FLITES TECHNICAL REPORT

Nick Polydorides & Hugh McCann
University of Edinburgh.

October 7, 2014

1 INTRODUCTION

The image reconstruction problem in two-dimensional chemical species tomography follows the classical paradigm of Radon transform inversion [1]. The aim is to produce a discrete image of an optically transparent medium from a finite set of path concentration integrals, also referred to as projection data. In this context, data acquisition typically entails a dense array of monochromatic laser radiation beams propagating through the medium at various angles spanning the half circle. Measuring the intensity of these beams at their sources and diametrically opposite detectors yields the attenuation information within a noise margin, and that relates linearly, under some assumptions, to the medium's chemical species concentration function. In situations where a dense sampling of the domain is feasible, implying a large number of beams per projection and a large number of angles; a reasonable discretisation of the domain suffices to regularise the inversion and enforce numerical stability against the noise in the data. In such circumstances, the image can then be reconstructed using Fourier transform-based methods like the popular filtered back-projection algorithm [2].

There are however some situations where measuring continuous Radon transform data is impractical, allowing only a sparse beam arrangement and a small number of projection angles within the half circle arc. This is the case known as limited data tomography and it is typically addressed in the context of algebraic image reconstruction [2]. In this work we focus on a rather extreme case where the number of projection data is substantially less than the dimension of the discrete image we seek to reconstruct, giving rise to a severely underdetermined inverse problem. Such problems are known to be ill-posed in the sense of Hadamard since they lack uniqueness, although a unique image can still be computed subject to enforcing some form of regularisation [5]. As the underlying optical attenuation model is linear upon neglecting beam hardening effects, the problem is well suited to the robust Tikhonov regularisation [11], [3], as well as iterative algorithms based on the Krylov and Landweber iterations, exploiting their regularising and convergence properties in solving linear ill-posed problems [12]. A drawback here is the computational inefficiency of these methods in handling large, dense, underdetermined systems. To see this consider having a small number m of projection measurements from which we wish to reconstruct the image at $N \times N$ pixel

resolution, where $m \leq N$. Implementing the direct method of Tikhonov regularisation or indeed a Krylov type iteration requires forming the normal equations resulting in a matrix with dimension $N^2 \times N^2$. Choosing for example, $m = N = 10^2$ yields an unknown with 10^4 degrees of freedom and a matrix with 10^8 elements to invert.

At the Flites experiment [9] we consider the problem of image reconstruction using attenuation data collected at the boundary of a square domain $\Omega = [-a, a] \times [-a, a]$, where a is in the order of a few meters. As beams of near infra red radiation propagate through the domain they get absorbed by intervening gases and they are assumed to attenuate with negligible scattering. Although the gases are in motion, high speed data acquisition provides static observation conditions for each data frame so here we simply consider the time-invariant system. If the amount of energy absorbed at a point is thought to be proportional to the gas density there, the attenuation of a beam with density $p > 0$ along a small segment $\mathrm{d}\ell$ is

$$\mathrm{d}p = -p\,x(\ell)\mathrm{d}\ell, \quad \mathrm{d}\ell \in \ell, \tag{1.1}$$

where $x \geq 0$ is the chemical species concentration function we seek to image. Let $\mathbf{r}_s$ denote the start of the ℓth beam; typically the position of the ℓth source, and $\mathbf{r}_d$ its end at the corresponding detector such as $|\ell| = |\mathbf{r}_s - \mathbf{r}_d|$ is the length of the beam. Further, let p_s to be the intensity of the beam leaving $\mathbf{r}_s$ and p_d the intensity of the beam arriving at $\mathbf{r}_d$, then dividing (1.1) by $-p$ and integrating over ℓ gives

$$\int_\ell \mathrm{d}\ell\; x(\ell) = -\int_{\mathbf{r}_s}^{\mathbf{r}_d} \mathrm{d}p\; \frac{1}{p} = \log\Big(\frac{p_s}{p_d}\Big), \tag{1.2}$$

where $p_s \geq p_d \geq 0$, and $\log p$ denotes the natural logarithm of p. The inverse imaging problem is then to estimate a bounded non-negative function $x : \Omega \to \Re_+$ from a set of m noise contaminated data $\Big\{\log\big(\frac{p_{s_1}}{p_{d_1}}\big), \ldots, \log\big(\frac{p_{s_m}}{p_{d_m}}\big)\Big\}$. Before focusing attention on the inverse problem, we make a brief remark on the impact of noise on the assembled logarithmic data compared to that on the actual measurements $\{p_{s_i}, p_{d_i}\}$. If p_s is known without uncertainty and p_m contains additive noise n, then $y^* = \log(p_s/p_d)$ denotes the exact data and $y = \log(p_s/p_d + n)$ the noisy. Subtracting the one from the other yields

$$|y^* - y| = \Big|\log\frac{p_d}{p_d + n}\Big| = \Big|\log(1 + \frac{n}{p_d})\Big| \leq \Big|\frac{n}{p_d}\Big|,$$

and therefore the data utilised in solving the inverse problem are affected less by the noise compared to the actual measurements from which they are formed. The presence of p_d in the denominator of the fraction above makes the case for using high-intensity radiation, and this is aligned with the general engineering principle for high signal to noise ratio. Populating the path concentration integrals in (1.2) for all beams over the range of angles in $[0, \pi)$ yields the linear operator equation

$$y = \mathscr{A}x + n, \tag{1.3}$$

where $\mathscr{A}x$ is the Radon transform of x, y is the data vector and n some additive noise corrupting the data. The mathematical problem of reconstructing x from y based on (1.3) is well studied and analysed in various textbooks, see for example [1] and [15]. Its theory postulates

that such a reconstruction is feasible and stable subject to the sufficiency of the data. In other words, the Radon operator has a continuous inverse that leads to a unique solution provided that the data suffice to resolve the degrees of freedom in the unknown. The following theorem, adapted from [15], formalises this idea.

Let $x \in L^2(B_a)$ be a continuous square integrable function supported completely within the ball $B_a \in \Omega$ centred at zero. Then there exists a constant $c > 0$ such that

$$\frac{1}{\sqrt{\pi}} \|x\|_{L^2} \leq \|\mathscr{A}x\|_{L^{2+1/2}} \leq c\|x\|_{L^2}. \tag{1.4}$$

This result postulates that in two dimensions, $\mathscr{A}x$ is half a degree smoother than x, manifesting that some information about x does not appear in $\mathscr{A}x$. In effect, trying to recover it from such data turns out to be problematic. However that loss tends to be small, especially when one considers finding a discrete form of x, thus in terms of stability the problem is said to be mildly ill-posed and thus subject to data sufficiency $\mathscr{A}$ has a continuous inverse. Unfortunately, in our setting the Radon transform is severely undersampled and such a stability is not feasible without some form of regularisation. In this paper we show that an effective re-parameterisation of the unknown in conjunction with a basic regularisation scheme can yield a stable, unique solution at a reduced computational cost. Our approach is fundamentally based on projection onto a low-dimensional image-feature subspace, an operation that also incurs some loss of information about x.

Discretising the operator equation (1.3) on an $N \times N$ grid of square pixels, with $N^2 \gg m$ yields an underdetermined system of linear equations

$$y = Ax + n, \tag{1.5}$$

where $A \in \Re^{m \times N^2}$, with $m = \text{rank}(A)$ and n is the additive noise of magnitude $\delta = \|n\|$. Without any loss of generality we assume the system normalised so that $\|A\| = 1$, while the singular value decomposition (SVD) of A yields $A = U\Sigma V'$ where $U \in \Re^{m \times m}$ and $V \in \Re^{N^2 \times N^2}$ are orthogonal matrices and $\Sigma \in \Re^{m \times N^2}$ is a diagonal holding the singular values of A in non-ascending order $1 \geq \sigma_2 \geq \ldots \sigma_m$. Expressing V and Σ like

$$V = [V_m \mid V_{N^2-m}], \quad V_m \in \Re^{N^2 \times m}, \quad \Sigma = [\Sigma_m | 0], \quad \Sigma_m \in \Re^{m \times m},$$

it is easy to see that

$$A = U\Sigma V' = U\Sigma_m V_m'.$$

The columns of U span the range of A, those of V_{N^2-m} its null space $\mathscr{N}(A)$ and $V_m \in \mathscr{N}^{\perp}(A)$. In our notation prime denotes transposition and $\|\cdot\|$ is the Euclidean norm of matrices and vectors. Despite being underdetermined, we remark that the condition number of A is small, thus let us denote $\kappa(A) = 1/\sigma_m = c$.

A suitable method for reconstructing a unique solution from the model (1.5) is by formulating the Tikhonov problem, using for example a smoothness imposing regularisation matrix $L \in \Re^{N^2 \times N^2}$ [11]. This is equivalent to solving the augmented least squares problem for a positive parameter λ

$$\min_x \left\| \begin{pmatrix} A \\ \sqrt{\lambda} L \end{pmatrix} x - \begin{pmatrix} y \\ 0 \end{pmatrix} \right\|^2 \tag{1.6}$$

whose solution

$$\hat{x}_\lambda = (A'A + \lambda L'L)^{-1} A'y, \tag{1.7}$$

can be computed directly by inverting an $N^2 \times N^2$ matrix. We will refer to this computationally expensive solution as a benchmark to compare the results obtained by implementing our methodology, that solves the problem without performing any computations in N^2 dimensions.

2 SUBSPACE PROJECTION

We embark our approach for image reconstruction by considering a suitable basis of functions where we could approximate the unknown image of the absorption function. The projection approach is quite popular in addressing various forms of large-scale inverse problems, see for example [5] for image processing, [14] for geophysical imaging and [19] for general linear inverse problems. The subspace projection relies on the hypothesis that there exist a potentially low-dimensional basis of functions that captures the dominant features of the desirable image, and as such this relies explicitly on the available prior information one possesses about the solution. Suppose $x^* \in \Re^{N^2}$ is the high-dimensional image we want to reconstruct. Further let $\Pi \in \Re^{N^2 \times N^2}$ be the orthogonal projection operator for functions in $\Re^{N^2}$ onto a low dimensional subspace

$$S \doteq \{\Phi r \,|\, r \in \Re^s\}, \quad s \ll N^2,$$

spanned by s linearly independent basis functions $\{\phi_1, \ldots, \phi_s\}$ comprising the columns of the tall matrix $\Phi \in \Re^{N^2 \times s}$. As we discuss in the following section, there are various options for choosing these functions, and although not it is not absolutely necessary, for stability reasons a choice of orthonormal basis is often preferred [5]. In these circumstances, adopted in this study, note that $\Phi'\Phi = I_s$ and $\Phi\Phi' = \Pi$. Further consider $r^* = \Phi' x^*$ to be the optimal coefficients in the approximation of x^* on S and $\hat{r}$ to be the computed low-dimensional solution of the projected inverse problem. Our methodology constitutes of two steps: the projection of the high-dimensional image in the low-dimensional subspace inducing r^* and the numerical solution of the projected inverse problem yielding $\hat{r}$. This process can be diagrammatically depicted as

$$x^* \quad \xrightarrow[\text{error}]{\text{approximation}} \quad \Pi x^* \quad = \quad \Phi r^* \quad \xrightarrow[\text{error}]{\text{computational}} \quad \Phi \hat{r}, \tag{2.1}$$

indicating the two sources of errors affecting the solution. In this context the total image error $\|x^* - \Phi\hat{r}\|$ will comprise the subspace approximation error, depending on the skill of the basis to express x^*, and the computational error which reflects the performance of the image reconstruction algorithm to approximate r^* given the numerical properties of the resulting low-dimensional inverse problem and the impact of noise in the data. To estimate this error consider that any bounded image in $\Re^{N^2}$ can be decomposed as

$$x = \Pi x + (I - \Pi)x, \tag{2.2}$$

where $\Pi x = \Phi r$ for a unique r. Denoting by $w_x = (I - \Pi)x$ the subspace approximation error encompassing the component of x that does not belong in S, the linear model for the

measurements (1.5) becomes

$$y = A(\Pi x + w_x) + n, \tag{2.3}$$

yielding the additional data error term Aw_x. To quantify this error further we project w_x onto $\mathcal{N}(A)$ using an orthogonal projection operator $P = V_{N^2-m} V'_{N^2-m}$ for the subspace $S_p \doteq \{V_{N^2-m} r \mid r \in \Re^{N^2-m}\}$, such that

$$w_x = P w_x + (I - P) w_x. \tag{2.4}$$

Since $P w_x \in \mathcal{N}(A)$ then the projected model for the low-dimensional variable $r \in \Re^s$ simplifies to

$$y = A\Phi r + Aq_x + n, \tag{2.5}$$

where $q_x = (I-P) w_x$. In terms of its actual magnitude, there is little to be said about this new error term, aside the rather obvious upper bound

$$\|Aq_x\| \le \|A\| \|q_x\| \le \|I - P\| \|w_x\| \le \|I - \Pi\| \|x\| \le \|x\|, \tag{2.6}$$

since $\|I - P\| = \|I - \Pi\| = 1$. Ultimately, to estimate the overall image error we have

$$x^* - \Phi\hat{r} = x^* - \Pi x^* + \Pi x^* - \Phi\hat{r} = x^* - \Pi x^* + \Phi r^* - \Phi\hat{r},$$

and using the triangle inequality and the orthogonality of Φ we obtain

$$\|x^* - \Phi\hat{r}\| \le \|w_x\| + \|r^* - \hat{r}\|. \tag{2.7}$$

This error bound is a mathematical expression equivalent to the schematic in (2.1) indicating how to account for the possible sources of error in the reconstructed image. Clearly the first component is the subspace approximation error and it relies entirely on the choice of basis Φ, which in term reflects the credibility of the prior information one has about x^*. To quantify the second term in (2.7) we must first specify $\hat{r}$ by formulating an appropriate inverse problem. For this we distinguish between two cases depending on the numerical properties of the matrix product $B = A\Phi$ assuming $m \le s \ll N^2$. For clarity in the notation let (2.5) be expressed as

$$y = Br + \varepsilon, \quad \text{where} \quad \varepsilon = Aq_x + n, \quad \text{and} \quad B = \widetilde{U}\widetilde{\Sigma}\widetilde{V}', \tag{2.8}$$

with $\widetilde{U}$, $\widetilde{\Sigma}$, and $\widetilde{V}$ the SVD factors of matrix B, such that for any index $1 \le k \le m$ these can be expressed as $\widetilde{U} = [\widetilde{U}_k | \widetilde{U}_{m-k}]$, $\widetilde{\Sigma} = [\widetilde{\Sigma}_k | \widetilde{\Sigma}_{s-k}]$ and $\widetilde{V} = [\widetilde{V}_k | \widetilde{V}_{s-k}]$. Further let $\widetilde{\Sigma}_{kk} \in \Re^{k\times k}$ to be the upper part of $\widetilde{\Sigma}_k$. Since $\|A\| = \|\Phi\| = 1$ then we can see immediately that $\|B\| \le \|A\|\|\Phi\| = 1$, while from [16], $\tilde{\sigma}_i \le \sigma_i \|\Phi\|, \quad i = 1, \ldots, m$, we can deduce that the singular values of B satisfy $1 > \tilde{\sigma}_1 \ge \tilde{\sigma}_2 \ge \ldots, \tilde{\sigma}_m > 0$, while similar arguments can be used to show that $\kappa(B) \ge \kappa(A)$ when $s < N^2$. In this context we define low-dimensional coefficients vector estimator

$$\hat{r} = \widetilde{V}_k \, \widetilde{\Sigma}_{kk}^{-1} \, \widetilde{U}'_k \, y, \quad k \le m, \tag{2.9}$$

where k depends on the numerical rank of B. Notice that since k is always less than s, therefore we distinguish between two cases. If B has full row rank then $k = m$ and the above $\hat{r}$ is simply the minimum norm estimator, that does not entail any regularisation. This estimator

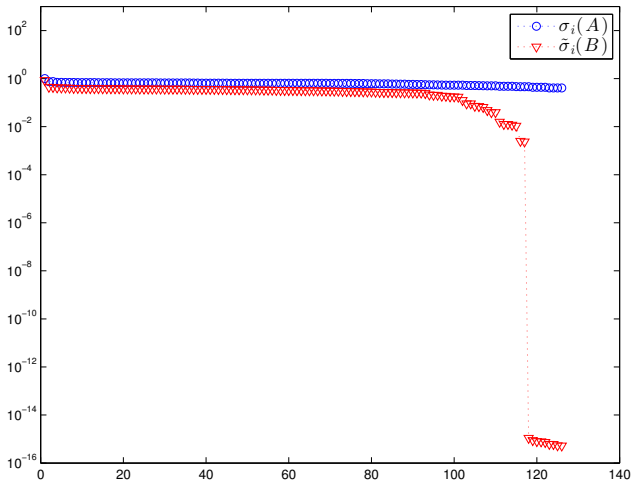

Figure 2.1: The singular values of matrices $A \in \Re^{m\times N^2}$ and $B = A\Phi$ with $\Phi \in \Re^{N^2\times s}$, $m = 126$, $s = 289$ and $N = 70$, indicating the small rate in the reduction of the singular values of A and the effective rank deficiency of B.

is seldom used, since most smooth feature bases will, at least in our experience, cause B to become rank deficient. In this, more general case, it is imperative to use $k < m$ hence the estimator $\hat{r}$ corresponds to a truncated SVD (TSVD) solution [6]. It must be emphasised that the ill-conditioning of $B \in \Re^{m\times s}$ and the need for regularisation arise as a consequence of a rank deficiency that is in turn introduced by the change of basis. Otherwise, as the figure 2.1 illustrates, the singular values of A decay very slowly, consistently to the analytical calculation of the SVD of the Radon operator [4], manifesting that the ill-posedness is merely due to the insufficiency of data. This is in stark contrast to severely ill-posed inverse problems [3] where the singular values of the operators involved decay at an exponential rate, and thus acquiring more data becomes pointless in the presence of noise, as asserted by the Picard criterion [6]. Effectively, the rank deficiency of B separates its m singular values in two clusters, one of which have values very closed to zero, which provides some insight into the choice of the truncation level k. More information to optimise the choice of k can be obtained by the error bounds of the solution we provide next. The figure 2.1 shows this rank deficiency of B for the case where the basis $\{\phi_1,\ldots,\phi_s\}$ consists of discrete cosine transform functions. The computational error $\|r^* - \hat{r}\|$ depends on how well we can solve the reduced inverse problem to approximate r^* in (2.8) given the impact of regularisation and the various types of additive noise embedded in ε. For the regularised estimator in (2.9), assuming the general rank deficient case where $\tilde{\sigma}_m \approx 0$

$$B_k^{-1} = \widetilde{V}_k \widetilde{\Sigma}_{kk}^{-1} \widetilde{U}_k',$$

and a given choice of index $k < m$ the estimator becomes

$$\hat{r} = B_k^{-1} B r^* + B_k^{-1}(Aq_x + n),$$

and by subtracting r^* from both sides and applying the triangle inequality we arrive at

$$\|\hat{r} - r^*\| \leq \|B_k^{-1} B - I\| \|r^*\| + \|B_k^{-1}\| \big(\|q_x\| + \|n\|\big). \tag{2.10}$$

Notice the presence of the so-called regularisation error $\|B_k^{-1} B - I\| \|r^*\|$, which constitutes the 'cost' of ensuring stability in inversion. With some simple algebra we can refine the above error bound using the chosen value of k and the singular values of B to get

$$\|\hat{r} - r^*\| \leq \|r^*\| + \frac{1}{\tilde{\sigma}_k} \big(\|q_x\| + \|n\|\big), \tag{2.11}$$

since $\|B_k^{-1} B - I\| = 1$ when $k < m$. In the special case where B has full row rank then $k = m$ and the inverse $B_m^{-1} = \widetilde{V} \widetilde{\Sigma}^{-1} \widetilde{U}$ exists,

$$\hat{r} = B_m^{-1} y = B_m^{-1}(B r^* + \varepsilon) = r^* + B_m^{-1} \varepsilon.$$

Rearranging and taking the norms

$$\|\hat{r} - r^*\| \leq \|B_m^{-1} A q_x\| + \|B_m^{-1} n\| \leq \|B_m^{-1}\| (\|q_x\| + \|n\|),$$

hence the respective error bound becomes

$$\|\hat{r} - r^*\| \leq \frac{1}{\tilde{\sigma}_m} (\|q_x\| + \|n\|) \tag{2.12}$$

with $\tilde{\sigma}_m > 0$. Since $\tilde{\sigma}_m < \tilde{\sigma}_1$ and $\tilde{\sigma}_1 \leq 1$, then some error amplification is included in the solution, albeit, the underlying assumption in this instance is that B is well conditioned for inversion. Comparing the two computational errors fort the TSVD and minimum norm estimators in (2.11) and (2.12) shows the first includes an additional term, the regularisation error, while the amplification of the noise is controlled by the inverse of the smallest singular value of B included in the solution. Notice however, that a small computational error $\|r^* - \hat{r}\|$ does not automatically imply a small overall error $\|x^* - \Phi \hat{r}\|$, as the next section demonstrates.

2.1 Choice of approximation basis and a special case

We have seen that projecting the unknown high-dimensional image into a low-dimensional subspace reduces the computational complexity by replacing the large matrix A with a smaller matrix B. This computational advantage causes an increase in the 'noise' in the data from n to ε, in manifestation of the approximation error w_x. Moreover, this projection typically leads to a rank deficient inverse problem, despite that the original large matrix has a small condition number. The choice of the appropriate basis of functions is thus instrumental in the image reconstruction process, in the sense that it controls the approximation error w_x, but it also has an impact on the computational error. One conclusion that becomes apparent even as early as this stage of the investigation is that the conventional local basis functions

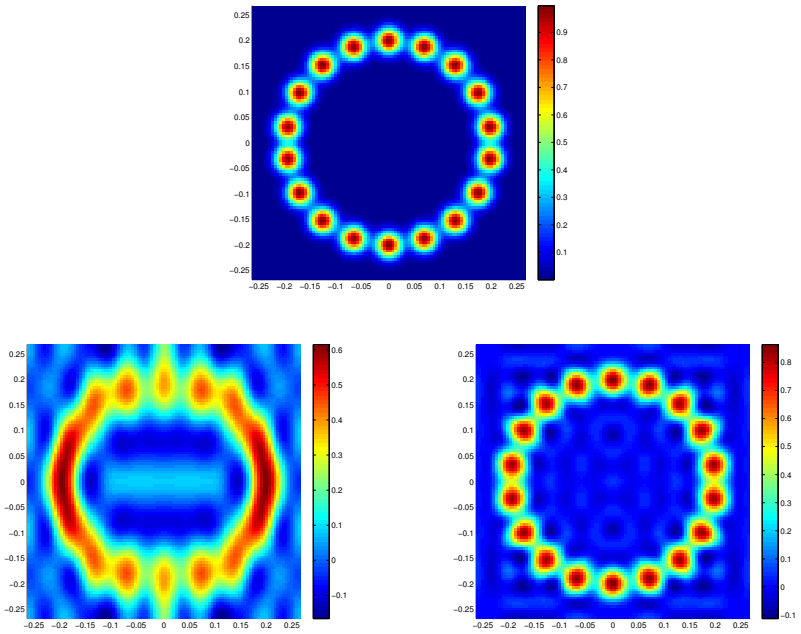

Figure 2.2: At the top, the high dimensional target image x^* and below its projections Πx^* in two distinct bases of discrete cosine transform functions. To the left, the projection on $s = 100$ functions which yields a relative approximation error $\|x^* - \Pi x^*\|/\|x^*\|$ of 0.50. To improve this crude approximation we enlarge the basis to include $s = 289$ functions yielding the projected image shown at the right that has a smaller relative error of 0.17. For comparison we also quote the condition number of the two B matrices at 2×10^9 and 1.5×10^{15} respectively.

with pixel-wise constant support; often adopted by default, is not necessarily the best basis to model the unknown image when having very few measurements. Also known as multi-resolution basis functions, this local basis is optimal for maximising the spatial resolution. However, projecting the unknown in a lower dimensional multi-resolution basis results in a severe image degradation since this is equivalent to using a coarser pixelation of the domain. Some alternative local basis functions for computed tomography are suggested in [17]. Alternatively, one could utilise a global basis of functions. If for example the solution is a priori known to be smooth, then a smooth basis is appropriate, involving for example some radial basis functions or wavelets [13]. For this study we utilise the orthonormal basis of functions corresponding to the discrete cosine transform suggested in [5], a smooth two-dimensional basis of gradually increasing frequency. Subject to a sufficiently large s, this basis can accommodate a large range of functions, including those that are not very smooth, see for example the projected images in figure 2.2.

2.1.1 THE MINIMUM NORM SOLUTION

Let us now look closer into computing the minimum norm solution for the high-dimensional model (1.5) formulated as

$$\hat{x} = \arg\min_{x \in \Re^{N^2}} \|x\| \quad \text{such that} \quad y = Ax.$$

Our intent is to show that this is a special case of the subspace projection method with $\Phi = V_m$ ($s = m$) in which case $\kappa(B) = \kappa(A) = c$, $\Pi = P^{\perp}$. Effectively,

$$Aq_x = A(I - P)w_x = A\Pi w_x = 0,$$

and hence the measurement error due to the subspace approximation error vanishes leaving $\varepsilon = n$. This unique estimator can be computed rather efficiently, and stably, by inverting only an $m \times m$ square matrix $B = AV_m$. Given that AA' is invertible by virtue of $\kappa(A) = c$ being small, then $\hat{x} = A'(AA')^{-1}y$ can alternatively be computed by solving the low-dimensional least squares problem

$$\hat{r} = \arg\min_{r \in \Re^m} \|AV_m r - y\|^2$$

If $\Sigma_m \in \Re^{m \times m}$ is the nonzero block of Σ that holds the singular values of A, then it is straight forward to show that $\hat{r} = \Sigma_m^{-1} U' y$ and therefore $\hat{x} = V_m \hat{r}$. Notice that in this case the computational error attains its minimum possible value $\|r^* - \hat{r}\| = c\|n\|$. To see this recall that $r^* = V_m' x^*$ and

$$\begin{aligned}
\hat{r} = \Sigma_m^{-1} U' y &= \Sigma_m^{-1} U' (AV_m r^* + A(I_{N^2} - P)w_x + \varepsilon) \\
&= \Sigma_m^{-1} U' (U\Sigma_m V_m' V_m r^* + U\Sigma_m V_m' (I_{N^2} - P)w_x + \varepsilon) \\
&= r^* + V_m'(I_{N^2} - P)w_x + \Sigma_m^{-1} U' \varepsilon \\
&= r^* + c\varepsilon = r^* + cn,
\end{aligned}$$

hence combining with the approximation error we have

$$\|x^* - V_m r^*\| \le \|x^* - \Pi x^*\| + \|V_m(\hat{r} - r^*)\| \le \|x^* - \Pi x^*\| + c\delta.$$

Choosing the basis $\Phi = V_m$ is thus 'optimal' for minimising the computational error, since $\|\hat{r} - r^*\| \to 0$ as $\delta \to 0$, and no regularisation error occurs. Unfortunately, this advantage is ruined by the fact that the basis V_m is unsuitable in reconstructing the anticipated smooth solutions, see for example projection of an image into this basis in figure 2.3, which results in a very large approximation error $\|x^* - \Pi x^*\|$.

3 A VARIANT ALGORITHM TO IMPOSE POSITIVITY

Having reduced the dimension of the inverse problem through the projection of the unknown image on a low-dimensional basis of feature functions we proceed to suggest a modification that precludes the solution from attaining non-positive values. In this way we can also preserve the feasibility of the images, as chemical concentration is by definition non-negative.

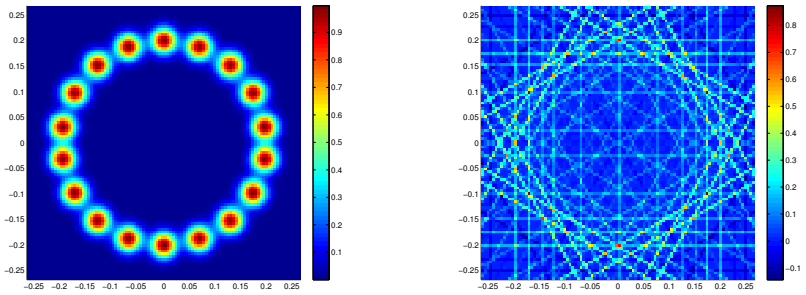

Figure 2.3: A target image x^* and its projection Πx^* in some bases of discrete cosine transform functions. The relative approximation error $\|x^* - \Pi x^*\|/\|x^*\|$ is 0.85 with $s = m = 126$ functions and the condition number of the square matrix B was 3.

The idea, originating from the geophysical inversion, is to reformulate the model and thereafter cast the inverse problem with respect to the logarithm of the unknown function [18]. In this context, let $z = \log x$ thus the original attenuation model now becomes nonlinear in the new unknown

$$y = A\mathrm{e}^z + n,$$

which can then be linearised around a (homogeneous) point $z_i \in S$ to

$$y = Ax_i + AX_i(z - z_i) + n,$$

where $X_i = \mathrm{diag}(x_i)$ and $x_i = \mathrm{e}^{z_i}$. Denoting $K(z_i) = AX_i$ and $y(z_i) = y - Ax_i + K(z_i)z_i$ then we cast a ith linear model for y at the iterate z_i

$$y(z_i) = K(z_i)z + n. \tag{3.1}$$

Note that the new unknown variable is still in the high-dimension, so similar to the unconstrained case we can project it onto S to get $z = \Pi z + w_z$ where now $\Pi z = \Phi r$, and $B(z_i) = K(z_i)\Phi$, yielding the low-dimensional model

$$y(z_i) = B(z_i)r + \varepsilon, \quad \varepsilon = K(z_i)w_z + n. \tag{3.2}$$

Assuming once again that the basis functions Φ spanning S are orthonormal then a positive subspace projected estimator of the image $\hat{x}^+$ can be obtained iteratively using the following algorithm:

1. Initialise $z_i \in S$ and compute $x_i = \log(z_i)$,
2. for $i = 1, 2, \ldots$
3. Compute matrix $B(z_i) = K(z_i)\Phi$ and vector $y(z_i)$,
4. Compute $\hat{r}_{i+1}$ using (2.9) with $B = B(z_i)$ and $y = y(z_i)$,
5. Update $z_{i+1} = \Phi\hat{r}_{i+1}$,
6. end

7. $\hat{x}^+ = e^{z_i}$

Although we do not derive error bounds for this positively constrained solution, we note that in principle there is some loss of information in the transformation from x to z, and therefore the bounds of the previous section no longer apply. In this case, only the positive part of the function x can be uniquely mapped to a corresponding function $z = \log x$, and consequently at the positively constrained estimator a zero concentration value will be assigned with a very small but still positive value.

4 NUMERICAL RESULTS

In this section we present some numerical results obtained with the projection-based algorithms presented in this paper. In particular, we consider reconstructing the two gas concentration functions, shown at the top left of figures 4.1 and 4.3. Notice that these target images, although consisting of Gaussian functions, they are not very smooth, thus we need to ensure that a sufficient number of high-frequency functions are included in the basis, in order to keep the approximation error as small as possible. In order to avoid committing an 'inverse crime' the $m = 126$ projection data are computed based on a fine square grid model with $N = 100$, while a coarser grid with $N = 70$ is utilised for solving the inverse problem. All data used in this study are infused with Gaussian noise of zero mean and variance equal to 5% of each projection's value. Although omitted from this report, our synthetic experiments replicated with 10% noise show only a marginal degradation from the results presented here.

In the first instance we consider reconstructing the image of eighteen narrow Gaussian functions arranged almost symmetrically on the periphery of a circle centred at the middle of Ω, as shown in figure 4.1 . Assuming 21 beams for 6 angles we compute 126 measurements on a fine grid and then infuse them with 5% Gaussian noise. Subsequently, we compute the model matrix $A \in \Re^{126\times 4900}$ on a coarser grid for which we also define a basis of $s = 289$ discrete cosine basis functions $\Phi \in \Re^{4900\times 289}$, and compute the projection of the target image in this basis $\Pi x \in \Re^{4900}$. Forming the low-dimensional matrix $B \in \Re^{126\times 289}$ yields $\kappa(B) \approx 10^{16}$ since the matrix comes out to be rank deficient, and solving the resulting problem for r using the TSVD estimator from (2.9) for $k = 100$ yields the image $\Phi\hat{r}$, also depicted in figure 4.1 along with the high-dimensional Tikhonov solution $\hat{x}_\lambda$ from (1.7) for $\lambda = 0.1$. This value of the regularisation parameter was optimised heuristically. In this case where $\|x^*\| = 17.3$, the approximation error was $\|w_x\| = 3$; a 17% relative error, the computational error $\|r^* - \hat{r}\| = 7.8$, while the error of the computed projected solution was $\|x^* - \Phi\hat{r}\| = 8.3$, confirming the error bound (2.7). Moreover, the cut-off singular value of B used in the solution $\tilde{\sigma}_{100} = 0.18$, $\|q_x\| = 1.18$ and $\|n\| = 0.49$, hence the error bound (2.11) is also confirmed. By comparison, the subspace projected solution looks quite similar to the Tikhonov solution whose error was evaluated at $\|x^* - \hat{x}_\lambda\| = 9.7$. It can thus be argued that a smooth Tikhonov-type regularised solution can be computed more efficiently using projection onto a space of smooth functions. Inspecting the reconstructed images indicates some overshoot of the pixel values into the negative range.

To eradicate the negative solution values we performed a few iterations of the positivity constrained algorithm maintaining the same level of truncation $k = 100$, and the same pro-

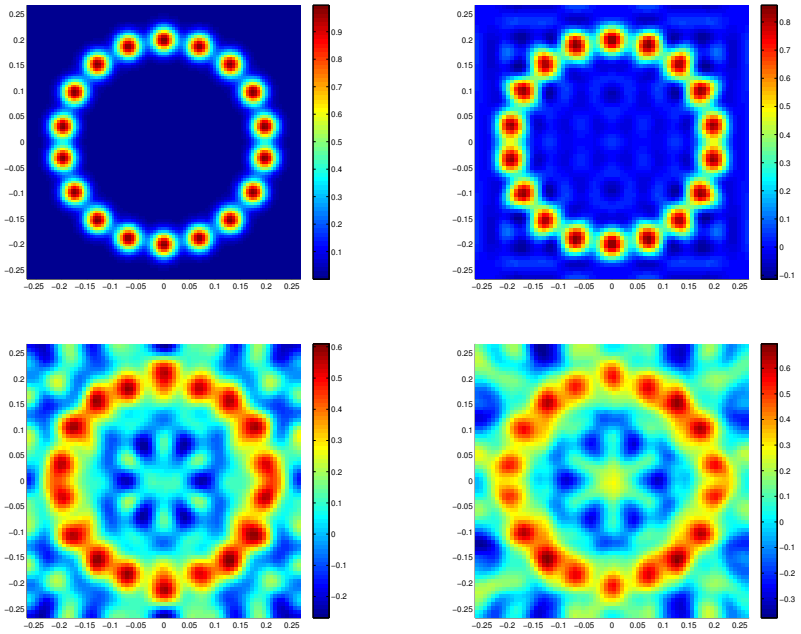

Figure 4.1: Top left, the target image x^* in resolution $N = 100$, and next its projection Πx^* into a subspace of $s = 289$ discrete cosine transform functions defined on a coarser grid with $N = 70$. Below left is the reconstructed image $\Phi\hat{r}$ with $\hat{r}$ the low-dimensional TSVD solution based on $k = 100$, and right the Tikhonov regularised solution $\hat{x}_\lambda$.

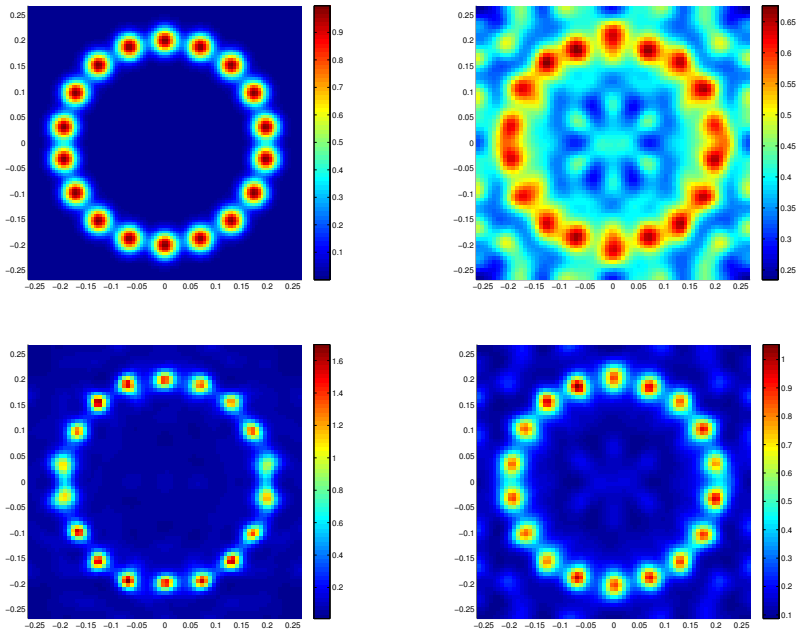

Figure 4.2: Clockwise, from top left: the target image in resolution $N = 100$, and the first three iterative solutions of the positivity imposing algorithm on a grid with $N = 70$. Note that the projected image becomes less smooth as the iterations progress.

jection subspace, starting at a guess $\hat{r}_0 = 0$ (equivalent to $\hat{x}_0^+ = 1$). The images corresponding to the first three iterations are illustrated in figure 4.2. The data and image errors recorded at these iterations were $\|y - A\exp(\Phi\hat{r}_i)\| = 17.8, 6.2, 2.6$, and $\|x^* - \exp(\Phi\hat{r}_i)\| = 23.8, 9.0, 6.1$, for $i = 1,2,3$ respectively. Contrary to the results obtained in the unconstrained case, apart from the imposed positivity these images appear to be less smooth. Comparing the errors in solving the nonlinear problem based on (3.2) against those in solving the linear problem based on (2.8) the results suggest that enforcing positivity has a significant impact on the solution's spatial resolution and overall error. This conjecture is consistent with the analysis on bound constrained nonlinear least-squares formulations, see for example the tomography related results in chapter 7 of [7]. We caution the readers that although this algorithm is based on the subspace projection too, the error bounds derived in the previous sections no longer apply. To see this notice that in this case $r^* = \Phi^\dagger \log x^*$ hence if x^* is sparse (c.f. the zero valued pixels in the target images in figures 4.1 and 4.3) then $\log x^*$ and r^* will have extremely large negative values. On the other hand, our formulation for $x = \exp(\Phi r)$ precludes x from attaining zero values, and the relevant elements of r_i effectively converge to a large negative value, albeit due to the impact of regularisation this value is sustained away from that of r^* and this causes $\|r^* - \hat{r}_i\|$ to remain quite large.

In the second example we test the performance of our algorithms on a simpler, yet not

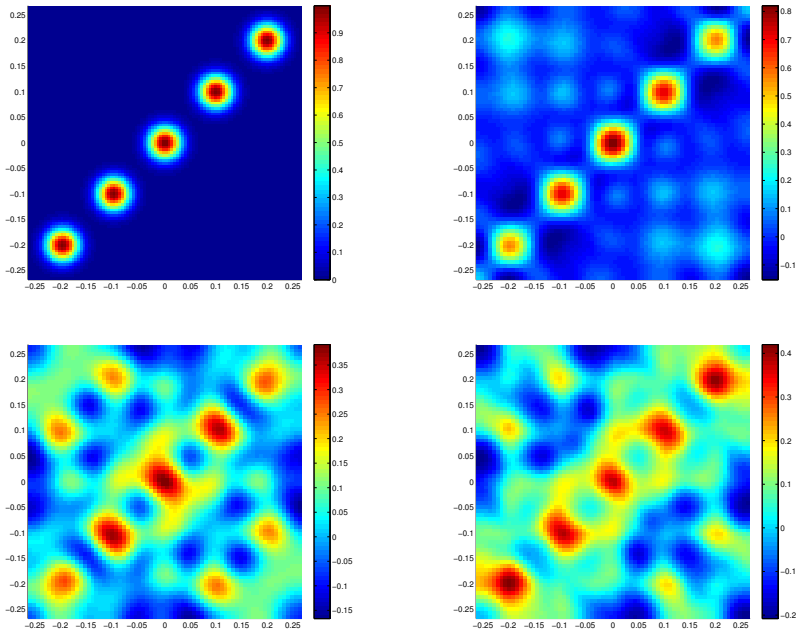

Figure 4.3: Top left, the target image x^* in resolution $N = 100$, and next its projection Πx^* into a subspace of $s = 144$ discrete cosine transform functions defined on a coarser grid with $N = 70$. Below left is the reconstructed image $\Phi\hat{r}$ with $\hat{r}$ the low-dimensional TSVD solution based on $k = 65$, and right the Tikhonov regularised solution $\hat{x}_\lambda$.

symmetric target image, as this appears at the top left of figure 4.3. The procedure to compute the measurements and projection matrices follows similar to the previous case, only this time we restrict the span of the subspace to the first $s = 144$ functions from the discrete cosine subspace. In this case, $\Phi \in \Re^{4900\times 144}$ yields a low dimensional, rank deficient matrix $B \in \Re^{126\times 144}$ with $\kappa(B) \approx 10^{16}$ and compute the truncated SVD estimator (2.9) with $k = 65$, for which $\tilde{\sigma}_{65} = 0.64$. Having used a smaller number of basis functions to reduce the dimensionality even further, the approximation error increases to $\|w_x\| = 5.3$ for a target with $\|x^*\| = 11.6$, constituting a relative error of 46%. The subspace projected image is shown in figure 4.3 on a grid with $N = 70$, along with the projected solution $\Phi\hat{r}$ and the Tikhonov solution $\hat{x}_\lambda$ for $\lambda = 0.1$. The corresponding image errors for these estimators were found to be $\|x^* - \Phi\hat{r}\| = 8.5$ and $\|x^* - \hat{x}_\lambda\| = 8.3$ respectively. Evaluating the computational error $\|r^* - \hat{r}\| = 6.5$, and $\|r^*\| = 10.27$, $\|q_x\| = 2.45$ and $\|n\| = 0.35$ verifies once again the error bounds (2.7) and (2.11).

Finally we reconstruct these measurements using the positivity enforcing iterative algorithm, assuming the same approximation basis and the same singular value truncation level $k = 65$, beginning once again with a homogeneous initial estimate of $\hat{r}_0 = 0$. Performing three iterations we obtain the reconstructed images illustrated in figure 4.4. Their corresponding

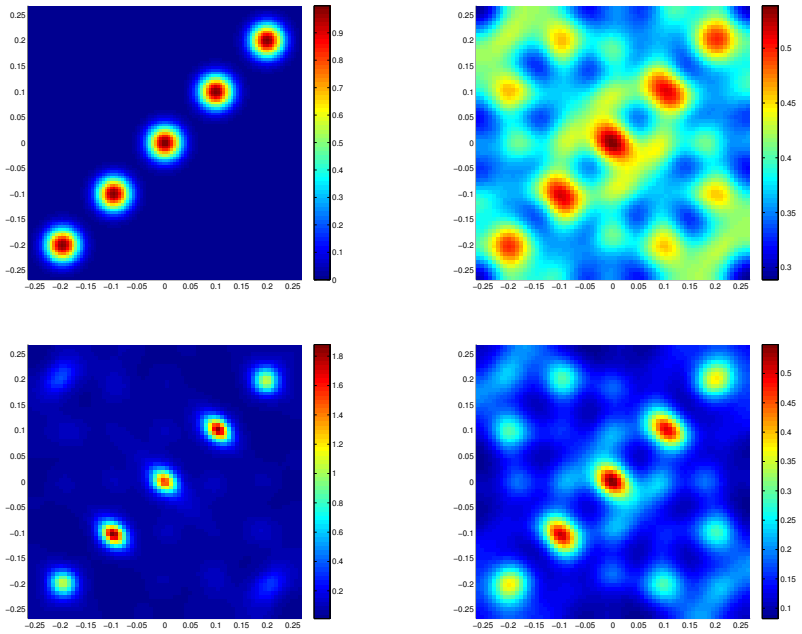

Figure 4.4: Clockwise, from top left: the target image in resolution $N = 100$, and the first three iterative solutions of the positivity imposing algorithm on a grid with $N = 70$. Note that the under-sampling image artifacts that profound in the unconstrained solutions in figure 4.3, are mitigated as the projected image becomes less smooth after each iteration.

data and image errors after the first three iterations are $\|y - A\exp(\Phi\hat{r}_i)\| = 19.8, 7.0, 3.5$, and $\|x^* - \exp(\Phi\hat{r}_i)\| = 25.3, 11.2, 7.2$, respectively. Similar to the first example, these results indicate a superiority of the constrained solution, yielding solutions with smaller errors, albeit at the expense of the smoothness of the image. The constrained images also seem to suppress the artefacts typically associated with limited data tomography, c.f. the noisy features symmetric to the diagonal line in figure 4.3.

5 Conclusions

This work addresses the computationally efficient solution of the inverse problem in attenuation tomography with undersampled data in the context of chemical species tomography. We have showed that projection of the unknown function in a smooth low-dimensional basis reduces drastically the dimensionality of the problem and produces images that are comparable if not better of those that would correspond to a large-scale Tikhonov formulation solution. Despite reducing the computational complexity of the problem the resulting matrices are typically rank-deficient rendering themselves available to regularisation methods based

on truncated singular value decomposition. This process incurs errors due to the subspace projection and the numerical solution of the resulted inverse problem under the impact of regularisation and noise in the data. Based on these errors we derive analytical bounds that quantify the overall image error, and verify their validity using numerical examples. As a variant of the suggested methodology we present an iterative algorithm that embodies a logarithmic transformation of the unknown to yield positively constrained images. This modification improves the image error and spatial resolution and was showed to be suitable in the cases where a sparse solution is sought.

REFERENCES

[1] A. C. Kak and M. Slaney, Principles of computerised tomographic imaging, SIAM, 2001.

[2] P. Kuchment, The Radon transform and medical imaging, SIAM, 2014.

[3] J. L. Mueller and S. Siltanen, Linear and nonlinear inverse problems with practical applications, SIAM, 2012.

[4] M. Bertero and P. Boccacci, Introduction to inverse problems in imaging, CRC Press, 1998.

[5] P. C. Hansen, Discrete inverse problems, insights and algorithms, SIAM, 2010.

[6] P. C. Hansen, Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion, SIAM, 1998.

[7] P. C. Hansen, V.Pereyra and G. Scherer, Least squares data fitting with applications, John Hopkins Press, 2013.

[8] G. T. Herman and R. Davidi, Image reconstruction from a small number of projections, Inverse Problems, 24, 2008.

[9] www.flites.eu. [or a paper on the Flites setup]

[10] M. G. Twynstra and K. J. Daun, Laser-absorption tomography beam arrangement optimisation using resolution matrices, Applied Optics, vol. 51(29), 7059-7068, 2012.

[11] K. J. Daun, Infrared species limited data tomography through Tikhonov reconstruction, Journal of Quantitative spectroscopy and Radiative Transfer, 111, 105-115, 2010.

[12] N. Terzija, J. L. Davidson, C. A. Garcia-Stewart, P. Wright, K. B. Ozanyan, S. Pegrum, T. J. Litt and H. McCann, Image optimisation for chemical species tomography with an irregular and sparse beam, Measurement Science and Technology, 19, 095007, 2008.

[13] M. Rantala, S. Vanska, S. Jarvenpaa, M. Kalke, M. Lassas, J. Moberg and S. Siltanen, Wavelet-based reconstruction for limited-angle X-ray tomography, IEEE Transactions on medical imaging, vol. 25, no. 2, 2006.

[14] D. W. Oldenburg, P. R. McGillvray and R. G. Ellis, Generalized subspace methods for large-scale inverse problems, Geophys. Journ. Intern., 114, 12-20, 1993.

[15] F. Natterer, The mathematics of computerized tomography, SIAM, 2001.

[16] A. S. Householder, The theory of matrices in numerical analysis, Dover, 1964.

[17] K. M. Hanson and G .W. Wecksung, Local basis-function approach to computed tomography, Applied Optics, 24(23), 4028-4039, 1985.

[18] M. Commer and G. A. Newman, New advances in three-dimensional controlled-source electromagnetic inversion, Geophys. J. Int. 172, 5130535, 2008.

[19] N. Polydorides, M. Wang and D. P. Bertsekas, A quasi Monte Carlo method for large scale inverse problems, Monte Carlo and Quasi-Monte Carlo Methods 2010, ed H. Wozniakowski and L. Plaskota, Springer-Verlag, 2012.