground truth (which is assumed to be known, as is the case with synthetic experiments). A recent work [25] used the L-curve method in which datafidelity residue is plotted against regularization norm. The parameter can then be selected based on the performance required for the application at hand.
- ⁹⁰ However, this method does not utilize the available information about noise statistics in low-dose imaging. In this work, we use the noise-model for the purpose of automated parameter selection.

3. Contributions

100

105

Our contributions are as follows:

- We propose two new statistically motivated cost functions for tomographic reconstruction from projections contaminated with Poisson-Gaussian noise: the Poisson-Gaussian convolution technique (Sec. 4.7), and the rescaled non-linear LASSO (Sec. 4.6).
 - 2. We propose a method for tomographic reconstruction from low-dose measurements (i.e. where Poisson-Gaussian noise dominates) of an object \boldsymbol{x} , which makes use of previous high-dosage reconstructions of similar objects. This is common in longitudinal studies where the same subject is scanned several times, for example in cancer imaging. Our technique (Sec. 5) detects new changes (i.e., differences between the test and templates) *directly* in the measurement space. This information is then used to adaptively infer weights to be applied to the previous template reconstructions and used in the current reconstruction. These weights are designed to be low in those regions of \boldsymbol{x} where there are new structural changes, and high in

those regions which have remained unchanged over time. The weights are obtained using sound statistical criteria.

- 3. Regions of low weight (see previous point) correspond to genuine anatomical changes and such regions are often small in area. Hence, such regions can be 're-irradiated' (Sec. 5.2) so as to improve the quality of finer structures within them, at the cost of just a small added amount of radiation. This concept of re-irradiation is a third major contribution.
- 4. Lastly, we present a technique (Sec. 6) for choice of the regularization parameter to relatively balance the contribution of the data fidelity term and the regularization term in tomographic reconstruction (as in all the different techniques presented in Sec. 4). To our best knowledge, there is no prior literature on this issue, in the context of tomography. Most papers tune this parameter manually or assuming that the true image is available to drive the choice of the optimal regularization parameter.

As far as the template based techniques are concerned, there is no prior literature (to our best knowledge) which looks at non-linear tomographic recon-¹²⁵ struction under Poisson-Gaussian noise and makes use of past reconstructions in a principled way. The existing template based reconstructions assume high dosage, or a single template [26], and most importantly do not use [18] the very important weighting scheme, which we have proposed in section (Sec. 5-Equation 23).

130 4. Reconstruction without prior

A good low-dose reconstruction technique should make optimal use of noise statistics as well as appropriate signal priors. Most techniques will involve minimizing a cost function of the following form: $J(\boldsymbol{x}; \boldsymbol{y}, \boldsymbol{\Phi}) = DF(\boldsymbol{y}|\boldsymbol{\Phi}\boldsymbol{x}) + \lambda R(\boldsymbol{x})$. Here the first term involves a data-fidelity cost, and may possibly (though not necessarily) be expressed by the negative log-likelihood of \boldsymbol{y} given $\boldsymbol{\Phi}$ and \boldsymbol{x} (i.e., by $-\log p(\boldsymbol{y}|\boldsymbol{\Phi}, \boldsymbol{x})$). Other alternatives could include a simple least squares term $\|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{x}\|_2^2$, or a weighted version of the same. In this section, we review several

110

115

such fidelity functions from the literature and propose two new ones. The second term $R(\boldsymbol{x})$ is a regularizer (weighted by the regularization parameter λ) representing prior knowledge about \boldsymbol{x} . This could be in the form of the well-known to-

- ¹⁴⁰ senting prior knowledge about \boldsymbol{x} . This could be in the form of the well-known total variation prior $TV(\boldsymbol{x}) = \sum_{i,j} \sqrt{(x(i+1,j)-x(i,j))^2 + (x(i,j+1)-x(i,j))^2}$ or penalty on the ℓ_1 norm of the coefficients $\boldsymbol{\theta}$ in a sparsifying basis $\boldsymbol{\Psi}$ where $\boldsymbol{x} = \boldsymbol{\Psi} \boldsymbol{\theta}$. Such cost functions are minimized by iterative shrinkage and thresholding algorithms such as ISTA. However, ISTA by itself is known to have sublinear convergence (as discussed in Sec.3 of [27]). Hence, faster methods such as the Fast Iterative Soft Thresholding Algorithm (FISTA) [27] may be used, which
 - have a quadratic rate of convergence, and hence are employed for the purpose of optimization in this paper. Below are some of the existing reconstruction methods, or intuitive variants thereof, and two new proposed techniques.

150 4.1. Post-log Compressed Sensing (CS)

A preliminary approach is to ignore the presence of Poisson noise and apply traditional CS reconstruction after linearizing the measurements [28]. The latter process is performed by computing the logarithm of the acquired measurements. The linearized measurements \boldsymbol{y}_0 are given by $\boldsymbol{y}_0 = -\log\left(\frac{\boldsymbol{y}+\boldsymbol{\epsilon}}{I_0}\right) = \boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}$, where $\boldsymbol{\epsilon}$ is a small positive constant added to the measurements to make them all positive and thus suitable for linearizing by applying a logarithm. For practical purposes, if min(\boldsymbol{y}) is zero or negative, $\boldsymbol{\epsilon}$ is set to $-\min(\boldsymbol{y}) + 0.001$. The cost function is given by

$$J_{PL-CS}(\boldsymbol{\theta}) = \|\boldsymbol{y}_{\boldsymbol{0}} - \boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}\|_{2}^{2} + \lambda \|\boldsymbol{\theta}\|_{1}, \text{ subject to } \boldsymbol{\Psi}\boldsymbol{\theta} \succeq \boldsymbol{0}$$
(2)

 J_{PL-CS} is minimized by l1 - ls solver [29]. This method is however not true to the Poisson-Gaussian statistics and suffers from an inherent statistical bias (as seen in Fig. 1) as it is a so-called 'post-log' method. The bias arises because for any non-negative random variable X, we have $\log(E[X]) \ge E(\log(X))$ as

per Jensen's inequality. Another way of viewing this is that the noise in y_0 (i.e. post-log) is being treated as if it were Gaussian with a constant variance (which is not true of Poisson or Poisson-Gaussian settings). This is not true except



Figure 1: Histogram of statistical bias in post-log methods. The bias is computed as $(y_0 - \Phi\Psi\theta)$, where y_0 refers to linearized post-log measurements. Here, the added Gaussian noise had a mean value of 0 and $\sigma = 0.01 \times$ average Poisson-corrupted projection value. The fact that every bin has a different bias, but is shifted by a constant ϵ is problematic. This results in poor reconstructions, as shown in a later Sec. 4.8.

at very high intensity (I_0) values. The adverse effects of computing post-log measurements is also discussed in [30].

160 4.2. Non-linear Least Squares with CS

An intuitive way to modify the previous cost J_{PL-CS} is by allowing the data fidelity cost to mimic the non-linearity inherent in the acquisition process. The cost function is then given by

$$J_{NL-CS} = \|\boldsymbol{y} - I_0 e^{-\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta}}\|_2^2 + \lambda \|\boldsymbol{\theta}\|_1, \text{ subject to } \boldsymbol{\Psi} \boldsymbol{\theta} \succeq \boldsymbol{0}$$
(3)

The FISTA routine [27] is used for this minimization. Since the attenuation constant of an object is never negative, a non-negativity constraint is imposed on $\Psi \theta$. It can be seen that this cost function is non-convex in θ .

4.3. Filtered Backprojection

In this technique, the classic filtered backprojection is applied on the linearized measurements: $y_0 = -\log \frac{y+\epsilon}{I_0} = \Phi x$. The slice or volume x is then reconstructed from the linearized measurements by filtered backprojection (FBP) in case of parallel beam projections or FeldKamp David Kress (FDK) algorithm [31] in case of cone beam projections. This method is called the post-log FBP. While it is computationally efficient, it suffers from a statistical bias for the same reasons as post-log CS, as described in 4.1. The performance of postlog FBP has been extensively compared with iterative schemes in [32],[33],[34] and the latter has been found to be well suited for low-dose reconstructions [35].

4.4. Negative Log Likelihood-Poisson with CS

This technique accounts for only the Poisson noise (ignoring the Gaussian part) and searches for a solution that minimizes the negative log-likelihood [36] of the observed measurements. Given m measurements, the likelihood of $\boldsymbol{\theta}$ is defined as

$$L(\boldsymbol{\theta}|\boldsymbol{y}) \coloneqq P_Y(\boldsymbol{Y} = \boldsymbol{y}|\boldsymbol{\theta}) = \prod_{i=1}^m \frac{e^{-a_i} a_i^{y_i}}{y_i!}$$
(4)

where $a_i = I_0 e^{-(\Phi \Psi \theta)_i}$. Thus, the negative log likelihood of θ is given by

$$-\log(P(\boldsymbol{y}|\boldsymbol{\theta})) = \sum_{i} (a_{i} - y_{i} \log a_{i} + \log(y_{i}!))$$

$$= \sum_{i} (I_{o}e^{-(\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_{i}} - y_{i}(\log(I_{0}) - (\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_{i}) + \log(y_{i}!))$$
(5)

The cost function combines the likelihood and the CS term as shown below:

$$J_{NLL-P}(\boldsymbol{\theta}) = \sum_{i} (I_o e^{-(\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i} - y_i (\log(I_0) - (\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i) + \lambda \|\boldsymbol{\theta}\|_1, \text{ subject to } \boldsymbol{\Psi} \boldsymbol{\theta} \succeq \boldsymbol{0}$$
(6)

175

This technique has been used in [30] for ultra-low-dose CT reconstruction. For the case of Poisson-Gaussian noise, a shifted form of the likelihood is used, where \boldsymbol{y}_i is replaced by $\boldsymbol{y}_i + \sigma^2$ and $(\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i$ is replaced by $(\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i + \sigma^2$.

4.5. Negative Log Likelihood-Poisson-Gaussian with CS

A natural extension of the earlier method is one wherein both the Poisson and Gaussian noise processes are accounted for in the design of the cost function. Here, given the measurements, the solution that minimizes the sum of negative likelihood terms of both Poisson and Gaussian noise models, is selected. Let V denote the Poisson random variable, i.e. $y = v + \eta$. As seen earlier, the Poisson likelihood of θ is given by

$$L(\boldsymbol{\theta}|\boldsymbol{v}) \coloneqq P_V(\boldsymbol{V}=\boldsymbol{v}|\boldsymbol{\theta}) = \prod_{i=1}^m \frac{e^{-a_i} a_i^{v_i}}{v_i!}$$
(7)

where $a_i = I_0 e^{-\Phi \Psi x}$. Poisson negative log-likelihood of θ is given by

$$-\log(P_V(\mathbf{V} = \mathbf{v}|\boldsymbol{\theta})) = \sum_i (a_i - v_i \log a_i + \log(v_i!))$$
$$= \sum_i (I_o e^{-(\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i} - v_i (\log(I_0) - (\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i) + \log(v_i!))$$
(8)

Next, if the assumed Gaussian noise has a variance of σ^2 , then Gaussian likelihood of σ is given by

$$L(\sigma|\boldsymbol{\eta}) \coloneqq P_E(\boldsymbol{E} = \boldsymbol{\eta}|\sigma) = P((\boldsymbol{y} - \boldsymbol{v})|\sigma) = \prod_{i=1}^m e^{-\frac{(y_i - v_i)^2}{2\sigma^2}}$$
(9)

The Gaussian negative log-likelihood of σ is given by

$$-\log(P(\boldsymbol{y}-\boldsymbol{v})|\sigma) = \sum_{i} \frac{(y_i - v_i)^2}{2\sigma^2}$$
(10)

We minimize the sum of the two negative log-likelihoods:

$$J_{PG-NLL}(\boldsymbol{\theta}, \boldsymbol{v}) = \sum_{i} (I_o e^{-(\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i} - v_i (\log I_0 - (\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i) + \log(v_i!) + \frac{(y_i - v_i)^2}{2\sigma^2}) + \lambda \|\boldsymbol{\theta}\|_1, \text{ subject to } \boldsymbol{\Psi} \boldsymbol{\theta} \succeq \boldsymbol{0}.$$
(11)

180

 θ and v are solved for alternately. Note that v is integer-valued, but a typical gradient-based method will not restrict v to remain in the domain of integers. For computational convenience, v needs to be 'softened' to real values. Consequently $\log(v_i!)$ must be replaced by the gamma function.

This cost function is non-convex. However it can be shown to be bi-convex, i.e., it is convex in θ if v is kept fixed and vice versa. Such a cost-function was used in [37] as a method of pre-processing/denoising of projections prior to tomographic reconstruction. In contrast, we directly use it as a data-fidelity term for tomographic reconstruction. This appears more principled because denoising of a projection induces some 'method noise' which cannot be accurately modelled and which may affect subsequent reconstruction quality.

190 4.6. Proposed Rescaled non-linear Least Squares (RNLLS) with CS

This new method integrates Poisson noise model into the technique described in Sec.4.2. Since, the variance of a Poisson random variable is proportional to its mean, the variance of \boldsymbol{y} is directly proportional to $I_0 \exp(-\Phi \Psi \boldsymbol{\theta})$. Hence the data-fidelity cost must be rescaled as shown below:

$$J_{RNLLS}(\boldsymbol{\theta}) = \sum_{i=1}^{m} \frac{(y_i - I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i})^2}{I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i}} + \lambda \|\boldsymbol{\theta}\|_1, \text{ subject to } \boldsymbol{\Psi}\boldsymbol{\theta} \succeq \boldsymbol{0}$$
(12)

Again, the cost is minimized using FISTA solver. This technique is in some sense similar to the Penalized Weighted Least Squares (PWLS) technique from [38] which seeks to minimize

$$J_{PWLS}(\boldsymbol{\theta}) = \|\boldsymbol{W}(\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})\|^2 + \lambda \|\boldsymbol{\theta}\|_1$$
(13)

where W is a diagonal matrix of weights which are explicitly set (prior to running the optimization) based on the values in y. This approach is heuristic in nature. Rather in RNLLS, the 'weights' are set to be equal to the underlying noiseless measurements, i.e. equal to $I_0 e^{(-\Phi\Psi\theta)}$, and are explicitly inferred on the fly. In fact, a major motivation for our proposed technique is based on the fact that

$$E\left(\frac{[y_i - I_0 \exp(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i]^2}{I_0 \exp(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i}\right) = Var\left(\frac{[y_i - I_0 \exp(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i]}{I_0 \exp(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i}\right) = 1$$
(14)

This technique can be used for the case of Poisson-Gaussian noise as well, as in

$$J_{RNLLS-PG}(\boldsymbol{\theta}) = \sum_{i=1}^{m} \frac{(y_i - I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i})^2}{I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i} + \sigma^2} + \lambda \|\boldsymbol{\theta}\|_1, \text{ subject to } \boldsymbol{\Psi}\boldsymbol{\theta} \succeq \boldsymbol{0} \quad (15)$$

We noticed that in [39], tomographic reconstruction was performed by minimizing the following cost function:

$$J_{RNLLS-PG-log}(\boldsymbol{\theta}) = \sum_{i=1}^{m} \frac{(y_i - I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i})^2}{I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i} + \sigma^2} + \langle \log(I_0 \exp(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i + \sigma^2), 1 \rangle$$
(16)

which is inspired by the approximation of Poisson(z) by $\mathcal{N}(z, z)$ and treating it as a maximum quasi-likelihood problem. On the other hand, the proposed method (RNLLS) can be interpreted as a weighted form of the well-known LASSO problem [40]. We also note that the cost function for RNLLS is convex ¹⁹⁵ in the case of Poisson noise, as shown in the supplemental material. In the case of Poisson-Gaussian noise, our numerical simulations reveal that the cost function is not convex in the worst case which does not often arise in practice. However, this non-convexity did not affect the numerical results significantly.

4.7. Proposed Poisson-Gaussian Convolution

This new technique models both the Poisson and Gaussian noise. It is based on the fact that if a random variable Q is the sum of two random variables R and S, then the density function of Q is given by the convolution of the density functions of R and S. This scheme has been used earlier [41] for image restoration from linear degradations such as blur, followed by Poisson-Gaussian corruption of the signal. In contrast, in CT, the measured signal is a non-linear function of the underlying image (i.e. its attenuation coefficients) as per Beer's law. Eq. 17 refers to the Beer's law along with the Poisson and Gaussian noise. The measurement is the sum of a Poisson random variable and a Gaussian random variable:

$$\boldsymbol{y} \sim \operatorname{Poisson}(\boldsymbol{a}) + \boldsymbol{\eta}$$
 (17)

where $\boldsymbol{a} = I_0 e^{-\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta}}$. The *i*th measurement is given as: $y_i \sim \text{Poisson}(a_i) + \eta_i$, where $a_i = I_0 e^{-[\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta}]_i}$. The probability density of the *i*th measurement y_i is given by the following convolution:

$$p_{y_i}(z_i) = \sum_{l=0}^{l=+\infty} \frac{e^{-a_i} a_i^l}{l!} \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(z_i-l)^2}{2\sigma^2}}$$
(18)

The running variable does not take on negative values because the Poisson is a counting process and hence the corresponding random variable is always positive. Because all the m measurements are independent (i.e., the noise in the sensor at any one pixel is independent of the noise at any other pixel on it), we have

$$p_{\boldsymbol{y}}(\boldsymbol{z}) = \prod_{i=1}^{i=m} \Big(\sum_{l=0}^{l=\infty} \frac{e^{-a_i} a_i^l}{l!} \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(z_i-l)^2}{2\sigma^2}} \Big)$$
(19)

The θ that maximizes the above probability needs to be computed. This is equivalent to minimizing the negative log-likelihood of the above probability.

Hence, our cost function J_{conv} is given by

$$\begin{aligned} J_{conv}(\boldsymbol{\theta}) &= -\log p_{\boldsymbol{y}}(\boldsymbol{z}) \\ &= \sum_{i=1}^{i=m} -\log \Big(\sum_{l=0}^{l=\infty} \frac{e^{-I_o e^{-[\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta}]_i}} (I_o e^{-[\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta}]_i})^l}{l!} \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(z_i-l)^2}{2\sigma^2}} \Big) \quad (20) \\ &+ \lambda \|\boldsymbol{\theta}\|_1, \text{ subject to } \boldsymbol{\Psi} \boldsymbol{\theta} \succeq \boldsymbol{0} \end{aligned}$$

Since l! is computationally intractable for large l, it has been approximated using Stirling's approximation: $l! \sim \sqrt{2\pi l} \left(\frac{l}{e}\right)^l$. Further, in order to make the optimization numerically feasible, the value that l takes for a particular measurement y_i is limited to the range $\max(0, y_i - K\sigma)$ to $y_i + K\sigma$ where K is an integer that is usually set to 3. It is assumed here that some estimate of the variance σ^2 of the Gaussian noise is already known. This is usually feasible by recording the values sensed by the detector during an empty scan (without any object), usually before the actual scan is taken.

4.8. Results on comparison of different methods

In order to compare the performance of various methods, 2D reconstructions of two datasets (Walnut[42] and Colon CT[43]) shown in Fig. 2 were computed for varying low-dose intensities. Reconstructions of two other datasets (Pelvis[44] and Shoulder CT[45]) are shown later in the supplemental material [46]. Following are the details of the datasets and the conditions used for simulating low-dose imaging: The size of the image from Walnut dataset was 156 × 156, and the size of image from Colon CT dataset was 154 × 154. The sum of the intensity values for the Walnut and Colon dataset images were 75 and 60 respectively. Measurements were simulated using equidistant angle sampling based on parallel beam geometry. The Cosine filter was applied for filtered backprojection. While the number of projection views was large (200

views for all datasets) and kept constant, the beam strength I_0 was varied as follows: $I_0 = 20, 40, 80, 160, 320$ and 620. Based on the intensity (attenuation coefficients) of the images, the above values of I_0 correspond to a Poisson noiseto-signal ratio (i.e. average value of $1/\sqrt{\kappa}$) of 25% for $I_0 = 20$, and 4.5% for $I_0 = 620$, for both the datasets. In addition, Gaussian noise of 0 mean and variance equal to 2% of average Poisson-corrupted measurement was added to measurements. The regularization parameter λ was chosen omnisciently. Among the methods discussed here, the ones that model both Poisson and Gaussian noise are non-convex. A few of the methods that model Poisson noise alone are convex and their convexity is proved in Sec.1 of [46].



Figure 2: Ground truth test slices used for comparison of low dose reconstruction techniques. A slice from (a) [42] dataset is of size 156×156 , (b) [43] dataset is of size 154×154

Sample reconstructions are shown in Figs. 3 and 4. The corresponding SSIM 230 values of the reconstructions are shown in Fig. 5. From these plots, the following can be inferred: the convolution method and the Poisson-Gaussian likelihood reconstructions were comparable and gave the best reconstructions for a majority of dose levels and datasets. The Poisson-Gaussian Likelihood and the Poissononly likelihood have very similar performance. However, at a theoretical level, 235 the former is a more principled method, and can deal with negative-valued measurements which have to be weeded out for the Poisson-only method. A shifted Poisson model as used in [30] for Poisson-Gaussian noise does not weed out measurements, but it matches the noise distribution for only the first two moments, and thus does not fully account for noise statistics. The non-linear least 240 squares method (Sec. 4.2) performed poorly. This is because the data-fidelity term assumes constant variance for all signal values. In reality, the variance of Poisson noise increases as signal intensity increases. The post-log linear least squares (Sec. 4.1) failed because the linear model fails to approximate the highly

14

non-linear low-dose acquisition. The post-log FBP yielded poor results, espe-



Convolution (Sec. 4.7)



Log-Likelihood Poisson-Gaussian (Sec. 4.5)



Rescaled Non-Linear Least Squares (Sec. 4.6)



Post-Log FBP (Sec. 4.3)

Figure 3: 2D Low-dose reconstructions of Walnut dataset for $I_0 = 20, 40, 80, 160, 320$ and 620 (from left to right). Gaussian noise of 0 mean and variance equal to 2% of average Poisson-corrupted measurement was added to simulate the low-dose acquisition. The SSIM values are shown in Fig. 5.

cially at slightly higher dose levels (for example at $I_0 = 620$ in Fig. 3). This could be due to the absence of iterative optimization when compared to the other methods and due to the post-log approximation. For all datasets except Walnut (Colon as discussed here, and Pelvis, Shoulder as discussed in [46]),

the performance of rescaled non-linear least squares (RNLLS) is inbetween the performance of likelihood-based methods and those of all other methods. For the Walnut dataset though, the RNLLS gives the best quality for many dosage levels. The performance of the above methods across multiple noise instances



Convolution (Sec. 4.7)



Log-Likelihood Poisson-Gaussian (Sec. 4.5)



Rescaled Non-Linear Least Squares (Sec. 4.6)



Post-Log FBP (Sec. 4.3)

Figure 4: 2D Low-dose reconstructions of Colon dataset for $I_0 = 20, 40, 80, 160, 320$ and 620 (from left to right). Gaussian noise of 0 mean and variance equal to 2% of average Poisson-corrupted measurement was added to simulate the low-dose acquisition. The SSIM values are shown in Fig. 5.

is discussed in Sec.2.1 of [46].

255

To summarize, among the techniques for which no templates are used, we have compared our techniques to recent ones such as [37] and [30]. The technique in [37] is the same as the one described in Sec. 4.5 and Equation 11. The work in [30] presents post-log (similar to the non-linear CS in Sec. 4.2) and pre-log techniques including the one in Sec. 4.4. Our rescaled nonlinear LASSO

technique from Sec. 4.6 is an improved version of the pre-log technique from [30], which sets the weights based on the noisy measurements in y. On the other



Figure 5: SSIM of the reconstructions for Walnut and Colon datasets shown in Fig. 4 for varying values of X-ray doses. A higher SSIM implies better reconstruction. Here, the reconstructions by Poisson-likelihood and Poisson-Gaussian likelihood methods were very similar. Hence, their SSIM plots (blue and yellow respectively) overlap.

hand, our technique sets these weights in a more principled fashion, as seen in Equation 15).

5. Reconstruction with prior

- As seen so far, principled data fidelity terms play a significant role in improving the reconstruction performance. However, when the x-ray dose is less, the performance can be further improved by incorporation of useful priors [26, 47]. These priors could be previous high-quality reconstructions of the same object in longitudinal studies, or high-quality reconstructions of similar objects. We refer to such prior data as templates. Here, our aim is to reconstruct an object from its *low-dose measurements*, using templates which are previous *high-dose reconstructions* of the same object in a longitudinal study. However, there is a danger of the templates overwhelming the current reconstruction and ad-
- 275 needs to be reconstructed from the current set of new tomographic projections) that are absent in any of the templates. In the case of reconstruction from few projection views, the above problem was tackled [48] by generating a map (known as 'weights-map') that shows an estimate of the regions of new changes and their magnitude. This map was then used to modulate the influence of

versely affecting reconstruction of new regions in the test (i.e., the object which

- the prior on the reconstruction of the test. The weights-map was computed based on the difference between the pilot reconstruction from the test measurements (acquired from a sparse set of projection angles) and its projection onto an eigenspace spanned by representative templates. However, in the lowdose case, this is not a preferable method because *all information about the*
- noise model is valid for the measurement space alone. The noise model (i.e., $\boldsymbol{y} \sim \text{Poisson}(I_0 \exp\{-\boldsymbol{\Phi} \boldsymbol{x}\}) + \boldsymbol{\eta})$ is not applicable to the spatial reconstructed image domain.

Hence, in this work, we propose a new algorithm to compute the weightsmap (i.e to detect differences between the test and the templates) directly in the measurement space. The aim is to identify those measurement bins which correspond to the new changes in the test. Following are the steps followed in order to accomplish this:

- 1. Let $x_{t_1}...x_{t_n}$ be *n* high quality template volumes, i.e. template volumes reconstructed from their standard dose measurements.
- 2. Simulate noiseless measurements from template volumes using the same I_0 used for imaging the test i.e. $y_{t_i} = I_0 \exp\{-\Phi x_{t_i}\}$, where $1 \le i \le n$.
 - 3. Let $y_{t_{i,j}}$ be the tomographic projection of the i^{th} template from the j^{th} angle, where $1 \leq j \leq Q$. Let $\{E_j\}_{j=1}^Q$ represent the set of eigenspaces, where E_j is the eigenspace built from the tomographic projections of each of the templates in the j^{th} angle, i.e. built from $\{y_{t_{i,j}}\}_{i=1}^n$
 - 4. Let y_j be the noisy tomographic projection of the test volume x from the jth angle. For each j ∈ {1,...,Q}, project y_j onto E_j, i.e., compute the eigen-coefficients α_j^m of the measurements y_j, along the set of eigenvectors V_j^m:

$$\boldsymbol{\alpha}_j^m = (\boldsymbol{V}_j^m)^T (\boldsymbol{y}_j - \boldsymbol{\mu}_j^m)$$
(21)

where $\boldsymbol{\mu}_{j}^{m}$ denotes the mean tomographic projection of all templates in the j^{th} angle. The *m* in the suffix denotes that the eigenspace $E_{j} :=$ $\{\boldsymbol{\mu}_{j}^{m}, \boldsymbol{V}_{j}^{m}\}$ is computed in the measurement space (We will contrast this with another eigen-space computed in image domain, used later in Eq. 23). Next, compute the resultant projection \boldsymbol{y}_{pj} , i.e.,

$$\boldsymbol{y}_{pj} = \boldsymbol{\mu}_j^m + \boldsymbol{V}_j^m \boldsymbol{\alpha}_j^m \tag{22}$$

5. Note that if a random variable $s \sim \text{Poisson}(\lambda) + \eta$, where $\eta \sim \mathcal{N}(0, \sigma^2)$, then $\sqrt{s + (3/8) + \sigma^2}$ is approximately distributed as $N(\sqrt{\lambda + (3/8) + \sigma^2}, 1/4)$. The quality of the approximation is known to improve as λ increases. In the absence of Gaussian noise (equivalent to the case where $\sigma = 0$), this transform is called the Anscombe transform [49, 50], and has been widely used in image processing. In the presence of Gaussian noise, it is referred to as the generalized Anscombe transform [51]. Now consider the k^{th} bin in the test measurement \boldsymbol{y} as well as in \boldsymbol{y}_{pj} , which we shall denote as y(k)

300

and $y_{p_j}(k)$ respectively. If y(k) represents the same underlying structure as in $y_{p_j}(k)$, barring the effect of Poisson-Gaussian noise, i.e. if the k^{th} bin in \boldsymbol{y} is not part of the 'new changes', then the following is true: $\sqrt{\boldsymbol{y}+3/8+\sigma^2}-\sqrt{\boldsymbol{y_p}+3/8+\sigma^2}\sim N(0,1/4).$

For bins falling in the regions of change in the test (compared to the template projections), the above hypothesis is false. The same argument can be extended for entire segments or 2D regions.

- 6. Based on the aforementioned fact, hypothesis testing is performed on √y + 3/8 + σ² - √y_p + 3/8 + σ² to detect bins corresponding to new changes in the measurement space. We use Z-test for hypothesis testing [52] on 2D patches in the measurement space (note that since the volume is in 3D, the measurement space is in 2D for every imaging view). This Z test computes the probability that the given sample is likely to be drawn from a population as specified by the null hypothesis. In this case, the null hypothesis is that the intensity values of small-sized patches taken from √y + 3/8 + σ² - √y_p + 3/8 + σ² are drawn from N(0, 1/4). The confidence level was set to 95%, i.e. for null hypothesis to be false, the probability p that the sample is drawn from Normal distribution must lie in the 2.5% tail-end of the Normal distribution on either side. A lower p-value denotes the presence of new changes i.e., presence of differences between the test and the templates in the measurement bins.
- 7. Once the new changes are detected in the measurement space, filtered backprojection of the vectors (containing p-values) resulting from the hypothesis test gives the location of the new changes (which we denote W_{inlier}) in the original (3D) spatial domain. The Cosine filter was used in the filtered backprojection process.
- 8. The final weights-map W^3 is computed from W_{inlier} by the following steps: (a) **Inversion**: $W = 1./(1 + (W_{inlier}).^2)$ where ./ and .² indi-

310

315

- 325

³An alternate method to compute a weights-map (a simpler binary weights-map) is discussed in Sec.3 of [46]

cate point-wise division and squaring, respectively. This step is just for inversion so that new regions get lower weight/intensity than prior-similar regions, (b) **Linear stretching**: Perform linear stretching on W so that the weights lie between 0 and 1.

Finally, the computed weights-map is used in a reconstruction optimization as follows:

$$J(\boldsymbol{\theta}, \boldsymbol{\alpha}) = \sum_{i=1}^{m} \frac{(\boldsymbol{y}_i - I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i})^2}{I_0 e^{(-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta})_i} + \sigma^2} + \lambda_1 \|\boldsymbol{\theta}\|_1 + \lambda_2 \|\boldsymbol{W}(\boldsymbol{\Psi}\boldsymbol{\theta} - (\boldsymbol{\mu} + \sum_{i=1}^{n-1} \boldsymbol{V}_i \alpha_i))\|_2^2$$
(23)

where the eigenvectors V and mean of the templates μ form the eigenspace which is built from the available high-dose *reconstructions* of the templates. Here, α is the vector of coefficients obtained by projecting the reconstruction of the test onto this eigenspace created from the high-quality templates. Information about the location and magnitude of new changes in the test is present in the weights-map W. Eq. 23 is solved by alternating minimization on θ and α until convergence is reached.

5.1. Reconstruction results

The above algorithm was validated by reconstructing a 3D volume from its low dose measurements. Fig. 6 shows a slice from each of the template and test volumes of the potato dataset. This dataset ⁴ consisted of four scans, each

⁴We are grateful to Dr. Andrew Kingston for facilitating data collection at the Australian National University.



Figure 6: Potato 3D dataset: One of the slices from template volumes (first four from the left) and test volume (extreme right). Size of each volume is $[150 \times 150 \times 20]$.

340



Figure 7: Prior-based low-dose reconstruction on 3D potato dataset. (a) Slice from test volume (b) Reconstruction using no prior (using RNLLS of Sec. 4.6); SSIM = 0.22 (c) Slice from unweighted prior (i.e. setting W to be the identity matrix in Eq. 23) reconstruction; SSIM = 0.42. The new change (highlighted as red RoI) is missing. (d) Slice from weighted prior reconstruction; SSIM = 0.69. The new change is detected here and its reconstruction is guided by the low-dose measurements. (e) Weights map showing the location and intensity of the new changes (darker regions indicate regions of change, coinciding with the red RoI). All SSIM values are averaged over 14 slices of the reconstructed volume in the red RoI region. The reconstructed volumes can be seen in [46].

acquired under high radiation dosage, of the humble potato, chosen for its simplicity. Measurements from each scan consisted of cone-beam projections from 900 views, each of size 150×150 . The corresponding size of the reconstructed

- volume is 150×150×20. While the first scan was taken of the undistorted potato, subsequent scans were taken of the same specimen, each time after drilling a new hole halfway into the potato. The ground truth consists of FDK reconstructions from the full set of acquired measurements from 900 equi-spaced projection views. Low dose cone-beam measurements were simulated from full-view FDK
- reconstructions of the test volume. I_0 was set to 4000, a value corresponding to Poisson noise of 1.5%. Mean of the added Gaussian noise was 0 and σ was set to 0.1% of the mean of Poisson-corrupted measurements. Fig 7 shows the same slice from each of the reconstructed volumes. A patch size of [5,5] was used for hypothesis testing and the location of new changes (marked in red RoI
- in test) was accurately detected in the weights-map as seen in Fig. 7e. The reconstructed volumes can be found in [46].

5.2. Re-irradiation to improve reconstruction

Once the regions of new changes are detected by the weights map, this information can be used to re-irradiate them with standard-dose rays and further improve the quality of their reconstruction. Following are the steps of the reirradiation process:

- 1. Let the X-rays passing through the new regions have their source points denoted by S_1 , and the corresponding bins at the detector be denoted by D_1 . Let the X-rays passing through the other regions (i.e. regions where the test and the templates are not structurally different) have their source points denoted by S_2 , and the corresponding bins at the detector be denoted by D_2 .
- 2. Block S_2 and re-irradiate the object by passing standard-dose rays from S_1 . This will generate measurements of high quality for regions of new changes. If the regions of new change are small in area, this process incurs only a small cost for the extra amount of radiation, since the latter is restricted to only specific regions.
- 3. In the measurement matrix captured for pilot reconstruction, replace all the bins in set D_1 by their new measurements. Therefore, the final measurement matrix consists of standard-dose measurements corresponding to new regions of the object and low-dose measurements corresponding to the other regions of the object.

Note that the original sampling pattern is uniform. Once the weights are obtained, the sampling pattern for re-irradiation is non-uniform and dependent on the location of the region of interest in the object. The new measurement model is: $\boldsymbol{y} \sim \text{Poisson}(\boldsymbol{I_0} \exp\{-\boldsymbol{\Phi}\boldsymbol{x}\}) + \boldsymbol{\eta}$. Here $\boldsymbol{I_0}$ now denotes a diagonal matrix (as opposed to a scalar quantity as in Eq. 1) with $\boldsymbol{I_0}(k,k)$ denoting the strength of the X-ray incident on the k^{th} bin of the sensor. Fig. 8 shows the templates and test images, and Fig. 9 shows the reconstructions and PSNR values illustrating the benefit of re-irradiation. Note that these reconstructions are from 360 (i.e.

dense) equi-spaced parallel-beam projections. The new changes within the RoI

375

380

370

are reconstructed very well after they are re-imaged with standard-dose X-rays. This is also reinforced by results on the sprouts data (Fig. 10), shown in Fig. 11. The selection of bins for re-irradiation and the choice of new X-ray intensity can also be chosen in a supervised manner by the physician or scientist based on the particular clinical or non-clinical setting.



Figure 8: Dataset for illustrating re-irradiation: Templates (first four from the left) and test (extreme right). Size of each slice is (310×310) . The RoI shows the region of difference between the test and the templates. (Also see Fig. 9.)



Figure 9: Improving reconstruction by re-irradiation in Okra 2D dataset (from Fig. 8). Measurements acquired were 360 equispaced parallel-beam projections. (a) test (b) pilot (PSNR=41.0 in the RoI, relative MSE = 0.24 in the RoI, relative MSE for full image = 0.40) (c) weights-map; the lower the intensity, the higher the magnitude of new changes. (d) weighted prior reconstruction (PSNR=49.0 in the RoI, relative MSE = 0.16 in the RoI, relative MSE for full image = 0.24); the quality of reconstruction of new regions is poor because it is guided by the measurements alone. (e) re-irradiated reconstruction (PSNR=64.7 in the RoI, relative MSE = 0.07 in the RoI, relative MSE for full image = 0.30); new measurements with twice the earlier low-dose X-ray intensity at 20% of the bins enable better reconstruction of new regions (as shown in RoI).



Figure 10: Sprouts Dataset for illustrating re-irradiation: Templates (first row) and test (second row). Size of each slice is (156×156) . The RoI shows the region of difference between the test and the templates. (Also see Fig. 11.)



Prior

re-irradiation

Figure 11: Improving reconstruction by re-irradiation in Sprouts 2D dataset (from Fig. 10). Measurements acquired were 350 equispaced parallel-beam projections. (a) test (b) pilot (PSNR= 39.3 in the RoI, relative MSE = 0.33 in the RoI, relative MSE for full image = 0.25) (c) weights-map; the lower the intensity, the higher the magnitude of new changes. (d) weighted prior reconstruction (PSNR=34.6 in the RoI, relative MSE = 0.42 in the RoI, relative MSE for full image = 0.22); the quality of reconstruction of new regions is poor because it is guided by the measurements alone. (e) re-irradiated reconstruction (PSNR=47.8 in the RoI, relative MSE = 0.22 in the RoI, relative MSE for full image = 0.17); new measurements with 8 times the earlier low-dose X-ray intensity at 25% of the bins enable better reconstruction of new regions (as shown in RoI).

6. Tuning of parameters

Two parameters were used in the techniques presented in this chapter: λ_1 : weight for CS term and λ_2 : weight for object-prior. Below are few of the ways to select these parameters.

6.1. Selection of weightage for CS term

In a large body of work on tomographic reconstruction [14], [53], the regularization parameter λ_1 in Eq. 23 is chosen in an "omniscient fashion". That is, the optimization problem is solved separately for many different values of

⁴¹⁰ λ_1 . The particular result which yields the least MSE with respect to a ground truth image is chosen to be the correct result. Such a method requires knowledge of the ground truth, and hence is infeasible in practice. Other alternatives include visual inspection or cross-validation. However none of these techniques are fully practical. Instead, we propose a method to choose λ_1 based on sound ⁴¹⁵ statistical principles pertaining to the Poisson or the Poisson-Gaussian noise model. The method is shown here in conjunction with the rescaled non-linear least squares method, however in principle, it can be used with any data fidelity term. For the Poisson-Gaussian noise model, the cost function is given by $J(\boldsymbol{\theta}) = \sum_{i=1}^{m} \frac{(y_i - I_0 e^{(-\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i})^2}{I_0 e^{(-\boldsymbol{\Phi} \boldsymbol{\Psi} \boldsymbol{\theta})_i + \sigma^2}} + \lambda_1 \|\boldsymbol{\theta}\|_1.$

Let *m* denote the total number of bins, θ_{opt} the reconstruction with optimal $\lambda_1 = \lambda_{1.opt}$. The measurements were based on equidistant angle sampling. Let $a_i \triangleq I_o e^{-[\Phi \Psi \theta_{opt}]_i}$. Clearly, we have $Var(y_i) = a_i + \sigma^2$. Hence we can state that $E[\sum_{i=1}^m (y_i - a_i)^2/(a_i + \sigma^2)] = m$. Furthermore, our simulations (Fig. 12) have shown that

$$E\Big(\|(\boldsymbol{y}-I_0e^{-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}_{opt}})\otimes\sqrt{I_0e^{-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}_{opt}}+\sigma^2}\|_2\Big)\approx\sqrt{m}$$
(24)

where \oslash denotes element-wise division. We also observed that the variance of the quantity $\|(\boldsymbol{y} - I_0 e^{-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}_{opt}}) \oslash \sqrt{I_0 e^{-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}_{opt}} + \sigma^2}\|_2$ is very small. This is illustrated in Fig. 12, which shows that the variance of $R = \sqrt{\sum_{i=1}^m \frac{(y_i - I_0 e^{-[\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}]_i})^2}{\sigma^2 + I_0 e^{-[\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}]_i}}}$ is very small compared to its mean. The expected value of R varies with the number of measurements (is equal to \sqrt{m}), and is independent of I_0 . Hence we conclude that the quantity R should be as close to \sqrt{m} as possible. Therefore, we consider

$$D = \operatorname{abs}\left(\left\|\left(\boldsymbol{y} - I_0 e^{-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}_{opt}}\right) \otimes \sqrt{\left(I_0 e^{-\boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{\theta}_{opt}} + \sigma^2\right)}\right\|_2 - \sqrt{m}\right)$$
(25)

and observe how D and relative MSE of reconstructions vary for different values of λ_1 . At a value of λ_1 close to the optimal one, D must achieve its minimum. The test image (154 × 154) and the reconstructions are shown in Figure 14. For these reconstructions, 410 projection views were chosen and Gaussian noise =



Figure 12: Mean and variance of the data-fidelity term $R = \sqrt{\sum_{i=1}^{m} \frac{(y_i - I_0 e^{-[\Phi \Psi \theta]_i})^2}{\sigma^2 + I_0 e^{-[\Phi \Psi \theta]_i}}}$ for different number of measurements (projection views) and beam strength I_0 . (a) Expected value of R exactly coincides with \sqrt{m} , (b) Variance of R is insignificant for any number of measurements, (c) mean of R is approximately independent of beam strength and very close to \sqrt{m} (here m was 8649), and (d) Variance of R is insignificant for all I_0 values.

0.3% was added to the measurements. The dose of X-rays resulted in a Poisson ⁴²⁵ NSR of 0.018. As shown in Fig. 13, the λ_1 for which D and relative MSE are minimum, are very close. In a real-life setting, when relative MSE cannot be computed because of absence of ground-truth, a brute force search needs to be done followed by selecting the value of λ_1 that minimizes D.

6.2. Selection of weightage for object-prior term

430

The weightage for the object prior, λ_2 term needs to be chosen omnisciently for every dataset. We observed that for a large range of values from 700 to 1200



Figure 13: A method to choose the parameter λ_1 in low-dose reconstruction: We expect D to be minimum at approximately the same λ_1 for which relative MSE is minimum. Here, the λ_1 for which D and relative MSE are minimum are very close. Refer to Fig. 14 to observe the reconstruction results for different values of λ_1 .

for okra dataset and from 400 to 700 for sprouts dataset, there was no significant effect on the reconstructions. Lower values indicate that the reconstructions are primarily guided by the measurements, and higher values will strengthen the effect of the prior.

7. Conclusions

435

In the low-dose CT imaging regime, the noise in the measurements becomes significant and needs to be accounted for during the reconstruction. Two new techniques: Poisson-Gaussian convolution and rescaled non-linear least squares (RNLLS) were presented and extensively compared with many of the existing methods. RNLLS was further used in low-dose reconstruction for longitudinal studies to specifically detect new regions in the test and simultaneously reduce noise in the other reconstructed regions. The results were validated on both 2D and 3D biological data. We demonstrated that the reconstructions of the regions

of new changes can be significantly improved by re-irradiating these specific regions by standard-dose X-rays. Further, different methods for choosing the parameters λ_1, λ_2 were also discussed, which has not been dealt with in the literature. Another interesting avenue of research is to consider the case of



Figure 14: Colon test data and its reconstructions for different values of λ_1 . *D* is minimum for $\lambda_1 = 1.2$, shown in green, with a relative MSE of 0.1691. The reconstruction for $\lambda_1 = 2$, shown in red, gives the minimum relative MSE of 0.1501.

tomographic reconstruction from a sparse set of projections (as opposed to the
dense angle sampling considered in this paper), all acquired under low dosage of
radiation. Our technique can possibly be extended to the case where templates
of a similar class of objects are available, as against previous scans of the same
object. This may further increase the utility of the technique in clinical settings.

References

- [1] D. J. Brenner, E. J. Hall, Computed tomography an increasing source of radiation exposure, New England Journal of Medicine 357 (22) (2007) 2277–2284.
 - [2] M. Howells, T. Beetz, H. Chapman, C. Cui, J. Holton, C. Jacobsen, J. Kirz, E. Lima, S. Marchesini, H. Miao, D. Sayre, D. Shapiro, J. Spence, D. Starodub, An assessment of the resolution limitation due to radiation-damage in X-ray

- diffraction microscopy, Journal of Electron Spectroscopy and Related Phenomena 170 (1) (2009) 4 – 12.
 - [3] J. C. Dainty, R. Shaw, Image Science: principles, analysis and evaluation of photographic-type imaging processes, Academic Press, 1974.
- [4] J. Xu, B. M. W. Tsui, Electronic noise modeling in statistical iterative reconstruction, IEEE Transactions on Image Processing 18 (6) (2009) 1228–1238.

465

460

- [5] H. Zhang, Y. Dong, Q. Fan, Wavelet frame based Poisson noise removal and image deblurring, Signal Processing 137 (2017) 363 – 372.
- [6] O. Barkan, J. Weill, S. Dekel, A. Averbuch, A mathematical model for adaptive computed tomography sensing, IEEE Transactions on Computational Imaging 3 (4) (2017) 551–565.
- [7] A. Fischer, T. Lasser, M. Schrapp, J. Stephan, P. B. Nol, Object specific trajectory optimization for industrial X-ray computed tomography, Scientific Reports 6 (19135).
- [8] A. Dabravolski, K. J. Batenburg, J. Sijbers, Dynamic angle selection in x-ray
 ⁴⁷⁵ computed tomography, Nuclear Instruments and Methods in Physics Research
 Section B: Beam Interactions with Materials and Atoms 324 (2014) 17 24, 1st
 International Conference on Tomography of Materials and Structures.
 - [9] X. Pan, E. Y. Sidky, M. Vannier, Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction?, Inverse problems 25 (12).
 - [10] Philips, Philips IMR offers new capabilities to simultaneously re-CTradiation duce and enhance image quality, https://www. philips.com/a-w/about/news/archive/standard/news/press/2013/ 20130617-Philips-IMR-offers-new-capabilities.html (news-article).
- [11] M. Fujita, T. Higaki, Y. Awaya, T. Nakanishi, Y. Nakamura, F. Tatsugami, Y. Baba, M. Iida, K. Awai, Lung cancer screening with ultra-low dose ct using full iterative reconstruction, Japanese Journal of Radiology 35 (4) (2017) 179–189.

470

- [12] K. Kim, G. El Fakhri, Q. Li, Low-dose CT reconstruction using spatially encoded nonlocal penalty, Medical physics 44 (10) (2017) e376–e390.
- ⁴⁹⁰ [13] Q. Lyu, D. Ruan, J. Hoffman, R. Neph, M. McNitt-Gray, K. Sheng, Iterative reconstruction for low dose CT using plug-and-play alternating direction method of multipliers (ADMM) framework, SPIE Medical Imaging 10949.
 - [14] X. Zheng, Z. Lu, S. Ravishankar, Y. Long, J. A. Fessler, Low dose CT image reconstruction with learned sparsifying transform, in: 2016 IEEE 12th Image, Widee and Markidian engine al Simul Discoursing Workshop 2016, pp. 1-5.
- ⁴⁹⁵ Video, and Multidimensional Signal Processing Workshop, 2016, pp. 1–5.
 - [15] X. Jia, Z. Bian, J. He, Y. Wang, J. Huang, D. Zeng, Z. Liang, J. Ma, Texturepreserved low-dose CT reconstruction using region recognizable patch-priors from previous normal-dose CT images, in: IEEE Nuclear Science Symposium, 2016, pp. 1–4.
- 500 [16] C. Zhang, T. Zhang, M. Li, C. Peng, Z. Liu, J. Zheng, Low-dose CT reconstruction via L1 dictionary learning regularization using iteratively reweighted least-squares, BioMedical Engineering OnLine 15 (1) (2016) 66.
 - [17] A. Andersen, A. Kak, Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm, Ultrasonic Imaging 6 (1) (1984) 81 - 94.
- 505
- [18] Q. Xu, H. Yu, X. Mou, L. Zhang, J. Hsieh, G. Wang, Low-dose X-ray CT reconstruction via dictionary learning, IEEE Transactions on Medical Imaging 31 (9) (2012) 1682–1697.
- [19] J. C. Park, B. Song, J. S. Kim, S. H. Park, H. K. Kim, Z. Liu, T. S. Suh, W. Y.
 Song, Fast compressed sensing-based CBCT reconstruction using barzilai-borwein formulation for application to on-line IGRT, Med. Phy. 39 (3) (2012) 1207–1217.
 - [20] J. Wu, F. Liu, L. Jiao, X. Wang, Multivariate pursuit image reconstruction using prior information beyond sparsity, Signal Processing 93 (6) (2013) 1662 – 1672, special issue on Machine Learning in Intelligent Image Processing.
- 515 [21] J. Wang, J. Ma, B. Han, Q. Li, Split Bregman iterative algorithm for sparse reconstruction of electrical impedance tomography, Signal Processing 92 (12) (2012) 2952 - 2961.

[22] J. Shtok, M. Elad, M. Zibulevsky, Learned shrinkage approach for low-dose reconstruction in computed tomography, International Journal of Biomedical Imaging 2013.

520

530

535

- [23] D. Wu, K. Kim, G. El Fakhri, Q. Li, Iterative low-dose CT reconstruction with priors trained by artificial neural network, IEEE Transactions on Medical Imaging 36 (12) (2017) 2479–2486.
- [24] H. Shan, A. Padole, F. Homayounieh, U. Kruger, R. D. Khera, C. Nitiwarangkul,
 M. K. Kalra, G. Wang, Competitive performance of a modularized deep neural network compared to commercial algorithms for low-dose CT image reconstruction, Nature Machine Intelligence 1 (6) (2019) 269–276.
 - [25] C. Gong, L. Zeng, Adaptive iterative reconstruction based on relative total variation for low-intensity computed tomography, Signal Processing 165 (2019) 149 – 162.
 - [26] G.-H. Chen, J. Tang, S. Leng, Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets, Medical Physics 35 (2) (2008) 660–663.
 - [27] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, SIAM Journal on Imaging Sciences 2 (1) (2009) 183–202.
 - [28] W. Hou, C. Zhang, A compressed sensing approach to low-radiation CT reconstruction, 2014 9th International Symposium on Communication Systems, Networks Digital Sign (CSNDSP) (2014) 793–797.
 - [29] K. Koh, S.-J. Kim, S. Boyd, l1-ls: Simple matlab solver for l1-regularized least squares problems, https://stanford.edu/~boyd/l1_ls/ (solver).
 - [30] L. Fu, et al., Comparison between pre-log and post-log statistical models in ultralow-dose CT reconstruction, IEEE transactions on medical imaging 36 (3) (2016) 707 – 720.
- [31] L. Feldkamp, L. C. Davis, J. Kress, Practical cone-beam algorithm, J. Opt. Soc.
 Am 1 (1984) 612–619.
 - 32

- [32] F. Pontana, A. Duhamel, J. Pagniez, T. Flohr, J.-B. Faivre, A.-L. Hachulla, J. Remy, M. Remy-Jardin, Chest computed tomography using iterative reconstruction vs filtered back projection (part 2): image quality of low-dose CT examinations in 80 patients, European Radiology 21 (3) (2011) 636–643.
- [33] H. Wang, B. Tan, B. Zhao, C. Liang, Z. Xu, Raw-data-based iterative reconstruction versus filtered back projection: image quality of low-dose chest computed tomography examinations in 87 patients, Clin. Imaging 37 (6) (2013) 1024 – 32.
 - [34] H. Koyama, Y. Ohno, M. Nishio, S. Matsumoto, N. Sugihara, T. Yoshikawa, S. Seki, K. Sugimura, Iterative reconstruction technique vs filter back projection: utility for quantitative bronchial assessment on low-dose thin-section MDCT in patients with/without chronic obstructive pulmonary disease, European Radiology 24 (8) (2014) 1860–1867.

555

560

565

570

575

- [35] M. J. Willemink, P. B. Noël, The evolution of image reconstruction for CT from filtered back projection to artificial intelligence, European Radiology 29 (5) (2019) 2185–2195.
- [36] Article, Likelihood, https://en.wikipedia.org/wiki/Likelihood_function (in Wikipedia).
- [37] Q. Xie, D. Zeng, Q. Zhao, D. Meng, Z. Xu, Z. Liang, J. Ma, Robust low-dose CT sinogram preprocessing via exploiting noise-generating mechanism, IEEE Transactions on Medical Imaging 36 (12) (2017) 2487–2498.
- [38] J. A. Fessler, Penalized weighted least-squares image reconstruction for positron emission tomography, IEEE Transactions on Med. Imaging 13 (2) (1994) 290–300.
- [39] Q. Ding, Y. Long, X. Zhang, J. A. Fessler, Statistical image reconstruction using mixed Poisson-Gaussian noise model for X-ray CT, submitted in Inverse Prob. and Imaging, axXivSubmitted.
- [40] T. Hastie, R. Tibshirani, M. Wainwright, Statistical Learning with Sparsity, CRC Press, Taylor & Francis Group, 2015.
- [41] E. Chouzenoux, A. Jezierska, J. Pesquet, H. Talbot, A convex approach for image restoration with exact poisson–gaussian likelihood, SIAM Journal on Imaging Sciences 8 (4) (2015) 2662–2682.

- [42] Walnut, Walnut CT dataset, https://www.uni-muenster.de/Voreen/download/ workspaces_and_data_sets.html (uni-muenster).
- [43] Colon, CT Colonography, https://idash.ucsd.edu/data-collections (idash.ucsd.edu).
- 580 [44] Pelvis, Pelvis CT dataset, https://medicine.uiowa.edu/mri/ facility-resources/images/visible-human-project-ct-datasets (medicine.uiowa.edu).
 - [45] Shoulder, Shoulder CT dataset, https://medicine.uiowa.edu/mri/ facility-resources/images/visible-human-project-ct-datasets

(medicine.uiowa.edu).

585

- [46] Supplementary material, videos, more results, https://www.cse.iitb.ac.in/ ~ajitvr/LowDose_Supplemental/ (video, proofs and comparison of methods).
- [47] E. A. Rashed, H. Kudo, Probabilistic atlas prior for CT image reconstruction, Computer Methods and Programs in Biomedicine 128 (2016) 119 – 136.
- ⁵⁹⁰ [48] P. Gopal, S. Chandran, I. D. Svalbe, A. Rajwade, Learning from past scans: Tomographic reconstruction to detect new structures, CoRR abs/1812.10998.
 - [49] F. J. Anscombe, The transformation of poisson, binomial and negative-binomial data, Biometrika 35 (3/4) (1948) 246–254.
 - [50] J. H. Curtiss, On transformations used in the analysis of variance, Ann. Math. Statist. 14 (2) (1943) 107–122.
 - [51] F. Murtagh, J. luc Starck, A. Bijaoui, Image restoration with noise suppression using a multiresolution support, Astronomy and Astrophysics, Suppl. Ser 112 (1995) 179–189.
 - [52] R. C. Sprinthall, Basic Statistical Analysis, Pearson Education, 2011.
- [53] J. Liu, Y. Hu, J. Yang, Y. Chen, H. Shu, L. Luo, Q. Feng, Z. Gui, G. Coatrieux, 3D feature constrained reconstruction for low dose CT imaging, IEEE Transactions on Circuits and Systems for Video Technology PP (99) (2016) 1–1.