Spectrum Reconstruction from Dose Measurements as a Linear Inverse Problem

Benjamin Armbruster†, Russell J Hamilton‡ and Arthur K Kuehl¶

† Department of Mathematics, University of Arizona, Tucson AZ 85721, USA
‡ Department of Radiation Oncology, University of Arizona, Tucson AZ 85724, USA
¶ Department of Materials Science and Engineering, University of Arizona, Tucson AZ 85721, USA

E-mail: rjh@email.arizona.edu

Abstract. There are three ways to determine the spectrum of a clinical photon beam: direct measurement, modeling the source, and reconstruction from ion-chamber measurements. We focus on reconstruction because the necessary equipment is readily available and it provides independent confirmation of source models for a given machine. Reconstruction methods involve measuring the dose in an ion chamber after the beam passes through an attenuator. We gain information about the spectrum from measurements using attenuators of differing compositions and thicknesses since materials have energy dependent attenuation. Unlike the procedures used in other papers, we do not discretize or parameterize the spectrum. With either of these two approximations, reconstruction is a least squares problem. The forward problem of going from a spectrum to a series of dose measurements is a linear operator, with the composition and thickness of the attenuators as parameters. Hence the singular value decomposition (SVD) characterizes this operator. The right singular vectors form a basis for the spectrum, and at first approximation, only those corresponding to singular values above a threshold are measurable. A more rigorous error analysis shows with what confidence different components of the spectrum can be measured. We illustrate this theory with simulations and an example utilizing six sets of dose measurements with water and lead as attenuators.

1. Introduction

Knowledge of the photon energy spectrum is necessary for accurate dose calculation. Sophisticated models of the accelerator head are used to deduce the emitted spectrum by considering interactions with the target, collimators, jaws, etc. The most accurate results are found using Monte Carlo simulation (c.f. Rogers et al. 1995). However, this is only a simulation. An accurate way to experimentally determine the spectrum is with a Compton scatter spectrometer (c.f. Landry & Anderson 1991). These however are prohibitively expensive for most clinics.

The reconstruction of a spectrum from a set of ion-chamber measurements is a third way to determine the photon energy spectrum. It complements physical models

of various accelerator types by providing simple means to estimate the spectrum for a particular linear accelerator and verify the model. Joseph (1975) was the first to address this problem. Approaches appearing in the literature include Laplace transforms (Baird 1981), simulated annealing (Nisbet et al. 1998), iterative least squares (Huang et al. 1982), matrix inversion (Francois et al. 1993) and parameterization (Baker et al. 1995, Hinson & Bourland 2002).

The contribution of this analysis, is to place spectrum reconstruction in the framework of linear inverse problems. The key is that a linear operator relates the spectrum to the measurements (even before any discretizations or approximations). Hence we can apply the theory directly by treating the spectrum as a vector in Hilbert space. Furthermore, we derive the sensitivity of the reconstruction from a model of measurement errors.

2. Reconstruction

2.1. Formal Problem Statement

Typically an ion chamber is placed on the central axis of the beam, a large distance from the source, with a known thickness and composition of attenuating material placed in the beam path. The *forward problem* relates the spectrum, $\Phi(E)$, to the chamber measurement, m, which has units of charge per monitor unit (MU):

$$m = \int \Phi(E)\kappa(E)\mathrm{d}E \tag{1}$$

where $\kappa(E) = N_D^{-1} E S_{a,w} \mu_{en}^{cap}(E) \rho_{cap}^{-1} \exp[-\mu(E)x - \mu^{cap}(E)x^{cap}]$. We define $\Phi(E)\mathrm{d}E$ to be the number of photons with energy in $[E, E+\mathrm{d}E]$ per monitor unit per cm^2 of the collimated beam's cross-sectional area. Here $\mu(E)$ is the attenuation coefficient for photons with energy E and x the thickness of the attenuator. μ^{cap} and x^{cap} refer to the attenuation coefficient and thickness of the buildup cap around the ion chamber while μ_{en}^{cap} refers to the absorption coefficient of the buildup cap. $S_{a,w}$ is the air to water stopping power ratio of secondary electrons passing through the chamber. Since, it is close to unity for photon beams with maximum energy greater than 4 MeV (Baker et al), it is a constant for a particular energy spectrum, and there is little variation of the photon beam spectrum with attenuator thickness, we set it to unity. ρ_{cap} is the density of the buildup cap. N_D is the chamber calibration factor relating the dose (energy deposited) to the measured charge. In general N_D is a function of E. However, because it varies by less than 5% (Almond et al. 1999) over the entire range of useful clinical beam energies (0.5-24MeV), is constant for any particular photon spectrum, and the spectrum varies only slightly with increasing attenuator thickness, N_D is considered a constant here. At these energies, $N_D \approx 2.4 \times 10^8$Gy/C for our chamber. For the purpose of this paper, the properties and thickness of the attenuator and buildup cap are known parameters (c.f. Hubbell & Seltzer 1997). We regard $\Phi(E)$ and $\kappa(E)$ as vectors, $\boldsymbol{\Phi}$ and

$\boldsymbol{\kappa}$, in a Hilbert space with a standard inner product:

$$\boldsymbol{a}^*\boldsymbol{b} = \int_0^{E_{max}} \bar{a}(E)b(E)\mathrm{d}E, \tag{2}$$

where E_{max} is the largest photon energy of interest (e.g. we might know that the spectrum is zero above 15MeV). In this notation, $\boldsymbol{a}^*$ is the Hermitian conjugate of $\boldsymbol{a}$ and $\bar{a}(E)$ is the complex conjugate of $a(E)$. We can then write (1) using (2) as $m = \boldsymbol{\kappa}^*\boldsymbol{\Phi}$. For n measurements, possibly using different depths and materials,

$$\boldsymbol{m} = A\boldsymbol{\Phi} = \begin{pmatrix} \boldsymbol{\kappa}_1^* \\ \boldsymbol{\kappa}_2^* \\ \vdots \\ \boldsymbol{\kappa}_n^* \end{pmatrix} \boldsymbol{\Phi}. \tag{3}$$

Here A is a discrete-by-continuous linear operator. The goal of the *inverse problem* is to find (or characterize) $\boldsymbol{\Phi}$ given $\boldsymbol{m}$. Tan and Fox (Tan & Fox n.d.) provide a nice introduction to linear inverse problems.

In contrast to previous papers, we are leaving $\Phi(E)$ as a general function. Previous work (c.f. Huang et al. 1982, Hinson & Bourland 2002) has either discretized or otherwise parameterized $\Phi(E)$ using a finite number of parameters. In the case where $\Phi(E)$ is discretized over l intervals, A is just a n x l matrix and $\boldsymbol{\Phi}$ an l dimensional vector. If $l < \text{rank } A \leq n$, determining $\boldsymbol{\Phi}$ from A is simply a least squares problem. The problem is that the solution is sensitive to the choice of discretization and ceases to be uniquely determined when $l > n$. For example consider the Elekta SLi 6MV spectrum calculated using Monte Carlo simulation (Fippel et al. 2003). Using this spectrum we simulated the forward problem assuming noise free measurements with 78 different attenuators (i.e., $n = 78$ in equation 3): lead attenuators from 1 to 10 cm thick in 0.5 cm increments and water attenuators from 1 to 30 cm thick in 0.5 cm increments. Figure 1 shows two vastly different reconstructions of $\Phi(E)$ from these simulated measurements, neither of which is close to the spectrum we assumed when simulating the measurements.

It is unclear how well the discretized, reconstructed spectrum, $\widetilde{\boldsymbol{\Phi}} \in \mathbb{R}^l$, can approximate $\Phi(E)$. Previous work estimates the accuracy of the reconstruction algorithms using simulation experiments (Catala et al. 1995) or by computing dosimetric beam properties (Nisbet et al. 1998). These methods do not provide an analysis of how the errors propagate. In order to obtain a more accurate reconstruction and an explicit error model, we will not discretize $\Phi(E)$. However, then we cannot ignore the underdetermined nature of the problem by choosing $l < n$. The inverse problem is inherently underdetermined since we are trying to reconstruct a *function* from a *finite* number of measurements.

2.2. Review of SVD in a Hilbert Space

Any matrix, B, has a singular value decomposition

$$B = \sum_i \boldsymbol{u}_i \sigma_i \boldsymbol{v}_i^*.$$

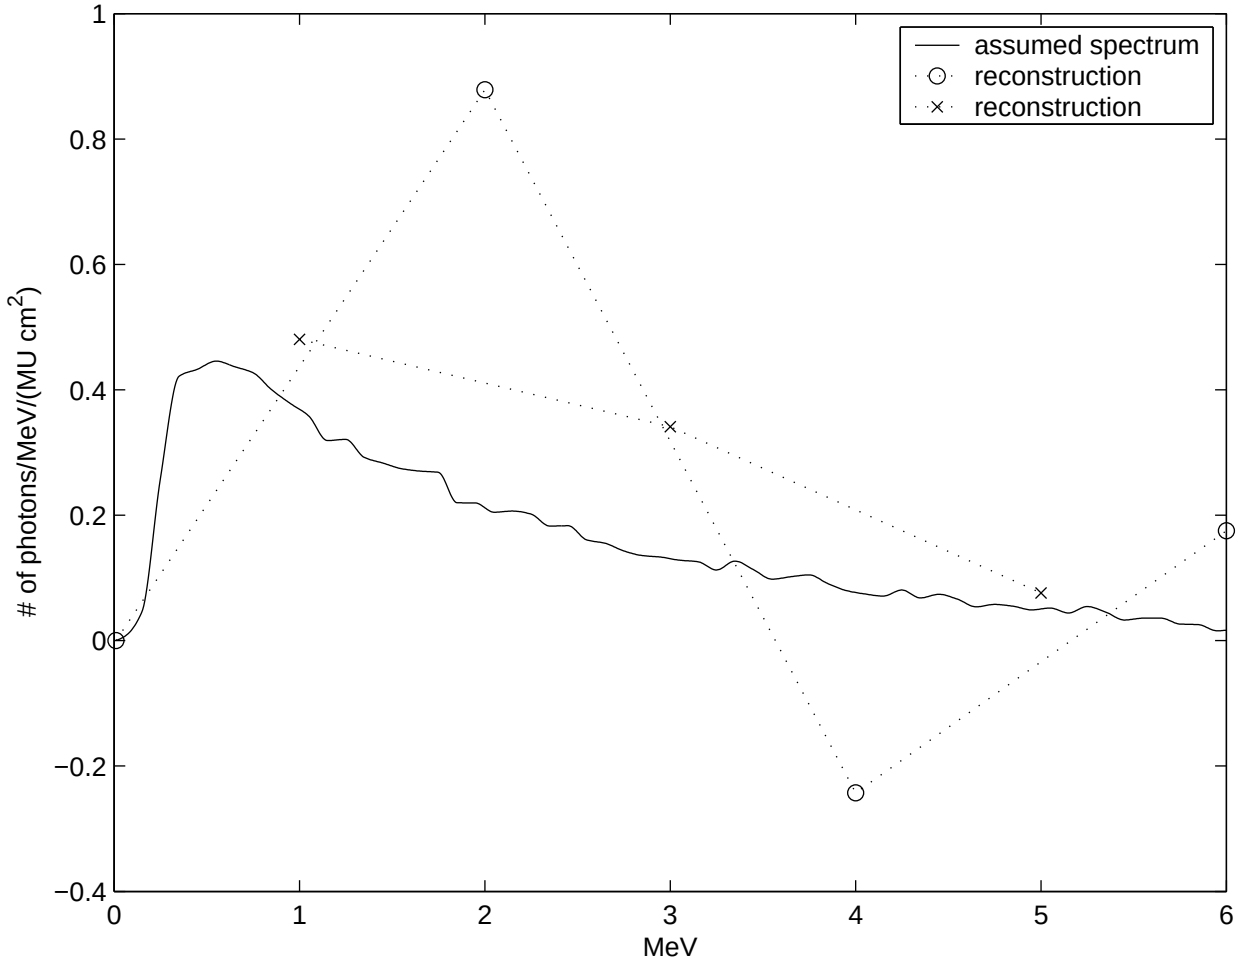

Figure 1. assumed spectrum (Elekta SLi 6MV from Fippel et al. (2003)) and reconstruction using evenly spaced grids of 3 and 4 points.

Here the *left singular vectors* $\{\boldsymbol{u}_i\}$ form an orthonormal basis of the range of B; the *right singular vectors* $\{\boldsymbol{v}_i\}$ form an orthonormal basis of the domain of B; and the singular values $\{\sigma_i\}$ are positive. The index i goes from 1 to the rank of B. It follows that

$$B\boldsymbol{v}_i = \boldsymbol{u}_i/\sigma_i. \tag{4}$$

Note it is easy to check that $BB^* = \sum_i \boldsymbol{u}_i\sigma_i^2\boldsymbol{u}_i^*$: the eigenvectors and eigenvalues of BB^* are $\{(\boldsymbol{u}_i, \sigma_i^2)\}$. In addition,

$$B^*\boldsymbol{u}_i = \sigma_i\boldsymbol{v}_i. \tag{5}$$

For our purposes, it is important that not just matrices but general linear operators (operating on function spaces) have an SVD. Consider our operator A. Note its rank is at most n. Hence its $\{\boldsymbol{v}_i\}$ do not span the entire space of functions. To find its SVD, we examine the n x n matrix AA^*:

$$(AA^*)_{ij} = \boldsymbol{\kappa}_i^*\boldsymbol{\kappa}_j = \int_0^{E_{max}} \kappa_i(E)\kappa_j(E)\mathrm{d}E.$$

(Note $\kappa_i(E)$ is real.) Its eigenvectors and eigenvalues give us the left singular vectors, $\{\boldsymbol{u}_i\}$, and the square of the singular values, $\{\sigma_i^2\}$‡. The right singular vectors (which form a basis of spectra) are $\boldsymbol{v}_i = A^*\boldsymbol{u}_i/\sigma_i = \sum_j \boldsymbol{\kappa}_j(\boldsymbol{u}_i)_j/\sigma_i$.

Note the null space of A has infinite dimension (another manifestation of the fact that this problem is underdetermined). That means for any reconstructed spectrum, $\widetilde{\boldsymbol{\Phi}}$ (such that $A\widetilde{\boldsymbol{\Phi}} = \boldsymbol{m}$), there is a whole family of solutions: $A(\widetilde{\boldsymbol{\Phi}} + \boldsymbol{r}) = \boldsymbol{m}$ where $A\boldsymbol{r} = 0$ (i.e., $\boldsymbol{r}$ is in the null space of A).

2.3. Naive Reconstruction

The general idea of spectrum reconstruction is to write the dose measurements in the basis $\{\boldsymbol{u}_i\}$ (these being the left singular vectors of A),§

$$\boldsymbol{m} = \sum_i c_i \boldsymbol{u}_i. \tag{6}$$

Then using (3) and (4), we divide the coefficients, c_i, by the singular values and use $\{\boldsymbol{v}_i\}$ as the basis for the spectrum:

$$\boldsymbol{\Phi} \approx \sum_i \boldsymbol{v}_i c_i/\sigma_i = \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^*\boldsymbol{m}}{\sigma_i}.$$

Ignoring measurement error for now, this would be exact if not for the null space of A:

$$\boldsymbol{\Phi} = \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^*\boldsymbol{m}}{\sigma_i} + \boldsymbol{r}$$

where $\boldsymbol{r}$ in the null space of A. Since we have no information about the residual $\boldsymbol{r}$ we assume $\boldsymbol{r} = 0$ and so our reconstructed $\boldsymbol{\Phi}$ is

$$\widetilde{\boldsymbol{\Phi}} := \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^*\boldsymbol{m}}{\sigma_i}. \tag{7}$$

These terms define the pseudoinverse of A, a continuous-by-discrete linear operator:

$$\widetilde{\boldsymbol{\Phi}} = A^+\boldsymbol{m} \quad \text{where} \quad A^+ := \sum_i \boldsymbol{v}_i \boldsymbol{u}_i^*/\sigma_i.$$

The relation between the reconstructed spectrum $\widetilde{\boldsymbol{\Phi}}$ and the actual spectrum $\boldsymbol{\Phi}$ is

$$\widetilde{\boldsymbol{\Phi}} = \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^* A\boldsymbol{\Phi}}{\sigma_i} = \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^* \sum_j \boldsymbol{u}_j \sigma_j \boldsymbol{v}_j^* \boldsymbol{\Phi}}{\sigma_i} = \sum_i \boldsymbol{v}_i \boldsymbol{v}_i^* \boldsymbol{\Phi}.$$

This means $\widetilde{\boldsymbol{\Phi}}$ is the projection of $\boldsymbol{\Phi}$ onto the subspace spanned by $\{\boldsymbol{v}_i\}$. Alternative assumptions on $\boldsymbol{r}$ are discussed later in the context of regularization. Figure 2 shows the set of basis functions $\{\boldsymbol{v}_i\}$ for an experiment involving no attenuator and attenuation by water with depths 3.03, 14.58, and 29.38 cm and lead of thicknesses 1.54 and 2.81 cm, using an $E_{max} = 6$MeV in equation 2. Figure 3 is a reconstruction of an assumed spectrum (the Monte Carlo simulated spectrum of Fippel et al. (2003)) by projection onto the subspace spanned by $\{\boldsymbol{v}_i\}$ of figure 2.

‡ For the remainder of the paper $\{\boldsymbol{u}_i\}$, $\{\boldsymbol{v}_i\}$, and $\{\sigma_i\}$ will refer to the SVD of A.

§ If $\{\boldsymbol{u}_i\}$ is not complete, then equation (6) is not an equality but a least squares problem. We ignore this as it does not change the theory nor occur in practice.

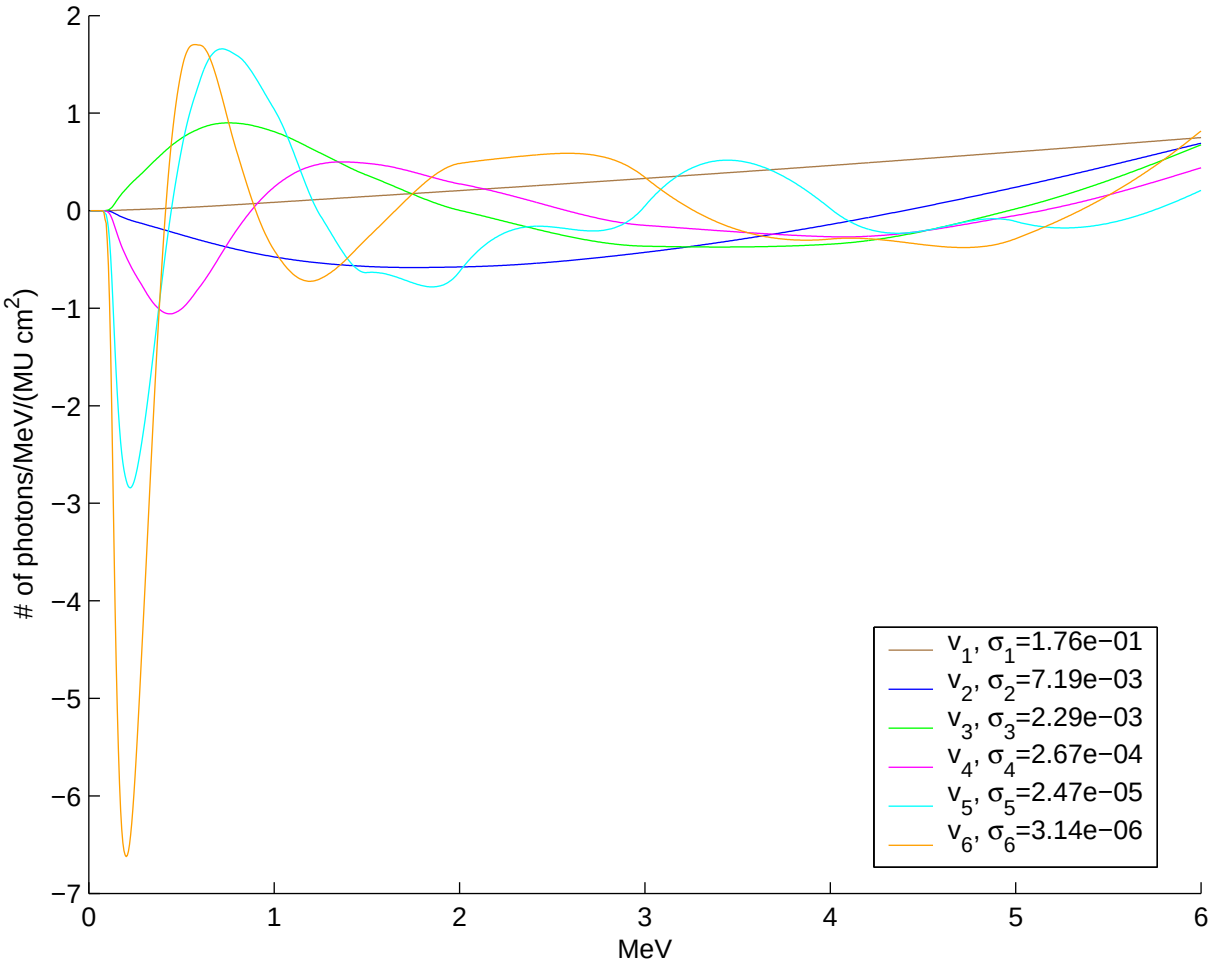

Figure 2. Basis functions.

2.4. Sensitivity Analysis

The projection of $\boldsymbol{\Phi}$ onto the subspace spanned by $\{\boldsymbol{v}_i\}$ in figure 3 gives a misleading impression of the accuracy attainable from the measurements. This is because it collapses the forward and inverse problems into an orthogonal projection, ignoring measurement and numerical error. The third curve in figure 3 reconstructs Φ using equation 7 from simulated measurements with relative error that is normally distributed with mean zero and standard deviation 0.005.

For a large number of measurements, many σ_i are very small. This means that the components of the spectrum $\boldsymbol{\Phi}$ along the corresponding $\boldsymbol{v}_i$ are "damped" to insignificance by the measurement process; they make almost no contribution to the measured values (along the corresponding $\boldsymbol{u}_i$). So when we reverse this process to reconstruct the spectrum from measurements and divide the coefficients of the measured values along $\boldsymbol{u}_i$ by tiny σ_i we only amplify measurement error. In the simulated reconstruction in figure 3, some of the σ_i are very small (see values given in figure 2). These terms are responsible for the poor reconstruction.

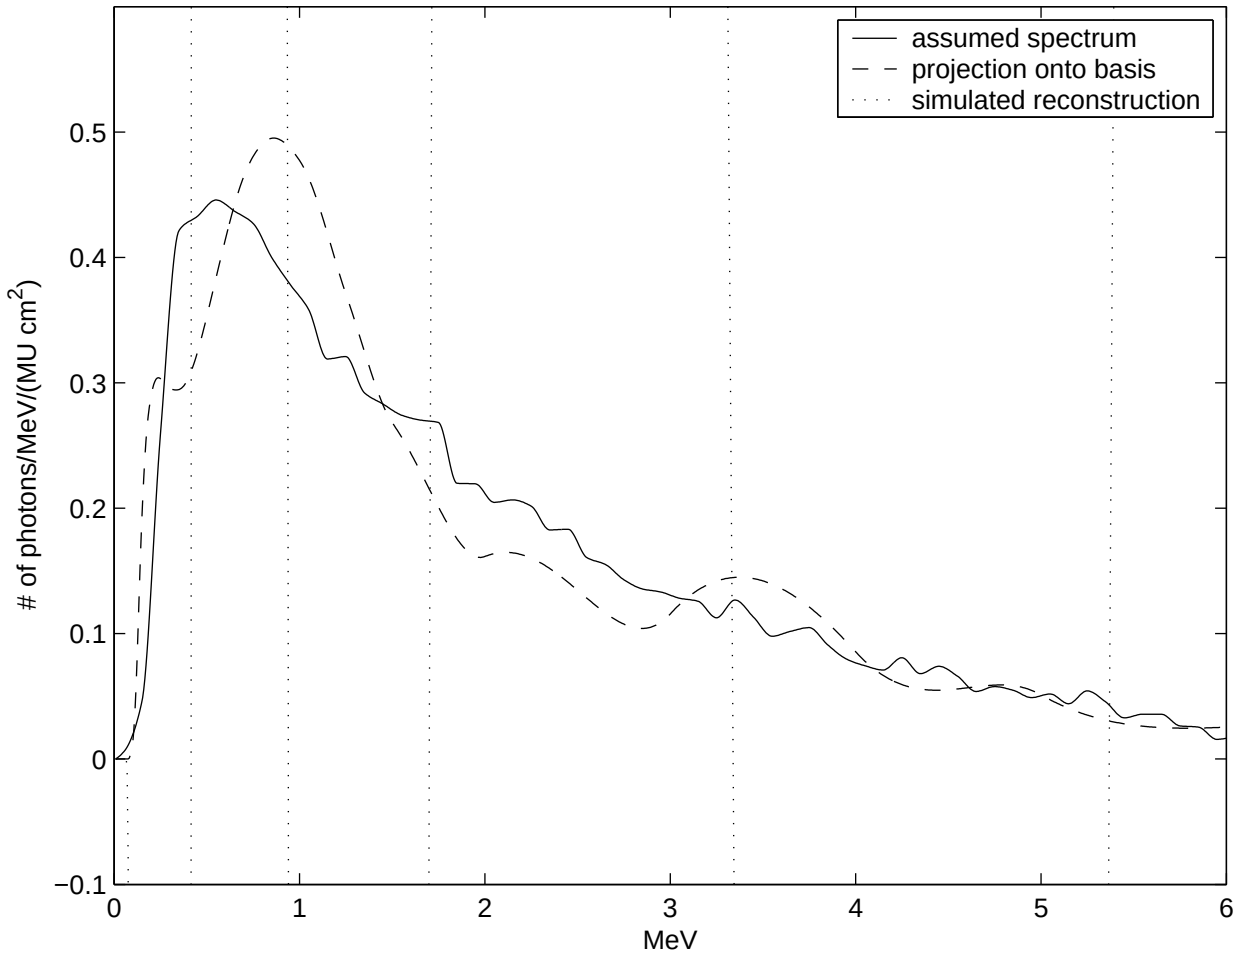

Figure 3. An assumed spectrum (Elekta SLi 6MV from Fippel et al. (2003)), its projection onto a basis defined by a set of measurements, $\{\boldsymbol{v}_i\}$, and a reconstruction from simulated measurements with 0.5% relative error.

The dependence of the reconstructed spectrum on the measurement errors can be made precise by looking at how $\widetilde{\boldsymbol{\Phi}}$ varies with respect to small changes in $\boldsymbol{m}$. This is basically the derivative (or more precisely the Jacobian) of the inverse problem. As our reconstruction is linear in $\boldsymbol{m}$ (7), the Jacobian is A^+. The rate at which $\widetilde{\Phi}(E)$ changes due to variations in measurement i, is A^+_{Ei} (i.e., "row" E and column i of A^+). We are often interested in the worst case: what is the largest change in $\widetilde{\boldsymbol{\Phi}}$ that can result from a small change in $\boldsymbol{m}$? This is the matrix norm $\|A^+\|$ (we will use the 2-norm). Due to properties of the SVD, $\|A^+\| = \sigma_n^{-1}$. If we are interested in relative changes, then the condition number of A^+ is appropriate, σ_1/σ_n. Generally, σ_n decays exponentially with n. So for large n, this condition number is huge and illustrates the ill-conditioned nature of the problem.

Therefore, if σ_1/σ_n is greater than the relative error of the measurements, then the reconstruction may have errors of order 1 (i.e., be completely useless). A solution is to only use in the reconstruction (7) the terms whose σ_i are sufficiently large. So a better

reconstruction would sum only those terms in equation (7) which have values of σ_i above some threshold, δ:

$$\widetilde{\boldsymbol{\Phi}} := \sum_{i\in\mathcal{I}} \boldsymbol{v}_i \frac{\boldsymbol{u}_i^* \boldsymbol{m}}{\sigma_i}$$
$$\mathcal{I} := \{i : \sigma_i > \delta\}. \tag{8}$$

We will denote this linear operator,

$$A_{\mathcal{I}}^{+} := \sum_{i\in\mathcal{I}} \boldsymbol{v}_i \boldsymbol{u}_i / \sigma_i,$$

as it is a generalization of A^+. Due to the above discussion, δ should be not be smaller than the relative error of the measurements. Choosing δ (or equivalently $\mathcal{I}$), is a tradeoff between the information and noise included in the terms summed in the reconstruction and a judgment of when the additional noise swamps or outweighs the additional information.

Figure 4 shows the reconstruction of an assumed spectrum using the same attenuators as in figures 2 and 3 from simulated measurements with 0.5% normally distributed error. The various colors correspond to reconstructions with different number of terms (i.e., different sizes of $\mathcal{I}$). Solid lines show the average reconstruction for a given $\mathcal{I}$ while the dashed lines show $\pm$ 1 standard deviations. (The average and standard deviation were calculated from 15 replicas of each simulated measurement.) Note how the dashed red lines are indistinguishable from the average reconstruction using $\mathcal{I} = \{1\}$. This shows how insensitive that reconstruction is compared to the last reconstruction using $\mathcal{I} = \{1, \ldots, 6\}$. At high photon energies some of the reconstructions in figure 4 are negative or increasing. This does not warrant concern as the large standard deviations signify the reconstruction is very poor there (due to the uniformly small size of the attenuation coefficients, μ, at high photon energies), but it does show the importance of the error analysis.

2.5. Error Model

To see how error is affecting our reconstruction, we must first understand the origin of the error and then incorporate it into our model of the measurement process (1). Errors are often characterized as either absolute or relative. We propose to extend the model of the measurement process to

$$m = (\boldsymbol{\kappa}^* \Phi + \boldsymbol{b}^* \boldsymbol{\Phi})(1 + \epsilon_r) + \epsilon_l \tag{9}$$

where ϵ_r is a relative error term, ϵ_l is an absolute error term, and $\boldsymbol{b}^*\boldsymbol{\Phi}$ is a bias term.

The absolute error term, ϵ_l models the *leakage current* (i.e., random current fluctuation) in the wire connected to the ion chamber. This is clearly an absolute error as the current fluctuations are present even with the beam off. Their magnitude is approximately 2×10^{-14}A. Modeling this as white noise (i.e., assuming independent increments), the error in the measured charge (obtained by integrating white noise)

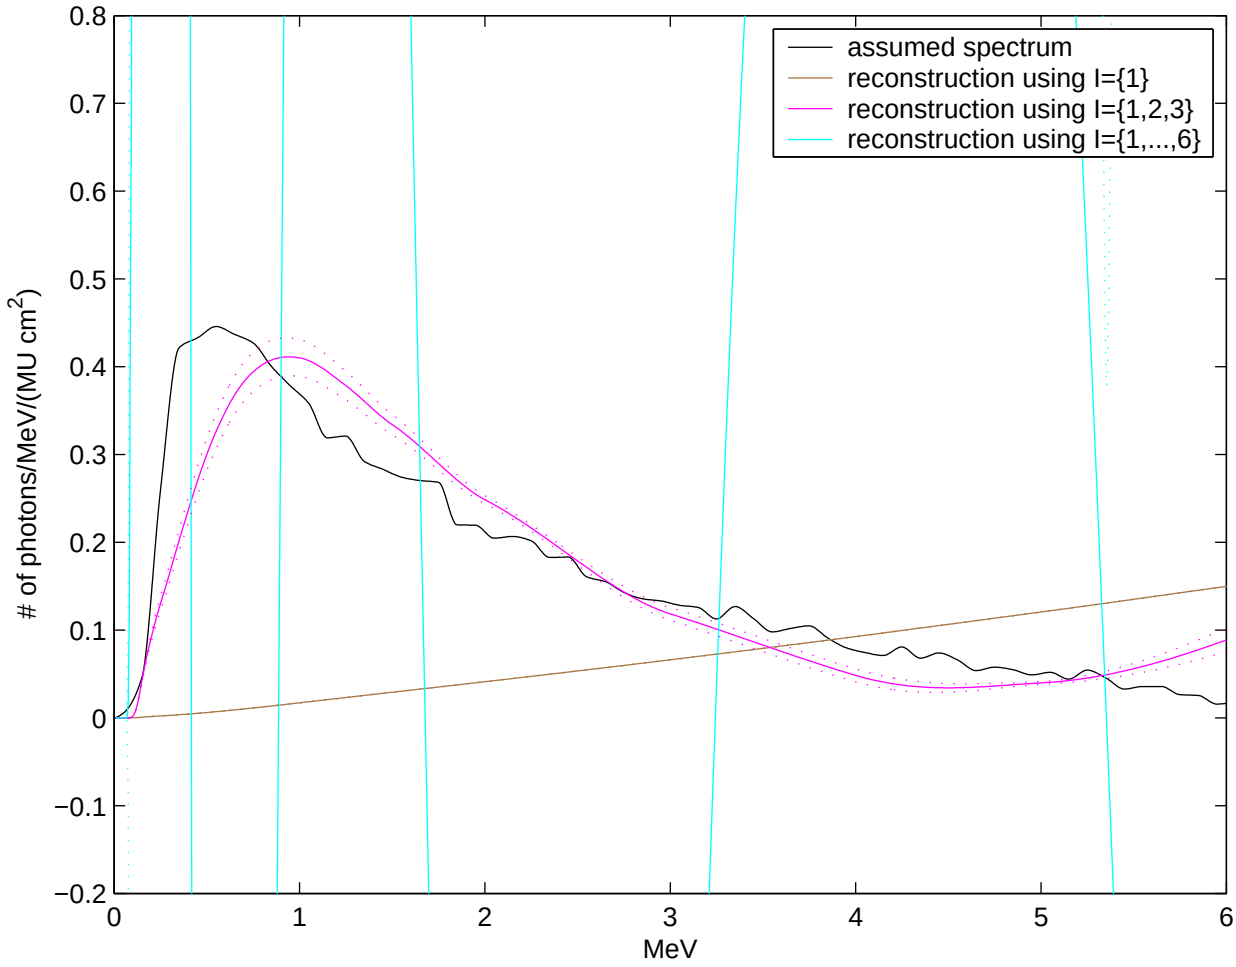

Figure 4. Reconstruction of a spectrum (Elekta SLi 6MV from Fippel et al. (2003)) from simulated measurements with 0.5% relative error using different numbers of basis functions. Dashed lines show ±1 standard deviation.

is a Brownian motion in time. Over t seconds, the error due to the leakage current is $2 \times 10^{-14}\sqrt{t}$ C. The source for our measurements, an Elekta SLi, had a delivery rate of approximately 500 MU/minute. So for a one minute measurement interval, $\epsilon_l \approx 3 \times 10^{-16}$C/MU. By extending the experiment time, ϵ_l can be made arbitrarily small since error increases with the square root of time, while the radiation (measured in MU) increases linearly with time. Nevertheless, we include this term in our model since for the one minute measurements we made, it is of comparable size to other terms, especially those with large attenuation.

The term $\boldsymbol{b}^*\boldsymbol{\Phi}$ in (9) models the *background*: the radiation reaching the detector that is not due to the primary beam. This is mainly radiation that leaks from the machine head. Clearly the background radiation is linear in the total number of photons, and hence linear in $\boldsymbol{\Phi}$. Therefore, we can model it as an inner product with some (Hilbert space) vector $\boldsymbol{b}$. This $\boldsymbol{b}$ may depend on any number of things: amount of head shielding, beam energy etc. Nevertheless, we can determine its effect from the

charge measured when the primary beam points away from the attenuators. We are modeling this deterministically. For the 6MV beam of an Elekta SLi, we measured $\boldsymbol{b}^*\boldsymbol{\Phi} \approx 1.8 \times 10^{-15}$C/MU.

In our measurements using 6 and 15 MV beams and attenuators between 0 and 30 cm, m ranged from 10^{-12} to 10^{-15} C/MU. For the larger m, the 0.1%–1% variation in repeated measurements of the same experimental setup cannot be explained by the absolute error term ϵ_l and justifies a relative error term, ϵ_r in equation 9, between 10^{-2} and 10^{-3}. These errors are due to phenomena such as thermal fluctuations in the machine, that produce variations in the number of photons produced per MU. Hence we assume ϵ_r is independent of the leakage current (and ϵ_l) and only multiplies $\boldsymbol{\kappa}^*\Phi + \boldsymbol{b}^*\boldsymbol{\Phi}$ in equation 9. So for measurements with little attenuation, the relative error ϵ_r dominates while for strongly attenuated measurements, the other two terms become more important (though ϵ_l can be reduced by increasing the time that the beam is on).

2.6. Error Analysis

The above error model (9) can be extended to a set of measurements (like equation 3):

$$\boldsymbol{m} = (I + R)(A + \mathbf{1}\boldsymbol{b}^*)\boldsymbol{\Phi} + \boldsymbol{\epsilon}. \tag{10}$$

We use $\mathbf{1}$ to denote a vector of ones and I to denote the identity matrix. Here R is an n by n diagonal matrix R whose diagonal entries, R_{11} to R_{nn}, are independently identically distributed (i.i.d.) with expectation 0 and represent ϵ_r for each of the n measurements. The components of the n dimensional vector $\boldsymbol{\epsilon}$ are also i.i.d. with expectation 0 and represent ϵ_l for the n measurements. Since R and $\boldsymbol{\epsilon}$ are random variables, the measurements $\boldsymbol{m}$ and the reconstruction $\widetilde{\boldsymbol{\Phi}}$ are too.

For our reconstruction to work we have to subtract out the background term, $\boldsymbol{b}^*\boldsymbol{\Phi}$ so that the expectation equals $A\boldsymbol{\Phi}$ and thereby matches the deterministic case (3). As mentioned in section 2.5, we can take a measurement (or an average of a series of measurements) of only the background,

$$m_b := \boldsymbol{b}^*\boldsymbol{\Phi}(1 + \epsilon_r) + \epsilon_l,$$

where ϵ_r and ϵ_l are random variables with expectation 0. Now we can define

$$\widetilde{\boldsymbol{m}} := \boldsymbol{m} - m_b\mathbf{1} = A\boldsymbol{\Phi} + R(A\boldsymbol{\Phi} + \mathbf{1}\boldsymbol{b}^*\boldsymbol{\Phi}) - \epsilon_r\mathbf{1}\boldsymbol{b}^*\boldsymbol{\Phi} + \boldsymbol{\epsilon} - \epsilon_l\mathbf{1}. \tag{11}$$

Clearly the expectation value of $\widetilde{\boldsymbol{m}}$, $E[\widetilde{\boldsymbol{m}}] = A\boldsymbol{\Phi}$ as desired. So if we redefine our reconstructed spectrum $\widetilde{\boldsymbol{\Phi}}$ to use $\widetilde{\boldsymbol{m}}$ instead of $\boldsymbol{m}$,

$$\widetilde{\boldsymbol{\Phi}} := A^+\widetilde{\boldsymbol{m}} = \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^*\widetilde{\boldsymbol{m}}}{\sigma_i},$$

by linearity of expectations,

$$E[\widetilde{\boldsymbol{\Phi}}] = A^+A\boldsymbol{\Phi} = \sum_i \boldsymbol{v}_i\boldsymbol{v}_i^*\boldsymbol{\Phi}.$$

as in the deterministic case.

We analyze the variance of $\widetilde{\boldsymbol{\Phi}}$ (i.e., the reliability of the reconstruction) to give a more complete picture than that provided by the sensitivity analysis of the inverse problem.

$$\text{Var } [\widetilde{\boldsymbol{\Phi}}] = \text{Var } [A^{+}\widetilde{\boldsymbol{m}}] = (A^{+})^2 \text{Var } [\widetilde{\boldsymbol{m}}]. \tag{12}$$

Here and in the remainder of the paper, squares of matrices, vectors, and operators (such as $(A^{+})^2$) refer to the square of their components. Now from equation 11,

$$\text{Var } [\widetilde{\boldsymbol{m}}] = \text{Var } [R](A\boldsymbol{\Phi} + \mathbf{1}\boldsymbol{b}^{*}\boldsymbol{\Phi})^2 + \text{Var } [\epsilon_r]\mathbf{1}(\boldsymbol{b}^{*}\boldsymbol{\Phi})^2 + \text{Var } [\boldsymbol{\epsilon}] + \text{Var } [\epsilon_l]\mathbf{1}$$

Note $\boldsymbol{m}$ is an estimator for $A\boldsymbol{\Phi} + \mathbf{1}\boldsymbol{b}^{*}\boldsymbol{\Phi}$, and m_b is an estimator for $\boldsymbol{b}^{*}\boldsymbol{\Phi}$. So,

$$\text{Var } [\widetilde{\boldsymbol{m}}] \approx \text{Var } [R]\boldsymbol{m}^2 + \text{Var } [\epsilon_r]\mathbf{1}m_b^2 + \text{Var } [\boldsymbol{\epsilon}] + \text{Var } [\epsilon_l]\mathbf{1}.$$

Substituting this into equation 12,

$$\text{Var } [\widetilde{\boldsymbol{\Phi}}] \approx (A^{+})^2(\text{Var } [R]\boldsymbol{m}^2 + \text{Var } [\epsilon_r]\mathbf{1}m_b^2 + \text{Var } [\boldsymbol{\epsilon}] + \text{Var } [\epsilon_l]\mathbf{1}).$$

Setting Var $[R_{ii}] = 0.005^2$ and ϵ_r, $\boldsymbol{\epsilon}$, ϵ_l to 0, we can reproduce the one standard deviation error bounds in figure 4 as expected.

2.7. Regularization

Recall that we arbitrarily chose the $\boldsymbol{r} = 0$ in our reconstruction (7). In fact we could have chosen any $\boldsymbol{r}$ as long as it is part of the null space of A. Regularization involves choosing $\boldsymbol{r}$ such that

$$\widetilde{\boldsymbol{\Phi}} = \sum_i \boldsymbol{v}_i \frac{\boldsymbol{u}_i^{*}\boldsymbol{m}}{\sigma_i} + \boldsymbol{r}$$

looks more plausible. Note that regularization does not reveal any additional information from our measurements, but is rather a way to include outside information in our reconstruction. For example if we happened to choose by accident or using outside knowledge, $\boldsymbol{r} := \boldsymbol{\Phi} - \sum_i \boldsymbol{v}_i\boldsymbol{v}_i^{*}\boldsymbol{\Phi}$ then our reconstruction would be perfect: $\widetilde{\boldsymbol{\Phi}} = \boldsymbol{\Phi}$. A more plausible scenario is that we make a guess, $\boldsymbol{\Phi}_0$, of $\boldsymbol{\Phi}$ and choose $\boldsymbol{r} := \boldsymbol{\Phi}_0 - \sum_i \boldsymbol{v}_i\boldsymbol{v}_i^{*}\boldsymbol{\Phi}_0$. Monte Carlo simulations of linear accelerators (Sheikh-Bagheri & Rogers 2002) show that the peak of the spectra, $\Phi(E)$, is usually at 1–3 MeV. So we could choose $\boldsymbol{\Phi}_0$ to have this property. This type of regularization is called *truncated singular value decomposition* (TSVD). Regularization is a broad topic and we have only scratched the surface (c.f. chapter 3 in Tan & Fox (n.d.) for other regularization methods). However, as we do not know of any significant outside information to use in the reconstruction our simple TSVD regularization should suffice. Figure 5 shows the reconstructions of figure 4 regularized using a plausible $\boldsymbol{\Phi}_0$, a piecewise linear function with a maximum at 1 MeV. Another way to make our reconstructions look more plausible is to constrain them to be nonnegative with a single maximum. $\boldsymbol{r}$ is then chosen accordingly.

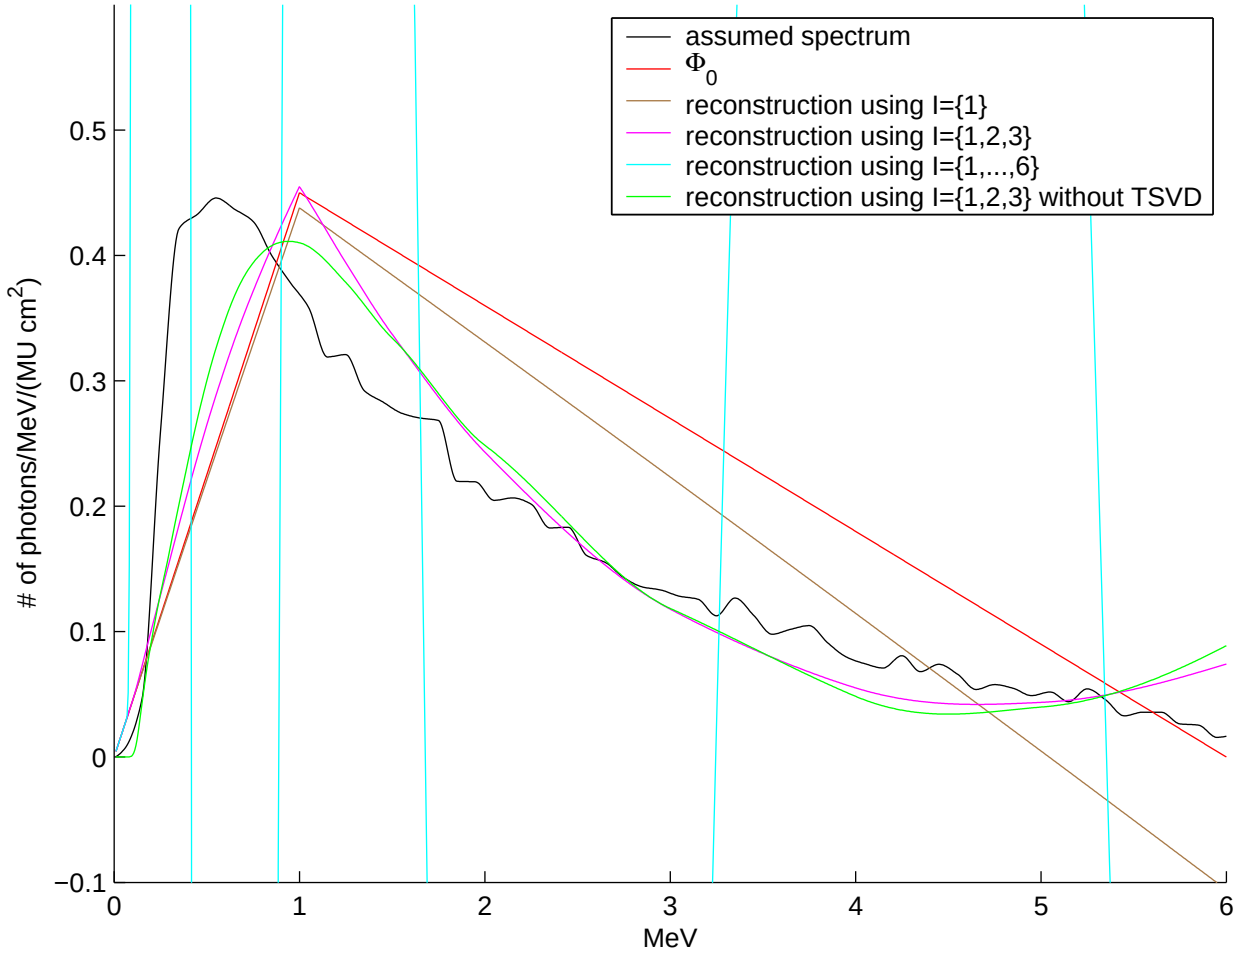

Figure 5. Reconstruction of a spectrum (Elekta SLi 6MV from Fippel et al. (2003)) with TSVD regularization.

2.8. Example

For the Elekta SLi we took two sets of measurements: one with the 6MV beam and one with the 15MV beam. The gantry was rotated 90 degrees. The attenuating material was placed on the couch 1m from the source (SSD=1m). A Scanditronix/Wellhöfer CC13 compact chamber (0.13 cm^3 volume, 5.8 mm active length) was placed perpendicular to the beam 478.5 cm from the source inside a "German Silver" buildup cap (density 8.9 g/cm^3, 57% Cu, 29% Zn, 12% Ni, 19% Pb, $\leq$ 0.7% Mn, $\leq$ 0.3% Fe) with a radius of 15.6mm for 6MV and 30mm for 15MV. A field size of 1cm $\times$ 1cm was used. This was sufficient for the beam to encompass the entire active area of the chamber. Exposures using between 300 and 1000 MU were performed at a dose rate of 500 MU/min. For 6MV, attenuators of 3.03 cm, 14.58 cm, and 29.38 cm of waters and 1.54 cm and 2.81 cm of lead were used. For 15MV, attenuators of 14.50 cm, 28.34 cm, and 3.25 cm of waters and 1.54 cm, 2.81 cm, 10.22 cm, and 5.05 cm of lead were used. In addition measurements with no attenuators were made for both the 6 MV and 15 MV beams. A

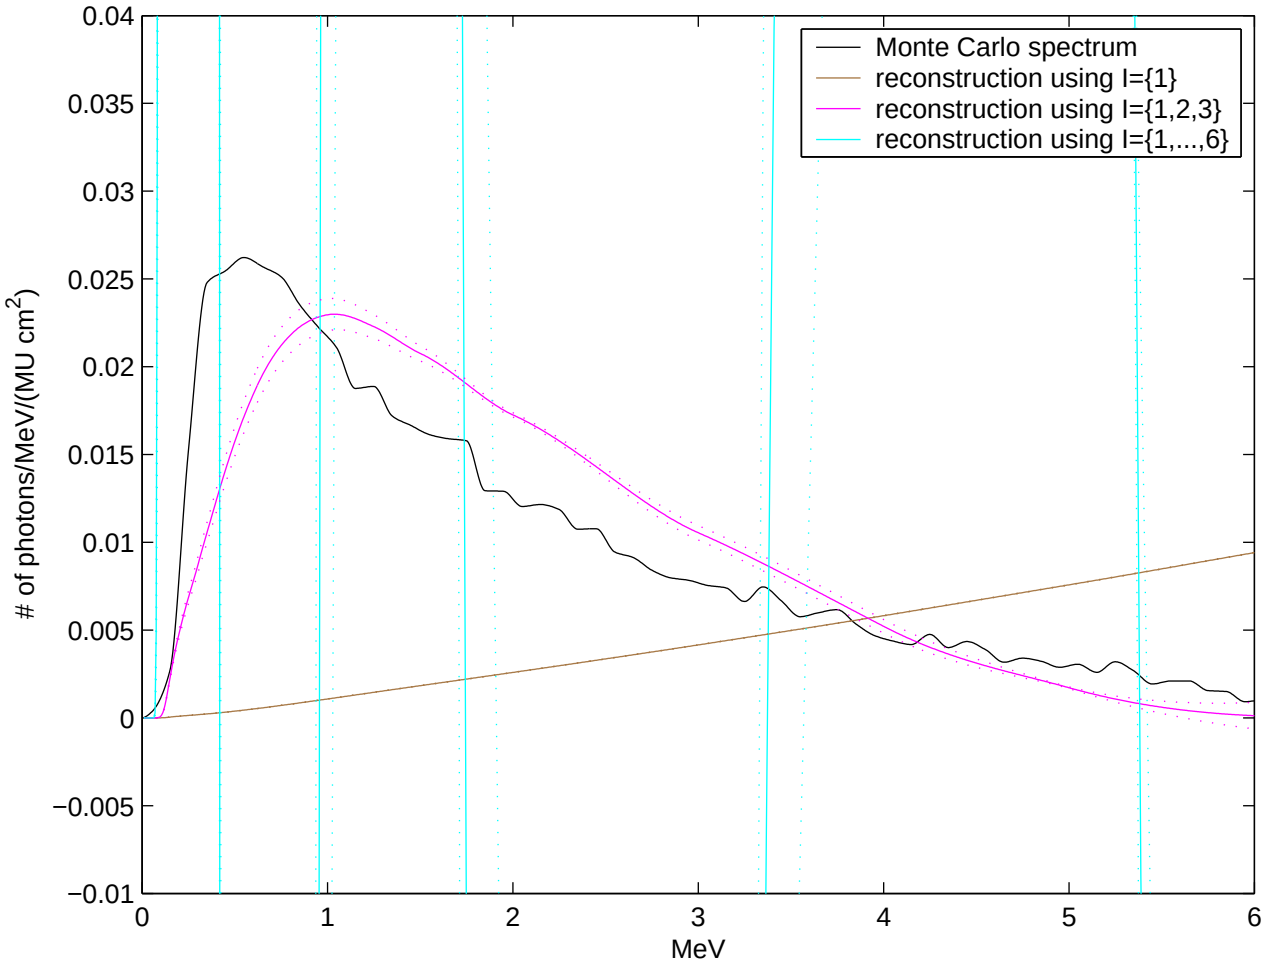

Figure 6. Reconstruction of the spectrum of a Elekta SLi 6 MV beam using different numbers of basis functions. Dashed lines show ±1 standard deviation. The Monte Carlo spectrum is from Fippel et al. (2003).

minimum of three measurements were made for each configuration. A leakage current of 2×10^{-14}A was found. The background radiation was measured to be 1.8×10^{-15}C and 7.3×10^{-16}C for 6 MV and 15 MV beams, respectively.

In figure 6 we compare the reconstruction from our measurements for the 6MV beam to an estimate of the spectrum for the same beam using Monte Carlo simulation of a physical model (Fippel et al. 2003). The spectrum calculated in Fippel et al. (2003) is normalized, so only the shape of the reconstruction can be judged. In figure 6 the curve labeled *Monte Carlo spectrum* is scaled to best match the reconstructions. Figure 7 tells a similar story for the 15 MV beam. Even though two more measurements were taken for the 15 MV beam than for the 6 MV beam, the singular values decay faster. A cutoff for the singular values (equation 8) of $\delta = 10^{-2}$ works well for both reconstructions.

These reconstructions result from only a few sets of measurements, six for the 6 MV beam and eight for the 15 MV beam. Some of the discrepancy between the

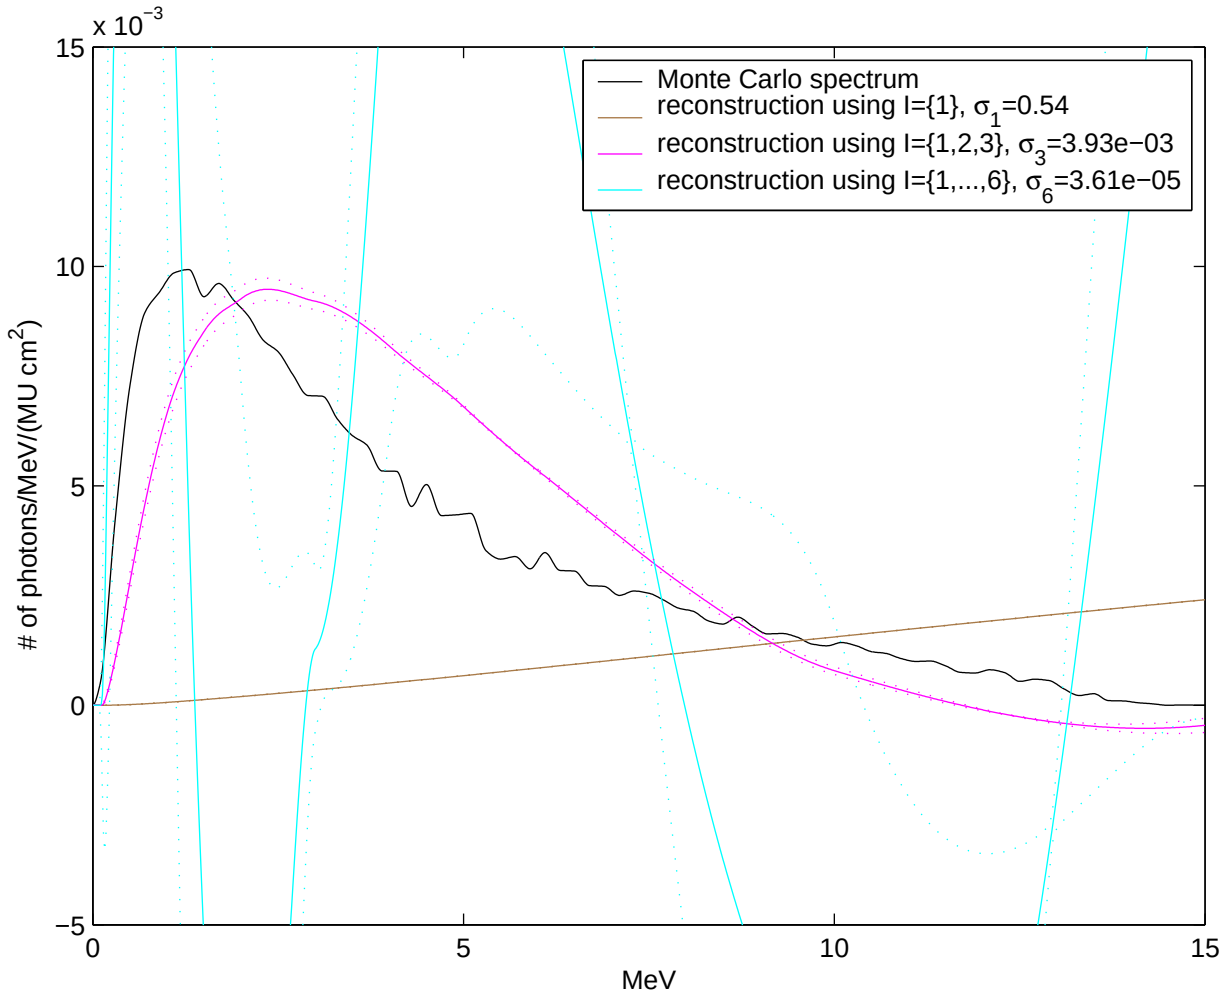

Figure 7. Reconstruction of the spectrum of a Elekta SLi 15 MV beam using different numbers of basis functions. Dashed lines show ±1 standard deviation. The Monte Carlo spectrum is from Fippel et al. (2003).

simulated spectrum and the reconstruction in figure 6 is due to the accuracy of the Monte Carlo simulation and not reconstruction error. This is evident from differences found in the calculation of the forward problem. Comparing the actual measurements (not the reconstruction) to simulated measurements calculated using equation 1 assuming the Monte Carlo spectrum of Fippel et al. (2003) represents our Elekta SLi, reveals a ±2.5% discrepancy for the 6MV beam and a ±5.7% discrepancy for the 15MV beam. This is after taking into account the normalization of the spectrum, and the background. While not large, these discrepancies are an order of magnitude larger than the measurement error (i.e., the variation among the replicas of each measurement). There are two possibilities. Either the spectrum of Fippel et al. (2003) is not identical to (or not an accurate enough approximation of) the true spectrum of our source or equation 1 does not accurately model the forward problem (e.g., perhaps the dependence of N_D on photon energy needs to be included). Further investigation is beyond the scope of this

work.

3. Discussion and Conclusions

Generically more measurements will not improve the reconstruction. This is because the reconstruction is limited by the number of large singular values. Generally the singular values of matrix decay exponentially. More replications will help as they reduce the noise and hence increase the range of singular values that are considered "large". Unfortunately, the noise decreases as $r^{-1/2}$ where r is the number of replications. Therefore, the gain from replications is often small (maybe one more singular value). This is a difficulty with the problem, rather than with the reconstruction procedure. One possible improvement in our reconstruction procedure is the use of regularization as discussed in section 2.7.

On the premise that the reconstruction is limited by the problem and not the algorithm, future research should focus on the optimal selection of measurements. With careful choices of attenuators, the singular values of A may decay more slowly.

Acknowledgments

The authors would like to acknowledge Matthias Fippel for emailing the Elekta SLi spectra in Fippel et al. (2003). Benjamin Armbruster would like to acknowledge some helpful discussions with Robert Indik and proofreading by David Lomen.

References

Almond P R, Biggs P J, Coursey B M, Hanson W F, Huq M S, Nath R & Rogers D W O 1999 *Medical Physics* **26**(9), 1847–1870.
*http://link.aip.org/link/?MPH/26/1847/1

Baird L C 1981 *Medical Physics* **8**.
*http://link.aip.org/link/?MPH/8/319/1

Baker C R, Ama'ee B & Spyrou N M 1995 *Physics in Medicine and Biology* **40**(4), 529–542.

Catala A, Francois P, Bonnet J & Scouarnec C 1995 *Medical Physics* **22**(1), 3–10.
*http://link.aip.org/link/?MPH/22/3/1

Fippel M, Haryanto F, Dohm O, Nusslin F & Kriesen S 2003 *Medical Physics* **30**.
*http://link.aip.org/link/?MPH/30/301/1

Francois P, Catala A & Scouarnec C 1993 *Medical Physics* **20**(6), 1695–1703.
*http://link.aip.org/link/?MPH/20/1695/1

Hinson W H & Bourland J D 2002 *Medical Physics* **29**.
*http://link.aip.org/link/?MPH/29/1789/1

Huang P H, Kase K R & Bjarngard B E 1982 *Medical Physics* **9**.
*http://link.aip.org/link/?MPH/9/695/1

Hubbell J H & Seltzer S M 1997 'Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients (version 1.03).'. [Online] Available: http://physics.nist.gov/xaamdi [2004, March 26]. Originally published as NISTIR 5632, National Institute of Standards and Technology, Gaithersburg, MD (1995).

Joseph P M 1975 *Medical Physics* **2**.
*http://link.aip.org/link/?MPH/2/201/1
Landry D J & Anderson D W 1991 *Medical Physics* **18**(3), 527–532.
Nisbet A, Weatherburn H, Fenwick J D & McVey G 1998 *Physics in Medicine and Biology* **43**(6), 1507–1521.
Rogers D, Faddegon B, Ding G, Ma C, We J & Mackie T 1995 *Medical Physics* **22**(5), 503–524.
Sheikh-Bagheri D & Rogers D W O 2002 *Medical Physics* **29**(3), 391–402.
*http://link.aip.org/link/?MPH/29/391/1
Tan S M & Fox C n.d. *Physics 707 Inverse Problems* University of Auckland.
*http://www.phy.auckland.ac.nz/Staff/smt/453707SC.html