

SVD-TAIL: A NEW LINEAR-SAMPLING
METHOD FOR INVERSE SCATTERING PROBLEMS

by M. Fares¹, S. Gratton² and Ph. L. Toint³

Report 09/09

19 March 2009

¹ CERFACS,
42, avenue Coriolis, 31057 Toulouse Cedex 01, France
email: fares@cerfacs.fr

² Centre National d'Etudes Spatiales (CNES),
18, avenue Edouard Belin, 31401 Toulouse, France
email: gratton@cerfacs.fr

¹ Department of Mathematics,
FUNDP-University of Namur,
61, rue de Bruxelles, B-5000 Namur, Belgium.
email: philippe.toint@fundp.ac.be

SVD-tail: a new linear-sampling reconstruction method for inverse scattering problems

M'Barek Fares, Serge Gratton and Philippe L. Toint

19 March 2009

Abstract

A new numerical procedure (SVD-tail) is proposed for the reconstruction of the shape and volume of unknown objects from measurements of their radiation in the far field. This efficient procedure is a variant of the linear sampling method and has a very acceptable computational load. At variance with previously published techniques where the solution is constructed as a regularized solution of the far-field equations using a variant of the Tikhonov-Morozov type, the new method uses a new eigenspace recovery method which exploits the combined presence of error in the operator and of eigenvalue clusters. The performance of the new technique on a battery of examples and its comparison with existing methods is shown to be promising.

Keywords: inverse scattering, linear sampling method, regularization heuristics, eigenspace recovery, numerical algorithms.

1 Introduction

The past decade has seen the rapid development of the Linear Sampling Method (LSM) for inverse scattering problems (see Arens, 2004, Cakoni, Colton and Haddar, 2002, Cakoni and Colton, 2003, Colton, Haddar and Monk, 2002, Colton, Haddar and Piana, 2003, Haddar and Monk, 2002 or Fares, Gratton and Toint, 2009, for instance). This interest is due to its effectiveness, especially in treating three-dimensional inverse problems, but also to its large spectrum of applications. We recall that this algorithm allows the reconstruction of the shape of an unknown obstacle (or a local inhomogeneity) from multistatic data at a fixed frequency and has the clear advantage of requiring no a priori knowledge on the physical properties of the scatterers. It is based on suitably solving a linear system of equations, known as the far-field equations.

This paper has its origins in the difficulties met by the authors in their efforts to apply different regularization methods to the far-field equations. While the Tikhonov-Morozov technique has been analyzed and reasonably successfully applied to inverse scattering problems (see Colton, Piana and Potthast, 1997, Collino, Fares and Haddar, 2003, Fares et al., 2009, for instance), other popular methods for regularization, such as the L-curve method by Hansen (1997), appear to give globally unsatisfactory results. It is suggested in this paper that this situation may be linked to the fact that the Picard coefficients associated with the problem do not converge to zero in practice, although this convergence is a recommended prerequisite for the successful application of the L-curve method and other similar techniques. It turns out however that the inner products between the singular vectors and the right-hand side of the far-field equations (which are normalized by the corresponding singular values to define the Picard coefficients) behave very differently when considering points interior or exterior to the object of interest. We first analyse a new characterization of these inner products which explains this behaviour and, in turn, allows us to propose a new general eigenspace

recovery technique. We then propose to apply this new technique to inverse scattering problems and define a non-standard technique for LSM inversion, called SVD-tail, where the point-wise solution of the far-field equation is never explicitly constructed. As a result, the method is computationally very efficient. Moreover, the visual quality of the reconstructed images compares well with that obtained using other existing inversion algorithms.

The paper is organized as follows. Section 2 gives a brief overlook of the LSM technique for recovering objects from far-field measurements. The new eigenspace recovery technique and the SVD-tail algorithm are then motivated and presented in Section 3, while numerical experiments are discussed in Section 4. A numerically efficient variant of SVD-tail is also presented in Section 5. Conclusions and perspectives are finally outlined in Section 6.

2 A brief description of the LSM

Suppose a bounded (sound-soft) domain $D \subset \mathbb{R}^3$ of class \mathcal{C}^2 is given. This domain is illuminated by a plane wave incident field $u^{\text{inc}}(x, d) = e^{ikx \cdot d}$, $x \in \mathbb{R}^3$, for some $d \in S$, where S is the unit sphere. The forward problem consist in finding a function $u^s \in \mathcal{C}^2(\mathbb{R}^3 \setminus \overline{D}) \cap \mathcal{C}(\mathbb{R}^3 \setminus D)$ that satisfies

$$\Delta u^s + k^2 u^s = 0 \quad \text{in } \mathbb{R}^3 \setminus \overline{D}, \quad (2.1)$$

$$u = u^{\text{inc}} + u^s, \quad (2.2)$$

$$u = 0 \quad \text{on } \partial D, \quad (2.3)$$

$$\lim_{r \rightarrow \infty} \left(\frac{\partial u^s}{\partial r} - iku^s \right) = 0, \quad (2.4)$$

where $k > 0$ is the wave number. The function u^s is called the scattered wave. In case of sound-hard obstacle, the boundary condition (2.3) is replaced by

$$\frac{\partial u}{\partial \nu} = 0 \quad \text{on } \partial D, \quad (2.5)$$

where ν is the normal to ∂D directed towards the exterior of D . The condition (2.4) characterizes outgoing waves and ensures uniqueness of the solution of the scattering problem. Any solution u^s to the forward problem (2.1) is called radiating. It can be shown (Colton and Kress, 1998) that every radiating solution u to the Helmholtz equation has the asymptotic behaviour of an outgoing spherical wave

$$u(x, d) = \frac{e^{ik\|x\|}}{\|x\|} \left\{ u^\infty(\hat{x}, d) + O\left(\frac{1}{\|x\|}\right) \right\}, \quad \|x\| \rightarrow \infty \quad (2.6)$$

uniformly in all directions $\hat{x} = x/\|x\|$, where u^∞ , defined on the unit sphere S is known as the far-field pattern of u and $\|x\|$ is the Euclidean norm of x .

We are now interested in the inverse problem consisting of the reconstruction of the domain D from the knowledge of $u^\infty(\hat{x}, \hat{d})$ for all $(\hat{x}, \hat{d}) \in S \times S$ for a fixed, known wave number k . The far-field pattern $u^\infty(\hat{x}, \hat{d})$ defines the far-field operator $\mathcal{F} : L^2(S) \rightarrow L^2(S)$ by

$$(\mathcal{F}g)(\hat{x}) = \int_S u^\infty(\hat{x}, d)g(d)ds(d). \quad (2.7)$$

The linear sampling method chooses a position $z \in \mathbb{R}^3$ and then looks for the solution $g = g(\cdot, z) \in L^2(S)$ of the *far-field equation*

$$(\mathcal{F}g)(\hat{x}) = \Phi_\infty(\hat{x}, z) = e^{-ik\hat{x} \cdot z} / 4\pi \quad (2.8)$$

where $\Phi_\infty(\hat{x}, z)$ is the far-field pattern associated with plane wave $e^{ik\hat{x}\cdot z}$. Notice first that since \mathcal{F} is compact, equation (2.8) is ill posed. Even if this equation does not have any solution in general (Colton, 2003), it turns out that one can prove the existence of a nearby solution $g_\varepsilon(\cdot, z) \in L^2(S)$, in the sense that it satisfies the inequality $\|\mathcal{F}g_\varepsilon(\cdot, z) - u^\infty(\cdot, z)\| \leq \varepsilon$, where ε is a small parameter independent of z . This solution is such that

$$\lim_{\substack{z \rightarrow \partial D \\ z \in D}} \|g_\varepsilon(\cdot, z)\|_{L^2(S)} = \infty. \quad (2.9)$$

Furthermore, (formally) $\|g_\varepsilon(\cdot, z)\| = \infty$ for $z \in \mathbb{R}^3 \setminus D$ (again, see Colton, 2003). Consequently the L^2 -norm of $g_\varepsilon(\cdot, z)$ can be used as an indicator of the domain D that we want to reconstruct. This is the central observation on which the LSM is based.

3 The new method

3.1 The discretized far-field equations

Following Colton, Giebermann and Monk (2000), we discretize the far-field equation (2.8) by considering a triangular meshing of the unit sphere S containing N vertices $(d_i)_{1 \leq i \leq N}$. These vertices serve as directions for the plane incident waves as well as degrees of freedom for the discrete solution of the far-field equation. As an empirical rule, we take $N \simeq (kR + 2 \log(kR + \pi))^2$ where R is the radius of a sphere containing the object. This rule can be inferred by studying the spectrum of the far field as functions of the object size.

For the numerical realization of the LSM, we construct a continuous approximation of the solution $g(\cdot, z)$ linear at each triangle, whose degrees of freedom are its values at the nodes $(d_\ell)_{1 \leq \ell \leq N}$. The nodal values are denoted by $(g_j(z))_{1 \leq j \leq N}$.

Assume now that an approximate far-field pattern $F_{\ell,j}^\infty \simeq u^\infty(d_\ell, d_j)$, $\ell, j = 1, \dots, N$ is known for N incident plane wave with directions d_ℓ and measured in the same directions. The integral equation (2.8) can then be transformed at the discrete level into the following linear system of N equations in N unknowns (g_j) :

$$\sum_{j=1}^N \omega_j F_{\ell,j} g_j = e^{-ikz d_\ell \cdot z}, \quad \ell = 1, \dots, N, \quad (3.1)$$

where the weights ω_j depend on the quadrature formulae used in evaluating the integrals over the mesh triangles.

For each z , the system (3.1) is then a discretized $N \times N$ linear system of the form

$$Fg(z) = b^\infty(z) \stackrel{\text{def}}{=} \left(e^{-ikd_1 \cdot z}, \dots, e^{-ikz d_N \cdot z} \right)^T \quad (3.2)$$

where F is $N \times N$ matrix independent of z , $g(z) = (g_1(z), \dots, g_N(z))^T$ is the unknown vector whose ℓ_2 norm is expected to be large when z is outside D and finally $b^\infty(z)$ is the right-hand side constructed from the far field of the plane waves $u^{\text{inc}}(z, d_j)$ for $j = 1, \dots, N$. Unfortunately, this system is ill-posed and its numerical solution requires the introduction of a suitable regularization scheme. A good example of such a scheme is given by the Tikhonov-Morozov technique, where one computes $g_\eta(z)$ solution of

$$(F^*F + \eta(z)I)g_\eta(z) = F^*b^\infty(z). \quad (3.3)$$

The regularization parameter $\eta(z)$ depending on z is then determined using the Morozov discrepancy principle (Colton et al., 1997), that is as the root of the discrepancy function

$$\|Fg_\eta(z) - b^\infty(z)\|^2 - \delta^2 \|g_\eta(z)\|^2 \quad (3.4)$$

where δ is an estimate of the error on the matrix F . This technique has been advocated by Colton et al. (1997) and since then successfully applied by a number of authors (see Collino et al., 2003 or Fares et al., 2009, amongst others). In practice, a box \mathcal{B} of \mathbb{R}^3 containing the object (more or less at its centre) is discretized using a regular cubic mesh and we denote by \mathcal{Z} the set of all these discretized points. The system (3.2) is then solved for each $z \in \mathcal{Z}$, using the singular-value (SVD) decomposition of the matrix F given by

$$F = U\Sigma V^* \quad (3.5)$$

where U and V are unitary and Σ is real diagonal with $\Sigma_{j,j} = \sigma_j$, $1 \leq j \leq N$. This decomposition is possibly truncated to ignore all singular values and vectors of index larger than $p \leq N$. The set of solutions $\{g_\eta(z)\}_{z \in \mathcal{B}}$ is then computed from

$$g_p(z) = V_p [\Sigma_p + \eta(z)I]^{-1} U_p^* b^\infty(z),$$

(where the quantities subscripted by p only contain the information related to the p largest singular values). A graphical tool is used to plot the level surfaces of the function $\|g_\eta(z)\|^{-1}$ for some suitably chosen level value, yielding a hopefully satisfactory representation of the unknown object.

3.2 Truncated singular values, L-curve and Picard coefficients

While regularization algorithms based on the Tikhonov-Morozov method have been successfully used for solving the system (3.2), their computational cost remains somewhat high, and it was argued in Fares et al., 2009 that a substantial fraction of this cost is caused by the need to compute a regularization parameter $\eta(z)$ which depends on the chosen point $z \in \mathcal{B}$. We are therefore interested in techniques where this dependence does not arise.

A first class of methods which would satisfy this requirement is that of truncated SVD schemes where the truncation level (i.e., the number p of singular vectors and values used for representing the signal) is kept constant for all values of $z \in \mathcal{B}$ and $\eta(z) = 0$ for all z . Determining this truncation level in an efficient manner which is also robust across a variety of examples is not an easy task. Inspiration can nevertheless be obtained from the existing literature on inverse problem regularization, such as Hansen, 1997. More specifically this book advocates the L-curve technique, which may be seen as a graphical tool where one exploits a plot of the value of $\|g_p(z)\|$ versus $\|Fg_p(z) - b^\infty(z)\|$ for all values of the regularization parameter p and a given z . This curve starts (for large p) with high values of the solution norm and small residual norm, and subsequently decreases with p until the solution norm is small and the residual norm large. Very often, this curve shows a sharp L-shaped corner for some value of the parameter. The L-curve regularization technique then amounts to choosing the value of p corresponding to the corner for regularizing the problem. It is based on the idea that this corresponds to the particular level of regularization where the noise (corresponding to small singular values) start dominating the signal. This observation is embodied in the so-called *discrete Picard condition*, which assumes that, for the numerically nonzero singular values σ_ℓ (taken in decreasing order), the corresponding sequence of *Picard coefficients* given by

$$\frac{u_\ell^* b^\infty(z)}{\sigma_\ell} \quad (3.6)$$

(for a particular value of $z \in \mathcal{Z}$) also decreases to zero, where the u_ℓ are the left singular vectors of the matrix F . This popular technique has been used in a number of contexts.

Unfortunately, our experience¹ with this regularization in inverse scattering problems has been mixed, and we observed in practice that the corner of the L-curve for our problem is commonly poorly determined, that it may significantly depend on z and that, even after hand-picking the most 'corner-like' point on the curve for selecting a regularization parameter, the quality of the representation obtained for the unknown object leaves much to be desired: the selection of a suitable p_\diamond is therefore typically far from obvious. As an alternative, we attempted to use the general cross-validation (GCV) technique², which is also related to the discrete Picard condition, without much more success.

Since neither of these methods seems to perform well, it is natural to consider whether the discrete Picard condition holds in our context. We therefore plotted the sequence of Picard coefficients for our application and obtained the topmost curves shown in Figure 1 for an exterior and an interior point: instead of decreasing, the sequence of Picard coefficients starts by increasing before it approximately stabilizes at a significantly nonzero level. Thus the discrete Picard condition cannot be assumed, which may explain at least partly the poor behaviour of the L-curve and GCV methods.




Figure 1: The values of $u_\ell^* b^\infty(z)/\sigma_\ell$ (top curve), $u_\ell^* b^\infty(z)$ (median curve) and σ_ℓ (bottom curve), as a function of ℓ for a point internal (left) or external (right) with respect to the considered object.

3.3 Operator perturbation and clustered eigenspace recovery

It is however the case that the information provided by the Picard coefficients can be exploited, although in a non-standard manner. A detailed analysis of the numerator in (3.6),

$$\theta_\ell(z) \stackrel{\text{def}}{=} u_\ell^* b^\infty(z),$$

indeed shows an interesting behaviour. While this sequence also stabilizes for large ℓ (i.e. for small singular values in the tail of their distribution), the level at which this stabilization occurs differs significantly depending on the position of the considered point z with respect to the object: this

¹Using Hansen's toolbox of regularization tools (Hansen, 1994).

²Also from Hansen's toolbox.

level is very close to zero for interior points but is significantly larger as soon as z is on the outside of the object. This behaviour is very noticeable in Figure 1. This phenomenon can be explained as follows.

Consider a general Hermitian matrix A and its eigenvalue decomposition $A = Q^* \Lambda Q$, where Q is the unitary matrix whose columns are the eigenvectors of A and where

$$\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n).$$

If we now define $\tilde{A} \stackrel{\text{def}}{=} A + tE$ for some bounded perturbation matrix E and some (small) real parameter t . From classical perturbation analysis for eigenvectors, we know that, for small enough t , the eigenvectors and eigenvalues of \tilde{A} are given by the matrices \tilde{Q} and $\tilde{\Lambda}$ defined by $\tilde{Q} = [q_1(t), \dots, q_n(t)]$ and $\tilde{\Lambda} = \text{diag}(\lambda_j(t))$ with $\tilde{A}\tilde{Q} = \tilde{Q}\tilde{\Lambda}$. Using the Landau little “ o ” notation, the first-order expansion of the eigenvalues and eigenvectors of \tilde{A} in the direction E are given (see Stewart, 2001, page 45) by

$$\lambda_\ell(t) = \lambda_\ell + t(q_\ell^* E q_\ell) + o(t) \quad \text{and} \quad \tilde{Q} = Q + \Delta Q(t) + o(t),$$

where the ℓ -th column of $\Delta Q(t)$ is defined by

$$\delta q_\ell(t) = t \underline{Q}_\ell (\lambda_\ell I_{n-1} - \underline{\Lambda}_\ell)^{-1} \underline{Q}_\ell^* E q_\ell, \quad (3.7)$$

with I_{n-1} being the $(n-1)$ -dimensional identity matrix, \underline{Q}_ℓ being the $n \times (n-1)$ matrix whose columns are those of Q , column ℓ excluded³, and $\underline{\Lambda}_\ell$ being the $(n-1) \times (n-1)$ diagonal matrix $\text{diag}(\lambda_j)_{j=1, \dots, n; j \neq \ell}$. When λ_ℓ belongs to a cluster of eigenvalues of A , the diagonal matrix $(\lambda_\ell I_{n-1} - \underline{\Lambda}_\ell)^{-1}$ has large diagonal terms corresponding to the positions of the clustered eigenvalues. In other words, (3.7) shows that, in the basis defined by the columns of Q , the components of an eigenvector of \tilde{A} along the eigenvectors associated with the cluster will dominate all other components. If we now consider an arbitrary combination of the eigenvectors of A , say $b = \sum_{j=1}^n \alpha_j q_j$, we then obtain that, for small t ,

$$\begin{aligned} \tilde{q}_\ell^* b &\approx \sum_{j=1}^n \alpha_j \left[q_\ell + t \underline{Q}_\ell (\lambda_\ell I_{n-1} - \underline{\Lambda}_\ell)^{-1} \underline{Q}_\ell^* E q_\ell \right]^* q_j \\ &= \alpha_\ell + t \sum_{j \neq \ell} \alpha_j q_\ell^* E^* \underline{Q}_\ell (\lambda_\ell I_{n-1} - \underline{\Lambda}_\ell)^{-1} e_{j,\ell}, \end{aligned} \quad (3.8)$$

where, using the unitary nature of Q , $e_{j\ell} = \underline{Q}_\ell^* q_j$ is a column of I_{n-1} . Thus $|\tilde{q}_\ell^* b|$ is expected to be largest whenever t is small, ℓ is the index of an eigenvalue in the cluster and b has a significant component along the eigenvectors of A corresponding to the cluster. As a result, we may thus use relation (3.8) to *recover an unknown eigenspace of A corresponding to an eigenvalue cluster by only considering inner products with eigenvectors of the perturbed \tilde{A}* . Remarkably, both the presence of the perturbation E and the clustering of the eigenvalues are crucial for this recovery to be possible.

We may now apply this technique to the far-field equations by identifying perturbed and unperturbed operators. We consider the unperturbed operator to be the matrix $F^\infty (F^\infty)^*$ where F^∞ is the *exact discretized* far-field operator. This operator is itself an approximation of the infinite-dimensional operator \mathcal{F} , and admits a cluster of near-zero eigenvalues associated with the discretized nullspace. However, F^∞ is unavailable in practice, and the computed F is typically contaminated by errors, because it results from physical measurements, themselves inaccurate, or is computed by a process involving approximate quadratures (see Section 3.1). The theory developed above therefore applies for the choices

$$\tilde{A} \leftarrow FF^* \quad \text{and} \quad tE \leftarrow FF^* - F^\infty (F^\infty)^*,$$

³That is $\underline{Q}_\ell = [Q(:, j)]_{j=1, \dots, n; j \neq \ell}$.

which implies that

$$\tilde{Q} \leftarrow U, \quad \text{and} \quad \tilde{q}_\ell^* b \leftarrow \theta_\ell(z).$$

Our reasoning then indicates that we may expect $|\theta_\ell(z)|$ to be significantly different from zero for large j provided that the error $\|FF^* - F^\infty(F^\infty)^*\|$ remains small and that b^∞ has a significant component along the (unknown) approximate nullspace of the far-field operator F^∞ . This is exactly what we observe in Figure 1, where the values of θ_ℓ (for large ℓ) are negligible for $b^\infty(z)$ corresponding to an internal z and significantly nonzero for $b^\infty(z)$ corresponding to an external z .

3.4 The SVD-tail method

In view of these observations, we propose to use the level at which the $\theta_\ell(z)$ stabilize for large ℓ as an indicator of the fact that z belongs to the unknown object. More specifically, our new technique, baptized SVD-tail, depends on a parameter d and consists of the following steps:

Algorithm 3.1: The SVD-tail method

Step 1: Select the dimension d of a subspace \mathcal{T} spanned (possibly approximately) by the d left singular vectors corresponding to the smallest d singular values of F .

Step 2: Compute $\{w_\ell\}_{\ell=1}^d$ a basis of \mathcal{T} , as well as

$$\vartheta_\ell(z) = w_\ell^* b^\infty(z) \quad \text{for} \quad \ell = 1, \dots, d. \quad (3.9)$$

Step 3: Define

$$\psi_d(z) \stackrel{\text{def}}{=} \|(\vartheta_1(z), \dots, \vartheta_d(z))^T\|^{-1}. \quad (3.10)$$

This procedure provides the function $\psi_d(z)$, whose values are expected to be small for z outside the object and significantly positive for z in its interior. It therefore plays a role entirely similar to that of $\|g_\eta(z)\|^{-1}$ in the more standard LSM approaches. In particular, a value ψ_* can be chosen such that the level curve

$$\mathcal{L} \stackrel{\text{def}}{=} \{z \in \mathcal{Z} \mid \psi_d(z) = \psi_*\} \quad (3.11)$$

is a suitable visual representation of the unknown object.

Observe that this technique makes use of the user-chosen dimension d and of the knowledge of a basis for \mathcal{T} . If the singular-value decomposition of F is known, then the choice $w_\ell = u_{N-\ell+1}$ is obviously adequate. Other techniques are however possible which may provide the necessary information at the fraction of the cost of the singular-value decomposition (as we discuss in Section 5). Observe also that, because of (3.9)-(3.10), the cost of computing $\psi_d(z)$ for a single z varies essentially linearly with d . Small values of d are therefore particularly advantageous when $\psi_d(z)$ must be computed for a large number of vectors z , i.e. when the discretization of the domain \mathcal{B} is fine.

We conclude this section by noting again that it is very remarkable that the actual error introduced in the discretized the far-field operator actually allows us to apply this inversion technique.

3.5 A simple isovalue heuristic

For the SVD-tail approach to be practical, we also need to select a suitable value ψ_* for the level curve (3.11) to be a good representation of the object. While fine hand-tuning is very often best for finalizing this representation, we tried, in line with our developments in Fares et al. (2009), to provide heuristics that provide a good starting point for this tuning.

We have first considered the technique developed for the adaptive Morozov regularization in Fares et al. (2009). This technique is based on the detection of large values of the gradient of $\psi_d(z)$, which are expected on the boundary of the unknown object. However, a simpler scheme turned out to be experimentally very efficient and is based on elementary statistics for the values of $\psi_d(z)$. Our choice of the isovalue is simply given by

$$\psi_* = \text{mean}_{z \in \mathcal{Z}}[\psi_d(z)] + 2 \text{std}_{z \in \mathcal{Z}}[\psi_d(z)] \quad (3.12)$$

where 'mean' and 'std' stand for the average and standard deviation, respectively. We refer to this technique as the *Global Mean and Standard Deviation* (GMSD) heuristic.

4 Preliminary numerical experience with SVD-tail

We now consider the performance of the SVD-tail method on a set of three-dimensional examples, and compare it with that of the continuous Morozov regularisation (see Colton et al., 1997, and Collino et al., 2003), its discrete variant (see Fares et al., 2009). In the notation of Section 3.1, the first of these techniques uses $p = N$ and varies $\eta(z)$ according to the Morozov discrepancy rule, while the second sets $\eta(z) = 0$ for all z but varies p with z according to a similar rule. A truncated SVD technique is also included in the comparison, for which the truncation parameter p is independent of z and is known *a priori*, thereby short-circuiting the difficulty to estimate it. Our objective is not only to compare the quality of the obtained images, but also to measure the sensitivity of the new method with respect to its parameter: the dimension d of the subspace \mathcal{T} . The adequacy of the isovalue heuristics is also of immediate interest.

The five objects used in our comparison are identical to those used in Fares et al. (2009): a cross, a teapot, a rocket, a plane and a Σ -shaped object. For each of these examples, we discretized the sphere in $N = 2252$ directions and discretized \mathcal{B} using $50 \times 50 \times 50$ points ($|\mathcal{Z}| = 125000$). Other parameters for these examples are given in Table 1, in which k is the wave number of the incident wave, $[\underline{x}, \bar{x}]$, $[\underline{y}, \bar{y}]$, and $[\underline{z}, \bar{z}]$ are the intervals (in x , y and z) defining the smallest box containing the object, and $[\underline{X}, \bar{X}]$, $[\underline{Y}, \bar{Y}]$ and $[\underline{Z}, \bar{Z}]$ are the intervals defining the scanned domain \mathcal{B} . The objects themselves are pictured in Figure 2 on page 14.

example	k	$[\underline{x}, \bar{x}]$	$[\underline{y}, \bar{y}]$	$[\underline{z}, \bar{z}]$	$[\underline{X}, \bar{X}]$	$[\underline{Y}, \bar{Y}]$	$[\underline{Z}, \bar{Z}]$
cross	10	$[-0.625, 0.625]$	$[-0.625, 0.625]$	$[-0.125, 0.125]$	$[-1, 1]$	$[-1, 1]$	$[-0.5, 0.5]$
teapot	31	$[-0.3, 0.34]$	$[-0.2, 0.2]$	$[0, 0.315]$	$[-0.5, 0.5]$	$[-0.5, 0.5]$	$[-0.1, 0.4]$
rocket	20	$[-0.165, 0]$	$[-0.241, 0.241]$	$[-0.241, 0.241]$	$[-2, 0.5]$	$[-0.5, 0.5]$	$[-0.5, 0.5]$
plane	12	$[-0.101, 0.101]$	$[-0.129, 0.75]$	$[-0.188, 0.416]$	$[-1.3, 1.3]$	$[-1.6, 1]$	$[-0.5, 0.5]$
Σ	30	$[-1, 0.8]$	$[-0.9, 0.9]$	$[-0.5, 0]$	$[-1.3, 1.1]$	$[-1.2, 1; 2]$	$[-0.8, 0.3]$

Table 1: Geometric parameters for the tested examples

In all our experiments, the measurements were generated using CASC, a solver for acoustic scattering problems developed at CERFACS. In this package, the sound-soft case is treated by

solving the integral equation,

$$S\varphi(x) = \int_{\partial D} \Phi(x, y)\varphi(y)ds(y) \quad x \in \partial D$$

whose unknown is $\varphi := [\partial_n \varphi] := \partial_n \varphi^+ - \partial_n \varphi^-$ and where $\Phi(x, y)$ is the fundamental solution to the Helmholtz equation. The numerical procedure is based upon a triangular meshing of the surface and uses finite elements of lowest degree. It leads to solving for φ_h such that,

$$\begin{cases} \int_{\partial D_h} \int_{\partial D_h} \Phi(x, y) (\varphi_h(y) \cdot \varphi_h^{test}(x)) ds(y)ds(x) = - \int_{\partial D_h} u^{inc}(x)\varphi_h^{test}(x)ds(x), \\ \text{for all } \varphi_h^{test}(x) \end{cases}$$

The numerical computation amounts to solving a set of linear systems with a dense symmetric non-hermitian matrix whose size is the number of nodes of the mesh. There are as many right-hand sides as the number of incident directions. Special attention has been paid to properly taking into account the singularity of the Green kernel during the assembly process. The LU decomposition of the matrix is then performed by means of a set of ScaLAPACK parallel routines. Once φ_h has been obtained, the associated far fields are easily deduced by applying the integral representation formulae

$$u(x) = u^{inc}(x) + S\varphi(x) \quad x \notin \partial D.$$

For improved accuracy, care is taken to ensure that the length of the longest edge in the discretization does not exceed a tenth of the wavelength. Once the matrix F was generated, it was then perturbed by random noise of relative magnitude 0.01.

We now consider applying the SVD-tail method in its most immediate form (that is when the SVD-decomposition (3.5) is computed and one chooses $w_q = u_{N-q+1}$ for $q = 1, \dots, d$), and investigate the sensitivity of the results to the choice of d , the dimension of the subspace \mathcal{T} corresponding to small singular values. We illustrate this sensitivity by considering the plane object and applying SVD-tail (with the GMSD heuristic) for values of $d = 5, 20, 35, 50, 65$ and 80 . The results are shown in Figure 3 on page 15. Examining this figure (and similar ones obtained for the other objects, not shown here), we see that increasing the dimension d is globally beneficial for the quality of the reconstructed image, but also that very low values of d can give very reasonable representations. In particular, choosing $d = 5$ does not give the best image, but, for a computing cost⁴ approximately 10 times smaller than for $d = 50$, the corresponding image remains broadly comparable with this latter case. The value $d = 50$ appears to be an adequate default value, and will be used in the rest of our study. The very limited sensitivity of the method on its parameter is clearly an advantage compared to other techniques.

We conclude this section with a comparison between SVD-tail(50), the adaptive the discrete and continuous Tikhonov-Morozov regularizations (DTM and CTM, respectively), and the truncated SVD method (TSVD) with *a priori* known truncation level $p = 150$. The level noise was set to 0.02 for the two Tikhonov-Morozov techniques, in accordance with the level noise added on our synthetic measurements. Figures 4 to 6 show that SVD-tail provides excellent images and compares well with TSVD⁵, DTM and even with CTM, albeit yielding slightly less smooth images.

We now turn to the computational cost of the methods discussed in this paragraph, which we split into two successive tasks:

⁴For computing $\{\psi_d(z)\}_{z \in \mathcal{Z}}$.

⁵In Figure 6, the isovalue for plotting the result of the TSVD method had to be tuned manually, the GMSD heuristic producing very poor results.

1. the time required for computing the necessary parts of the singular value decomposition,
2. the time required for evaluating $\psi(z)$ for all $z \in \mathcal{Z}$,

We report these times in Tables 2 to 4 for a Fortran 90 implementation of Algorithms CTM, DTM and TSVD(150) (see Fares et al., 2009, for details) and SVD-tail(50). They were obtained on a single 1.6 GHz “Itanium 2” processor of a Bull Novascale 3045 machine, using the MKL 64 library for linear algebra kernels. The time necessary for the operations other than described above, including computing the isovalue heuristic, was negligible (below 0.005 sec).

	cross				teapot			
	CTM	DTM	TSVD	SVD-tail(50)	CTM	DTM	TSVD	SVD-tail(50)
SVD	1138.61	43.99	43.99	1138.61	1094.29	47.30	47.30	1094.29
eval. ψ	2084.44	100.66	130.92	59.60	1986.09	92.60	128.43	56.86
total	3223.05	144.65	174.91	1198.21	3080.38	139.90	175.73	1151.15
percent.	100.00%	4.5%	5.4%	37.2%	100.00%	4.5%	5.7%	37.4%

Table 2: CPU times (in secs.) for the different computational tasks and the four algorithms, applied on the cross and the teapot

	rocket				plane			
	CTM	DTM	TSVD	SVD-tail(50)	CTM	DTM	TSVD	SVD-tail(50)
SVD	1092.77	48.39	43.39	1092.77	1147.75	44.30	44.30	1147.75
eval. ψ	1970.75	80.07	128.06	58.30	1950.52	95.45	127.36	56.15
total	3063.52	128.46	171.45	1151.07	3098.27	139.75	171.66	1203.90
percent.	100.00%	4.2%	5.6%	37.6%	100.00%	4.5%	5.5%	38.9%

Table 3: CPU times (in secs.) for the different computational tasks and the four algorithms, applied on the rocket and the plane

	Σ			
	CTM	DTM	TSVD	SVD-tail(50)
SVD	1033.19	241.95	241.95	1033.19
eval. ψ	1944.66	92.00	127.54	57.61
total	2977.85	333.95	369.49	1090.80
percent.	100.00%	11.2%	12.4%	36.6%

Table 4: CPU times (in secs.) for the different computational tasks and the four algorithms, applied on the Σ

Because SVD-tail (in this form) still relies on the SVD decomposition, we should not expect substantial reduction in CPU-time with respect to methods that do not use the full SVD, and this is apparent in the table. On the other hand, the relatively low dimension of the subspace \mathcal{T} makes the computation of the values of ψ substantially faster than for other methods who use more singular vectors, because the computational effort needed in SVD-tail to obtain $\psi_d(z)$ of each $z \in \mathcal{Z}$ is linear in $d = \dim(\mathcal{T})$. Moreover, the gains are expected to be even more substantial if larger problems are considered and finer space discretizations can thus be accommodated without excessive cost.

The next section presents a variant of SVD-tail which does not use the SVD decomposition, and is therefore globally more efficient.

5 ISVD-tail: a SVD-tail variant without SVD

As we have indicated in Section 3.4, any basis of \mathcal{T} is suitable, and there is no need to choose $w_q = u_{N-q+1}$ as we have done so far. We thus propose to use subspace iteration (see Chapter 6 in Stewart, 1998) to compute a basis of \mathcal{T} . The method now depends on two parameters (d and $n_{\mathcal{T}}$) and is stated as Algorithm 5.1 below.

Algorithm 5.1: Iterative SVD-tail algorithm (ISVD-tail)

Step 1: The dimension d of \mathcal{T} and a number of iterations $n_{\mathcal{T}}$ are given.

Choose a random $N \times d$ matrix W .

Step 2: Compute a LU factorization $F = LU$.

Step 3: Perform $n_{\mathcal{T}}$ iterations of the following process:

Step 3.1: Solve the lower triangular system $LY_1 = W$ for Y_1 .

Step 3.2: Solve the upper triangular system $UY_2 = Y_1$ for Y_2 .

Step 3.3: Solve the lower triangular system $U^*Y_3 = Y_2$ for Y_3 .

Step 3.4: Solve the upper triangular system $L^*Y_4 = Y_3$ for Y_4 .

Step 3.5: Overwrite W using the orthogonal factorization $Y_4 = WR$.

The columns of the final matrix W provide the desired approximate basis of \mathcal{T} . Note that Steps 3.1 to 3.4 of Algorithm 5.1 amount to solving the system

$$FF^*Y_4 = W. \quad (5.1)$$

Also observe that the cost of one execution of Step 3 remains modest, given that W and the Y_q ($q = 1, \dots, 4$) are $N \times d$ matrices, where $d \ll N$. If $n_{\mathcal{T}}$ remains small, the overall cost is then much smaller than that of computing the SVD decomposition (3.5).

In our experience, a value of $n_{\mathcal{T}}$ as small as 10 is very often satisfactory, the obtained images being visually undistinguishable from those obtained with SVD-tail which uses the more expensive choice $w_q = u_q$ ($q = 1, \dots, d$). As a consequence, all comments made above on the relative visual performance of SVD-tail apply to ISVD-tail without modification. The cost of executing ISVD-tail is however again considerably smaller than that of the previously described methods. This is illustrated in Table 5, which should be compared to Tables 2-4 above.

	cross	teapot	rocket	plane	Σ
\mathcal{T} basis	12.92	15.41	16.62	16.85	15.41
eval. ψ	57.79	54.28	57.60	58.44	54.28
total	70.71	69.69	74.22	75.29	69.69
percent.	2.2%	2.3%	2.4%	2.4%	2.3%

Table 5: CPU times (in secs.) for ISVD-tail(50,10) (percentage computed w.r.t. CTM)

This table shows how fast ISVD-tail can be compared to the other methods: a speed-up factor of 50 with respect to CTM, of 2.5 with respect to DTM, of 2.8 with respect to TSVD(150) and of 16 with respect to (the crude) SVD-tail are typical, without altering the quality of the visual representation significantly.

6 Conclusion and perspectives

We have introduced a new technique for reconstructing images using the Linear Sampling Method, whose computing cost is extremely modest and which produces visually satisfactory visual representations. This technique is based on identifying the value of the signal in a low-dimensional subspace associated with small singular values of the measurement matrix. This identification is efficiently conducted by exploiting a new eigenspace recovery technique, itself capitalizing on the combined presence of error in the data and clustered eigenvalues. Two variants of the new approach have been detailed, with varying computational requirements. We have also described a new heuristic for automatically determining an isovalue for the visual representation from which finer hand-tuning can be initiated if necessary. The combination of these new proposals leads to a cheap and automatic image reconstruction algorithm, hopefully paving the way for real-time image reconstruction.

The authors are well aware that several questions remain open at this point. In particular, further applications of the new eigenspace recovery principle in other areas seems of interest. Continued experimentation with the SVD-tail methods on a wider battery of meaningful examples is also necessary to assess their true potential. Further work on the automatic determination of the isovalue is finally desirable and can be inspired by existing ideas in the field of computer graphics.

Acknowledgments

The authors are grateful to F. Collino for his interest in this work and for his questioning, which ultimately lead to the ideas discussed in Section 3.3. They would like to thank the C.I.C.T. (Centre Interuniversitaire de Calcul de Toulouse) for providing access to the Grid'MIP machine for producing the synthetic data.

References

- T. Arens. Why linear the sampling method works. *Inverse Problems*, **20**, 163–173, 2004.
- F. Cakoni and D. Colton. A uniqueness theorem for an inverse electromagnetic scattering problem in inhomogeneous anisotropic media. *Proc. Edinburgh Math. Soc.*, **46**, 293–314, 2003.
- F. Cakoni, D. Colton, and H. Haddar. The linear sampling method for anisotropic media. *J. Comput. Appl. Math.*, **146**, 285–299, 2002.
- F. Collino, M. Fares, and H. Haddar. Numerical and analytical studies of the linear sampling method in electromagnetic scattering problems. *Inverse Problems*, **19**, 1279–1298, 2003.
- D. Colton. Inverse acoustic and electromagnetic scattering theory. *Inverse Problems*, **47**, 67–109, 2003.
- D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer Verlag, 2nd edn, 1998.
- D. Colton, K. Giebermann, and P. Monk. A regularized sampling method for solving three-dimensional inverse scattering problems. *Sci. Comput.*, **21**(6), 2316–2330, 2000.

- D. Colton, H. Haddar, and P. Monk. The linear sampling method for solving the electromagnetic inverse scattering problem. *SIAM Journal on Scientific Computing*, **24**, 719–731, 2002.
- D. Colton, H. Haddar, and M. Piana. The linear sampling method in inverse electromagnetic scattering theory. *Inverse problems*, **19**, S105–S137, 2003.
- D. Colton, M. Piana, and R. Potthast. A simple method using Morozov's discrepancy principle for solving inverse scattering problems. *Inverse Problems*, **13**, 1477–1493, 1997.
- M. Fares, S. Gratton, and Ph. L. Toint. Fast regularized linear sampling for inverse scattering problems. Technical Report TR09/08, Department of Mathematics, FUNDP - University of Namur, Namur, Belgium, 2009.
- H. Haddar and P. Monk. The linear sampling method for solving the electromagnetic inverse medium problem. *Inverse Problems*, **18**, 891–906, 2002.
- P. C. Hansen. Regularization tools: A Matlab package for analysis and solution of discrete ill-posed problems. *Numerical Algorithms*, **6**, 1–35, 1994.
- P. C. Hansen. *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. SIAM, Philadelphia, USA, 1997.
- G. W. Stewart. *Matrix Computations: Basic Decompositions*. SIAM, Philadelphia, USA, 1998.
- G. W. Stewart. *Matrix Algorithms. Volume II: Eigensystems*. SIAM, Philadelphia, USA, 2001.

Appendix : the figures




Figure 2: The five objects of the study




Figure 3: The plane for $d = 5, 20, 35, 50, 65, 80$ (from top to bottom and left to right)

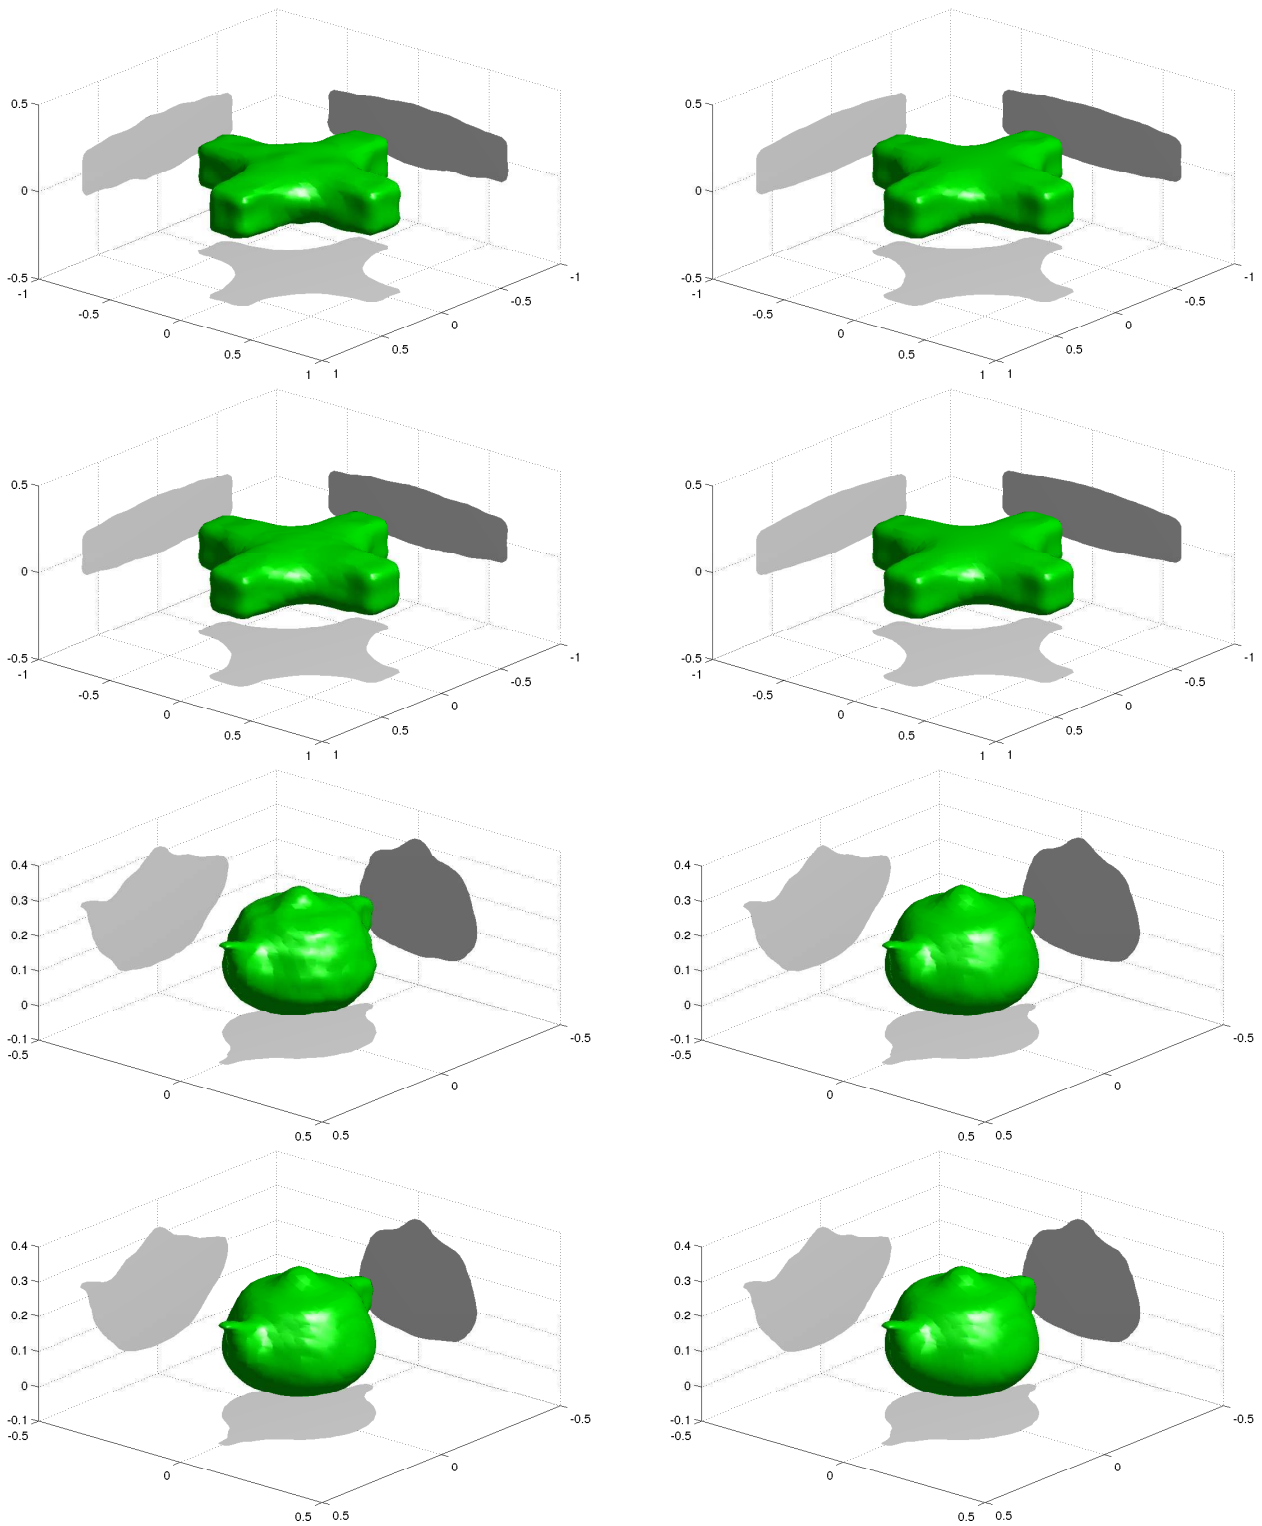


Figure 4: The cross and the teapot reconstructed by SVD-tail(50) (top left), TSVD (top right), DTM (bottom left) and CTM (bottom right)




Figure 5: The rocket and the plane reconstructed by SVD-tail(50) (top left), TSVD (top right), DTM (bottom left) and CTM (bottom right)




Figure 6: The Σ reconstructed by SVD-tail(50) (top left), TSVD (top right, manually tuned), DTM (bottom left) and CTM (bottom right)